PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Hyun		Lee	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Ladera Ranch	State/Province	CA	Country of Residence i	US

Mailing Address of Inventor:

Address 1	21 Thalia		
Address 2			
City	Ladera Ranch	State/Province	CA
Postal Code	92694	Country i	US

Inventor 2 Remove

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Jayesh	R.	Bhakta	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Cerritos	State/Province	CA	Country of Residence i	US

Mailing Address of Inventor:

Address 1	21 Thalia		
Address 2			
City	Cerritos	State/Province	CA
Postal Code	92694	Country i	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. Add

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

☐ An Address is being provided for the correspondence Information of this application.			
Customer Number	79141		
Email Address	jz@jzpatent.com	Add Email	Remove Email

Application Information:

Title of the Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION		
Attorney Docket Number	NT003US	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Suggested Class (if any)		**Sub Class (if any)**	
Suggested Technology Center (if any)			
Total Number of Drawing Sheets (if any)	26	**Suggested Figure for Publication (if any)**	2A

Publication Information:

☐	Request Early Publication (Fee required at time of Request 37 CFR 1.219)
☒	**Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	79141		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending		Remove
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
	non provisional of	61676883	2012-07-27

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button. Add

Foreign Priority Information:

This section allows for the applicant to claim benefit of foreign priority and to identify any prior foreign application for which priority is not claimed. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(a).

Remove

Application Number	Country i	Filing Date (YYYY-MM-DD)	Priority Claimed
			(•) Yes ( ) No

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button. Add

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date of filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Applicant 1 Remove

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

Clear

(•) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. [X]

Organization Name	Netlist, Inc.		
Mailing Address Information:			
Address 1	51 Discovery, #150		
Address 2			
City	Irvine	**State/Province**	CA
Country i	US	Postal Code	92618
Phone Number	949-435-0025	Fax Number	
Email Address	gsasaki@netlist.com		

Additional Applicant Data may be generated within this form by selecting the Add button. Add

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

Remove

If the Assignee is an Organization check here. [ ]

PTO/AIA/14 (08-12)
Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Prefix	**Given Name**	Middle Name	**Family Name**	Suffix

Mailing Address Information:

Address 1			
Address 2			
City		**State/Province**	
Country i		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button. Add

Signature:

Remove

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications

Signature	/Jamie J. Zheng/			Date (YYYY-MM-DD)	2013-07-27
First Name	Jamie	Last Name	Zheng	Registration Number	51167

Additional Signature may be generated within this form by selecting the Add button. Add

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether the Freedom of Information Act requires disclosure of these records.

2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.

3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.

4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).

5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.

6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).

7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.

8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspections or an issued patent.

9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

CLAIM OF PRIORITY

[0001] The present application claims priority to U.S. Provisional App. No. 61/676,883, filed on July 27, 2012, which is incorporated by reference herein in its entirety.

CROSS REFERENCE TO RELATED APPLICATIONS

[0002] The present application is related to commonly-owned U.S. Pat. App. No. 12/504,131, filed on July 16, 2009, now U.S. Pat. No. 8,417,870, U.S. Pat. App. No. 12/761,179, filed on April 15, 2010, U.S. Pat. App. No. 13/287,042, filed on November 1, 2011, and U.S. Pat. App. No. 13/287,081, filed on November 1, 2011, each of which is incorporated herein by reference in its entirety.

FIELD

[0003] The disclosure herein is related generally to memory modules, and more particularly to multi-rank memory modules and methods of operation.

BACKGROUND

[0004] With recent advancement of information technology and widespread use of the Internet to store and process information, more and more demands are placed on the acquisition, processing, storage and dissemination of vocal, pictorial, textual and numerical information by microelectronics-based combination of computing and communication means. In a typical computer or server system, memory modules are used to store data or information. A memory module usually includes multiple memory devices, such as dynamic random access memory devices (DRAM) or synchronous dynamic random access

Micron Technology Inc. et al.
Ex. 1010, p. 7

memory devices (SDRAM), packaged individually or in groups, and/or mounted on a printed circuit board (PCB). A processor or a memory controller accesses the memory module via a memory bus, which, for a single-in-line memory module (SIMM), can have a 32-bit wide data path, or for a dual-in-line memory module (DIMM), can have a 64-bit wide data path.

[0005] The memory devices of a memory module are generally organized in ranks, with each rank of memory devices 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 "x64" or "by 64" organization. Similarly, a memory module having 72-bit-wide ranks is described as having an "x72" or "by 72" organization.

[0006] The memory capacity or memory density of a memory module increases with the number of memory devices on the memory module. 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.

[0007] In certain conventional memory modules, the ranks are selected or activated by control signals from a processor or memory controller during operation. 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 a limited number of ranks per memory module, which limits the memory density of the memory modules that can be used in these computer and server systems.

[0008] For memory devices in such as a memory module to be properly accessed, distribution of control signals and a control clock signal in the memory module is subject to strict constraints. In some conventional memory modules, control wires are routed so there is an equal length to each memory component, in order to eliminate variation of the timing of the control signals and the control clock signal between different memory devices in the memory modules. The balancing of the length of the wires to each memory devices compromises system performance, limits the number of memory devices, and complicates

Micron Technology Inc. et al.
Ex. 1010, p. 8

their connections.

[0009] In some conventional memory systems, the memory controllers include leveling mechanisms for write and/or read operations to compensate for unbalanced wire lengths and memory device loading on the memory module. As memory operating speed and memory density continue to increase, however, such leveling mechanisms are also insufficient to insure proper timing of the control and/or data signals received and/or transmitted by the memory modules.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] FIG. 1 is a diagram illustrating a memory system including at least one memory module according to one embodiment.

[0011] FIGS. 2A-2D are each a diagrams illustrating interactions among components in a a memory module according to certain embodiments.

[0012] FIG. 3 is a diagram illustrating one of a plurality of data buffers in a memory module according to one embodiment.

[0013] FIGS. 4A-4B are each a diagram illustrating data and data strobe signal lines coupled to memory devices in a memory module according to certain embodiments.

[0014] FIGS. 5A-5B are diagrams illustrating different numbers of memory devices that can be coupled to each data buffer in a memory module according to certain embodiments.

[0015] FIG. 6 is a diagram illustrating a control circuit in a data buffer according to certain embodiments.

[0016] FIG. 7 is a diagram illustrating control signals from a module control device to a plurality of data buffers in a memory module according to certain embodiments.

[0017] FIG. 8 is a timing diagram illustrating alignment of module control signals with respect to module clock signals.

Micron Technology Inc. et al.
Ex. 1010, p. 9

[0018] FIG. 9 is a diagram illustrating a metastability detection circuit and signal adjustment circuit in a data buffer according to certain embodiments.

[0019] FIGS. 10A-10C are diagrams illustrating a metastability detection circuit according to certain embodiments.

[0020] FIGS. 10D is a diagram illustrating a signal adjustment circuit according to certain embodiments.

[0021] FIGS. 11A-11B are diagrams illustrating a metastability detection circuit and signal adjustment circuit, respectively, according to certain embodiments.

[0022] FIGS. 12A-12B are a timing diagrams illustrating a write operation and a read operation, respectively, performed by a memory module according to one embodiment.

[0023] FIG. 13 is a diagram illustrating a delay control circuit in a data buffer according to certain embodiments.

[0024] FIG. 14 is a diagram illustrating a DQ or DQS routing circuit in a data buffer according to an embodiment.

[0025] FIG. 15 a diagram illustrating a DQS routing circuit having a delay circuit in a data buffer according to an embodiment.

[0026] FIG. 16 a diagram illustrating a DQ routing circuit having a delay circuit in a data buffer according to an embodiment.

[0027] FIG. 17 is a diagram illustrating a delay circuit in a DQ or DQS routing circuit according to an embodiment.

[0028] FIG. 18 is a flowchart illustrating a method for data edge alignment according to embodiments.

[0029] FIG. 19 is a diagram illustrating a control circuit in a data buffer according to certain embodiments.

Micron Technology Inc. et al.
Ex. 1010, p. 10

DESCRIPTION OF EMBODIMENTS

[0030] A memory module according to one embodiment includes memory devices organized in groups, a module control device, and data buffers (DB). The data buffers are sometimes referred to herein as buffer circuits, isolation devices (I.D.) or load reduction devices. The memory module is operable to perform memory operations in response to memory commands (e.g., read, write, refresh, precharge, etc.), each of which is represented by a set of control/address (C/A) signals transmitted by the memory controller to the memory module. The C/A signals may include, for example, a row address strobe signal (/RAS), a column address strobe signal (/CAS), a write enable signal (/WE), an output enable signal (/OE), one or more chip select signals, row/column address signals, and bank address signals. The memory controller may also transmit a system clock signal to the memory module. In one embodiment, the C/A signals and the system clock signal are received by the module control device, which generates a set of module command signals and a set of module control signals in response to each memory command from the memory controller. The module command signals are transmitted by the module control device to the memory devices via module C/A signal lines, and the module control signals (referred sometimes herein as module control signals) are transmitted by the module control device to the buffer circuits via module control signal lines.

[0031] The buffer circuits are associated with respective groups of memory devices and are distributed across the memory module at positions corresponding to the respective groups of memory devices. Thus, during certain high speed operations, each module control signal may arrive at different buffer circuits at different points of time across more than one clock cycle of the system clock. Also, each buffer circuit associated with a respective group of memory devices is in the data paths between the respective group of memory devices and the memory controller. Thus, the memory controller does not have direct control of the memory devices. In one embodiment, each group of memory devices include at least two subgroups, each subgroup including at least one memory device. Each buffer circuit is configured to select a subgroup in the respective group of memory devices to communicate data with the memory controller in response to the module control signals.

Micron Technology Inc. et al.
Ex. 1010, p. 11

Thus, the memory module can have more ranks of memory devices than what is supported by the memory controller.

[0032] In one embodiment, each buffer circuit includes metastability detection circuits to detect metastability condition in the module control signals and signal adjustment circuits to adjust the module control signals and/or a module clock signal to mitigate any metastability condition in the module control signals.

[0033] Further, in one embodiment, each buffer circuit includes signal alignment circuits that determine, during a write operation, a time interval between a time when one or more module control signals are received from the module control circuit and a time when a strobe or data signal is received from the memory controller. This time interval is used during a subsequent read operation to time transmission of read data to the memory controller, such that the read data arrives at the memory controller within a time limit in accordance with a read latency parameter associated with the memory system.

[0034] FIG. 1 shows a system 100 including a memory controller (MCH) 101 and one or more memory modules 110 coupled to the MCH by a memory bus 105, according to one embodiment. As shown, the memory bus includes C/A signal lines 120 and groups of system data/strobe signal lines 130. Also as shown, each memory module 110 has a plurality of memory devices 112 organized in a plurality of ranks 114. Each memory module 110 further includes a module control circuit (module controller or module control device) 116 coupled to the MCH 101 via the C/A signal lines 120, and a plurality of buffer circuits or isolation devices 118 coupled to the MCH 101 via respective groups of system data/strobe signal lines 130. In one embodiment, the memory devices 112, the module control circuit 116 and the isolation devices 118 can be mounted on a same side or different sides of a printed circuit board (module board) 119.

[0035] In the context of the present description, a rank refers to a set of memory devices that are selectable by a same chip select signal from the memory controller. The number of ranks of memory devices in a memory module 110 may vary. For example, as shown, each

Micron Technology Inc. et al.
Ex. 1010, p. 12

memory module 110 may include four ranks of memory devices 112. In another embodiment, the memory module 110 may include 2 ranks of memory devices. In yet another embodiment, the memory module may include six or more ranks of memory devices 112.

[0036] In the context of the present description, a memory controller refers to any device capable of sending instructions or commands, or otherwise controlling the memory devices 112. Additionally, in the context of the present description, a memory bus refers to any component, connection, or groups of components and/or connections, used to provide electrical communication between a memory module and a memory controller. For example, in various embodiments, the memory bus 105 may include printed circuit board (PCB) transmission lines, module connectors, component packages, sockets, and/or any other components or connections that provide connections for signal transmission.

[0037] Furthermore, the memory devices 112 may include any type of memory devices. For example, in one embodiment, the memory devices 112 may include dynamic random access memory (DRAM) devices. Additionally, in one embodiment, each memory module 110 may include a dual in-line memory module (DIMM).

[0038] Referring to FIG. 2A, which illustrates one memory module 110 according to an embodiment, the module control device 116 receives system memory commands represented by a set of system control/address (C/A) signals from the MCH 101 via signal lines 120 and generates module command signals and module control signals based on memory commands from the system. The module control device 116 also received a system clock MCK and generates a module clock signal CK in response to the system clock signal MCK. The MCK signal may include a pair of complementary clock signals, MCK and $\overline{\text{MCK}}$, and the module clock signal may include a pair of complementary clock signals CK and $\overline{\text{CK}}$.

[0039] Examples of the system C/A signals include, but are not limited to, Chip Select (or /CS) signal, which is used to select a rank of memory devices to be accessed during a

Micron Technology Inc. et al.
Ex. 1010, p. 13

memory (read or write) operation; Row Address Strobe (or /RAS) signal, which is used mostly to latch a row address and to initiate a memory cycle; Column Address Strove (or /CAS) signal, which is used mostly to latch a column address and to initiate a read or write operation; address signals, including bank address signals and row/column address signals, which are used to select a memory location on a memory device or chip; Write Enable (or /WE) signal, which is used to specify a read operation or a write operation, Output Enable (or /OE) signal, which is used to prevent data from appearing at the output until needed during a read operation, and the system clock signal MCK.

[0040] Examples of module command signals include, but are not limited to module /CS signals, which can be derived from the system /CS signals and one or more other system C/A signals, such as one or more bank address signals and/or one or more row/column address signals; a module /RAS signal, which can be, for example, a registered version of the system /RAS signal; a module /CAS signal, which can be, for example, a registered version of the system /CAS signal; module address signals, which can be, for example, registered versions of some or all of the address signals; a module /WE signal, which can be, for example, a registered version of the system /WE signal; a module /OE signal, which can be, for example a registered version of the system /OE signal. In certain embodiments, the module command signals may also include the module clock signal CK.

[0041] Examples of module control signals include, but are not limited to a mode signal (MODE), which specifies a mode of operation (e.g., test mode or operating mode) for the isolation devices 118; one or more enable signals, which are used by an isolation device to select one or more subgroups of memory devices to communicate data with the memory controller; and one or more ODT signals, which are used by the isolation devices to set up on-die termination for the data/strobe signals. In one embodiment, the module control signals are transmitted to the isolation devices 118 via respective module control signal lines 230. Alternatively, the module control signals can be packetized before being transmitted to the isolation devices 118 via the module control signal lines and decoded/processed at the isolation devices.

Micron Technology Inc. et al.
Ex. 1010, p. 14

[0042] Module control device 116 transmits the module command signals to the memory devices 112 via module C/A signal lines 220. The memory devices 112 operate in response to the module command signals to receive write data or output read data as if the module command signals were from a memory controller. The module control device transmits the module control signals together with the module clock signal CK to the isolation devices 118 via module control signal lines 230. As shown in FIG. 2, at least some of the memory devices in a same rank share a same set of module C/A signal lines 220, and at least some of the isolation devices 118 share a same set of module control signal lines 230.

[0043] As shown n FIGS. 2A and 2B, each rank 114 includes *N* memory devices, where *N* is an integer larger than one. For example, a first rank includes memory devices M_{11}, ..., M_{i1}, $M_{i+1,1}$, ..., M_N, a second rank includes memory devices M_{12}, ..., M_{i2}, $M_{i+1,2}$, ..., $M_{N,2}$, and so on. In one embodiment, the memory devices 112 are also organized in groups or sets, with each group corresponding to a respective group of system data/strobe signal lines 130 and including at least one memory device from each rank. For example, memory devices M_{11}, M_{12}, M_{13}, and M_{14} form a first group of memory devices, memory devices M_{i1}, M_{i2}, M_{i3}, and M_{i4} form an ith group of memory devices, and so on.

[0044] As shown, the isolation devices 118 are associated with respective groups of memory devices and are coupled between respective groups of system data/strobe signal lines 130 and the respective groups of memory devices. For example, isolation device ID-1 among the isolation devices 118 is associated with the first group of memory devices M_{11}, M_{12}, M_{13}, and M_{14} and is coupled between the group of system data/strobe signal lines 130-1 and the first group of memory devices, isolation devices ID-*i* among the isolation devices 118 is associated with the ith group of memory devices M_{i1}, M_{i2}, M_{i3}, and M_{i4} and is coupled between the group of system data/strobe signal lines 130-i and the *i^{th} group* of memory devices, and so on.

[0045] In one embodiment, each group or sets of memory devices are coupled to the associated isolation device 118 via a set of module data/strobe lines 210. Each group or set of memory devices is organized in subgroups or subsets, with each subgroup or subset

Micron Technology Inc. et al.
Ex. 1010, p. 15

including at least one memory device. The subgroups in a group of memory devices may be coupled to the associated isolation device 118 via a same set of module data/strobe lines 210 (as shown in FIG. 2A) or via respective subsets of module data/strobe lines 210 (as shown in FIG. 2B). For example, as shown in FIG. 2B, in the first group of memory devices, memory devices M_{11} and/or M_{13} form a first subgroup, and memory devices M_{12} and/or M_{14} form a second subgroup; in the ith group of memory devices, memory devices M_{i1} and/or M_{i3} form a first subgroup, and memory devices M_{i2} and/or M_{i4} form a second subgroup; and so on. The first subgroup of at least one memory device in each group of memory devices is coupled to the associated isolation device 118 via an associated first subset of module data/strobe lines YA, and the second subgroup of at least one memory device in each group of memory devices is coupled to the associated isolation device via an associated second subset of module data/strobe lines YB, as shown. For example, memory devices M_{11} and/or M_{13} form the first subgroup are/is coupled to the isolation device ID-1 via the corresponding first subset of module data/strobe lines YA-1, and memory devices M_{12} and/or M_{14} form the second subgroup are/is coupled to the isolation device ID-1 via the corresponding second subset of module data/strobe lines YA-2.

[0046] In one embodiment, the isolation devices 118 are in the data paths between the MCH 101 and the memory module 110 and include data buffers between the MCH 101 and the respective groups of memory devices. In one embodiment, each isolation device 118 is configured to select a subgroup in the respective group of memory devices to communicate data with the MCH 101 in response to the module control signals, such that the memory module can include more ranks than what is supported by the MCH 101. Further, each isolation devices 118 is configured to isolate unselected subgroup(s) of memory devices from the MCH 101 during write operations, so that the MCH sees a load on each data line that is less than a load associated with the respective group of memory devices. In one embodiment, the MCH sees only a load associated with one memory device on each data/strobe signal line during write operations.

[0047] In one embodiment, the isolation devices 118 are distributed across the memory module 110 or the module board 119 in positions corresponding to the respective groups of

Micron Technology Inc. et al.
Ex. 1010, p. 16

memory devices. For example, isolation device ID-1 is disposed in a first position corresponding to the first group of memory devices M_{11}, M_{12}, M_{13}, and M_{14}, and isolation device ID-*i* is disposed in an ith position separate from the first position and corresponding to the ith group of memory devices M_{i1}, M_{i2}, M_{i3}, and M_{i4}. In one embodiment, the first position is between the first group of memory devices and an edge 201 of the module board 119 where connections (not shown) to the data/strobe signal lines 130 are placed, and i^{th} position is between the ith group of memory devices and the edge 201 of the module board 119. In one embodiment, the isolation devices 118 are distributed along the edge 201 of the memory module 110. In one embodiment, each isolation device 118 is a separate integrated circuit device packaged either by itself or together with at least some of the respective group of memory devices. In one embodiment, the module data/strobe signal lines 210, the module C/A signal lines 220, and the module control signal lines 230 include signal traces formed on and/or in the module board 119.

[0048] As an option, memory module 110 may further include a serial-presence detect (SPD) device 240, which may include electrically erasable programmable read-only memory (EEPROM) for storing data that characterize various attributes of the memory module 110. Examples of such data include a number of row addresses, a number of column addresses, a data width of the memory devices, a number of ranks on the memory module 110, a memory density per rank, a number of memory device on the memory module 110, and a memory density per memory device, etc. A basic input/output system (BIOS) of system 100 can be informed of these attributes of the memory module 110 by reading from the SPD 240 and can use such data to configure the MCH 101 properly for maximum reliability and performance.

[0049] In certain embodiments, the SPD 240 and/or the control circuit 116 store module configuration information, such as: memory space translation code, memory address mapping function code, input and output signals timing control information for the control circuit 116, input and output signals electrical and logical level control information for the control circuit 116, etc. In certain embodiments, the SPD 240 contains a system view of the module 110 which can be different from an actual physical construction of the module 110.

Micron Technology Inc. et al.
Ex. 1010, p. 17

For example, the SPD 240 stores at least one memory operation parameter that is different from a correspondng memory operation parameter in a system memory controller setting. The SPD 240 may also store at least on data buffer operation parameter that is different from a corresponding parameter in the system memory controller setting.

[0050] Thus, in certain embodiment, in the memory module 110, C/A signals representing a memory command are received and buffered by the module control circuit 116, so that the MCH sees only the module control circuit 116 as far as the C/A signals are concerned. Write data and strobe signals from the controller are received and buffered by the isolation devices 118 before being transmitted to the memory devices 112 by the isolation devices 118. On the other hand, read data and strobe signals from the memory devices are received and buffered by the isolation devices before being transmitted to the MCH via the system data/strobe signal lines 130. Thus, MCH 101 does not directly operate or control the memory devices 112. As far as data/strobe signals are concerned, the MCH 101 mainly sees the isolation devices 118, and the system 100 depends on the isolation devices 118 to properly time the transmission of the read data and strobe signals to the MCH 101.

[0051] In certain embodiments, the memory module 110 is a dual in-line memory module (DIMM) and the memory devices are double data rate (DDR) dynamic random access memory devices (DRAM). In certain embodiments, the control circuit 116 includes a DDR register, and logic for memory space translation between a system memory domain and a module level physical memory domain. Such translation may produce address mapping, proper interface timing for the control signals to the module level physical memory domain, and a proper interface electrical and logical level for the control signals to the module level physical memory domain.

[0052] As shown in FIG. 2C, in certain embodiments, the control circuit 116 transmits registered C/A and clock signals to the memory devices 112, and transmits module control signals and a registered clock signal (or module clock signal) to the isolation devices 118, in a fly-by configuration. As the speed of memory operations increase, issues can arise with respect to signal alignment for input, output delay variation due process, voltage and

Micron Technology Inc. et al.
Ex. 1010, p. 18

temperature (PVT) variations, synchronization with system memory controller interface, and phase drift accumulation during operation, etc. Electrical interface calibration drift during operation due to charge build up and timing interface calibration drift during operation due to environment change can also create issues.

[0053] For example, load reduction mechanism in the isolation devices 118 would provide a single data bus interface to the respective set of memory devices, which is hidden from the system memory controller 101. Thus, a long sequence of interface timing training may be required due to limited controllability of the system memory controller 101 over the interface between the memory devices 112 and the associated isolation devices 118. Furthermore, interface signal alignment-drift after the initial training would not be easily detected by the system memory controller 101, which may cause silent system failure.

[0054] Moreover, clock skew amongst the memory devices 112 and the associated isolation devices 118 due to the distributed architecture of the memory module 110 can cause synchronization issues. As the speed of memory operation increase, data period can become very close to the signal propagation delay time. Thus, such issues cannot simply be addressed by pipelining the data paths, as variation of the signal propagation time through I/Os becomes a very significant portion of a data period.

[0055] To address at least some of the above issues, in certain embodiments, as shown in FIG. 2D, the control circuit 116 transmits registered C/A signals to the memory devices 112, and transmits the module control signals and the module clock signal to the data buffers 118, in a fly-by arrangement. The memory devices 112 do not receive the module clock signal from the control circuit 116. Instead, each data buffer 118 regenerates the clock that is used by the respective set of memory devices 112. Each Data buffer 118 is thus responsible for providing a correct data timing interface between the respective set of memory devices 112 and the system memory controller 101. Each data buffer 118 is also responsible for providing the correct control signal timing between the control circuit 116 and the respective set of memory devices 112.

Micron Technology Inc. et al.
Ex. 1010, p. 19

[0056] Thus, the memory module 110 in FIG. 2D allows a locally synchronized operation for each respective set of memory devices 112, which can correspond to a nibble or a byte of a DDR data bus between the memory module 110 and the system memory controller 101. Also, signal interface between each data buffer 118 and the respective set of memory devices 112 can be synchronized. In one embodiment, each data buffer 118 has a set of configurable operations, including, for example: programmable phase relationship between the clock it receives and the clock it regenerates, programmable phase adjustment for the data and data-strobe signals coupled to the memory devices 112, programmable phase adjustment for the data and data-strobe signals coupled to the system memory controller 101, programmable phase adjustment related to at least one control signal that is coupled to the control circuit 116. The locally synchronized operation also makes it easier for each data buffer 118 to perform self-testing of the associated set of memory devices 112, independent of the self-testing of other sets of memory devices performed by the other data buffers, as disclosed in commonly-owned U.S. Patent 8,001,434, entitled "Memory Board with Self-Testing Capability," which is incorporated herein by reference in its entirety.

[0057] In certain embodiments, operations of the isolation devices 118 are controlled by the module control signals from the module control circuit 116, which generates the module control signals according to the C/A signals received from the MCH. Thus, the module control signals need to be properly received by the isolation devices 118 to insure their proper operation. In one embodiment, the module control signals are transmitted together with the module clock signal CK, which is also generated by the module control circuit 116 based on the system clock signal MCK. The isolation circuits 118 buffers the module clock signal, which is used to time the sampling of the module control signals. Since the isolation devices 118 are distributed across the memory module, the module control signal lines 230 can stretch across the memory module 110, over a distance of several centimeters. As the module control signals travel over such a distance, they can become misaligned with the module clock signal, resulting in metastability in the received module control signals. Therefore, in one embodiment, the isolation circuits 118 includes metastability detection circuits to detect metastability condition in the module control

Micron Technology Inc. et al.
Ex. 1010, p. 20

signals and signal adjustment circuits to adjust the module control signals and/or the module clock signal to mitigate any metastability condition in the module control signals, as explained in further detail below.

[0058] Because the isolation devices 118 are distributed across the memory module 110, during high speed operations, it may take more than one clock cycle time of the system clock MCK for the module control signals to travel along the module control signals lines 230 from the module control device 116 to the farthest positioned isolation devices 118, such as isolation device ID-1 and isolation device ID-(n-1) in the exemplary configuration shown in FIG. 2. In other words, a same set of module control signals may reach different isolation devices 118 at different times across more than one clock cycle of the system clock. For example, when the clock frequency of the system clock is higher than 800 MHz, the clock cycle time is less than about 1.2 ns. With a signal travel speed of about 70ps per centimeter of signal line, a module control signal would travel about 15 cm during one clock cycle. When the clock frequency increases to 1600 MHz, a module control signal would travel less than 8 cm during one clock cycle. Thus, a module control signal line can have multiple module control signals on the line at the same time, i.e., before one module control signal reaches an end of the signal line, another module control signal appear on the signal line.

[0059] With the isolation devices 118 receiving module control signals at different times across more than one clock cycle, the module control signals alone are not sufficient to time the transmission of read data signals to the MCH 101 from the isolation devices 118. In one embodiment, each isolation devices includes signal alignment circuits that determine, during a write operation, a time interval between a time when one or more module control signals are received from the module control circuit 116 and a time when a write strobe or write data signal is received from the MCH 101. This time interval is used during a subsequent read operation to time the transmission of read data to the MCH 101, such that the read data follows a read command by a read latency value associated with the system 100, as explained in more detail below.

Micron Technology Inc. et al.
Ex. 1010, p. 21

[0060] More illustrative information will now be set forth regarding various optional configurations, architectures, and features with which the foregoing framework may or may not be implemented, per the desires of the user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the exclusion of other features described.

[0061] In one embodiment, as shown in FIG. 3, each group of signal lines 130 include a set of *n* data (DQ) signal lines 322 each for transmitting one of a set of data signals DQ_0, DQ_1, ..., DQ_{n-1}, and at least one strobe (DQS) signal line 324 for transmitting at least one strobe signal DQS. Each set of module data/strobe lines Y include a set of *n* module data signal lines Y_0, Y_1, ..., Y_{n-1} and at least one module strobe signal line Y_{DQS}. When the subsets of memory devices are coupled to the associated isolation device 118 via respective subsets of memory devices, each set of module data/strobe lines Y may include multiple subsets of module data/strobe lines, such as the subsets of module data/strobe lines YA and YB shown in FIG. 2B. Each subset of module data/strobe lines YA include a set of *n* first module data lines YA_0, YA_1, ..., YA_n and at least one first module strobe signal line YA_{DQS}; and each subset of module data/strobe lines YB include a set of *n* second module data lines YB_0, YB_1, ..., YB_n and at least one second module strobe signal line YB_{DQS}.

[0062] Each isolation device 118 includes a set of DQ routing circuits 320 coupled on one side to respective ones of the set of *n* DQ signal lines 322, and on another side to respective ones of the respective set of *n* module data lines, or respective ones of the respective subsets of module data lines, such as the first module data lines YA_0, YA_1, ..., YA_n and the second module data lines YB_0, YB_1, ..., YB_n. Each isolation device 118 further includes an ID control circuit 310 coupled on one side to the at least one DQS signal line 324, on another side to the one or more module strobe signal lines Y_{DQS}, or the first module strobe signal line YA_{DQS} and second module strobe signal line YB_{DQS}. The ID control circuit 310 also receives the module clock signal CK and the module control signals via the module control signal lines 230, and outputs ID control signals 330 to the DQ routing circuits 320, including, for example, one or more enable signals ENA and/or ENB, and some or all of

Micron Technology Inc. et al.
Ex. 1010, p. 22

the other received, decoded, and/or otherwise processed module control signals, a delay signal DS, a read DQS signal RDQS, a write DQS signal WDQS, and a buffer clock signal CK0. Each DQ routing circuit 320 is configured to enable data communication between the respective DQ signal line 322 with a selected subgroup of one or more memory devices in response to the module control signals, as explained in more detail below.

[0063] In certain embodiments, the ID control circuit 310 also provides a delay signal DS, which is used by the DQ routing circuits 320 to align read data output by the isolation device 118 with read data output by the other isolation devices 118, as explained in further detail below. In certain embodiments, the ID control circuit 310 regenerates a clock signal from the module clock signal CK, which can have a programmable delay from the module clock signal. The regenerated clock signal is used as the clock signal CK0 and a clock signal CKM that is provided to the corresponding set of memory devices, as explained in more detail below.

[0064] The memory devices 112 are coupled to the isolation devices 118 via a same set of module data/strobe signal lines or different subsets of module data/strobe signal lines. For example, as shown in FIG. 4A, memory devices M_{11}, M_{12}, M_{13}, and M_{14} in the first group of memory devices can be coupled to the isolation device ID-1 via a same set of module data lines $Y\text{-}1_0$, $Y\text{-}1_1$, ..., $Y\text{-}1_{n-1}$ and module strobe line $Y\text{-}1_{DQS}$. In such embodiment, a subgroup in the group of memory devices can be selected by the isolation devices to communicated data with the MCH based on the phases of the data/strobe signals, which can be different with respect to different subgroups of memory devices.

[0065] Alternatively, as shown in FIG. 4B, memory devices M_{11} and M_{13}, which form a subgroup in the first group of memory devices, are coupled to the isolation device ID-1 via the module data lines $YA\text{-}1_0$, $YA\text{-}1_1$, ..., $YA\text{-}1_n$ and module strobe line $YA\text{-}1_{DQS}$ and memory devices M_{12} and M_{14}, which form another subgroup in the first group of memory devices, are coupled to the isolation device ID-1 via the module data lines $YB\text{-}1_0$, $YB\text{-}1_1$, ..., $YB\text{-}1_n$ and module strobe line $YB\text{-}1_{DQS}$. Memory devices coupled to the same isolation devices can be disposed on a same side or different sides of the memory board

Micron Technology Inc. et al.
Ex. 1010, p. 23

119. Memory devices coupled to the same isolation devices may be placed side by side, on opposite sides of the module boards 119, or stacked over each other, and/or over the associated isolation device.

[0066] Multiple memory devices having a data width that is less than a data width of the isolation devices 118 may be used in place of one of the memory devices 112, which has the same data width as that of the isolation devices. For example, as shown in FIG. 5A, two memory devices $M_{11\text{-}1}$ and $M_{11\text{-}2}$ may be used in place of the memory device M_{11}. Each of the two memory devices $M_{11\text{-}1}$ and $M_{11\text{-}2}$ has a data width of 4, and together they act like a memory device M_{11} of a data width of 8. Thus, memory device $M_{11\text{-}1}$ is coupled to the isolation device ID-1 via module data lines $YA\text{-}1_0$, ..., $YA\text{-}1_3$ and module strobe line $YA\text{-}1_{DQS\text{-}1}$ while memory circuit $M_{11\text{-}2}$ is coupled to the isolation device ID-1 via module data lines $YA\text{-}1_4$, ..., $YA\text{-}1_7$ and module strobe line $YA\text{-}1_{DQS\text{-}2}$.

[0067] In another embodiment, as shown in FIG. 5B, four memory devices $M_{11\text{-}1}$ to $M_{11\text{-}4}$ may be used as the memory device M_{11}. Each of the four memory devices $M_{11\text{-}1}$ to $M_{11\text{-}4}$ has a data width of 4, and together they act like a memory device M_{11} of a data width of 16. Thus, memory device $M_{11\text{-}1}$ is coupled to the isolation device ID-1 via module data lines $YA\text{-}1_0$, ..., $YA\text{-}1_3$ and module strobe line $YA\text{-}1_{DQS\text{-}1}$ while memory device $M_{11\text{-}2}$ is coupled to the isolation device ID-1 via module data lines $YA\text{-}1_4$, ..., $YA\text{-}1_7$ and module strobe line $YA\text{-}1_{DQS\text{-}2}$, and so on.

[0068] FIG. 6 illustrates the ID control circuit 310 in an isolation device 118. As shown, the ID control circuit 310 includes a clock buffer 610 to receive the module clock signal CK from the module control device 116, and to output a module clock signal CK0. The ID control circuit 310 further includes a strobe routing circuit 620 that are coupled on one side to the corresponding system DQS signal line 324 and on another side to the corresponding module DQS signal lines YA_{DQS} and YB_{DQS}. The ID control circuit 310 further includes a receiver circuit 630 with respect to each of at least some of the module control signals (MCS) to receive a respective one of the module control signals. The ID control circuit 310 further includes a command processing circuit 640 that provides the received, decoded,

Micron Technology Inc. et al.
Ex. 1010, p. 24

and/or otherwise processed module control signals 330 to the DQ routing circuits 320 and the strobe routing circuit 620 either directly or after further processing, if needed. The received/decoded/processed module control signals may include, for example, one or more enable signals ENA and/or ENB that are used by the DQ routing circuits 320 and the strobe routing circuit 620 to selectively enabling data communication between the MCH 101 and one of the subgroups in the respective group of memory devices, with which the isolation device is associated.

[0069] The strobe routing circuit 620 also buffers strobe signals received from either the MCH 101 or the memory devices 112, and output either a write strobe WDQS or read strobe RDQS to the DQ routing circuits 320. In one embodiment, the ID control circuit 310 further includes a delay control circuit 650 that receives one of the module control signals and either a data signal or a strobe signal and determines a delay amount to be used by the DQ routing circuit 320 and the strobe routing circuit 620. The delay amount is provided to the DQ routing circuit 320 and the strobe routing circuit in a delay signal DS.

[0070] In a receiver circuit 630, the respective MCS is received in accordance with the module clock signal CK0. In one embodiment, receiver circuit 630 samples the respective MCS using rising (or falling) edges of the module clock CK0. Since the isolation devices 118 are distributed across the memory module 110 at positions corresponding to the respective groups of memory devices, the module control signal lines 230 that carry the MCS to the isolation devices can stretch over a distance of more than 10 centimeters, as shown in FIG. 7. As the MCS and CK0 travel along their respective module control signal lines 710 and 720, they can become misaligned with each other when they reach the input pins 730 of an isolation device 118.

[0071] For example, a module control signal, like the MCS 810 shown in FIG. 8, can be perfectly aligned with the module clock signal CK, with a rising edge 801 of the module clock signal CK being at a center of a data eye 802, when the MCS signal and the clock signal leave the module control circuit 116. When the module control signal and the module clock signal reach an isolation device, however, their alignment can become shifted

Micron Technology Inc. et al.
Ex. 1010, p. 25

like the MCS 820 with respect to the CK signal, i.e., the rising edge 801 of the clock signal is near a left edge of a data eye of the MCS 820, barely providing enough set up time for proper sampling of the module control signal. Or, the module control signal, like the MCS 830, can be shifted with respect to the module clock signal such that a rising edge 801 of the clock signal is near a right edge of a data eye of the MCS, barely providing enough hold time for proper sampling of the module control signal. Or, ever worse, the module control signal, like the MCS 840, can be so shifted with respect to the module clock signal such that a rising edge 801 of the clock signal falls in the glitches 803 at the edge of a data eye of the MCS, meaning that the sampled results could be metastable.

[0072] In one embodiment, as shown in FIG. 9, a receiver circuit 630 includes a metastability detection circuit (MDC) 910 to determine a metastability condition in a corresponding module control signal MCS0. In one embodiment, the MDC 910 generates at least one delayed version of the module clock signal CK and at least one delayed version of the corresponding MCS0. The MDC 910 also generates one or more metastability indicators and outputs the one or more metastability indicators via lines 912 and/or 914.

[0073] The receiver circuit 630 further includes a signal selection circuit 920 that receives the module clock CK and the at least one delayed version of the module clock via signal lines 916. The signal selection circuit 920 also receives the corresponding MCS and the at least one delayed version of the corresponding MCS via signal lines 918. The signal selection circuit 920 selects a clock signal CK_i from among the module clock CK and the at least one delayed version of the module clock based on one or more of the metastability indicators. The signal selection circuit 920 may also select an MCS signal MCS_i from among the corresponding MCS and the at least one delayed version of the corresponding MCS based on at least one other metastability indicator.

[0074] The receiver circuit 630 further includes a sampler or register circuit 930 that samples the selected module control signal MCS_i according to the selected clock signal CK_i and outputs the sampled signal as the received module control signal, which is provided to the command processing circuit 640 for further processing (if needed) before being

Micron Technology Inc. et al.
Ex. 1010, p. 26

provided to the DQ routing circuits 320 and DQS routing circuit 620.

[0075] FIG. 10A illustrates an MDC 910 according to one embodiment. As shown, the MDC 910 includes a delay circuit 1012 that generates a delayed version MCS1 of the corresponding MCS0 by adding a predetermined amount of delay (e.g., 10ps) to MCS0. MDC 910 also includes a delay circuit 1016 that generates a delayed version CK1 of the clock signal CK0 by adding a predetermined amount of delay to CK0. In one embodiment, CK1 is delayed from CK0 by about 1/10th of a clock cycle, e.g., 50-70 ps for an operating frequency of about 1600 MHz. The MDC 910 further includes a sampler circuit 1042 that samples MCS1 according to CK0 and outputs a sampled result A, a sampler circuit 1044 that samples MCS0 according to CK0 and outputs a sampled result B, and a sampler circuit 1046 that samples MCS0 according to CK1 and outputs a sampled result C. The MDC 910 further includes a logic circuit (e.g., a majority decision circuit) that generates metastability indicators Z1 and Z2 based on the sampled results A, B, and C.

[0076] In one embodiment, Z1 is the result of a logic operation (e.g., an XNOR operation) on the sampled result, e.g., $Z1 = \overline{A \oplus B}$, and Z2 is the result of another logic operation on the sampled results, e.g., $Z2 = \overline{B \oplus C}$. Thus, as shown in FIG. 10B and Table 1 below, when a metastability condition of insufficient hold time occurs, i.e., a rising clock edge 1061 of CK0 is close to the right side of a data eye where gliches at the edges of the data eyes can make C unpredicatable, A and B can be in agreement (i.e., Z1 is true) while B and C are likely not in agreement (i.e., Z2 is false). FIG. 10 C illustrates a metastability condition when there is insufficient set-up time. As shown in FIG. 10C and Table 1 below, a rising clock edge 1061 of CK0 is close to the left side of a data eye where gliches at the edges of the data eyes can make A unpredicatable. Thus, A and B can be in disagreement so Z1 is false while B and C can be in agreement so Z2 is true. Not shown in the figures is the situation that all A, B, and C are in agreement, meaning that both the rising clock edge 1061 of CK0 and the rising clock edge 1062 of CK1 are near the middle of an MCS0 data eye so there is no metastability issues and both Z1 and Z2 are true, as shown in Table 1.

[0077] FIG. 10D illustrates a signal selection circuit 920 according to an embodiment. As

Micron Technology Inc. et al.
Ex. 1010, p. 27

shown, in one embodiment, the signal selection circuit 920 includes a first multiplexor 1071 that selects between CK0 and CK1 based on the metastability indicator Z1, and a second multiplexor 1072 that selects between MCS0 and MCS1 based on the metastability indicator Z2. Thus, as shown in Table 1, where a metastability condition of insufficient hold time occurs, Z1=1 and Z2=0, and MCS1 is output from multiplexor 1071 while CK0 is output from multiplexor 1072. Sampler 930 thus samples MCS1 according to the rising edges of CK0. Thus, more hold time is provided to mitigate the metastability condition since MCS1 is shifted from MCS0 toward the right.

[0078] On the other hand, where a metastability condition of insufficient set-up time occurs, Z1=0 and Z2=1, and CK1 is output from multiplexor 1071 while MCS0 is output from multiplexor 1072. Sampler 930 thus samples MCS0 according to the rising edges of CK1. Since CK1 is shifted from CK0 toward the right, more set-up time is provided to mitigate the metastability condition.

[0079] In the case when no metastability is detected, Z1=1 and Z2=1, and CK0 is output from multiplexor 1071 while MCS0 is output from multiplexor 1072. So, the unshifted module control signal is sampled according to the unshifted module clock signal.

Table 1 Metastability Detection and Signal Selection

Sampler Output			MS Indicators		MS Condition	Signal Selection	
A	B	C	Z1	Z2		CK	MCS
D1	D1	D2	1	0	insufficient hold time	CK0	MCS1
D1	D2	D2	0	1	insufficient set-up time	CK1	MCS0
D1	D1	D1	1	1	no metastability	CK0	MCS0

[0080] FIGS. 10A-10D illustrate a relatively simple implementation of the metastability detection circuit (MDC) 910 where only three different sample points are provided to detect metastability condition in the module control signal. In general, the MDC 910 may generate more delayed versions of the module clock signal CK0 and/or the corresponding module control signal MCS0, and may include more sampler circuits to sample any

Micron Technology Inc. et al.
Ex. 1010, p. 28

additional delayed versions of the module control signal according to either the module clock signal or one of the delayed versions of the module clock signal. For example, as shown in FIG. 11A, the MDC 910 can include a plurality of delay circuits 1102 that generate *m* delayed versions of MCS0, e.g., MCS1, MCS2, ... MCS*m*, and *m* delayed versions of CK0, e.g., CK1, CK2, ..., CK*m*. The MDC 910 can include sampler circuits 1104 that sample MCS0 according to CK0, CK1, ..., CK*m*, respectively, and sampler circuits 1104 that sample MCS0, MCS1, MCS2, ... MCS*m* according to CK0, respectively. The outputs of the samplers 1104 are provided to a logic circuit 1120, which determines a metastability condition in MCK0 based on the sampler outputs using, for example, a majority decision logic. The logic circuit 1120 outputs a first metastability indicator on line(s) 912 and a second metastability indicator on line(s) 914.

[0081] FIG. 11B illustrates a signal selection circuit 920 according to an embodiment. As shown, in one embodiment, the signal selection circuit 920 includes a first multiplexor 1171 that selects between CK0, CK1, ..., CK*m* based on the metastability indicator provided on line(s) 912, and a second multiplexor 1172 that selects between MCS0, MCS1, ..., MCS*m* based on the metastability indicator provided on line(s) 914, such that the rising edges of the selected clock signal, e.g., Ck*i*, are close to the middle of the respective data eyes in the selected module control signal, e.g., MCS*i*. The selected signals MCS*i* and Ck*i* are provided to the sampler 930, which samples MCS*i* according to the rising edges of CK*i*.

[0082] As stated above, in certain embodiments, since the isolation devices 118 are in the data paths between the MCH 101 and the respective groups of memory devices 112, the MCH 101 does not have direct control of the memory devices 112. Thus, conventional read/write leveling techniques are not sufficient for managing read/write data timing. In one embodiment, the isolation devices 118 includes signal alignment mechanism to time the transmission of read data signals based on timing information derived from a prior write operation, as discussed further below.

[0083] FIG. 12A is a timing diagram for a write operation according to one embodiment.

Micron Technology Inc. et al.
Ex. 1010, p. 29

As shown, after a write command W/C associated with the write operation is received by the module control circuit 116 at time t1, the module control circuit 116 outputs one or more enable signals EN at time t2 in response to the write commands. The one or more enable signals are received by an isolation device 118 at time t3, which afterwards receives one or more strobe signal DQS from the MCH 101 at time t4. Note that the same enable signal may be received by another isolation device 118 at time t3′, which can be in a different cycle of the system clock MCK from the cycle which t3 is in. The time interval between t4 and t1 is consistent with a write latency W.L. associated with the system 100, and is controllable by the MCH 101 and knowable to the isolation device 118. The time interval between t4 and t3, referred to hereafter as an enable-to-write data delay EWD, can be determined by the isolation device 118 since both these signals are received by the isolation device. Based on such determination, the isolation device 118 can have knowledge of the time interval between t3 and t1, referred to hereafter as a command-to-enable delay CED, which can be used by the isolation device 118 to properly time transmission of read data to the MCH, as explained further below.

[0084] FIG. 12B is a timing diagram for a read operation according to one embodiment. As shown, after a read command R/C associated with the read operation is received by the module control circuit 116 at time t5, the module control circuit 116 outputs one or more enable signals EN at time t6 in response to the read commands. The one or more enable signals are received by an isolation device 118 at time t7, which afterwards receives at time t8 read data signals (not shown) and one or more strobe signal DQS from the respective group of memory devices. Note that the same enable signal may be received by another isolation device 118 at time t3′, which can be in a different cycle of the system clock MCK from the cycle which t3 is in. Thus, the enable signals alone cannot be used to time the transmission of the read signals by the isolation devices 118.

[0085] With knowledge of the time interval between t7 and t5, which should be about the same as the time interval between t3 and t1, i.e., the command-to-enable delay CED, in certain embodiments, the isolation device can add a proper amount of delay to the read data signals and the one or more DQS signal such that the read data signals and the one or more

Micron Technology Inc. et al.
Ex. 1010, p. 30

DQS signal are transmitted at time t9 by the isolation device to the MCH 101 via the respective group of data/strobe signal lines 130, with the time interval between t9 and t5 being consistent with a read latency R.L. associated with the system 100.

[0086] The time interval between t4 and t3, i.e., the enable to write data delay EWD, is determined by the delay control circuit 650 in the ID control circuit 310, as shown in FIG. 6. According to one embodiment, as shown in FIG. 13, the delay control circuit 650 includes a peramble detector 1310 to detect a write preamble in the DQS, a flip-flop circuit 1320 having an enable input EN receiving one of the module control signals and a clock input CK receiving the buffered module clock signal CK0, and a counter circuit 1330 having a Start input receiving the one of the module control signals, a Stop input receiving an output of the flip-flop circuit 1320. Thus, the output of the counter circuit, i.e., the delay signal DS, would indicate a time interval from when the write preamble is detected and when the one of the module control signal is received.

[0087] FIG. 14 illustrates a DQ or DQS routing circuit 320 or 620 according to an embodiment. As shown, the DQ/DQS routing circuit 320/620 includes a DQ/DQS pin 1401 that is coupled to the corresponding DQ/DQS signal line 322/324, a set of one or more DQS pins 1402 that is coupled to a corresponding module DQ/DQS line(s) Y/Y_{DQS}, or YA/YA_{DQS} and YB/YB_{DQS}. The DQ/DQS routing circuit 320/620 further includes a write strobe buffer 1410 that buffers write data/strobe, and a write data/strobe receiver 1420 that samples the write data/strobe. The DQ/DQS routing circuit 320/620 further includes a plurality of write paths 1430 that are selectable or can be selectively enabled by one or more of the module control signals, such as the enable signals ENA and ENB.

[0088] The DQS routing circuit further includes a plurality of read paths 1450 that are selectable by the one or more of the module control signals. Output from the selcted read path is delayed in a delay circuit 1460 by an amount controlled by the delay signal DS, and sampled by a sampler circuit 1470. The sampled read data/strobe is transmitted by transmitter 1480 onto the corresponding data/strobe signal line 322/324 via the DQ/DQS pin 1401.

Micron Technology Inc. et al.
Ex. 1010, p. 31

[0089] FIG. 15 illustrates a DQS routing circuit 620 according to an embodiment. As shown, the DQS routing circuit 620 includes a first DQS pin 1501 that is coupled to a corresponding DQS signal line 324, a second DQS pin 1502A that is coupled to a corresponding module DQS line YA_{DQS}, a third DQS pin 1502B that is coupled to a corresponding module DQS line YB_{DQS}. The DQS routing circuit 620 further includes a first write strobe path coupled between the first DQS pin 1501 and the second DQS pin 1502A and a second write strobe path coupled between the first DQS pin 1501 and the third DQS pin 1502B. The first write strobe path includes a write strobe buffer 1510 that buffers a write strobe, a write strobe receiver 1520 that samples the write strobe according to the buffered module signal CK0. The sampled write strobe is provided to the DQ routing circuits 320 as the write strobe WDQS. The first write strobe path further includes a first write strobe transmitter 1530A that transmits the write strobe to one or more memory devices 112 coupled to the module strobe line YA_{DQS}. The second write strobe path includes the write strobe buffer 1510, the write strobe receiver 1520, and a second write strobe transmitter 1530B that transmits the write strobe to one or more memory devices 112 coupled to the module strobe line YB_{DQS}. The first and second write strobe transmitters, 1530A and 1530B, are controlled by two enable signals, ENA and ENB, respectively, such that the first write strobe path and the second write strobe path can be selectively enabled/disabled by the enable signals, ENA and ENB.

[0090] The DQS routing circuit further includes a read strobe path coupled between the first DQS pin 1501 and a selected one of the second and third DQS pins 1502A and 1502B. In the read strobe path, a select circuit 1550 (e.g., a multiplexor) selects either a read strobe signal received via DQS pin 1502A or a read strobe signal received via DQS pin 1502B based on one or both of the enable signals ENA or ENB. The selected read strobe signal is delayed in a delay circuit 1560 by an amount controlled by the delay signal DS, and sampled by a sampler circuit 1570 according to the buffered module clock signal CK0. The sampled read strobe is provided to the DQ routing circuits 320 as the read strobe RDQS and is transmitted by transmitter 1580 onto the corresponding strobe signal line 324 via the first DQS pin 1501.

Micron Technology Inc. et al.
Ex. 1010, p. 32

[0091] FIG. 16 illustrates a DQ routing circuit 320 according to an embodiment. As shown, the DQ routing circuit 320 includes a first DQ pin 1601 that is coupled to a corresponding DQ signal line 130, a second DQ pin 1602A that is coupled to a corresponding module DQ line YA_{DQ}, a third DQ pin 1602B that is coupled to a corresponding module DQ line YB_{DQ}. The DQ routing circuit 320 further includes a first write data path coupled between the first DQ pin 1601 and the second DQ pin 1602A and a second write data path coupled between the first DQ pin 1601 and the third DQ pin 1602B. The first write data path includes a write data buffer 1610, a write data receiver 1620 that samples write data according to the write strobe WDQS from the DQS routing circuit 620, and a first write data transmitter 1630A that transmits the write data to one or more memory devices 112 coupled to the module data line YA_{DQ}. The second write data path includes the write data buffer 1610, the write data receiver 1620, and a second write data transmitter 1630B that transmits the write data to one or more memory devices 112 coupled to the module data line YB_{DQ}. The first and second write data transmitters, 1530A and 1530B, are controlled by two enable signals, ENA and ENB, respectively. Thus, the first write data path and the second write data path can be selectively enabled/disabled by the enable signals, ENA and ENB.

[0092] The DQ routing circuit further includes a read data path coupled between the first DQ pin 1601 and a selected one of the second and third DQ pins 1602A and 1602B. In the read data path, a select circuit 1650 (e.g., a multiplexor) selects either a read data signal received via DQ pin 1602A or a read data signal received via DQ pin 1602B based on one or both of the enable signals ENA or ENB. The selected read data signal is delayed in a delay circuit 1660 by an amount controlled by the delay signal DS. The delayed read data signal is then sampled by a receiver circuit 1670 according to the read strobe RDQS from the DQS routing circuit 620, and transmitted by transmitter 1680 onto the corresponding data signal line 130 via the first DQ pin 1601.

[0093] FIG. 17 illustrate a delay circuit 1560 or 1660 according to an embodiment. As shown, the delay circuit 1560 or 1660 includes a plurality of delay stages, such as delay stages 1710, 1720, and 1730, each delaying a read data or read strobe signal from the select

Micron Technology Inc. et al.
Ex. 1010, p. 33

circuit 1550/1650 by a predetermined amount. The delay circuit 1560 or 1660 further includes a select circuit 1740 (e.g., a multiplexor) that selects from among the read data or read strobe signal and the outputs from the delay stages according to the delay signal DS. The output of the select circuit 1740, is provided to the sampler circuit 1570 or 1670, either directly or after being buffered by a buffer circuit 1750.

[0094] Thus, as shown in FIG. 18, in one embodiment, a memory module 110 operates in the memory system 100 according to a method 1800. In the method, during a write operation, one or more module control signals are received by an isolation device 118 from a module control circuit or module controller 116 (1810). The module controller 116 generates the one or more module control signals in response to C/A signals representing a write command from the MCH 101. The one or more module control signals are used to control the isolation device 118. For example, the one or more module control signals may include one or more first enable signals to enable a write path to allow write data be communicated to a selected subgroup of memory devices among the group of memory devices coupled to the isolation device 118. After a time interval from receiving the one or more first enable signals, write data DQ and write strobe DQS are received by the isolation device 118 from the MCH 101 (1820). In one embodiment, upon receiving the one or more first enable signal, a counter is started, which is stopped when the write data DQ or write strobe DQS is received. Thus, a time interval EWD between receiving the one or more first enable signals and receiving the write strobe signal DQS is recorded.

[0095] Since the time interval between the arrival of the command signals from the MCH 101 and the arrival of the write data/strobe signal DQ/DQS from the MCH 101 is a set according to a write latency parameter associated with the system 100, the time interval EWD can be used to ascertain a time interval CED between the time when a command signal is received by the memory module 110 and the time when the one or more enable signals are received by the isolation device 118. The time interval CED can be used by the isolation device 118 to properly time the transmission of read data to the MCH 101, as described above and explained further below.

Micron Technology Inc. et al.
Ex. 1010, p. 34

[0096] As shown in FIG. 18, a delay signal DS is generated according to the time interval EWD (1830). Concurrent to receiving the write strobe signal DQS, the isolation device 118 also receives a set of write data signals DQ (1840). The received write data signals are transmitted to the subgroup of memory devices (1850), which are selected from the group of memory devices coupled to the isolation device 118 by the one or more first enable signals.

[0097] During a read operation, another set of module control signals including, for example, one or more second enable signals, are received by the isolation device 118 from the module controller 116 (1860). The one or more second enable signals are generated by the module controller 116 in response to read command signals received from the MCH 101, and are used by the isolation device 118 to select a subgroup of memory devices from which to receive read data. Afterwards, a read strobe signal DQS and a set of read data signal DQ are received from the selected subgroup of memory devices (1870). To properly time the transmission of the DQS and DQ signals to the MCH 101, the DQS and DQ signals are adjusted (e.g., delayed) according to the delay signal DS, such that the DQS and DQ signals follow a read command by a time interval consistent with a read latency parameter associated with the system 100.

[0098] In certain embodiments, especially the embodiments shown in FIG. 2D, the delay circuits 1560 and 1660 shown in FIGS. 15 and 16 are not needed to provide alignment of the read data. As shown in FIG. 19, the ID control circuit 310 includes a clock regeneration circuit 1920 that regenerates the clock signal CK received from the control circuit 116, according to the delay signal DS. The regenerated clock signals CK0 and CKM each includes a proper amount of delay as compared to the clock signal CK. The clock CK0 is provided to the strobe routing circuit 620 so that the strobe signals are properly timed to result in proper data alignment. The regenerated clock signal CKM is provided to the respective set of memory devices so that the respective data buffer 118 and the respective set of memory devices are locally synchronized.

Micron Technology Inc. et al.
Ex. 1010, p. 35

WE CLAIM:

1. A memory module to operate in a memory system with a memory controller, the memory system operating according to a system clock, comprising:

a module control device to receive command signals from the memory controller and to output module command signals and module control signals;

memory devices organized in groups, each group including at least one memory device, the memory devices receiving the module command signals from the module control device; and

a plurality of buffer circuits to receive the module control signals, a respective buffer circuit corresponding to a respective group of memory devices, the respective buffer circuit including data paths for communicating data between the memory controller and the respective group of memory devices, the data paths being controlled by at least one of the module control signals; and

wherein the plurality of buffer circuits are distributed across a surface of the memory module such that each module control signal arrives at the plurality of buffer circuits at different points in time, and

wherein the plurality of buffer circuits are configured to align read data signals received from the respective groups of memory devices such that the read data signals are transmitted to the memory controller from the memory module substantially aligned with each other and in accordance with a read latency parameter of the memory system.

2. The memory module of claim 1, wherein each module control signal arrives at the plurality of buffer circuits at different points in time across more than one clock cycle of the system clock, and wherein the plurality of buffer circuits are configured to align the read data signals such that they are transmitted by the buffer circuits to the memory controller within one system clock cycle.

3. The memory module of claim 1, wherein each buffer circuit includes receiver circuits to receive corresponding ones of the module control signals, each receiver circuit including a metastability detection circuit to detect metastability condition in the

Micron Technology Inc. et al.
Ex. 1010, p. 36

corresponding module control signal, and wherein each buffer circuit further includes at least one signal adjustment circuit to adjust one or more the module control signals to mitigate any metastability condition in the module control signals.

4. The memory module of claim 1, wherein the module control device further receives the system clock from the system memory controller and transmits a registered clock signal to the plurality of buffer circuits, and wherein each respective buffer circuit regenerates one or more clock signals from the registered clock signal and provides a regenerated clock signal to the respective group of memory devices.

5. The memory module of claim 1, wherein each buffer circuit comprises:

a time interval determination circuit to receive, during a write operation, a first signal from the module controller and a second signal from the memory controller and to generate a delay signal indicating a time interval between the first signal and the second signal;

a delay circuit to delay a signal in the buffer circuit according to the delay signal indicating the time interval, the signal in the buffer circuit being one of a read data signal, a read strobe signal, and a regenerated clock signal.

6. A method of operating a memory module in a memory system having a memory controller, comprising:

receiving a read command from the memory controller;

generating module command signals and module control signals in response to the read command;

producing read data from the memory devices in response to the module command signals;

controlling a plurality of data paths in a plurality of data buffers using the module control signals; and

outputting read data from the plurality of data paths to the memory controller, wherein at least one signal in each of the plurality of data buffers is delayed such that the read data output from the plurality of data paths are aligned with each other and in

Micron Technology Inc. et al.
Ex. 1010, p. 37

accordance with a read latency parameter of the memory system.

7. The method of claim 6, further comprising:

determining a metastability condition associated with a respective module control signal at each of the plurality of data buffers; and

adjusting the respective module control signal based on the metastability condition.

8. The method of claim 6, further comprising, in the each data buffer:

receiving at least one first module control signal generated in the memory module in response to a write command;

determining a time interval between receiving the at least one first module control signal and receiving a write signal associated with the write command from the memory controller, the write signal being one of a write data signal and a write strobe signal; and

delaying the at least one signal in the respective data buffer according to the time interval.

9. The method of claim 8 wherein the at least one signal is a clock signal.

10. The method of claim 8 wherein the at least one signal is at least one of a read data signal and a read strobe signal.

11. The method of claim 8, wherein the write signal is a write strobe signal, the method further comprising:

detecting a preamble of the write strobe signal to determine the time interval.

12. A buffer circuit for use in a memory module coupled to a memory controller, the memory module including a module controller, comprising:

a time interval determination circuit to receive, during a write operation, a first signal from the module controller and a second signal from the memory controller and to generate a delay signal indicating a time interval between the first signal and the second signal; and

a delay circuit to delay at least one signal in the buffer circuit according to the delay

Micron Technology Inc. et al.
Ex. 1010, p. 38

signal indicating the time interval.

13. The buffer circuit of claim 12 wherein the at least one signal is a clock signal.

14. The buffer circuit of claim 12 wherein the at least one signal is at least one of a read data signal and a read strobe signal.

15. The buffer circuit of claim 12, wherein the first signal is an enable signal to enable a write data path in the buffer circuit.

16. The buffer circuit of claim 15, wherein the second signal is a write strobe signal.

17. The buffer circuit of claim 15, wherein the second signal has a preamble, and wherein the time interval determination circuit includes a preamble detector to detect the preamble of the second signal.

18. The buffer circuit of claim 18, wherein the time interval determination circuit further includes a counter to count a number of clock cycles between receipt of the first signal and detection of the preamble of the second signal.

19. The buffer circuit of claim 12, further comprising:

a clock regeneration circuit to receive a clock signal from the module controller and to regenerate a regenerated clock signal according to the delay signal indicating the time interval.

20. The buffer circuit of claim 12, wherein the buffer circuit is coupled to a respective set of memory devices on the memory module, and wherein the buffer circuit further comprises a clock regeneration circuit to receive a clock signal from the module controller and to provide a regenerated clock signal to the respective memory devices.

Micron Technology Inc. et al.
Ex. 1010, p. 39

MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

ABSTRACT OF THE DISCLOSURE

A memory module is operatable in a memory system with a memory controller. The memory module comprises a module control device to receive command signals from the memory controller and to output module command signals and module control signals. The module command signals are provided to memory devices organized in groups, each group including at least one memory device, while the module control signals are provided to a plurality of buffer circuits to control data paths in the buffer circuits. The plurality of buffer circuits are associated with respective groups of memory devices and are distributed across a surface of the memory module such that each module control signal arrives at the plurality of buffer circuits at different points in time. The plurality of buffer circuits are configured to align read data signals received from the memory devices such that the read data signals are transmitted to the memory controller from the memory module substantially aligned with each other and in accordance with a read latency parameter of the memory system.

Micron Technology Inc. et al.
Ex. 1010, p. 40

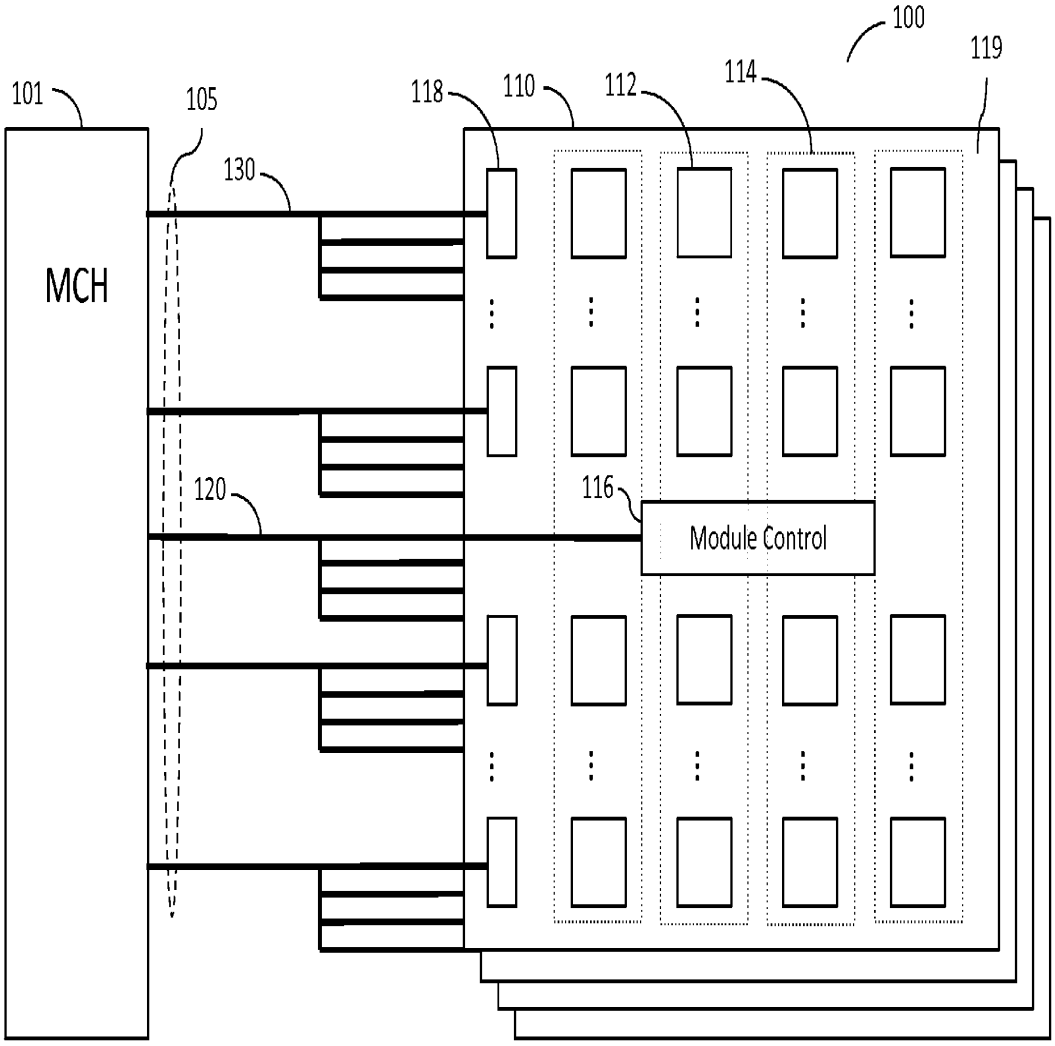


FIG. 1

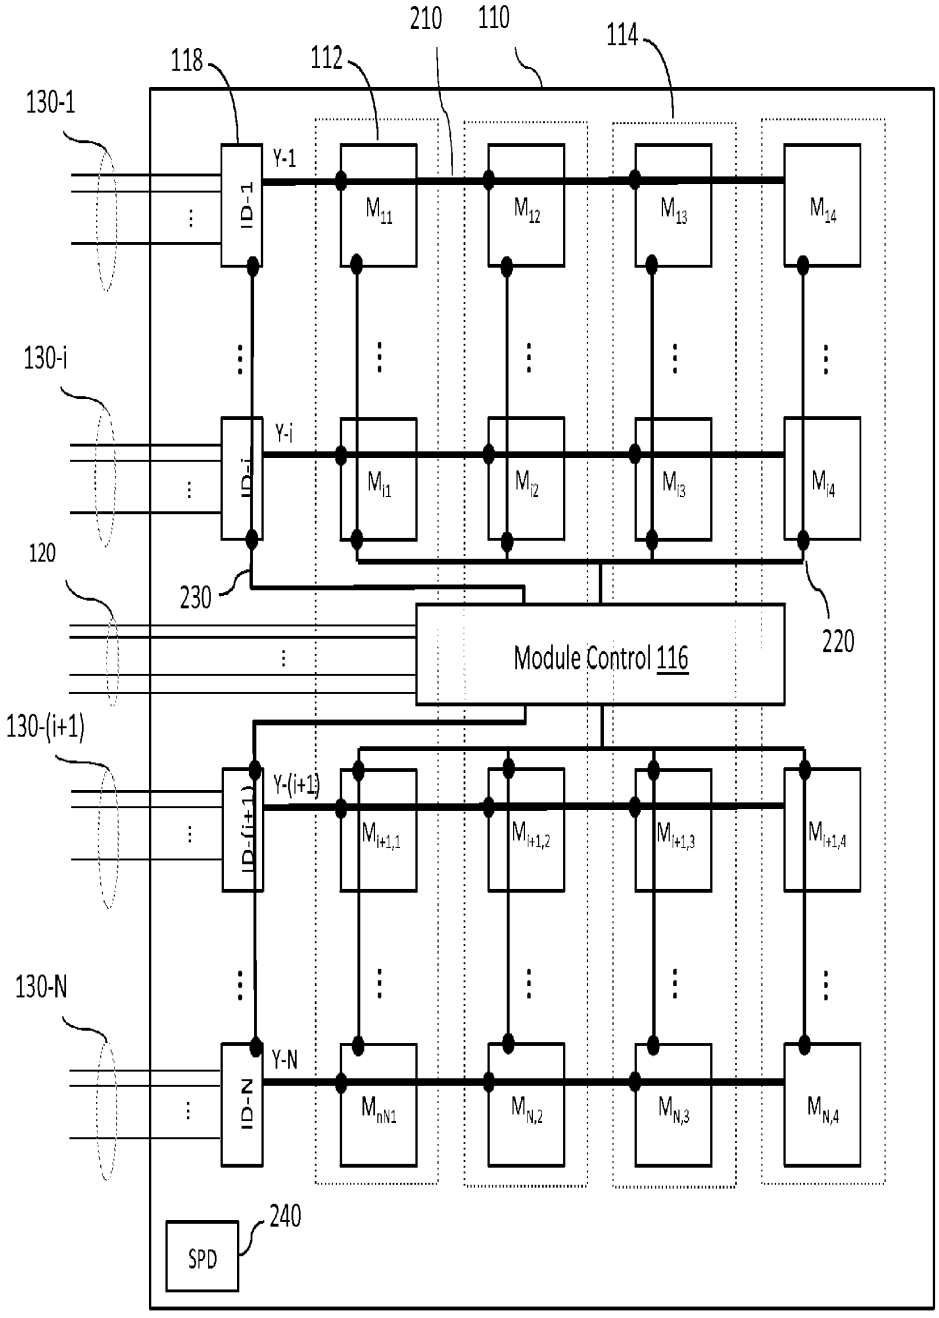

210
110
114
118
112
130-1
Y-1
ID-1
M11
M12
M13
M14
130-i
Y-i
ID-i
Mi1
Mi2
Mi3
Mi4
120
230
220
Module Control 116
130-(i+1)
Y-(i+1)
ID-(i+1)
Mi+1,1
Mi+1,2
Mi+1,3
Mi+1,4
130-N
Y-N
ID-N
MnN1
MN,2
MN,3
MN,4
240
SPD

FIG. 2A

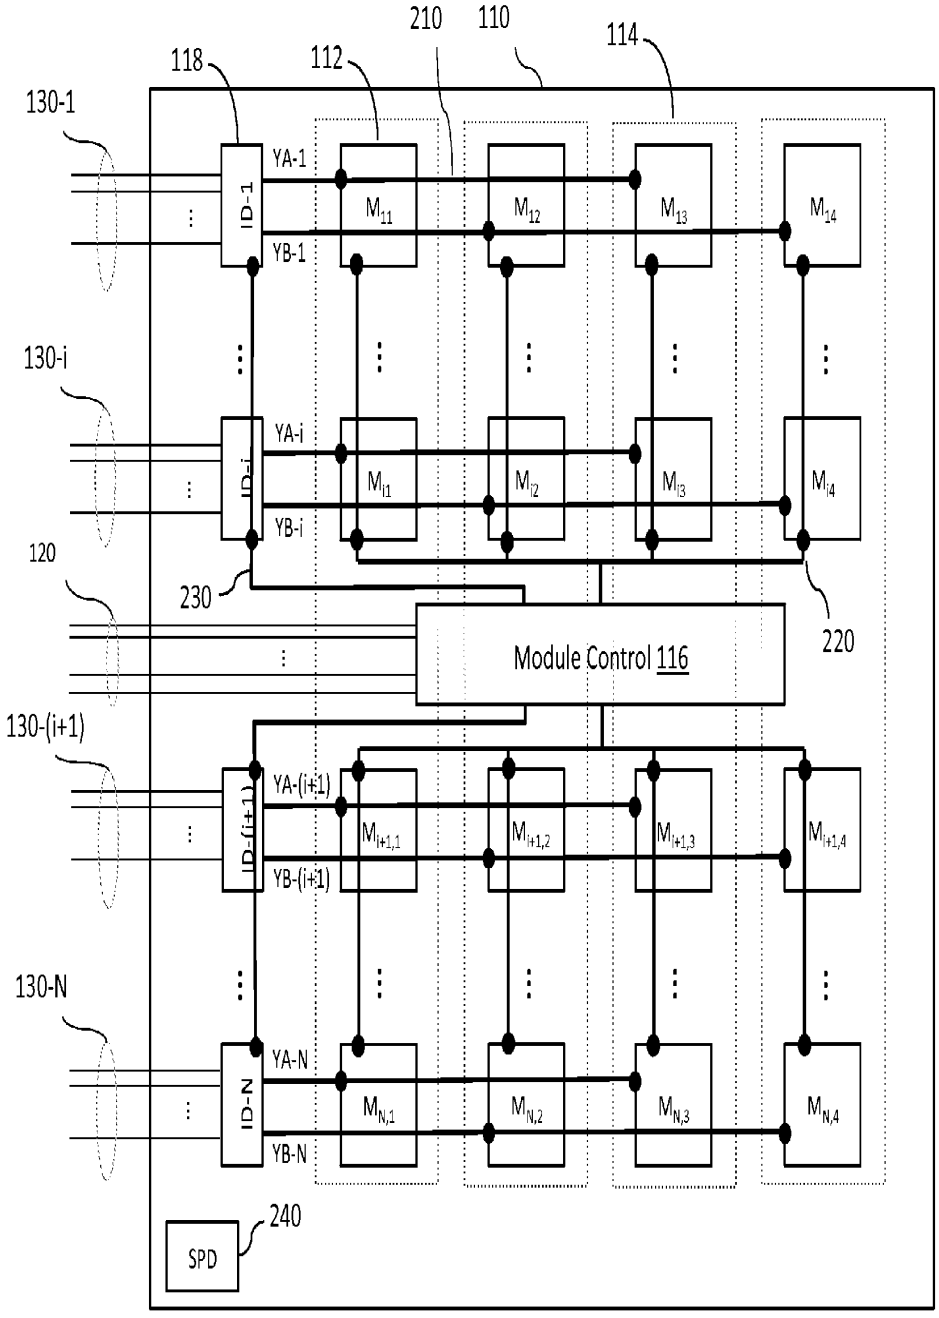

210
110
114
118
112
130-1
YA-1
ID-1
YB-1
M11
M12
M13
M14
130-i
YA-i
ID-i
YB-i
Mi1
Mi2
Mi3
Mi4
120
230
220
Module Control 116
130-(i+1)
YA-(i+1)
ID-(i+1)
YB-(i+1)
Mi+1,1
Mi+1,2
Mi+1,3
Mi+1,4
130-N
YA-N
ID-N
YB-N
MN,1
MN,2
MN,3
MN,4
240
SPD

FIG. 2B

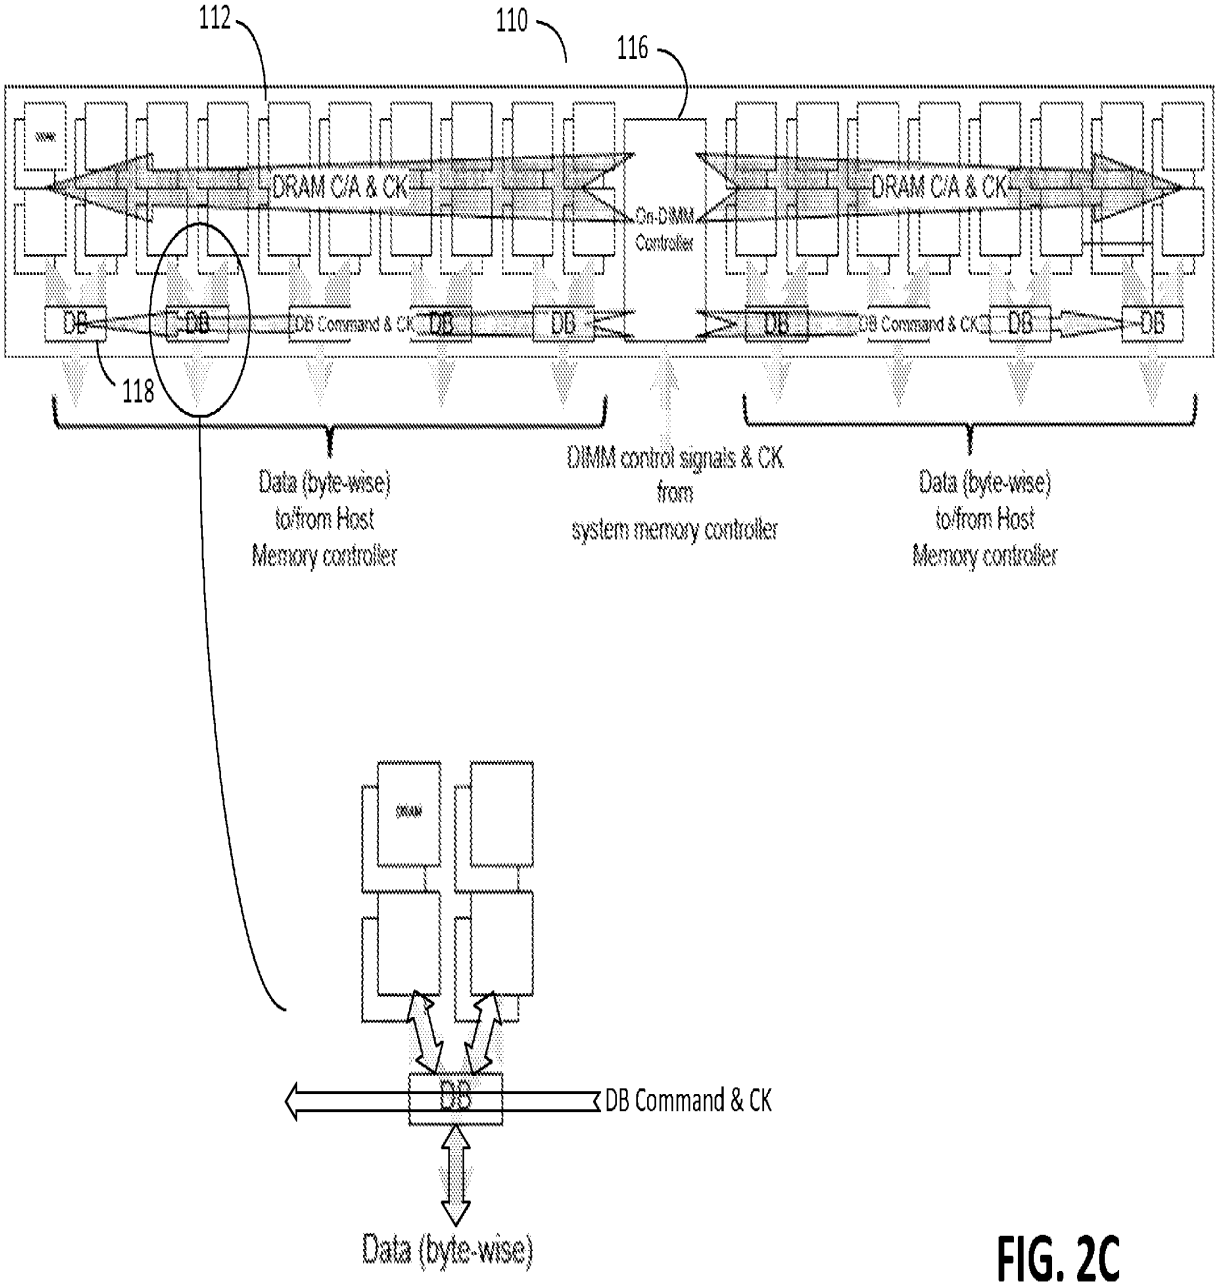

112
110
116
DRAM C/A & CK
On-DIMM Controller
DRAM C/A & CK
DB
DB Command & CK
118
Data (byte-wise) to/from Host Memory controller
DIMM control signals & CK from system memory controller
Data (byte-wise) to/from Host Memory controller
DB
DB Command & CK
Data (byte-wise)

FIG. 2C

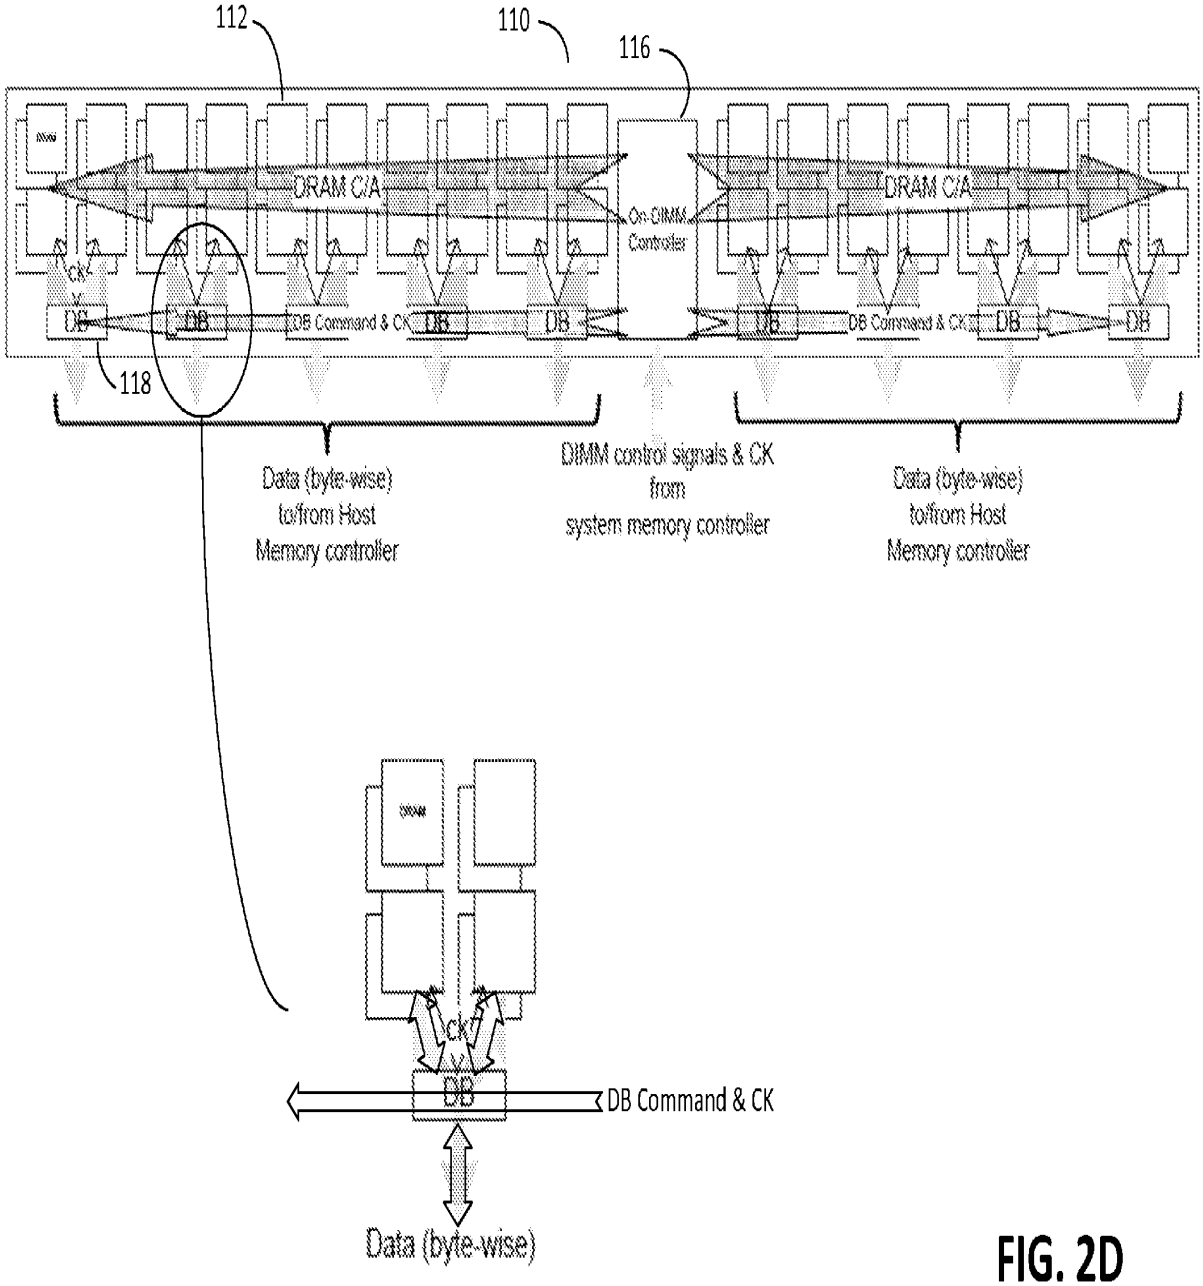

112
110
116
DRAM C/A
On-DIMM Controller
DRAM C/A
CK
DB
DB Command & CK
DB
DB
DB
DB Command & CK
DB
DB
118
Data (byte-wise) to/from Host Memory controller
DIMM control signals & CK from system memory controller
Data (byte-wise) to/from Host Memory controller
CK
DB
DB Command & CK
Data (byte-wise)

FIG. 2D

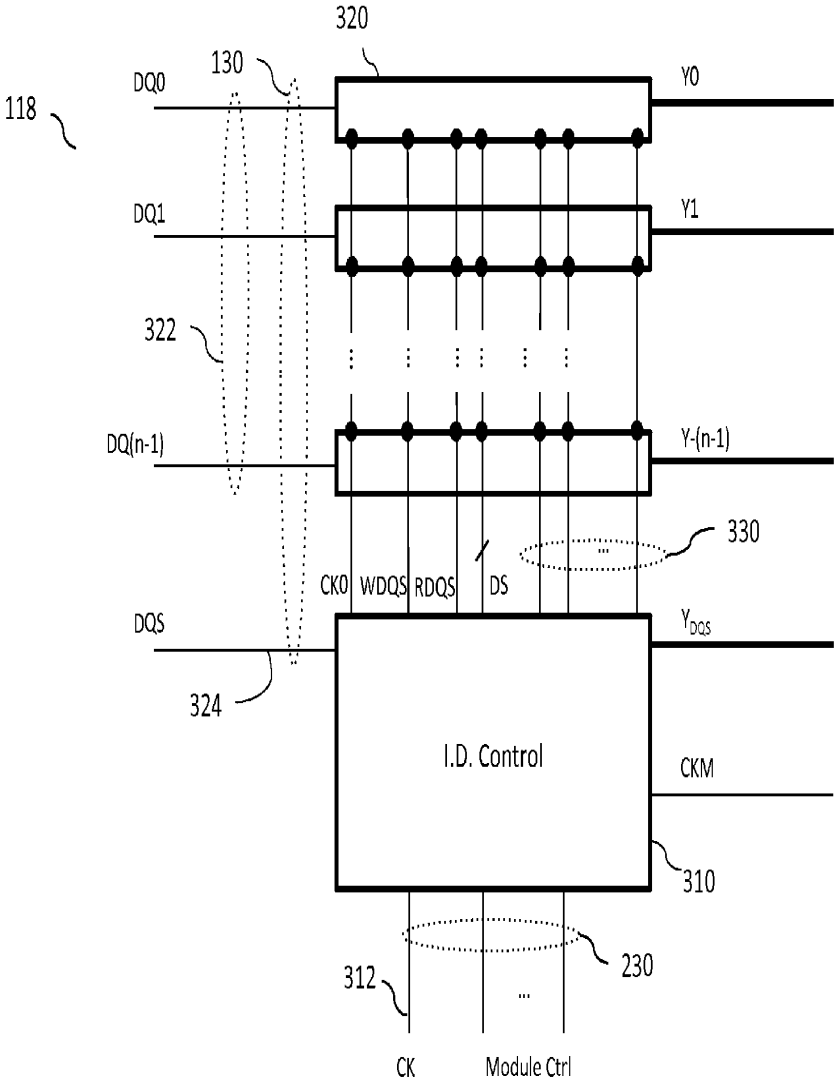

320
130
118
DQ0
Y0
DQ1
Y1
322
DQ(n-1)
Y-(n-1)
330
CK0
WDQS
RDQS
DS
DQS
Y_{DQS}
324
I.D. Control
CKM
310
230
312
CK
Module Ctrl

FIG. 3

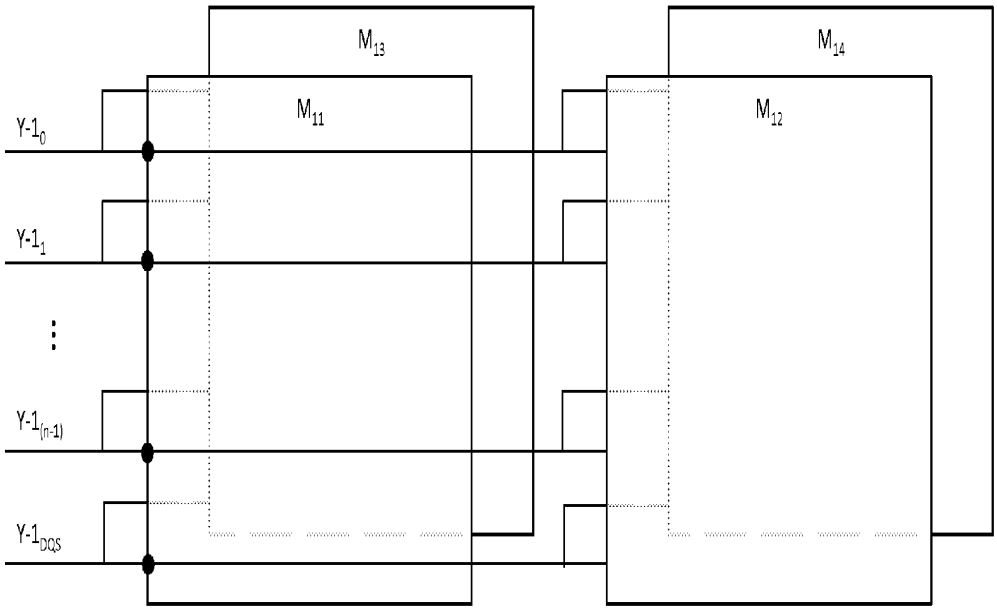


FIG. 4A

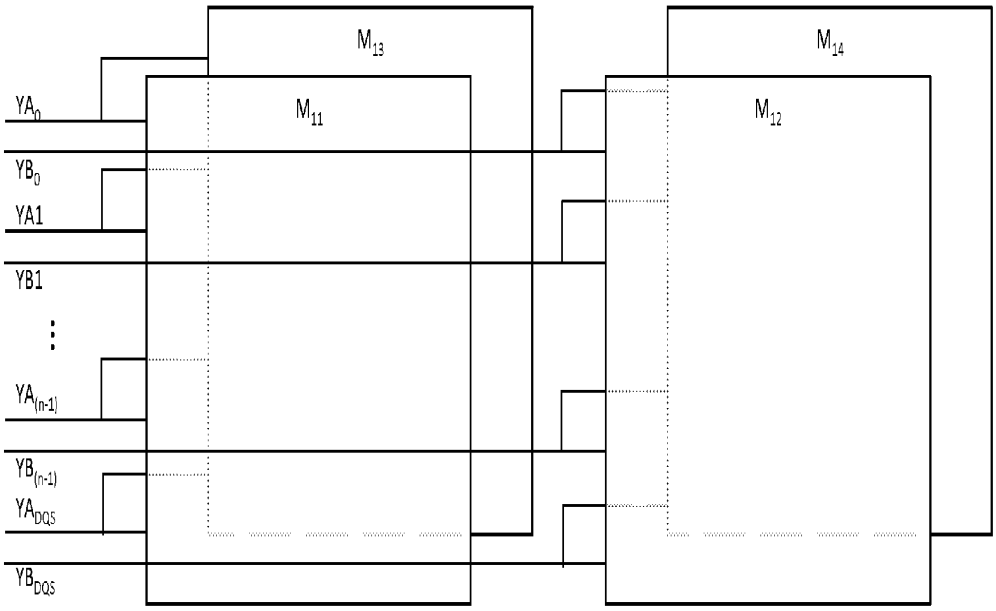


FIG. 4B

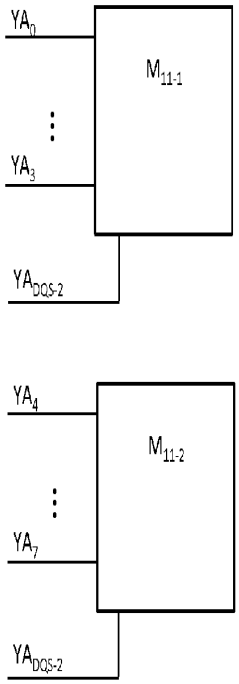

YA0
YA3
YADQS-2
M11-1
YA4
YA7
YADQS-2
M11-2

FIG. 5A

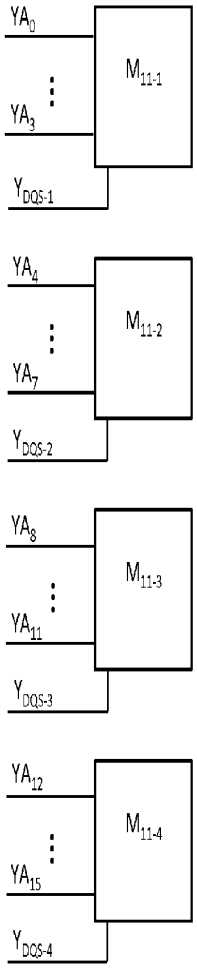

YA0
YA3
YDQS-1
M11-1
YA4
YA7
YDQS-2
M11-2
YA8
YA11
YDQS-3
M11-3
YA12
YA15
YDQS-4
M11-4

FIG. 5B

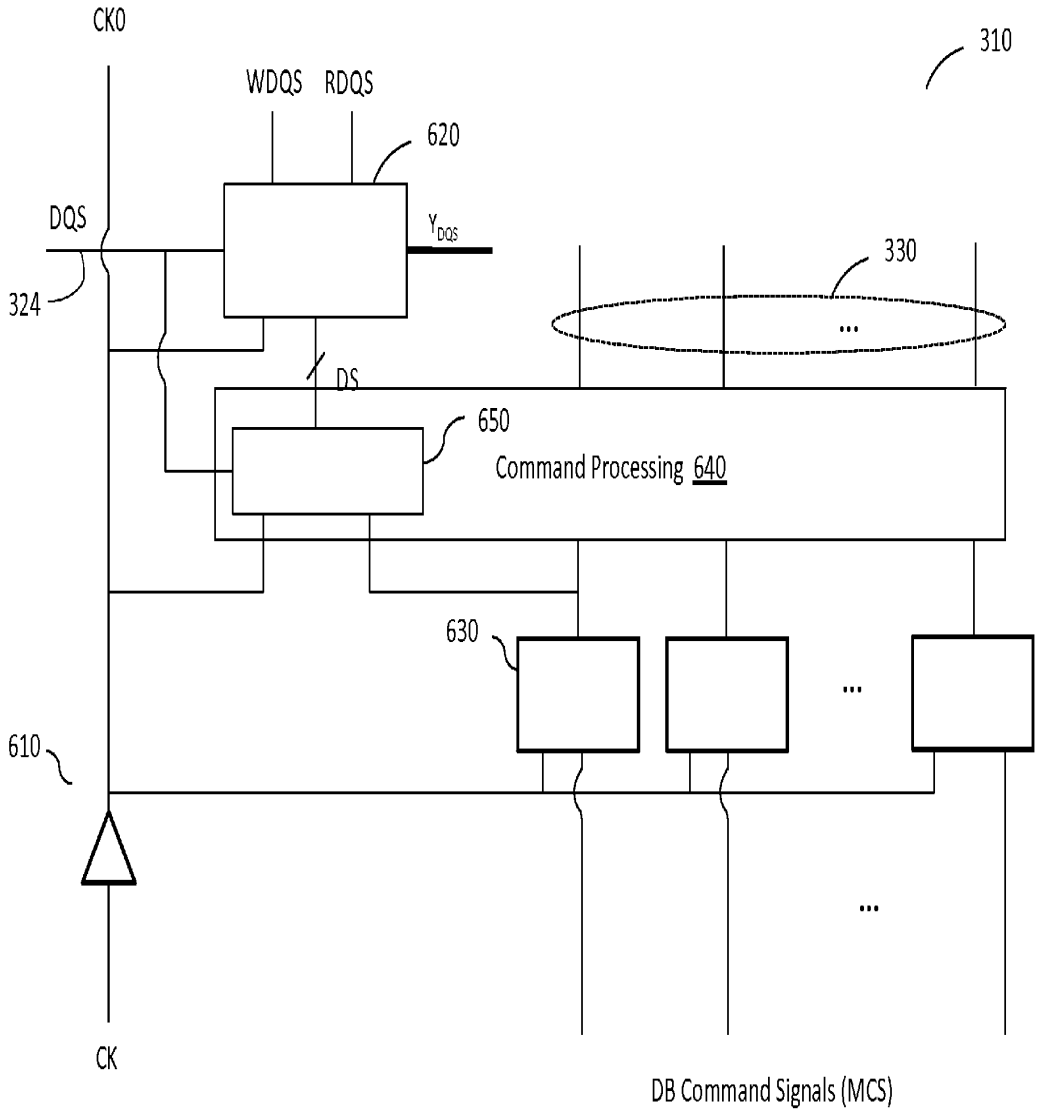

CK0
310
WDQS
RDQS
620
DQS
Y_{DQS}
330
324
DS
650
Command Processing 640
630
610
CK
DB Command Signals (MCS)

FIG. 6

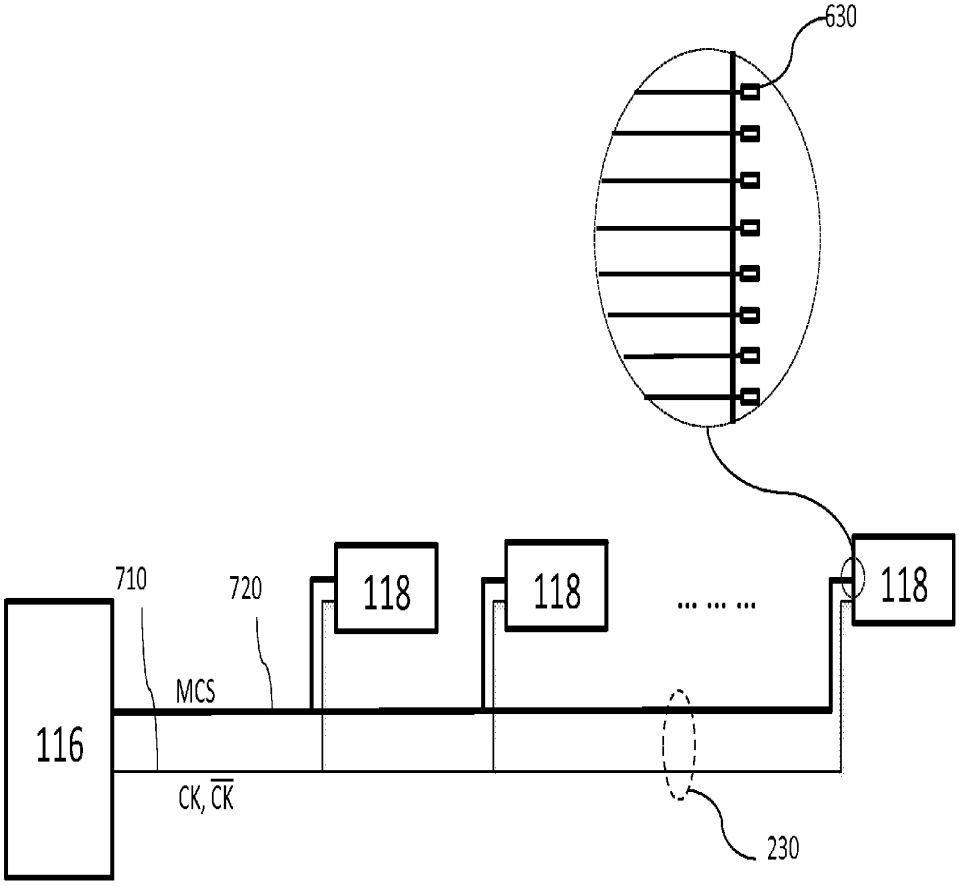


FIG. 7

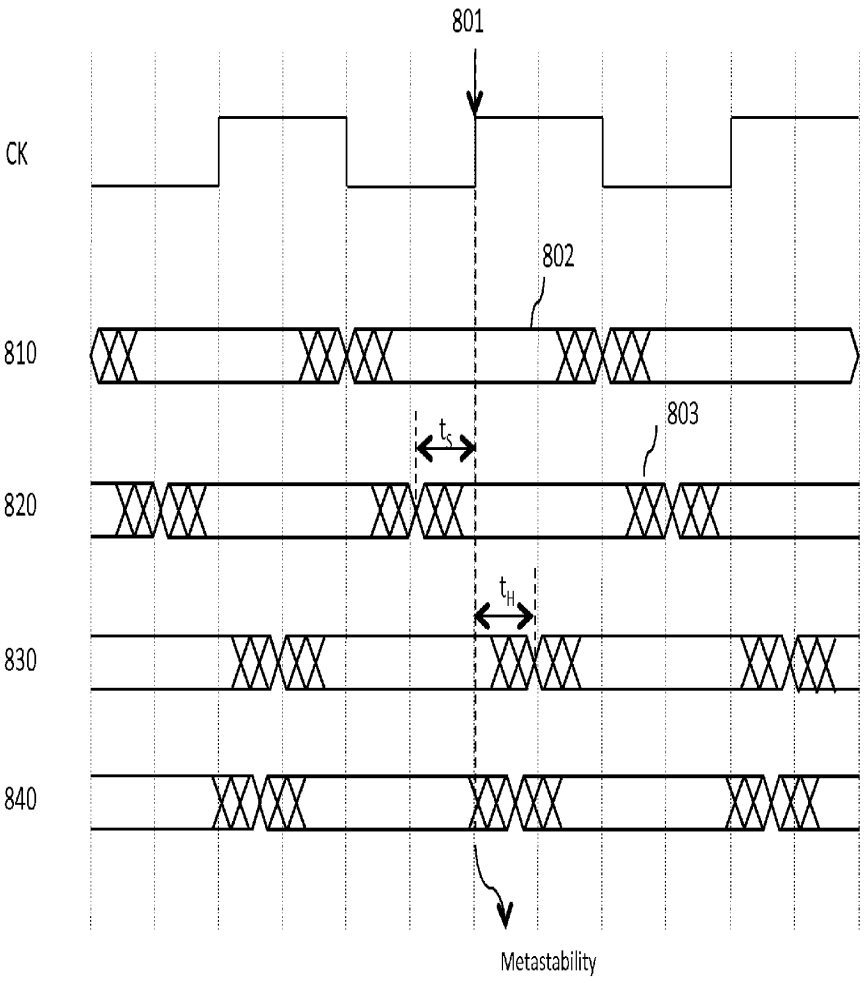


FIG. 8

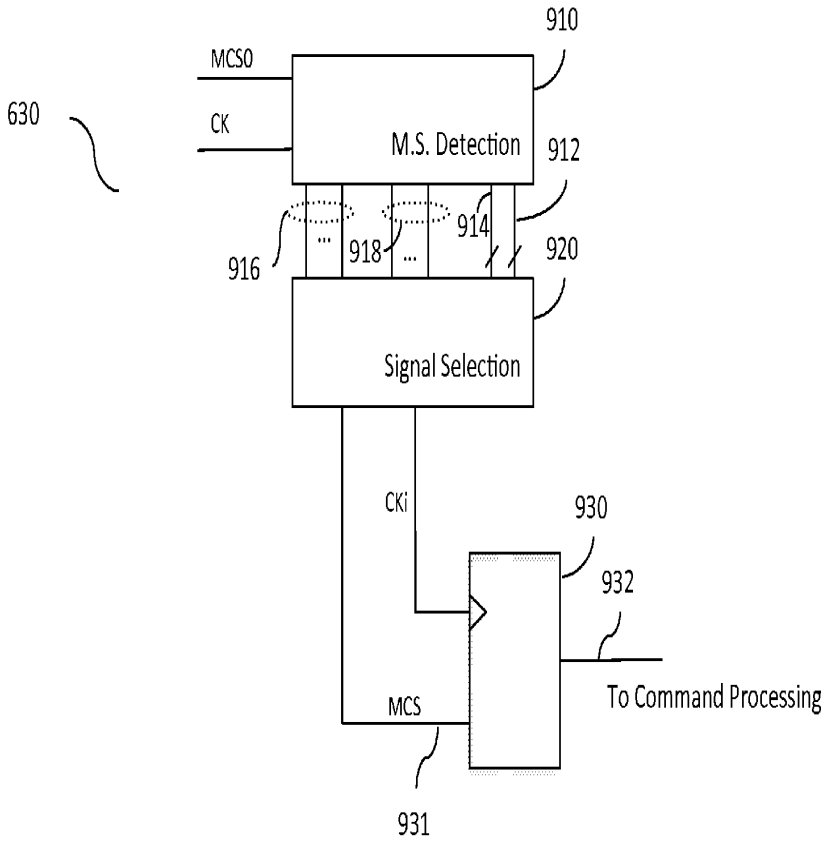

630
MCS0
CK
910
M.S. Detection
912
914
916
918
920
Signal Selection
CKi
930
932
To Command Processing
MCS
931

FIG. 9

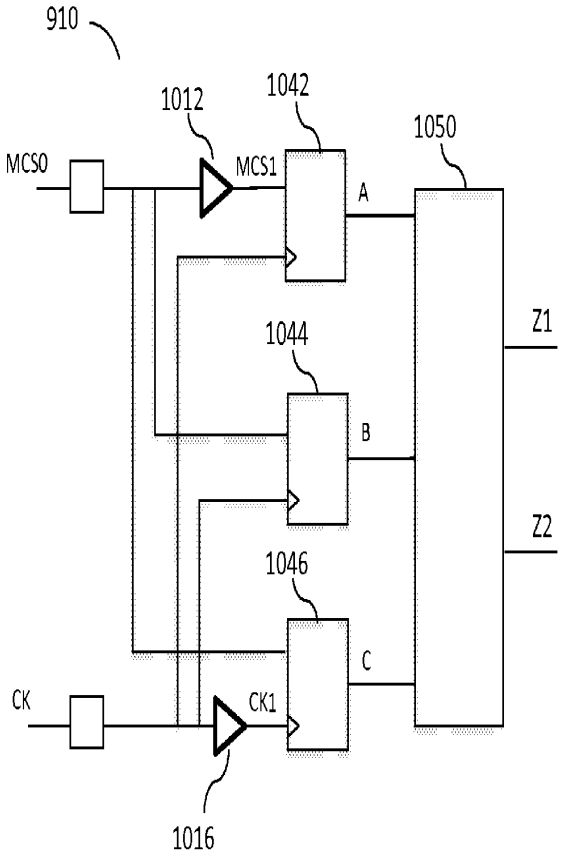


FIG. 10A

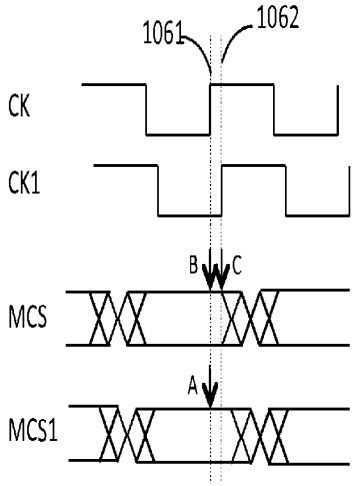


FIG. 10B

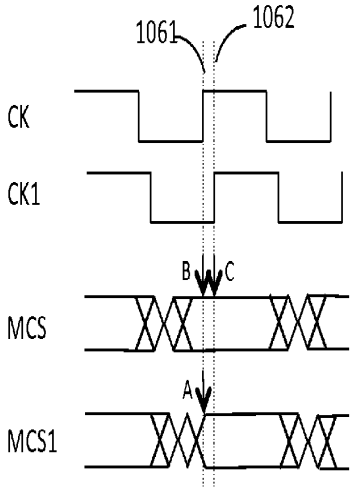


FIG. 10C

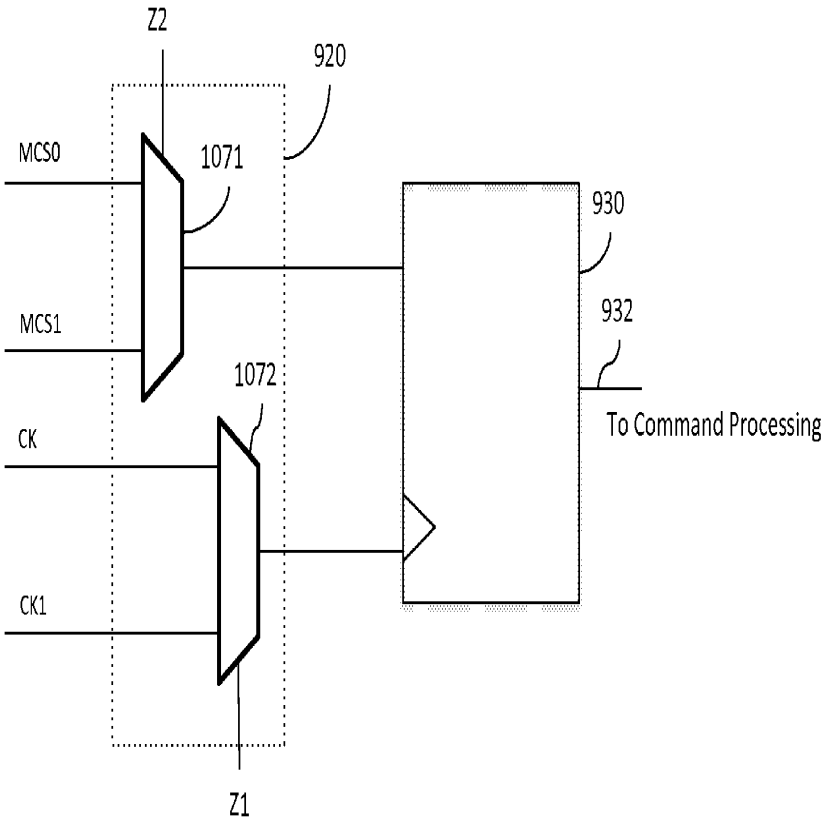

Z2
920
MCS0
1071
930
MCS1
932
1072
To Command Processing
CK
CK1
Z1

FIG. 10D

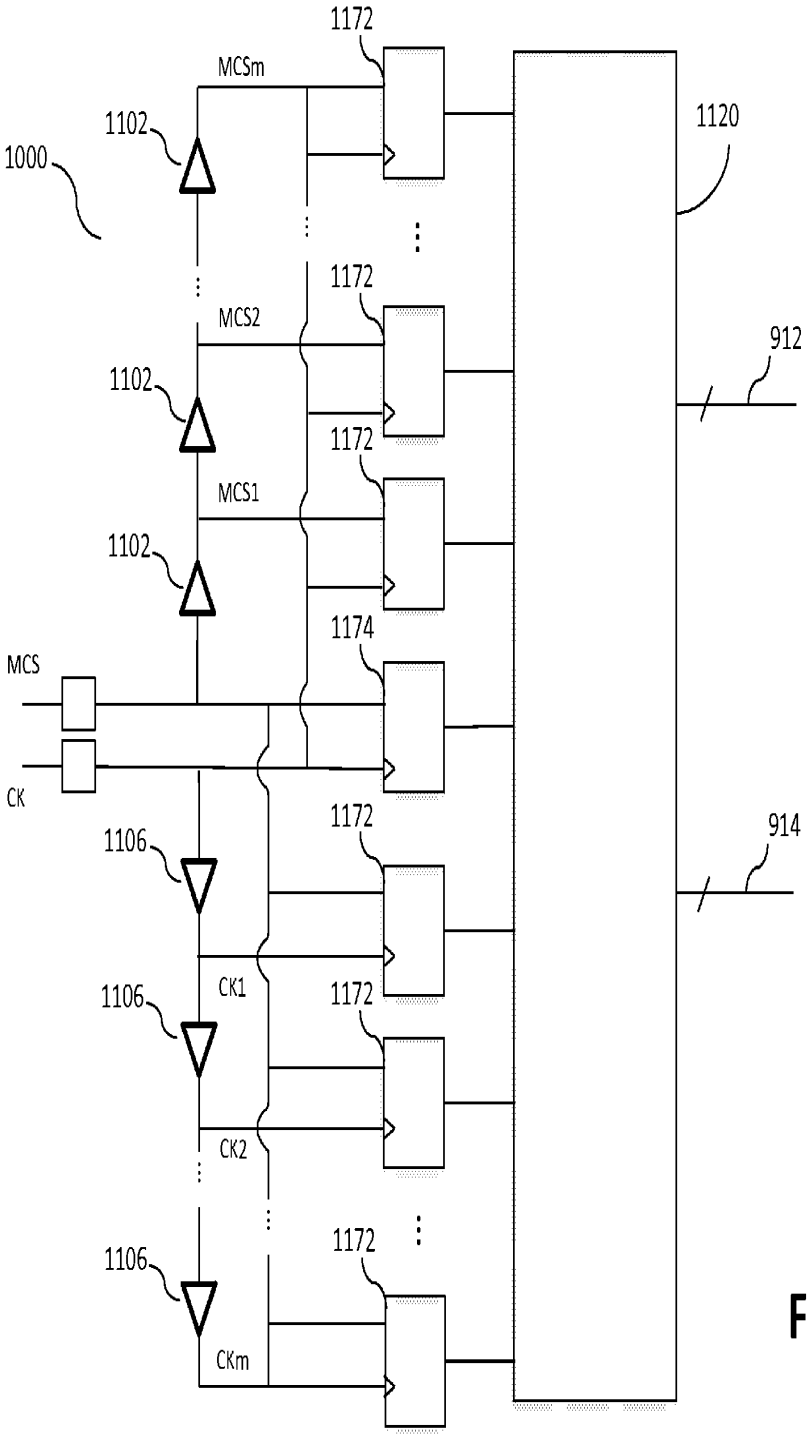

1172
MCSm
1102
1000
1120
1172
MCS2
912
1102
1172
MCS1
1102
1174
MCS
CK
1172
914
1106
CK1
1106
1172
CK2
1172
1106
CKm

FIG. 11A

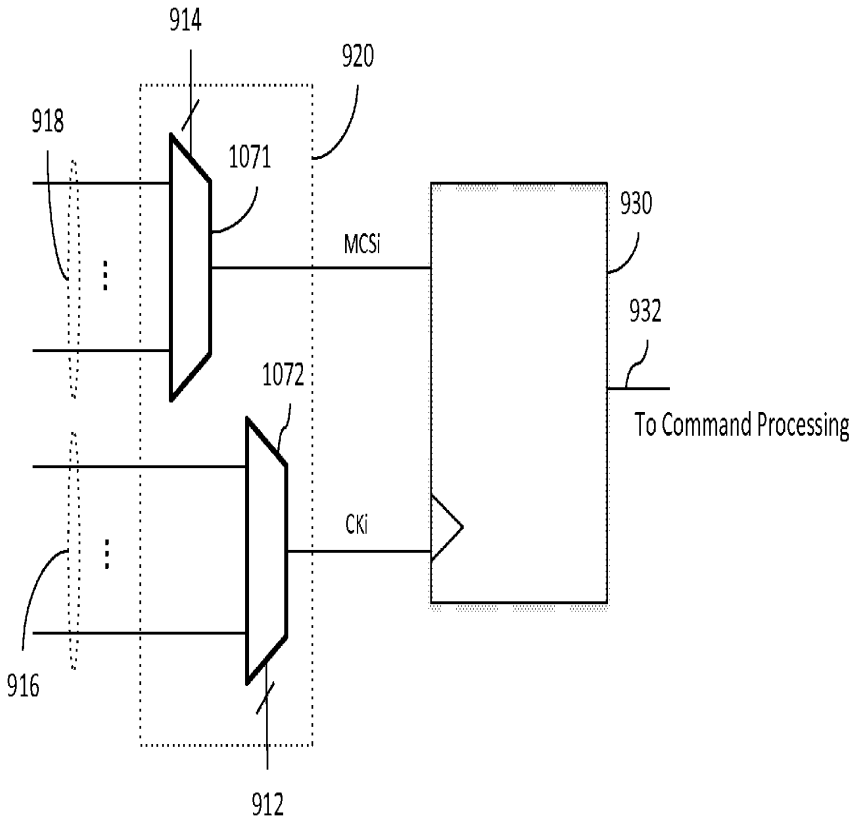

914
920
918
1071
930
MCSi
932
1072
To Command Processing
CKi
916
912

FIG. 11B

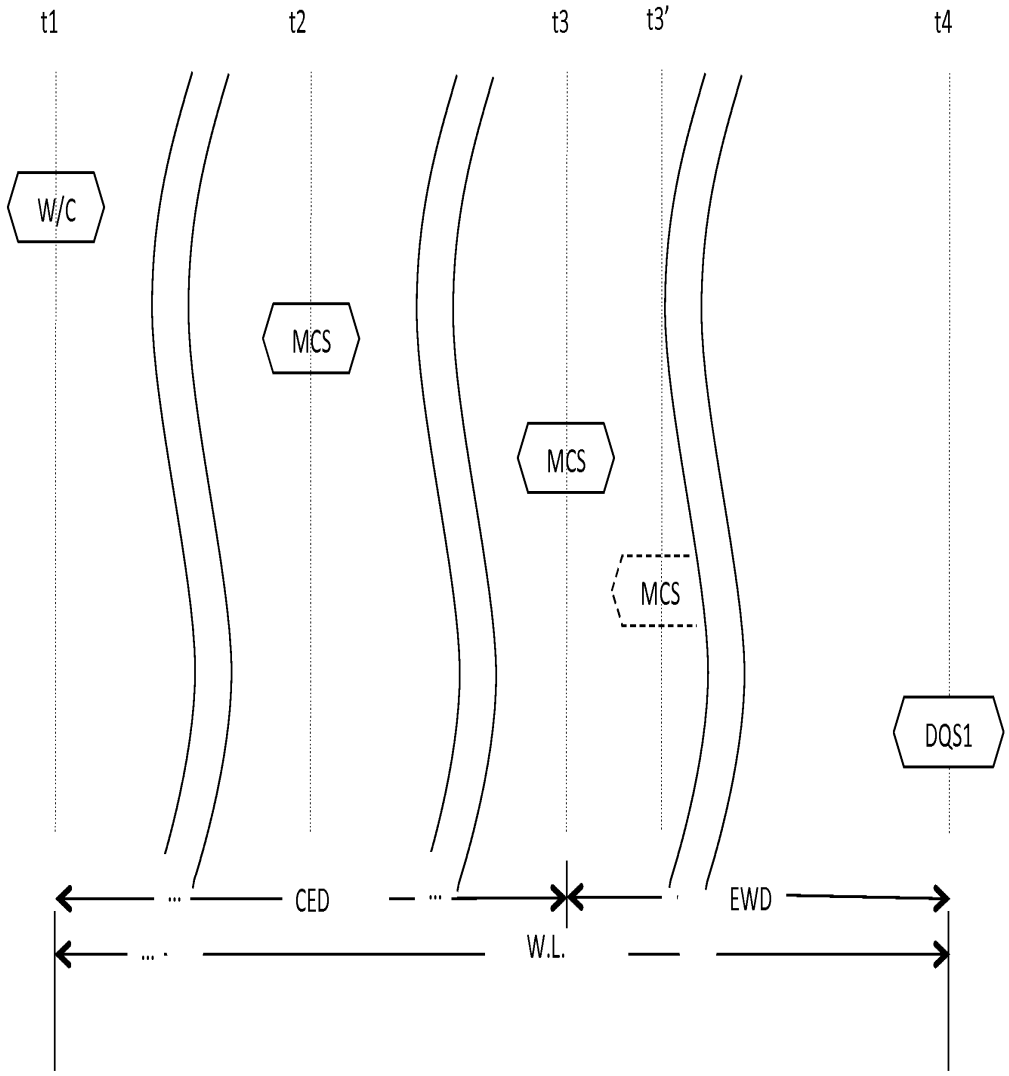


FIG. 12A

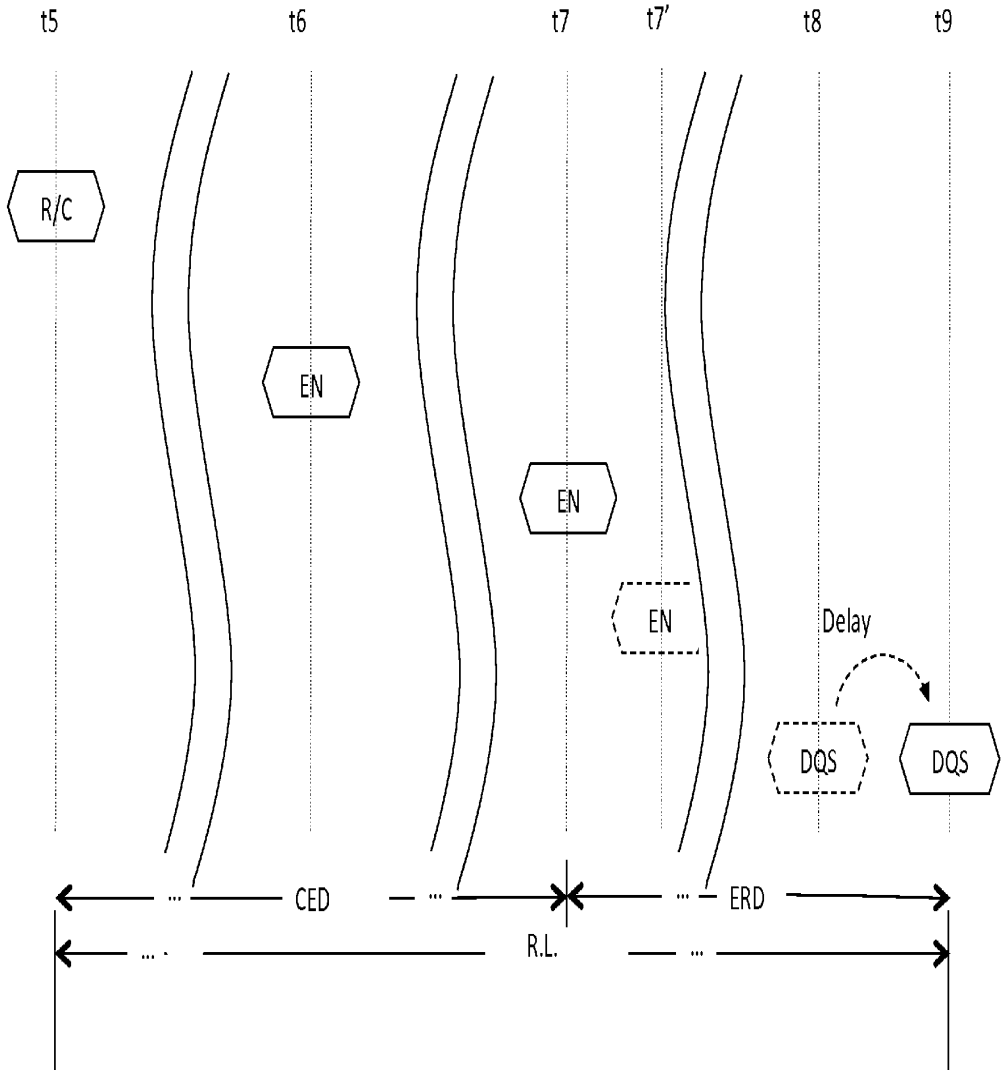


FIG. 12B

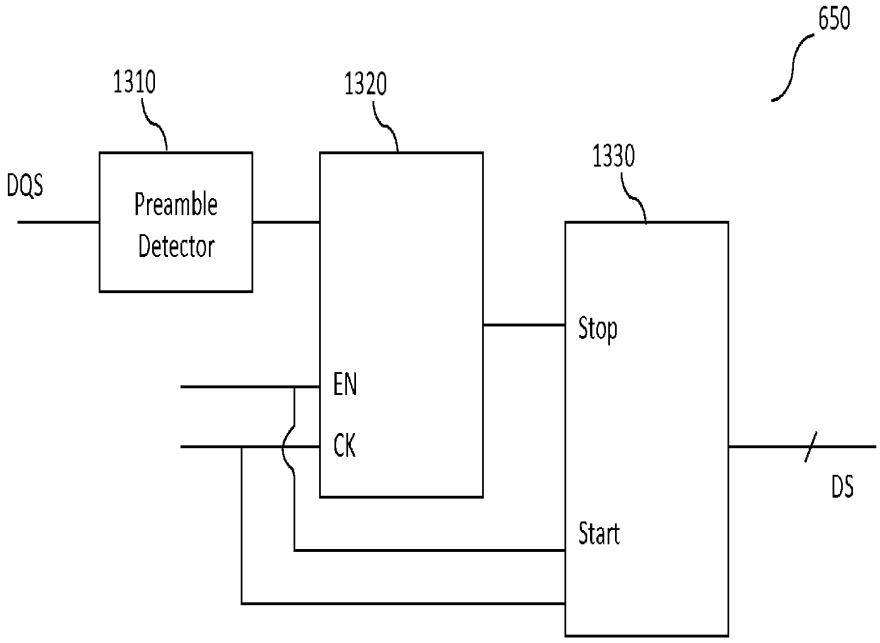


FIG. 13

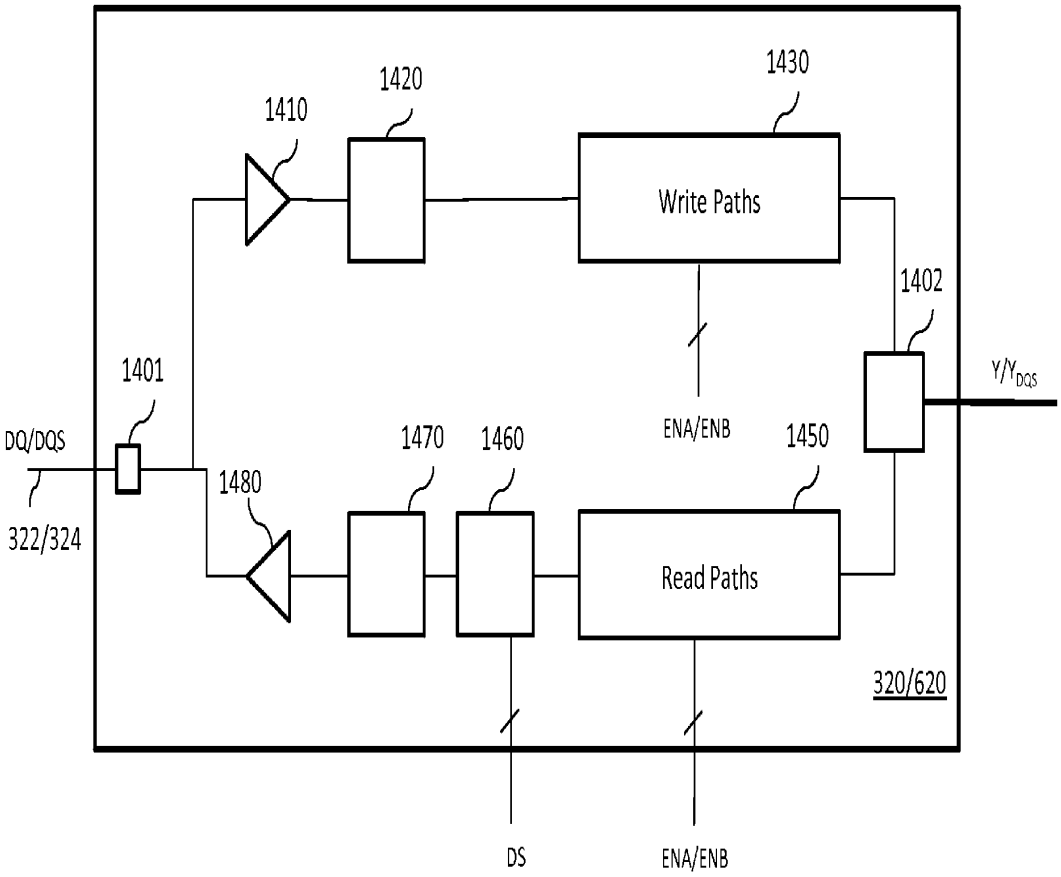

1410
1420
1430
Write Paths
1402
1401
DQ/DQS
322/324
1480
1470
1460
1450
ENA/ENB
Read Paths
320/620
DS
ENA/ENB

FIG. 14

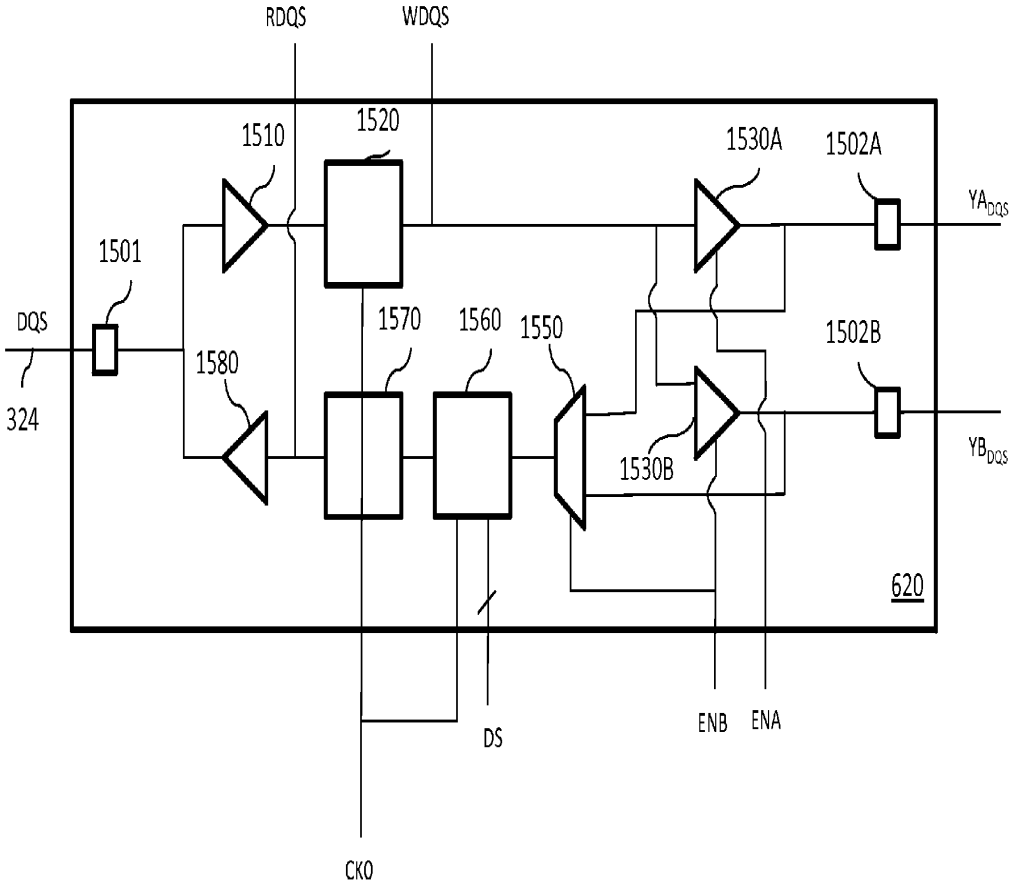

RDQS
WDQS
1510
1520
1530A
1502A
YA_DQS
1501
DQS
324
1580
1570
1560
1550
1502B
YB_DQS
1530B
620
DS
ENB
ENA
CKO

FIG. 15

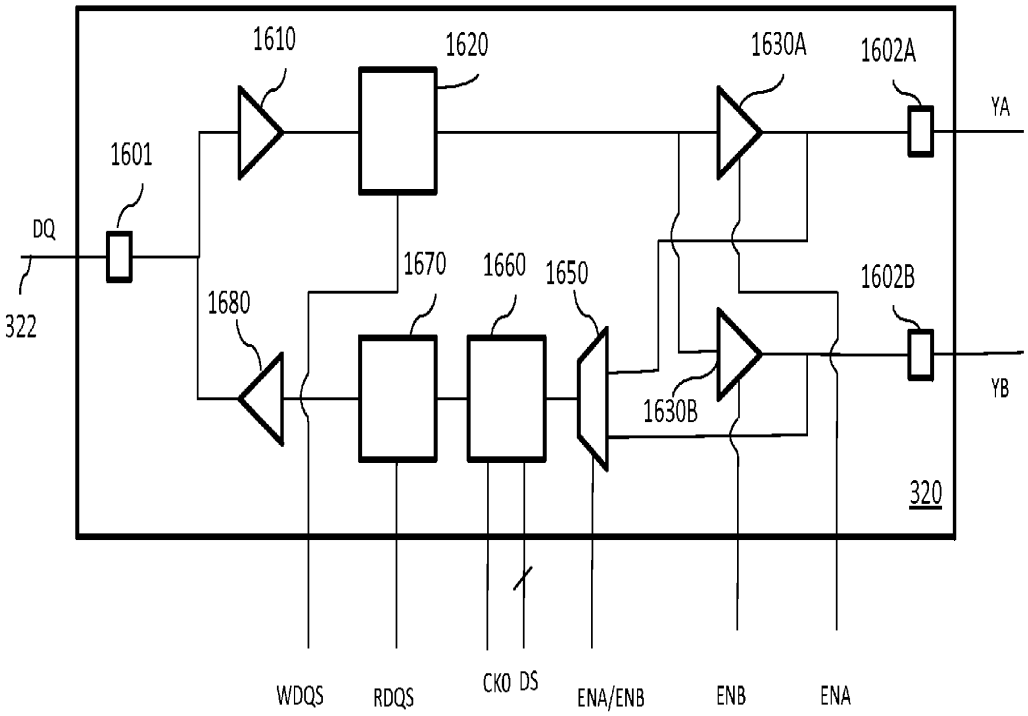

1610
1620
1630A
1602A
YA
1601
DQ
322
1670
1660
1650
1602B
1680
1630B
YB
320
WDQS
RDQS
CK0
DS
ENA/ENB
ENB
ENA

FIG. 16

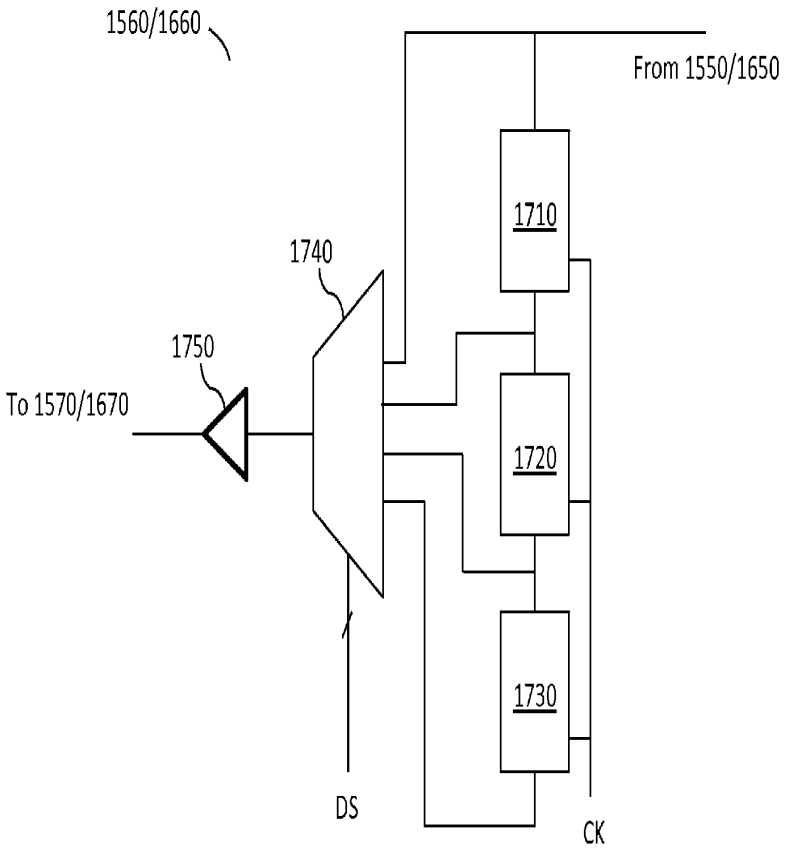

1560/1660
From 1550/1650
1710
1740
1750
To 1570/1670
1720
1730
DS
CK

FIG. 17

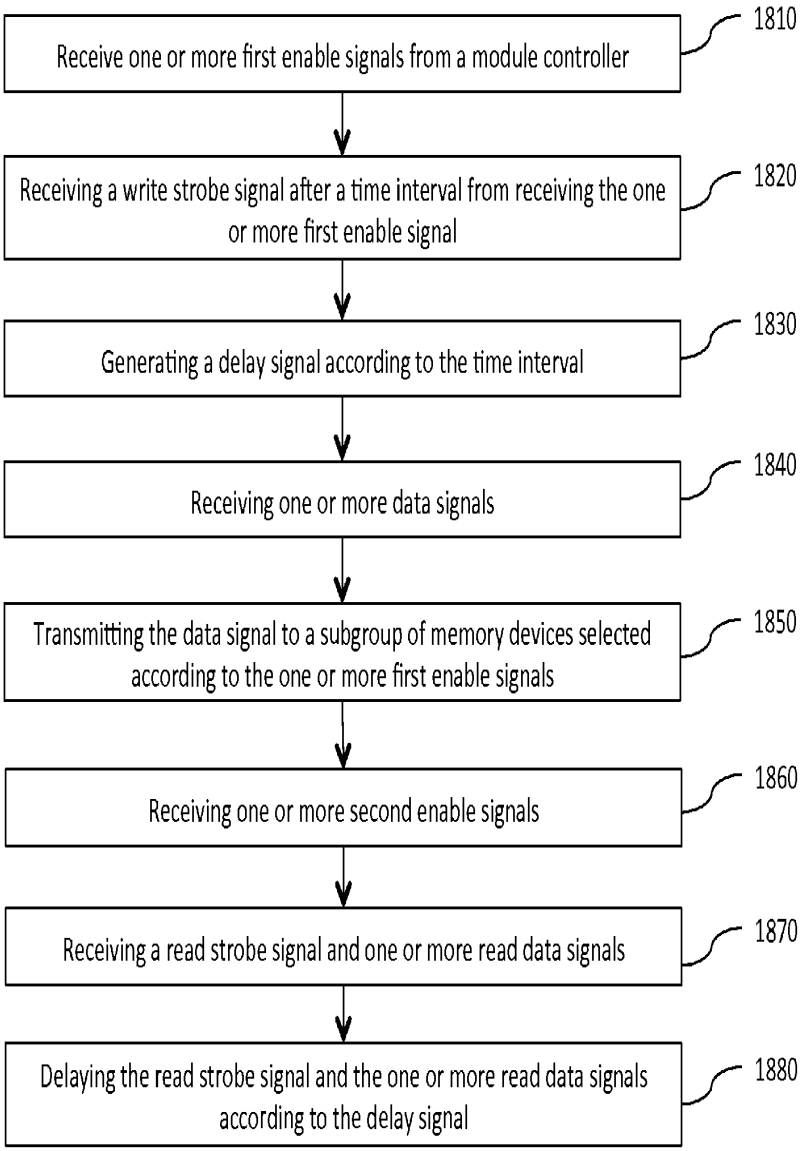

Receive one or more first enable signals from a module controller
1810
Receiving a write strobe signal after a time interval from receiving the one or more first enable signal
1820
Generating a delay signal according to the time interval
1830
Receiving one or more data signals
1840
Transmitting the data signal to a subgroup of memory devices selected according to the one or more first enable signals
1850
Receiving one or more second enable signals
1860
Receiving a read strobe signal and one or more read data signals
1870
Delaying the read strobe signal and the one or more read data signals according to the delay signal
1880

FIG. 18

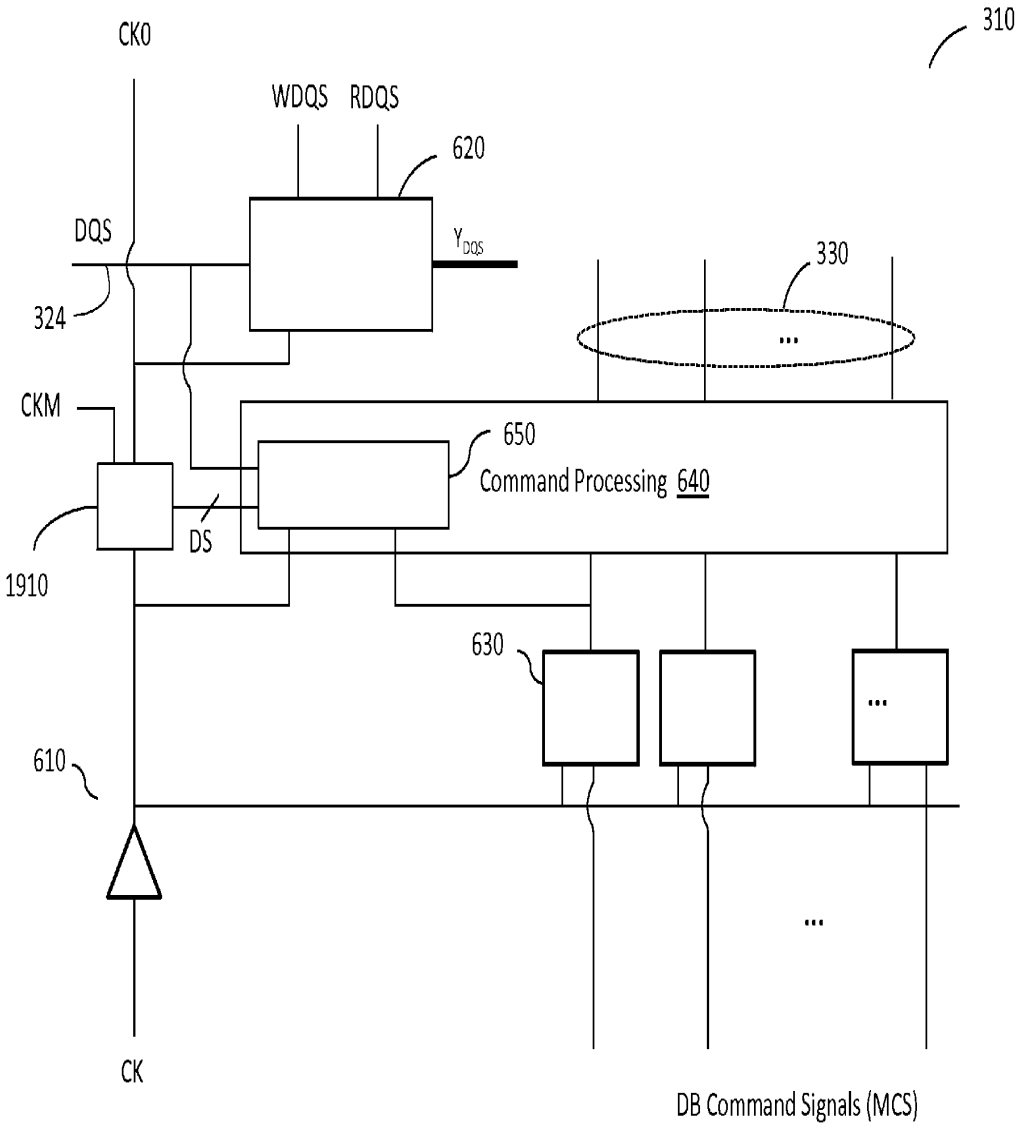

CK0
310
WDQS
RDQS
620
DQS
Y_{DQS}
330
324
CKM
650
Command Processing 640
DS
1910
630
610
CK
DB Command Signals (MCS)

FIG. 19

Electronic Patent Application Fee Transmittal	
Application Number:	
Filing Date:	
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Filer:	Jamie Jie Zheng
Attorney Docket Number:	NT003US
Filed as Small Entity	

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Utility filing Fee (Electronic filing)	4011	1	70	70
Utility Search Fee	2111	1	300	300
Utility Examination Fee	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		**730**

Electronic Acknowledgement Receipt	
EFS ID:	16434407
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	27-JUL-2013
Filing Date:	
Time Stamp:	17:22:05
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Application Data Sheet	2013-07-27_ApplicationDataSheet_NT003US.pdf	1437143 b57cd8746db0d11120620b555565845a6e8e4bff	no	6
Warnings:					
Information:					

2		2013-07-27_Specification_NT003US.pdf	354802 1f72aaca03558fcb9d8b3371938fd7904a5c93d21	yes	34

	Multipart Description/PDF files in .zip description		
	Document Description	**Start**	**End**
	Specification	1	29
	Claims	30	33
	Abstract	34	34

Warnings:

Information:

3	Drawings-only black and white line drawings	2013-07-27_Drawings_NT003US.pdf	339521 928257330c13e3f746dba792c8501ed691ecd833	no	26

Warnings:

The page size in the PDF is too large. The pages should be 8.5 x 11 or A4. If this PDF is submitted, the pages will be resized upon entry into the Image File Wrapper and may affect subsequent processing

Information:

4	Fee Worksheet (SB06)	fee-info.pdf	32774 953205ab8223e56a615a5daf21baf5affd487d9d	no	2

Warnings:

Information:

Total Files Size (in bytes):	2164240

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

<u>New Applications Under 35 U.S.C. 111</u>
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

<u>National Stage of an International Application under 35 U.S.C. 371</u>
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

<u>New International Application Filed with the USPTO as a Receiving Office</u>
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

DocCode - SCORE

SCORE Placeholder Sheet for IFW Content

Application Number: 13952599 Document Date: 07/27/2013

The presence of this form in the IFW record indicates that the following document type was received in electronic format on the date identified above. This content is stored in the SCORE database.

- Drawings – Other than Black and White Line Drawings

Since this was an electronic submission, there is no physical artifact folder, no artifact folder is recorded in PALM, and no paper documents or physical media exist. The TIFF images in the IFW record were created from the original documents that are stored in SCORE.

To access the documents in the SCORE database, refer to instructions developed by SIRA.

At the time of document entry (noted above):

- Examiners may access SCORE content via the eDAN interface.
- Other USPTO employees can bookmark the current SCORE URL (http://es/ScoreAccessWeb/).
- External customers may access SCORE content via the Public and Private PAIR interfaces.

Form Revision Date: February 8, 2006

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 13/952,599

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	70		N/A	
SEARCH FEE (37 CFR 1.16(k), (l), or (m))	N/A	N/A	N/A	300		N/A	
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	360		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	20 minus 20=	*	x 40 =	0.00	OR		
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*	x 210 =	0.00			
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			0.00			
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				0.00			
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	730		TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE(S)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE(S)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/952,599	07/27/2013	Hyun Lee	NT003US

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

FORMALITIES LETTER

OC000000063191960

Date Mailed: 08/19/2013

NOTICE TO FILE MISSING PARTS OF NONPROVISIONAL APPLICATION

FILED UNDER 37 CFR 1.53(b)

Filing Date Granted

Items Required To Avoid Abandonment:

An application number and filing date have been accorded to this application. The item(s) indicated below, however, are missing. Applicant is given **TWO MONTHS** from the date of this Notice within which to file all required items below to avoid abandonment. Extensions of time may be obtained by filing a petition accompanied by the extension fee under the provisions of 37 CFR 1.136(a).

- The statutory basic filing fee is missing.
  *Applicant must submit $**70** to complete the basic filing fee for a small entity.*

The applicant needs to satisfy supplemental fees problems indicated below.

The required item(s) identified below must be timely submitted to avoid abandonment:

- A surcharge (for late submission of the basic filing fee, search fee, examination fee or inventor's oath or declaration) as set forth in 37 CFR 1.16(f) of $ **70** for a small entity in compliance with 37 CFR 1.27, must be submitted.

SUMMARY OF FEES DUE:

Total fee(s) required within **TWO MONTHS** from the date of this Notice is $ **800** for a small entity
- $ **70** Statutory basic filing fee.
- $ **70** Surcharge.
- The application search fee has not been paid. Applicant must submit $ **300** to complete the search fee.
- The application examination fee has not been paid. Applicant must submit $ **360** to complete the examination fee for a small entity in compliance with 37 CFR 1.27.

Items Required To Avoid Processing Delays:

Applicant is notified that the above-identified application contains the deficiencies noted below. No period for reply is set forth in this notice for correction of these deficiencies. However, if a deficiency relates to the inventor's oath or declaration, the applicant must file an oath or declaration in compliance with 37 CFR 1.63, or a substitute statement in compliance with 37 CFR 1.64, executed by or with respect to each actual inventor no later than the expiration of the time period set in the "Notice of Allowability" to avoid abandonment. See 37 CFR 1.53(f).

- A properly executed inventor's oath or declaration has not been received for the following inventor(s):

Hyun Lee
Jayesh R. Bhakta

Applicant may submit the inventor's oath or declaration at any time before the Notice of Allowance and Fee(s) Due, PTOL-85, is mailed.

Replies must be received in the USPTO within the set time period or must include a proper Certificate of Mailing or Transmission under 37 CFR 1.8 with a mailing or transmission date within the set time period. For more information and a suggested format, see Form PTO/SB/92 and MPEP 512.

Replies should be mailed to:

Mail Stop Missing Parts
Commissioner for Patents
P.O. Box 1450
Alexandria VA 22313-1450

Registered users of EFS-Web may alternatively submit their reply to this notice via EFS-Web.
https://sportal.uspto.gov/authenticate/AuthenticateUserLocalEPF.html

For more information about EFS-Web please call the USPTO Electronic Business Center at **1-866-217-9197** or visit our website at http://www.uspto.gov/ebc.

If you are not using EFS-Web to submit your reply, you must include a copy of this notice.

/mkibret/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

Micron Technology Inc. et al.
Ex. 1010, p. 74

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
13/952,599	07/27/2013	2189	0.00	NT003US	20	3

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

FILING RECEIPT

OC000000063191959

Date Mailed: 08/19/2013

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Hyun Lee, Ladera Ranch, CA;
Jayesh R. Bhakta, Cerritos, CA;

Applicant(s)

Netlist, Inc., Irvine, CA

Assignment For Published Patent Application

Netlist, Inc., Irvine, CA

Power of Attorney: None

Domestic Priority data as claimed by applicant

This appln claims benefit of 61/676,883 07/27/2012

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

If Required, Foreign Filing License Granted: 08/12/2013

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 13/952,599**

Projected Publication Date: Request for Non-Publication Acknowledged

Non-Publication Request: Yes

Early Publication Request: No
**** SMALL ENTITY ****

Micron Technology Inc. et al.
Ex. 1010, p. 75

Title

MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

Preliminary Class

711

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

Micron Technology Inc. et al.
Ex. 1010, p. 76

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

Electronic Patent Application Fee Transmittal

Application Number:	13952599
Filing Date:	27-Jul-2013
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Filer:	Jamie Jie Zheng
Attorney Docket Number:	NT003US

Filed as Small Entity

Utility under 35 USC 111(a) Filing Fees

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Utility filing Fee (Electronic filing)	4011	1	70	70
Utility Search Fee	2111	1	300	300
Utility Examination Fee	2311	1	360	360
Pages:				
Claims:				
Miscellaneous-Filing:				
Late Filing Fee for Oath or Declaration	2051	1	70	70
Petition:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				
Miscellaneous:				
		Total in USD ($)		800

Electronic Acknowledgement Receipt	
EFS ID:	17174418
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	19-OCT-2013
Filing Date:	27-JUL-2013
Time Stamp:	01:48:48
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Credit Card
Payment was successfully received in RAM	$800
RAM confirmation Number	8590
Deposit Account	
Authorized User	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)

1	Fee Worksheet (SB06)	fee-info.pdf	36797 066b7b5960a0935134aef7247dffb109023678ad	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	36797		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

Customer No. 79141
Attorney Docket No.: NT003US

COMBINED DECLARATION AND ASSIGNMENT FOR UTILITY OR DESIGN APPLICATION

Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

DECLARATION

As a below named inventor, I hereby declare that:

This declaration is directed to:

☐ The attached application, or

■ United States application or PCT international application number 13/952,599 filed on July 27, 2013.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

Note to Inventor: 37 C.F.R. § 1.63(c) states: "A person may not execute an oath or declaration for an application unless that person has reviewed and understands the contents of the application, including the claims, and is aware of the duty to disclose to the Office all information known to the person to be material to patentability as defined in § 1.56."

ASSIGNMENT

In consideration of the sum of One Dollar ($1.00) and other good and valuable consideration paid, I hereby sell and assign to

Netlist, Inc.
(Name of Assignee)

of **51 Discovery, Suite #150, Irvine, CA 92618**
(Address of Assignee)

(hereinafter designated as the Assignee) my entire right, title and interest for the United States as defined in 35 USC 100, for any invention set forth in the above-identified application.

I agree to execute all papers necessary in connection with this application for patent in the USPTO for the invention, and any continuation, divisional or reissue applications thereof and also to execute separate assignments in connection with such applications as the Assignee may deem necessary or expedient.

I agree to execute all papers necessary in connection with any post-grant proceeding which may occur in connection with this application or continuation, divisional or reissue thereof and to cooperate with the Assignee in every way possible in obtaining evidence and going forward with such post-grant proceeding.

I agree to execute all papers and documents and perform any act which may be necessary in connection with claims or provisions of the International Convention for Protection of Industrial Property or similar agreements.

I agree to perform all affirmative acts which may be necessary to obtain a grant of a valid United States patent to the Assignee.

I hereby authorize and request the USPTO to issue any and all Letters Patents of the United States resulting from the application or any continuation, divisional or reissue applications thereof to the Assignee, as Assignee of my entire interest, and covenant that I have not executed, and will not execute, any agreement in conflict herewith.

First Inventor (Legal Name): Hyun Lee

Signature: _____ Date: 10/9/2013

Second Inventor (Legal Name): Jayesh R. Bhakta

Signature: Jayesh Bhakta Date: 10/14/2013

Micron Technology Inc. et al.
Ex. 1010, p. 83

Electronic Acknowledgement Receipt	
EFS ID:	17174432
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	19-OCT-2013
Filing Date:	27-JUL-2013
Time Stamp:	03:15:48
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Oath or Declaration filed	2013-10-15_Declaration_Signed_NT003US.pdf	201957 077e7a3e75ea46a46ebcf4c6382d25028db4b7b3	no	2

Warnings:

The page size in the PDF is too large. The pages should be 8.5 x 11 or A4. If this PDF is submitted, the pages will be resized upon entry into the Image File Wrapper and may affect subsequent processing	
Information:	
Total Files Size (in bytes):	201957

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PATENT APPLICATION FEE DETERMINATION RECORD

Substitute for Form PTO-875

Application or Docket Number: 13/952,599

APPLICATION AS FILED - PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	SMALL ENTITY RATE($)	SMALL ENTITY FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY FEE($)
BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	70		N/A	
SEARCH FEE (37 CFR 1.16(k), (l), or (m))	N/A	N/A	N/A	300		N/A	
EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	360		N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	20 minus 20=	*	x 40 =	0.00	OR		
INDEPENDENT CLAIMS (37 CFR 1.16(h))	3 minus 3 =	*	x 210 =	0.00			
APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			0.00			
MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				0.00			
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	730		TOTAL	

APPLICATION AS AMENDED - PART II

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE(S)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT A	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	SMALL ENTITY RATE(S)	SMALL ENTITY ADDITIONAL FEE($)	OR	OTHER THAN SMALL ENTITY RATE($)	OTHER THAN SMALL ENTITY ADDITIONAL FEE($)
AMENDMENT B	Total (37 CFR 1.16(i))	*	Minus	**	=	x =		OR	x =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	x =		OR	x =	
	Application Size Fee (37 CFR 1.16(s))									
	FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))							OR		
						TOTAL ADD'L FEE		OR	TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest found in the appropriate box in column 1.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
13/952,599	07/27/2013	2189	800	NT003US	20	3

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

UPDATED FILING RECEIPT

OC00000004604370

Date Mailed: 10/29/2013

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Hyun Lee, Ladera Ranch, CA;
Jayesh R. Bhakta, Cerritos, CA;

Applicant(s)

Netlist, Inc., Irvine, CA

Assignment For Published Patent Application

Netlist, Inc., Irvine, CA

Power of Attorney: None

Domestic Priority data as claimed by applicant

This appln claims benefit of 61/676,883 07/27/2012

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

If Required, Foreign Filing License Granted: 08/12/2013

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 13/952,599**

Projected Publication Date: Request for Non-Publication Acknowledged

Non-Publication Request: Yes

Early Publication Request: No
**** SMALL ENTITY ****

Micron Technology Inc. et al.
Ex. 1010, p. 87

Title

MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

Preliminary Class

711

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

Micron Technology Inc. et al.
Ex. 1010, p. 88

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

The application data sheet is part of the provisional or nonprovisional application for which it is being submitted. The following form contains the bibliographic data arranged in a format specified by the United States Patent and Trademark Office as outlined in 37 CFR 1.76.
This document may be completed electronically and submitted to the Office in electronic format using the Electronic Filing System (EFS) or the document may be printed and included in a paper filed application.

Secrecy Order 37 CFR 5.2

☐ Portions or all of the application associated with this Application Data Sheet may fall under a Secrecy Order pursuant to 37 CFR 5.2 (Paper filers only. Applications that fall under Secrecy Order may not be filed electronically.)

Inventor Information:

Inventor 1 [redacted]

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Hyun		Lee	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Ladera Ranch	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	21 Thalia
Address 2	

City	Ladera Ranch	State/Province	CA
Postal Code	92694	Country	US

Inventor 2 [redacted]

Legal Name

Prefix	Given Name	Middle Name	Family Name	Suffix
	Jayesh	R.	Bhakta	

Residence Information (Select One) ◉ US Residency ○ Non US Residency ○ Active US Military Service

City	Cerritos	State/Province	CA	Country of Residence	US

Mailing Address of Inventor:

Address 1	~~21 Thalia~~ 12220 Rose St.
Address 2	

City	Cerritos	State/Province	CA
Postal Code	~~92694~~ 90703	Country	US

All Inventors Must Be Listed - Additional Inventor Information blocks may be generated within this form by selecting the **Add** button. [redacted]

Correspondence Information:

Enter either Customer Number or complete the Correspondence Information section below.
For further information see 37 CFR 1.33(a).

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

☐ **An Address is being provided for the correspondence Information of this application.**

Customer Number	79141	
Email Address	jz@jzpatent.com	[redacted]

Application Information:

Title of the Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION		
Attorney Docket Number	NT003US	**Small Entity Status Claimed**	☒
Application Type	Nonprovisional		
Subject Matter	Utility		
Suggested Class (if any)		**Sub Class (if any)**	
Suggested Technology Center (if any)			
Total Number of Drawing Sheets (if any)	26	**Suggested Figure for Publication (if any)**	2A

Publication Information:

☐ Request Early Publication (Fee required at time of Request 37 CFR 1.219)

☒ **Request Not to Publish.** I hereby request that the attached application not be published under 35 U.S.C. 122(b) and certify that the invention disclosed in the attached application **has not and will not** be the subject of an application filed in another country, or under a multilateral international agreement, that requires publication at eighteen months after filing.

Representative Information:

Representative information should be provided for all practitioners having a power of attorney in the application. Providing this information in the Application Data Sheet does not constitute a power of attorney in the application (see 37 CFR 1.32).
Either enter Customer Number or complete the Representative Name section below. If both sections are completed the customer Number will be used for the Representative Information during processing.

Please Select One:	◉ Customer Number	○ US Patent Practitioner	○ Limited Recognition (37 CFR 11.9)
Customer Number	79141		

Domestic Benefit/National Stage Information:

This section allows for the applicant to either claim benefit under 35 U.S.C. 119(e), 120, 121, or 365(c) or indicate National Stage entry from a PCT application. Providing this information in the application data sheet constitutes the specific reference required by 35 U.S.C. 119(e) or 120, and 37 CFR 1.78.

Prior Application Status	Pending	[redacted]	
Application Number	Continuity Type	Prior Application Number	Filing Date (YYYY-MM-DD)
13/952,599	non provisional of	61676883	2012-07-27

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Additional Domestic Benefit/National Stage Data may be generated within this form by selecting the **Add** button.

Foreign Priority Information:

This section allows for the applicant to claim benefit of foreign priority and to identify any prior foreign application for which priority is not claimed. Providing this information in the application data sheet constitutes the claim for priority as required by 35 U.S.C. 119(b) and 37 CFR 1.55(a).

Application Number	Country[i]	Filing Date (YYYY-MM-DD)	Priority Claimed
			(•) Yes ( ) No

Additional Foreign Priority Data may be generated within this form by selecting the **Add** button.

Authorization to Permit Access:

☐ Authorization to Permit Access to the Instant Application by the Participating Offices

If checked, the undersigned hereby grants the USPTO authority to provide the European Patent Office (EPO), the Japan Patent Office (JPO), the Korean Intellectual Property Office (KIPO), the World Intellectual Property Office (WIPO), and any other intellectual property offices in which a foreign application claiming priority to the instant patent application is filed access to the instant patent application. See 37 CFR 1.14(c) and (h). This box should not be checked if the applicant does not wish the EPO, JPO, KIPO, WIPO, or other intellectual property office in which a foreign application claiming priority to the instant patent application is filed to have access to the instant patent application.

In accordance with 37 CFR 1.14(h)(3), access will be provided to a copy of the instant patent application with respect to: 1) the instant patent application-as-filed; 2) any foreign application to which the instant patent application claims priority under 35 U.S.C. 119(a)-(d) if a copy of the foreign application that satisfies the certified copy requirement of 37 CFR 1.55 has been filed in the instant patent application; and 3) any U.S. application-as-filed from which benefit is sought in the instant patent application.

In accordance with 37 CFR 1.14(c), access may be provided to information concerning the date of filing this Authorization.

Applicant Information:

Providing assignment information in this section does not substitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Applicant 1

If the applicant is the inventor (or the remaining joint inventor or inventors under 37 CFR 1.45), this section should not be completed. The information to be provided in this section is the name and address of the legal representative who is the applicant under 37 CFR 1.43; or the name and address of the assignee, person to whom the inventor is under an obligation to assign the invention, or person who otherwise shows sufficient proprietary interest in the matter who is the applicant under 37 CFR 1.46. If the applicant is an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest) together with one or more joint inventors, then the joint inventor or inventors who are also the applicant should be identified in this section.

[redacted]

(●) Assignee	( ) Legal Representative under 35 U.S.C. 117	( ) Joint Inventor
( ) Person to whom the inventor is obligated to assign.	( ) Person who shows sufficient proprietary interest	

If applicant is the legal representative, indicate the authority to file the patent application, the inventor is:

Name of the Deceased or Legally Incapacitated Inventor :

If the Applicant is an Organization check here. ☒

Organization Name	Netlist, Inc.		
Mailing Address Information:			
Address 1	~~51 Discovery, #150~~ 175 Technology, Suite 150		
Address 2			
City	Irvine	**State/Province**	CA
Country	US	Postal Code	92618
Phone Number	949-435-0025	Fax Number	
Email Address	gsasaki@netlist.com		

Additional Applicant Data may be generated within this form by selecting the Add button.

Non-Applicant Assignee Information:

Providing assignment information in this section does not subsitute for compliance with any requirement of part 3 of Title 37 of CFR to have an assignment recorded by the Office.

Assignee 1

Complete this section only if non-applicant assignee information is desired to be included on the patent application publication in accordance with 37 CFR 1.215(b). Do not include in this section an applicant under 37 CFR 1.46 (assignee, person to whom the inventor is obligated to assign, or person who otherwise shows sufficient proprietary interest), as the patent application publication will include the name of the applicant(s).

If the Assignee is an Organization check here. ☐

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it contains a valid OMB control number.

Application Data Sheet 37 CFR 1.76	Attorney Docket Number	NT003US
	Application Number	
Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION	

Prefix	Given Name	Middle Name	Family Name	Suffix

Mailing Address Information:

Address 1	
Address 2	

City		**State/Province**	
Country[i]		Postal Code	
Phone Number		Fax Number	
Email Address			

Additional Assignee Data may be generated within this form by selecting the Add button.

Signature:

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications

Signature	/Jamie J. Zheng/			Date (YYYY-MM-DD)	2014-07-29
First Name	Jamie	Last Name	Zheng	Registration Number	51167

Additional Signature may be generated within this form by selecting the Add button.

This collection of information is required by 37 CFR 1.76. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 23 minutes to complete, including gathering, preparing, and submitting the completed application data sheet form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

Electronic Acknowledgement Receipt	
EFS ID:	19721499
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	30-JUL-2014
Filing Date:	27-JUL-2013
Time Stamp:	01:12:42
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Application Data Sheet	2014_07_29_Supplemental_ADS_NT003US.pdf	1297862 8827c0fb1cc58f9af476a1f4a3854e3c10c3d6eb	no	5
Warnings:					
Information:					

This is not an USPTO supplied ADS fillable form	
Total Files Size (in bytes):	1297862

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
13/952,599	07/27/2013	Hyun Lee	NT003US	6784

79141 7590 07/31/2014
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

EXAMINER	
SUN, MICHAEL	
ART UNIT	PAPER NUMBER
2184	

NOTIFICATION DATE	DELIVERY MODE
07/31/2014	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

jz@jzpatent.com

Office Action Summary	Application No. 13/952,599	Applicant(s) LEE ET AL.	
	Examiner MICHAEL SUN	Art Unit 2184	AIA (First Inventor to File) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE *3* MONTHS FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *7/27/2013*.

☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on _____; the restriction requirement and election have been incorporated into this action.

4)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims*

5)☒ Claim(s) *1-20* is/are pending in the application.

5a) Of the above claim(s) _____ is/are withdrawn from consideration.

6)☐ Claim(s) _____ is/are allowed.

7)☒ Claim(s) *1-20* is/are rejected.

8)☐ Claim(s) _____ is/are objected to.

9)☐ Claim(s) _____ are subject to restriction and/or election requirement.

* If any claims have been determined allowable, you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.

Application Papers

10)☐ The specification is objected to by the Examiner.

11)☒ The drawing(s) filed on *7/27/2013* is/are: a)☒ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:

a)☐ All b)☐ Some** c)☐ None of the:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____.

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

** See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☐ Information Disclosure Statement(s) (PTO/SB/08a and/or PTO/SB/08b) Paper No(s)/Mail Date _____.

3) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

4) ☐ Other: _____.

The present application is being examined under the pre-AIA first to invent provisions.

DETAILED ACTION

Status of the Application

This Office Action is in response to Applicant's Application filed on 7/27/2013.

Claims 1-20 are pending for this examination.

Claim Objections

Claims 1-5 are objected to because of the following informalities:

In claim 1, line 18, "substantially" is a term that is vague and ambiguous. There is no requisite level of understanding to determine what Applicants view is "substantially" aligned, and as such this term can mean just about anything would fall under the definition of "substantially". Examiner recommends removing the term "substantially" or replacing it with a definitive term.

Appropriate correction is required.

Claim Rejections - 35 U.S.C. § 102

The following is a quotation of the appropriate paragraphs of pre-AIA 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

> A person shall be entitled to a patent unless –
>
> (e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the

international application designated the United States and was published under Article 21(2) of such treaty in the English language.

Claims 1-20 are rejected under pre-AIA 35 U.S.C. 102(e) as being anticipated by Manohararajah et al. (US 8,565,033), herein referred to as Manohararajah '033.

Referring to **claim 1,** Manohararajah '033 teaches a memory module (see Figs. 2 and 3, memory module 22) to operate in a memory system with a memory controller (see Fig. 2, memory system with a memory controller 28), the memory system operating according to a system clock (see Col. 5, lines 8-20, wherein the signal clock may serve as the system reference clock), comprising: a module control device to receive command signals from the memory controller and to output module command signals and module control signals (see Col. 5, lines 1-20); memory devices organized in groups, each group including at least one memory device, the memory devices receiving the module command signals from the module control device (see Fig. 3, memory groups 52 of the memory module 22; also see Col. 6, lines 6-19); and a plurality of buffer circuits to receive the module control signals, a respective buffer circuit corresponding to a respective group of memory devices, the respective buffer circuit including data paths for communicating data between the memory controller and the respective group of memory devices, the data paths being controlled by at least one of the module control signals (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group); and wherein the plurality of buffer circuits are distributed across a surface of the memory module such that each module control signal arrives at the plurality of buffer circuits at different points in time (see Col. 6, lines 60-62, wherein buffer circuitry is a part of the memory interface,

in which Examiner points out that the position of the buffer circuitry being on the memory module itself or on the interface would be an obvious variant as the utility function is the same; also see Col. 6, lines 33-48, wherein in a non-zero layout skew, control signals may arrive at each of the memory groups at different times), and wherein the plurality of buffer circuits are configured to align read data signals received from the respective groups of memory devices such that the read data signals are transmitted to the memory controller from the memory module substantially aligned with each other and in accordance with a read latency parameter of the memory system (see Col. 8, lines 18-34, wherein buffer circuitry can be used to synchronize buffers so that read data is output simultaneously in parallel from the multiple memory groups).

As to **claim 2,** Manohararajah '033 teaches the memory module (see Figs. 2 and 3, memory module 22) wherein each module control signal arrives at the plurality of buffer circuits at different points in time across more than one clock cycle of the system clock (see Col. 6, lines 33-48, wherein in a non-zero layout skew, control signals may arrive at each of the memory groups at different times), and wherein the plurality of buffer circuits are configured to align the read data signals such that they are transmitted by the buffer circuits to the memory controller within one system clock cycle (see Col. 8, lines 18-34, wherein buffer circuitry can be used to synchronize buffers so that read data is output simultaneously in parallel from the multiple memory groups).

As to **claim 3,** Manohararajah '033 teaches the memory module (see Figs. 2 and 3, memory module 22) wherein each buffer circuit includes receiver circuits to receive corresponding ones of the module control signals, each receiver circuit including a metastability detection circuit to detect metastability condition in the corresponding module control signal, and

wherein each buffer circuit further includes at least one signal adjustment circuit to adjust one or more the module control signals to mitigate any metastability condition in the module control signals (Examiner points out that circuits are inherently designed to detect for metastability, or the stability of energy levels, to make a decision on whether the signal being received is considered as a logic '0' or logic '1').

As to **claim 4,** Manohararajah '033 teaches the memory module (see Figs. 2 and 3, memory module 22) wherein the module control device further receives the system clock from the system memory controller and transmits a registered clock signal to the plurality of buffer circuits, and wherein each respective buffer circuit regenerates one or more clock signals from the registered clock signal and provides a regenerated clock signal to the respective group of memory devices (see Col. 5, lines 21-30, wherein the memory interface circuit generates the timing information to meet the timing requirements of the protocols being used, i.e. taking the system clock and turning it into timing information for each individual memory group, which would use that information, i.e. a regenerated clock, to time read and writes).

As to **claim 5,** Manohararajah '033 teaches the memory module (see Figs. 2 and 3, memory module 22) wherein each buffer circuit comprises: a time interval determination circuit to receive, during a write operation, a first signal from the module controller and a second signal from the memory controller and to generate a delay signal indicating a time interval between the first signal and the second signal; a delay circuit to delay a signal in the buffer circuit according to the delay signal indicating the time interval, the signal in the buffer circuit being one of a read data signal, a read strobe signal, and a regenerated clock signal (see Col. 7, lines 42-53, wherein

the buffers can be used to delay signals; also see Col. 8, lines 51-67, Col. 9, lines 1-18; also see Col. 9, lines 27-40).

Referring to **claim 6,** Manohararajah '033 teaches a method of operating a memory module (see Abstract; see Figs. 2 and 3, memory module 22) in a memory system having a memory controller (see Fig. 2, memory system with a memory controller 28) comprising: receiving a read command from the memory controller (see Col. 5, lines 1-20); generating module command signals and module control signals in response to the read command (see Col. 5, lines 8-20; also see Col. 5, lines 46-60); producing read data from the memory devices in response to the module command signals (see Col. 6, lines 49-62); controlling a plurality of data paths in a plurality of data buffers using the module control signals (see Col. 5, lines 8-20); and outputting read data from the plurality of data paths to the memory controller (see Col. 6, lines 49-62), wherein at least one signal in each of the plurality of data buffers is delayed such that the read data output from the plurality of data paths are aligned with each other and in accordance with a read latency parameter of the memory system (see Col. 8, lines 18-34, wherein buffer circuitry can be used to synchronize buffers so that read data is output simultaneously in parallel from the multiple memory groups).

As to **claim 7,** Manohararajah '033 teaches the method (see Abstract; see Figs. 2 and 3, memory module 22) further comprising: determining a metastability condition associated with a respective module control signal at each of the plurality of data buffers; and adjusting the respective module control signal based on the metastability condition (Examiner points out that

circuits are inherently designed to detect for metastability, or the stability of energy levels, to make a decision on whether the signal being received is considered as a logic '0' or logic '1').

As to **claim 8,** Manohararajah '033 teaches the method (see Abstract; see Figs. 2 and 3, memory module 22) further comprising, in the each data buffer: receiving at least one first module control signal generated in the memory module in response to a write command; determining a time interval between receiving the at least one first module control signal and receiving a write signal associated with the write command from the memory controller, the write signal being one of a write data signal and a write strobe signal; and delaying the at least one signal in the respective data buffer according to the time interval (see Col. 7, lines 42-53, wherein the buffers can be used to delay signals; also see Col. 8, lines 51-67, Col. 9, lines 1-18; also see Col. 9, lines 27-40).

As to **claim 9,** Manohararajah '033 teaches the method (see Abstract; see Figs. 2 and 3, memory module 22) wherein the at least one signal is a clock signal (see Col. 5, lines 8-20).

As to **claim 10,** Manohararajah '033 teaches the method (see Abstract; see Figs. 2 and 3, memory module 22) wherein the at least one signal is at least one of a read data signal and a read strobe signal (see Col. 5, lines 1-20).

As to **claim 11,** Manohararajah '033 teaches the method (see Abstract; see Figs. 2 and 3, memory module 22) wherein the write signal is a write strobe signal, the method further comprising: detecting a preamble of the write strobe signal to determine the time interval (see Col. 5, lines 1-20).

Referring to **claim 12,** Manohararajah '033 teaches a buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group)for use in a memory module (see Figs. 2 and 3, memory module 22) coupled to a memory controller (see Fig. 2, memory system with a memory controller 28), the memory module including a module controller (see Fig. 2, memory system with a memory controller 28), comprising: a time interval determination circuit to receive, during a write operation (see Col. 5, lines 1-20), a first signal from the module controller and a second signal from the memory controller and to generate a delay signal indicating a time interval between the first signal and the second signal (see Col. 7, lines 42-53, wherein the buffers can be used to delay signals; also see Col. 8, lines 51-67, Col. 9, lines 1-18; also see Col. 9, lines 27-40); and a delay circuit to delay at least one signal in the buffer circuit according to the delay signal indicating the time interval (see Col. 7, lines 42-53, wherein the buffers can be used to delay signals; also see Col. 8, lines 51-67, Col. 9, lines 1-18; also see Col. 9, lines 27-40).

As to **claim 13,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the at least one signal is a clock signal (see Col. 5, lines 8-20).

As to **claim 14,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to

equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the at least one signal is at least one of a read data signal and a read strobe signal (see Col. 5, lines 1-20).

As to **claim 15,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the first signal is an enable signal to enable a write data path in the buffer circuit (see Col. 7, lines 42-53).

As to **claim 16,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the second signal is a write strobe signal (see Col. 5, lines 1-20).

As to **claim 17,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the second signal has a preamble, and wherein the time interval determination circuit includes a preamble detector to detect the preamble of the second signal (see Col. 5, lines 1-20).

As to **claim 18,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the time interval determination circuit further includes a counter to count a number of clock cycles between receipt of the first signal and detection of the preamble of the second signal (see Col. 5, lines 1-20).

As to **claim 19,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) further comprising: a clock regeneration circuit to receive a clock signal from the module controller and to regenerate a regenerated clock signal according to the delay signal indicating the time interval (see Col. 5, lines 21-30, wherein the memory interface circuit generates the timing information to meet the timing requirements of the protocols being used, i.e. taking the system clock and turning it into timing information for each individual memory group, which would use that information, i.e. a regenerated clock, to time read and writes).

As to **claim 20,** Manohararajah '033 teaches the buffer circuit (see Col. 6, lines 60-62, wherein the buffer circuitry can be a part of the memory interface circuitry that is used to equalize the skew across the multiple memory groups, i.e. since each memory group is individual from each other, there would inherently be a respective buffer for each memory group) wherein the buffer circuit is coupled to a respective set of memory devices on the memory module, and

wherein the buffer circuit further comprises a clock regeneration circuit to receive a clock signal from the module controller and to provide a regenerated clock signal to the respective memory devices (see Col. 5, lines 21-30, wherein the memory interface circuit generates the timing information to meet the timing requirements of the protocols being used, i.e. taking the system clock and turning it into timing information for each individual memory group, which would use that information, i.e. a regenerated clock, to time read and writes).

Contact Information

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael Sun whose telephone number is 571-270-1724. The examiner can normally be reached on Mondays to Thursdays from 6:30 am to 5:00 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Dr. Henry Tsai, can be reached at (571) 272-4176. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/MICHAEL SUN/

Primary Examiner, Art Unit 2184

Notice of References Cited	Application/Control No. 13/952,599	Applicant(s)/Patent Under Reexamination LEE ET AL.	
	Examiner MICHAEL SUN	Art Unit 2184	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	Classification
*	A	US-8,565,033	10-2013	Manohararajah et al.	365/193
*	B	US-7,464,225	12-2008	Tsern, Ely	711/115
*	C	US-2009/0198924	08-2009	Shaeffer et al.	711/154
*	D	US-2010/0228891	09-2010	Talbot, Gerald R.	710/14
*	E	US-2007/0064462	03-2007	Matsui, Yoshinori	365/063
*	F	US-2004/0105292	06-2004	Matsui, Yoshinori	365/063
*	G	US-6,754,746	06-2004	Leung et al.	710/100
*	H	US-6,799,252	09-2004	Bauman, Mitchell A.	711/149
	I	US-			
	J	US-			
	K	US-			
	L	US-			
	M	US-			

FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	

*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Micron Technology Inc. et al.
Ex. 1010, p. 109

Search Notes	Application/Control No.	Applicant(s)/Patent Under Reexamination
	13952599	LEE ET AL.
	Examiner	Art Unit
	MICHAEL SUN	2184

CPC- SEARCHED		
Symbol	Date	Examiner

CPC COMBINATION SETS - SEARCHED		
Symbol	Date	Examiner

US CLASSIFICATION SEARCHED			
Class	Subclass	Date	Examiner

SEARCH NOTES		
Search Notes	Date	Examiner
EAST Search - see attached	7/23/2014	MS
Inventor Search - see attached EAST Search	7/23/2014	MS
limited EAST Search of class/subclasses 710/5, 6, and 7	7/23/2014	MS
limited EAST Search of class/subclasses 711/1, 2, 3, 4, 154, and 155	7/23/2014	MS

INTERFERENCE SEARCH			
US Class/ CPC Symbol	US Subclass / CPC Group	Date	Examiner

	/MICHAEL SUN/ Primary Examiner, Art Unit 2184

EAST Search History

EAST Search History (Prior Art)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
S1	2728	(710/5).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:54
S2	716	(710/6).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:54
S3	418	(710/7).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:54
S4	809	(711/1).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S5	505	(711/2).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S6	1012	(711/3).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S7	1152	(711/4).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S8	7947	(711/154).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S9	555	(711/155).CCLS.	US-PGPUB; USPAT; USOCR	OR	OFF	2014/07/23 11:56
S10	1144	(S1 or S2 or S3 or S4 or S5 or S6 or S7 or S8 or S9) and memory module and command	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:57
S11	826	S10 and buffer	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:57
S12	773	S11 and read and write	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:57
S13	599	S12 and memory controller	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:57
S14	345	S13 and delay	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:58
S15	0	S14 and (align$4 with read data) and (sequential with write)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:59
S16	11	S14 and (align$4 with read data)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 11:59
S17	3203	(lee, hyun).in.	US-PGPUB; USPAT;	ADJ	ON	2014/07/23

			USOCR; EPO; DERWENT; IBM_TDB			12:00
S18	83	(bhakta, jayesh).in.	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 12:00
S19	3699	memory module and buffer and memory device and memory controller and control signal and command	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 13:28
S20	797	S19 and system clock	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 13:29
S21	19	S20 and (write with (different near4 time))	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 13:29
S22	1	S21 and (align$4 with read)	US-PGPUB; USPAT; USOCR; EPO; DERWENT; IBM_TDB	ADJ	ON	2014/07/23 13:29

EAST Search History (Interference)

< This search history is empty>

7/ 25/ 2014 12:18:57 PM
C:\ Users\ msun\ Documents\ EAST\ Workspaces\ 13952599.wsp

Index of Claims	Application/Control No.	Applicant(s)/Patent Under Reexamination
	13952599	LEE ET AL.
	Examiner	Art Unit
	MICHAEL SUN	2184

✓	Rejected	-	Cancelled	N	Non-Elected	A	Appeal
=	Allowed	÷	Restricted	I	Interference	O	Objected

☐ Claims renumbered in the same order as presented by applicant ☐ CPA ☐ T.D. ☐ R.1.47

CLAIM		DATE								
Final	Original	07/25/2014								
	1	✓								
	2	✓								
	3	✓								
	4	✓								
	5	✓								
	6	✓								
	7	✓								
	8	✓								
	9	✓								
	10	✓								
	11	✓								
	12	✓								
	13	✓								
	14	✓								
	15	✓								
	16	✓								
	17	✓								
	18	✓								
	19	✓								
	20	✓								

Doc Code: PA..

Document Description: Power of Attorney

PTO/AIA/82A (07-12)

Approved for use through 11/30/2014. OMB 0651-0035

U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

TRANSMITTAL FOR POWER OF ATTORNEY TO ONE OR MORE REGISTERED PRACTITIONERS

NOTE: This form is to be submitted with the Power of Attorney by Applicant form (PTO/AIA/82B or equivalent) to identify the application to which the Power of Attorney is directed, in accordance with 37 CFR 1.5. If the Power of Attorney by Applicant form is not accompanied by this transmittal form or an equivalent, the Power of Attorney will not be recognized in the application.

Application Number	13/952,599
Filing Date	7/27/2013
First Named Inventor	Hyun Lee
Title	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
Art Unit	2184
Examiner Name	SUN, MICHAEL
Attorney Docket Number	NT003US

SIGNATURE of Applicant or Patent Practitioner

Signature	/Jamie J. Zheng/	Date	July 31, 2014
Name	Jamie J. Zheng	Telephone	650-273-6008
Registration Number	51,167		

NOTE: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications.

☐ *Total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32 and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Doc Code: PA..

Document Description: Power of Attorney

PTO/AIA/82B(07-12)
Approved for use through 11/30/2014. OMB 0651-0035
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

POWER OF ATTORNEY BY APPLICANT

I hereby revoke all previous powers of attorney given in the application identified in the attached transmittal letter.

☐ I hereby appoint Practitioner(s) associated with the following Customer Number as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent):

OR

☐ I hereby appoint Practitioner(s) named below as my/our attorney(s) or agent(s), and to transact all business in the United States Patent and Trademark Office connected therewith for the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent):

Name	Registration Number	Name	Registration Number

Please recognize or change the correspondence address for the application identified in the attached transmittal letter to:

☐ The address associated with the above-mentioned Customer Number.

OR

☐ The address associated with Customer Number:

OR

☐ Firm or Individual Name					
Address					
City		State		Zip	
Country					
Telephone		Email			

I am the Applicant:

☐ Inventor or Joint Inventor

☐ Legal Representative of a Deceased or Legally Incapacitated Inventor

☐ Assignee or Person to Whom the Inventor is Under an Obligation to Assign

☐ Person Who Otherwise Shows Sufficient Proprietary Interest (*e.g.*, a petition under 37 CFR 1.46(b)(2) was granted in the application or is concurrently being filed with this document)

SIGNATURE of Applicant for Patent

Signature		Date	
Name		Telephone	
Title and Company			

NOTE: Signature - This form must be signed by the applicant in accordance with 37 CFR 1.33. See 37 CFR 1.4 for signature requirements and certifications. Submit multiple forms for more than one signature, see below *.

☐ *Total of _____ forms are submitted.

This collection of information is required by 37 CFR 1.31, 1.32 and 1.33. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 3 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

PTO/SB/36 (07-09)
Approved for use through 07/31/2012. OMB 0651-0031
U.S. Patent and Trademark Office; U. S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

RESCISSION OF PREVIOUS NONPUBLICATION REQUEST
(35 U.S.C. 122(b)(2)(B)(ii))
AND, IF APPLICABLE,
NOTICE OF FOREIGN FILING
(35 U.S.C. 122(b)(2)(B)(iii))

Send completed form to:
Mail Stop PG Pub
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450
FAX: (571) 273-8300

Application Number	13/952,599
Filing Date	July 27, 2013
First Named Inventor	Hyun Lee
Title	MEMORY MODULE WITH DISTRIBUTED DAT
Atty Docket Number	NT003US
Art Unit	2184
Examiner	SUN, MICHAEL

A request that the above-identified application not be published under 35 U.S.C. 122(b) (nonpublication request) was included with the above-identified application on filing pursuant to 35 U.S.C. 122(b)(2)(B)(i).
I hereby **rescind** the previous nonpublication request.

If a notice of foreign or international filing is or will be required by 35 U.S.C. 122(b)(2)(B)(iii) and 37 CFR 1.213(c), I hereby provide such notice. This notice is being provided no later than forty-five (**45**) days after the date of such foreign or international filing.

*If a notice of subsequent foreign or international filing required by 35 U.S.C. 122(b)(2)(B)(iii) and 37 CFR 1.213(c) was not filed within forty-five (**45**) days after the date of filing of the foreign or international application, the application is ABANDONED, and a petition to revive under 37 CFR 1.137(b) is required. See 37 CFR 1.137(f).*

/Jamie J. Zheng/	July 31, 2014
Signature	Date
Jamie J. Zheng	51167
Typed or printed name	Registration Number, if applicable
650-273-6008	
Telephone Number	

This request must be signed in compliance with 37 CFR 1.33(b).

If information or assistance is needed in completing this form, please contact the Pre-Grant Publication Division at (703)605-4283 or by e-mail at PGPub@USPTO.gov.

CERTIFICATE OF MAILING OR TRANSMISSION

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to: Mail Stop PG Pub, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, or facsimile transmitted to the U.S. Patent and Trademark Office on the date shown below.

Signature

Name (Print/Type) | Date

This collection of information is required by 37 CFR 1.213(b). The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 6 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Mail Stop PG Pub, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (*i.e.*, GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Electronic Acknowledgement Receipt

EFS ID:	19748564
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	01-AUG-2014
Filing Date:	27-JUL-2013
Time Stamp:	01:35:27
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	no

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1	Power of Attorney	2014-07-31_AIA82A_transmittal_POA_NT003US.pdf	195893 d8ea99615344b78e205e8e486ccc1871c76a43d9	no	3
Warnings:					
Information:					

2	Power of Attorney	General_POA.pdf	133545 f8258889ee9edc63b1cc1705b3158faaed49c9c	no	1
Warnings:					
The page size in the PDF is too large. The pages should be 8.5 x 11 or A4. If this PDF is submitted, the pages will be resized upon entry into the Image File Wrapper and may affect subsequent processing					
Information:					
3	Rescind Nonpublication Request for Pre Grant Pub	2014-08-01_Rescind_Non_Pub_Req_NT003US.pdf	259424 5bdae459e9d5d67a3c28e194272f9230d00f823f	no	2
Warnings:					
Information:					
		Total Files Size (in bytes):	588862		

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

GENERAL POWER OF ATTORNEY BY APPLICANT FOR UNITED STATES PATENT

I hereby revoke all previous powers of attorney given in the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent).

I hereby appoint the practitioner associated with the following Customer Number :

79141

as my/our attorney or agent, and to transact all business in the United States Patent and Trademark Office in connection with the application referenced in the attached transmittal letter (form PTO/AIA/82A or equivalent).

Please recognize or change the **correspondence address** for the application referenced in the attached transmittal letter to the address associated with the above-mentioned **Customer Number.**

I am the Applicant:

- ☐ Inventor or Joint Inventor
- ☐ Legal Representative of a Deceased or Legally Incapacitated Inventor
- ■ Assignee or Person to Whom the Inventor is Under an Obligation to Assign
- ☐ Person Who Otherwise Shows Sufficient Proprietary Interest (e.g., a petition under 37 CFR 1.46(b)(2) was granted in the application or is currently being filed in this document)

SIGNATURE of Applicant for Patent

Signature		Date: January 9, 2013
Name	**Gail Sasaki**	
Title	Vice President and Chief Financial Officer	
Company	Netlist, Inc.	

Note: Signature - this form must be signed by the Applicant or a person authorized to act on behalf of Applicant.

Submit multiple forms for more than one signature
Total of __ forms are submitted.

PATENT ASSIGNMENT COVER SHEET

Electronic Version v1.1
Stylesheet Version v1.2

EPAS ID: PAT2963893

SUBMISSION TYPE:	NEW ASSIGNMENT
NATURE OF CONVEYANCE:	ASSIGNMENT

CONVEYING PARTY DATA

Name	Execution Date
HYUN LEE	10/09/2013
JAYESH R BHAKTAH	10/14/2013

RECEIVING PARTY DATA

Name:	NETLIST, INC.
Street Address:	51 DISCOVERY, SUITE 150
City:	IRVINE
State/Country:	CALIFORNIA
Postal Code:	92618

PROPERTY NUMBERS Total: 1

Property Type	Number
Application Number:	13952599

CORRESPONDENCE DATA

Fax Number: (650)852-9432

Correspondence will be sent to the e-mail address first; if that is unsuccessful, it will be sent using a fax number, if provided; if that is unsuccessful, it will be sent via US Mail.

Phone: 650-273-6008
Email: jz@jzpatent.com
Correspondent Name: THE LAW OFFICE OF JAMIE ZHENG, PH.D ESQ.
Address Line 1: P.O. BOX 60573
Address Line 4: PALO ALTO, CALIFORNIA 94306

ATTORNEY DOCKET NUMBER:	NT003US
NAME OF SUBMITTER:	JAMIE J. ZHENG
SIGNATURE:	/Jamie J. Zheng/
DATE SIGNED:	08/01/2014
	This document serves as an Oath/Declaration (37 CFR 1.63).

Total Attachments: 2
source=2013-10-15_Declaration_Signed_NT003US#page1.tif
source=2013-10-15_Declaration_Signed_NT003US#page2.tif

Customer No. 79141
Attorney Docket No.: NT003US

COMBINED DECLARATION AND ASSIGNMENT FOR UTILITY OR DESIGN APPLICATION

Title of Invention	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

DECLARATION

As a below named inventor, I hereby declare that:

This declaration is directed to:

☐ The attached application, or

■ United States application or PCT international application number 13/952,599 filed on July 27, 2013.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

Note to Inventor: 37 C.F.R. § 1.63(c) states: "A person may not execute an oath or declaration for an application unless that person has reviewed and understands the contents of the application, including the claims, and is aware of the duty to disclose to the Office all information known to the person to be material to patentability as defined in § 1.56."

ASSIGNMENT

In consideration of the sum of One Dollar ($1.00) and other good and valuable consideration paid, I hereby sell and assign to

Netlist, Inc.
(Name of Assignee)

of **51 Discovery, Suite #150, Irvine, CA 92618**
(Address of Assignee)

(hereinafter designated as the Assignee) my entire right, title and interest for the United States as defined in 35 USC 100, for any invention set forth in the above-identified application.

I agree to execute all papers necessary in connection with this application for patent in the USPTO for the invention, and any continuation, divisional or reissue applications thereof and also to execute separate assignments in connection with such applications as the Assignee may deem necessary or expedient.

I agree to execute all papers necessary in connection with any post-grant proceeding which may occur in connection with this application or continuation, divisional or reissue thereof and to cooperate with the Assignee in every way possible in obtaining evidence and going forward with such post-grant proceeding.

I agree to execute all papers and documents and perform any act which may be necessary in connection with claims or provisions of the International Convention for Protection of Industrial Property or similar agreements.

I agree to perform all affirmative acts which may be necessary to obtain a grant of a valid United States patent to the Assignee.

I hereby authorize and request the USPTO to issue any and all Letters Patents of the United States resulting from the application or any continuation, divisional or reissue applications thereof to the Assignee, as Assignee of my entire interest, and covenant that I have not executed, and will not execute, any agreement in conflict herewith.

First Inventor (Legal Name): Hyun Lee

Signature: [signature] Date: 10/9/2013

Second Inventor (Legal Name): Jayesh R. Bhakta

Signature: Jayesh Bhakta Date: 10/14/2013

Micron Technology Inc. et al.
Ex. 1010, p. 124

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/952,599	07/27/2013	Hyun Lee	NT003US

CONFIRMATION NO. 6784

NEW OR REVISED PPD NOTICE

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

OC00000070059617

NOTICE OF NEW OR REVISED PROJECTED PUBLICATION DATE

The above-identified application has a new or revised projected publication date. The current projected publication date for this application is 11/13/2014. If this is a new projected publication date (there was no previous projected publication date), the application has been cleared by Licensing & Review or a secrecy order has been rescinded and the application is now in the publication queue.

If this is a revised projected publication date (one that is different from a previously communicated projected publication date), the publication date has been revised due to processing delays in the USPTO or the abandonment and subsequent revival of an application. The application is anticipated to be published on a date that is more than six weeks different from the originally-projected publication date.

More detailed publication information is available through the private side of Patent Application Information Retrieval (PAIR) System. The direct link to access PAIR is currently http://pair.uspto.gov. Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Questions relating to this Notice should be directed to the Office of Data Management, Application Assistance Unit at (571) 272-4000, or (571) 272-4200, or 1-888-786-0101.

Micron Technology Inc. et al.
Ex. 1010, p. 125

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/952,599	07/27/2013	Hyun Lee	NT003US

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

NONPUBLICATION RESCISSION LETTER



Date Mailed: 08/08/2014

Communication Regarding Rescission Of Nonpublication Request and/or Notice of Foreign Filing

Applicant's rescission of the previously-filed nonpublication request and/or notice of foreign filing is acknowledged. The paper has been reflected in the Patent and Trademark Office's (USPTO's) computer records so that the earliest possible projected publication date can be assigned.

The projected publication date is 11/13/2014.

If applicant rescinded the nonpublication request before or on the date of "foreign filing,"[1] then no notice of foreign filing is required.

If applicant foreign filed the application after filing the above application and before filing the rescission, and the rescission did not also include a notice of foreign filing, then a notice of foreign filing (not merely a rescission) is required to be filed within 45 days of the date of foreign filing. See 35 U.S.C. § 122(b)(2)(B)(iii), and Clarification of the United States Patent and Trademark Office's Interpretation of the Provisions of 35 U.S.C. § 122(b)(2)(B)(ii)-(iv), 1272 Off. Gaz. Pat. Office 22 (July 1, 2003).

If a notice of foreign filing is required and is not filed within 45 days of the date of foreign filing, then the application becomes abandoned pursuant to 35 U.S.C. § 122(b)(2)(B)(iii). In this situation, applicant should either file a petition to revive or notify the Office that the application is abandoned. See 37 CFR 1.137(f). Any such petition to revive will be forwarded to the Office of Petitions for a decision. Note that the filing of the petition will not operate to stay any period of reply that may be running against the application.

Questions regarding petitions to revive should be directed to the Office of Petitions at (571) 272-3282.

[1] Note, for purpose of this notice, that "foreign filing" means "filing an application directed to the same invention in another country, or under a multilateral international agreement, that requires publication of applications 18 months after filing".

/hsarwari/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

Micron Technology Inc. et al.
Ex. 1010, p. 126

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING or 371(c) DATE	GRP ART UNIT	FIL FEE REC'D	ATTY.DOCKET.NO	TOT CLAIMS	IND CLAIMS
13/952,599	07/27/2013	2184	800	NT003US	20	3

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

REPLACEMENT FILING RECEIPT

OC00000007000630l

Date Mailed: 08/08/2014

Receipt is acknowledged of this non-provisional patent application. The application will be taken up for examination in due course. Applicant will be notified as to the results of the examination. Any correspondence concerning the application must include the following identification information: the U.S. APPLICATION NUMBER, FILING DATE, NAME OF APPLICANT, and TITLE OF INVENTION. Fees transmitted by check or draft are subject to collection. Please verify the accuracy of the data presented on this receipt. **If an error is noted on this Filing Receipt, please submit a written request for a Filing Receipt Correction. Please provide a copy of this Filing Receipt with the changes noted thereon. If you received a "Notice to File Missing Parts" for this application, please submit any corrections to this Filing Receipt with your reply to the Notice. When the USPTO processes the reply to the Notice, the USPTO will generate another Filing Receipt incorporating the requested corrections**

Inventor(s)

Hyun Lee, Ladera Ranch, CA;
Jayesh R. Bhakta, Cerritos, CA;

Applicant(s)

Netlist, Inc., Irvine, CA

Assignment For Published Patent Application

Netlist, Inc., Irvine, CA

Power of Attorney: The patent practitioners associated with Customer Number 79141

Domestic Priority data as claimed by applicant

This appln claims benefit of 61/676,883 07/27/2012

Foreign Applications for which priority is claimed (You may be eligible to benefit from the **Patent Prosecution Highway** program at the USPTO. Please see http://www.uspto.gov for more information.) - None.
Foreign application information must be provided in an Application Data Sheet in order to constitute a claim to foreign priority. See 37 CFR 1.55 and 1.76.

If Required, Foreign Filing License Granted: 08/12/2013

The country code and number of your priority application, to be used for filing abroad under the Paris Convention, is **US 13/952,599**

Projected Publication Date: 11/13/2014

Non-Publication Request: No

Early Publication Request: No
**** SMALL ENTITY ****

Micron Technology Inc. et al.
Ex. 1010, p. 127

Title

MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

Preliminary Class

710

Statement under 37 CFR 1.55 or 1.78 for AIA (First Inventor to File) Transition Applications: No

PROTECTING YOUR INVENTION OUTSIDE THE UNITED STATES

Since the rights granted by a U.S. patent extend only throughout the territory of the United States and have no effect in a foreign country, an inventor who wishes patent protection in another country must apply for a patent in a specific country or in regional patent offices. Applicants may wish to consider the filing of an international application under the Patent Cooperation Treaty (PCT). An international (PCT) application generally has the same effect as a regular national patent application in each PCT-member country. The PCT process **simplifies** the filing of patent applications on the same invention in member countries, but **does not result** in a grant of "an international patent" and does not eliminate the need of applicants to file additional documents and fees in countries where patent protection is desired.

Almost every country has its own patent law, and a person desiring a patent in a particular country must make an application for patent in that country in accordance with its particular laws. Since the laws of many countries differ in various respects from the patent law of the United States, applicants are advised to seek guidance from specific foreign countries to ensure that patent rights are not lost prematurely.

Applicants also are advised that in the case of inventions made in the United States, the Director of the USPTO must issue a license before applicants can apply for a patent in a foreign country. The filing of a U.S. patent application serves as a request for a foreign filing license. The application's filing receipt contains further information and guidance as to the status of applicant's license for foreign filing.

Applicants may wish to consult the USPTO booklet, "General Information Concerning Patents" (specifically, the section entitled "Treaties and Foreign Patents") for more information on timeframes and deadlines for filing foreign patent applications. The guide is available either by contacting the USPTO Contact Center at 800-786-9199, or it can be viewed on the USPTO website at http://www.uspto.gov/web/offices/pac/doc/general/index.html.

For information on preventing theft of your intellectual property (patents, trademarks and copyrights), you may wish to consult the U.S. Government website, http://www.stopfakes.gov. Part of a Department of Commerce initiative, this website includes self-help "toolkits" giving innovators guidance on how to protect intellectual property in specific countries such as China, Korea and Mexico. For questions regarding patent enforcement issues, applicants may call the U.S. Government hotline at 1-866-999-HALT (1-866-999-4258).

Micron Technology Inc. et al.
Ex. 1010, p. 128

LICENSE FOR FOREIGN FILING UNDER

Title 35, United States Code, Section 184

Title 37, Code of Federal Regulations, 5.11 & 5.15

GRANTED

The applicant has been granted a license under 35 U.S.C. 184, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" followed by a date appears on this form. Such licenses are issued in all applications where the conditions for issuance of a license have been met, regardless of whether or not a license may be required as set forth in 37 CFR 5.15. The scope and limitations of this license are set forth in 37 CFR 5.15(a) unless an earlier license has been issued under 37 CFR 5.15(b). The license is subject to revocation upon written notification. The date indicated is the effective date of the license, unless an earlier license of similar scope has been granted under 37 CFR 5.13 or 5.14.

This license is to be retained by the licensee and may be used at any time on or after the effective date thereof unless it is revoked. This license is automatically transferred to any related applications(s) filed under 37 CFR 1.53(d). This license is not retroactive.

The grant of a license does not in any way lessen the responsibility of a licensee for the security of the subject matter as imposed by any Government contract or the provisions of existing laws relating to espionage and the national security or the export of technical data. Licensees should apprise themselves of current regulations especially with respect to certain countries, of other agencies, particularly the Office of Defense Trade Controls, Department of State (with respect to Arms, Munitions and Implements of War (22 CFR 121-128)); the Bureau of Industry and Security, Department of Commerce (15 CFR parts 730-774); the Office of Foreign AssetsControl, Department of Treasury (31 CFR Parts 500+) and the Department of Energy.

NOT GRANTED

No license under 35 U.S.C. 184 has been granted at this time, if the phrase "IF REQUIRED, FOREIGN FILING LICENSE GRANTED" DOES NOT appear on this form. Applicant may still petition for a license under 37 CFR 5.12, if a license is desired before the expiration of 6 months from the filing date of the application. If 6 months has lapsed from the filing date of this application and the licensee has not received any indication of a secrecy order under 35 U.S.C. 181, the licensee may foreign file the application pursuant to 37 CFR 5.15(b).

SelectUSA

The United States represents the largest, most dynamic marketplace in the world and is an unparalleled location for business investment, innovation, and commercialization of new technologies. The U.S. offers tremendous resources and advantages for those who invest and manufacture goods here. Through SelectUSA, our nation works to promote and facilitate business investment. SelectUSA provides information assistance to the international investor community; serves as an ombudsman for existing and potential investors; advocates on behalf of U.S. cities, states, and regions competing for global investment; and counsels U.S. economic development organizations on investment attraction best practices. To learn more about why the United States is the best country in the world to develop technology, manufacture products, deliver services, and grow your business, visit http://www.SelectUSA.gov or call +1-202-482-6800.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.C. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/952,599	07/27/2013	Hyun Lee	NT003US

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

POA ACCEPTANCE LETTER

OC000000070006283

Date Mailed: 08/08/2014

NOTICE OF ACCEPTANCE OF POWER OF ATTORNEY

This is in response to the Power of Attorney filed 08/01/2014.

The Power of Attorney in this application is accepted. Correspondence in this application will be mailed to the above address as provided by 37 CFR 1.33.

/hsarwari/

Office of Data Management, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

UNITED STATES PATENT AND TRADEMARK OFFICE
COMMISSIONER FOR PATENTS
P.O.BOX 1450
ALEXANDRIA VA 22313-1451

PRESORTED
FIRST-CLASS MAIL
U.S. POSTAGE PAID
POSTEDIGITAL
NNNNN

The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

Courtesy Reminder for
Application Serial No: 13/952,599

Attorney Docket No: NT003US
Customer Number: 79141
Date of Electronic Notification: 08/08/2014

This is a courtesy reminder that new correspondence is available for this application. If you have not done so already, please review the correspondence. The official date of notification of the outgoing correspondence will be indicated on the form PTOL-90 accompanying the correspondence.

An email notification regarding the correspondence was sent to the following email address(es) associated with your customer number:
jz@jzpatent.com

To view your correspondence online or update your email addresses, please visit us anytime at **https://sportal.uspto.gov/secure/myportal/privatepair.** If you have any questions, please email the Electronic Business Center (EBC) at EBC@uspto.gov or call 1-866-217-9197.

UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NUMBER	FILING OR 371(C) DATE	FIRST NAMED APPLICANT	ATTY. DOCKET NO./TITLE
13/952,599	07/27/2013	Hyun Lee	NT003US

CONFIRMATION NO. 6784

79141
The Law Office of Jamie Zheng, Ph.D Esq.
P.O. Box 60573
Palo Alto, CA 94306

PUBLICATION NOTICE

OC00000071837370

Title:MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION

Publication No.US-2014-0337539-A1
Publication Date:11/13/2014

NOTICE OF PUBLICATION OF APPLICATION

The above-identified application will be electronically published as a patent application publication pursuant to 37 CFR 1.211, et seq. The patent application publication number and publication date are set forth above.

The publication may be accessed through the USPTO's publically available Searchable Databases via the Internet at www.uspto.gov. The direct link to access the publication is currently http://www.uspto.gov/patft/.

The publication process established by the Office does not provide for mailing a copy of the publication to applicant. A copy of the publication may be obtained from the Office upon payment of the appropriate fee set forth in 37 CFR 1.19(a)(1). Orders for copies of patent application publications are handled by the USPTO's Office of Public Records. The Office of Public Records can be reached by telephone at (703) 308-9726 or (800) 972-6382, by facsimile at (703) 305-8759, by mail addressed to the United States Patent and Trademark Office, Office of Public Records, Alexandria, VA 22313-1450 or via the Internet.

In addition, information on the status of the application, including the mailing date of Office actions and the dates of receipt of correspondence filed in the Office, may also be accessed via the Internet through the Patent Electronic Business Center at www.uspto.gov using the public side of the Patent Application Information and Retrieval (PAIR) system. The direct link to access this status information is currently http://pair.uspto.gov/. Prior to publication, such status information is confidential and may only be obtained by applicant using the private side of PAIR.

Further assistance in electronically accessing the publication, or about PAIR, is available by calling the Patent Electronic Business Center at 1-866-217-9197.

Office of Data Managment, Application Assistance Unit (571) 272-4000, or (571) 272-4200, or 1-888-786-0101

Micron Technology Inc. et al.
Ex. 1010, p. 132

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant	:	Netlist, Inc.
App. No.	:	13/952,599
Filed	:	July 27, 2013
For	:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
Examiner	:	SUN, MICHAEL
Art Unit	:	2184
Conf. No.	:	6784
Docket No.	:	NT003US

CERTIFICATE OF EFS WEB TRANSMISSION

I hereby certify that this correspondence and any other attachment noted on the Automated Acknowledgement Receipt, is being transmitted from within the Pacific Time zone to the Commissioner for Patents via the EFS Web Server on:

December 31, 2014

(Date)

/Jamie J. Zheng/

Jamie J. Zheng, Reg. No. 51167

AMENDMENT RESPONSE TO OFFICE ACTION DATED APRIL 3, 2014

Mail Stop Amendment
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Dear Sir:

In response to the Non-Final Office Action dated July 31, 2014, Applicant hereby submits the following amendment and remarks.

The Commissioner is hereby authorized to charge any required fee(s) to Deposit Account No. 50-5963.

Amendments to the Claims: Begin on page 2 of this paper.

Remarks: Begin on page 7 of this paper.

AMENDMENTS TO THE CLAIMS

A complete listing of all claims is presented below with insertions underlined (e.g., insertion), and deletions struckthrough or in double brackets (e.g., ~~deletion~~ or [[deletion]]):

1. (Currently Amended) A memory module to operate in a memory system with a memory controller, the memory system operating according to a system clock, the memory system including a memory bus coupling the memory module to the memory controller, the memory bus including a set of control/address signal lines and a plurality of sets of data/strobe signal lines, the memory module comprising:

a module control device to receive memory command signals from the memory controller and to output module command signals and module control signals in response to each of the memory command signals;

memory devices organized in groups, each group including at least one memory device, the memory devices receiving the module command signals from the module control device and performing one or more memory operations in accordance with the module command signals; and

a plurality of buffer circuits to receive the module control signals, ~~a~~ each respective buffer circuit corresponding to a respective group of memory devices and coupled between the respective group of memory devices and a respective set of the plurality of sets of data/strobe signal lines, the respective buffer circuit including data paths for communicating data between the memory controller and the respective group of memory devices, the data paths being controlled by at least one of the module control signals; and

wherein the plurality of buffer circuits are distributed across a surface of the memory module in positions corresponding to respective sets of the plurality of sets of data/strobe signal lines such that each module control signal arrives at the plurality of buffer circuits at different points in time, and

wherein ~~the plurality of buffer circuits are configured to align~~ the each respective buffer circuit is configured to determine a respective time interval based on signals received by the each respective buffer circuit during a memory write operation and is further configured to time transmission of a respective set of read data signals received from the respective group ~~groups~~ of memory devices ~~such that the read data signals are transmitted to the memory controller from the memory module substantially aligned with each other and~~ in accordance with the time interval

Micron Technology Inc. et al.
Ex. 1010, p. 134

and a read latency parameter of the memory system during a memory read operation.

2. (Currently Amended) The memory module of claim 1, wherein each module control signal arrives at the plurality of buffer circuits at different points in time across more than one clock cycle of the system clock, and wherein the ~~plurality of buffer circuits are~~ each respective buffer circuit is configured to align the respective set of read data signals such that ~~they~~ different sets of read data signals corresponding to the memory read operation are transmitted by the different buffer circuits onto ~~to~~ different sets of data/strobe signal lines in the memory bus ~~controller~~ within one system clock cycle.

3. (Original) The memory module of claim 1, wherein each buffer circuit includes receiver circuits to receive corresponding ones of the module control signals, each receiver circuit including a metastability detection circuit to detect metastability condition in the corresponding module control signal, and wherein each buffer circuit further includes at least one signal adjustment circuit to adjust one or more the module control signals to mitigate any metastability condition in the module control signals.

4. (Currently Amended) The memory module of claim 1, wherein the module control device further receives the system clock from the system memory controller and transmits a registered clock signal to the plurality of buffer circuits, and wherein the each respective buffer circuit regenerates one or more clock signals from the registered clock signal and provides a regenerated clock signal to the respective group of memory devices.

5. (Currently Amended) The memory module of claim 1, wherein the respective ~~each~~ buffer circuit comprises:

a time interval determination circuit to receive, during a write operation, a first signal from the module controller and a second signal from the memory ~~controller~~ bus and to generate a delay signal indicating a time interval between the first signal and the second signal;

a delay circuit to delay the respective set of read data signals received from the respective group of memory devices ~~a signal in the buffer circuit~~ according to the delay signal ~~indicating the time interval, the signal in the buffer circuit being one of a read data signal, a read strobe signal, and a regenerated clock signal~~.

Micron Technology Inc. et al.
Ex. 1010, p. 135

6. (Currently Amended) A method of operating a memory module in a memory system having a memory controller, the memory module being coupled to the memory controller via a memory bus, the memory bus including a set of control/address signal lines and a plurality of sets of data/strobe signal lines, the memory module including memory devices and a module controller coupled to the memory controller via the set of control/address signal lines, the method comprising:

receiving a read command from the memory controller via the set of control/address signal lines;

generating module command signals and module control signals in response to the read command;

producing read data from the memory devices in response to the module command signals;

controlling a plurality of sets data paths in a plurality of data buffers using the module control signals such that each respective set of the plurality of sets of data paths receive and conduct a respective portion of the read data; and

outputting read data from the plurality of data paths ~~to~~ onto the memory bus, wherein a first set of the plurality of data paths output a first portion of the read data onto a first set of data/strobe signal lines in accordance with a first time interval, wherein a second set of the plurality of data paths output a second portion of the read data onto a second set of data/strobe signal lines in accordance with a second time interval different from the first time interval, the first time interval and the second time interval having been determined during a prior write operation performed in the memory system ~~controller, wherein at least one signal in each of the plurality of data buffers is delayed such that the read data output from the plurality of data paths are aligned with each other and in accordance with a read latency parameter of the memory system~~.

7. (Original) The method of claim 6, further comprising:

determining a metastability condition associated with a respective module control signal at each of the plurality of data buffers; and

adjusting the respective module control signal based on the metastability condition.

8. (Currently Amended) The method of claim 6, further comprising, in the each data buffer:

receiving at least one first module control signal generated in the memory module in

Micron Technology Inc. et al.
Ex. 1010, p. 136

response to a write command associated with the prior write operation;

determining a time interval between receiving the at least one first module control signal and receiving a write signal associated with the write command from the memory controller, the write signal being one of a write data signal and a write strobe signal; and

delaying ~~the~~ at least one signal in the respective data buffer according to the time interval.

9. (Original) The method of claim 8 wherein the at least one signal is a clock signal.

10. (Original) The method of claim 8 wherein the at least one signal is at least one of a read data signal and a read strobe signal.

11. (Original) The method of claim 8, wherein the write signal is a write strobe signal, the method further comprising:

detecting a preamble of the write strobe signal to determine the time interval.

12. (Currently Amended) A buffer circuit for use in a memory module coupled to a memory controller via a memory bus, the memory bus including a set of control/address signal lines and a plurality of sets of data/strobe signal lines, the memory module including a plurality of memory devices organized in groups and a module controller coupled to the memory controller via the set of control/address signal lines, comprising:

a time interval determination circuit to ~~receive, during a write operation, a first signal from the module controller and a second signal from the memory controller and to generate a delay signal indicating~~ determine a time interval between receiving a first signal from the module controller and receiving a second signal from the memory bus ~~the first signal and the second signal~~; and

a delay circuit to ~~delay at least one signal in the buffer circuit~~ time transmission of data/strobe signals received from a respective group of the plurality of memory devices to the memory controller via a respective set of the plurality of sets of data/strobe signal lines according to ~~the delay signal indicating~~ the time interval.

13. (Currently Amended) The buffer circuit of claim 12 wherein the buffer further comprises receiver circuits to receive corresponding ones of the module control signals, each receiver circuit including a metastability detection circuit to detect metastability condition in the corresponding module control signal ~~the at least one signal is a clock signal~~.

Micron Technology Inc. et al.
Ex. 1010, p. 137

14. (Currently Amended) The buffer circuit of claim 12 wherein ~~the at least one signal is at least one of a read data signal and a read strobe signal~~ the buffer circuit further comprises at least one signal adjustment circuit to adjust one or more the module control signals to mitigate metastability condition in the module control signals.

15. (Original) The buffer circuit of claim 12, wherein the first signal is an enable signal to enable a write data path in the buffer circuit.

16. (Original) The buffer circuit of claim 15, wherein the second signal is a write strobe signal.

17. (Original) The buffer circuit of claim 15, wherein the second signal has a preamble, and wherein the time interval determination circuit includes a preamble detector to detect the preamble of the second signal.

18. (Currently Amended) The buffer circuit of claim 15 ~~18~~, wherein the time interval determination circuit further includes a counter to count a number of clock cycles between receipt of the first signal and detection of the preamble of the second signal.

19. (Currently Amended) The buffer circuit of claim 12, further comprising:

a clock regeneration circuit to receive a clock signal from the module controller and to regenerate a regenerated clock signal according to the ~~delay signal indicating the~~ time interval.

20. (Currently Amended) The buffer circuit of claim 12, wherein the buffer circuit is coupled to a respective ~~set~~ group of the plurality of memory devices on the memory module, and wherein the buffer circuit further comprises a clock regeneration circuit to receive a clock signal from the module controller and to provide a regenerated clock signal to the respective group of the plurality of memory devices.

Micron Technology Inc. et al.
Ex. 1010, p. 138

REMARKS

This amendment responds to the non-final Office Action dated July 31, 2014, in which the Examiner:

- objected to certain informalities in claims 1-5;
- rejected claims 1-20 under pre-AIA 35 U.S.C. 102(e) as being anticipated by Manohararajah et al. (US 8,565,033), hereafter referred to as Manohararajah.

I. CLAIM AMENDMENT

Claims 1-2, 4-6, 8, 12-14, and 18-20 are amended. No new matter is added.

II. OBJECTION TO INFORMALITIES IN CLAIMS 1-5

Claim 1 has been amended and the term "substantially" has been removed. Therefore, the objection is now moot.

III. REJECTION OF CLAIMS 1-20 UNDER PRE-AIA 35 U.S.C. 102(E) AS BEING ANTICIPATED BY MANOHARARAJAH

Applicant disagrees with the Examiner's assertion that Manohararajah teaches a memory module comprising, among other things, "a plurality of buffer circuits to receive the module control signals from the module control device, a respective buffer circuit corresponding to a respective group of memory devices, the respective buffer circuit including data paths for communicating data between the memory controller and the respective group of memory devices, the data paths being controlled by at least one of the module control signals." The interface circuitry 24 in Manohararajah, which the Examiner asserts to include the "buffer circuits", is part of the memory controller 10 (Manohararajah, FIG. 3). Thus, whatever "buffer circuits" the Examiner speculates could be part of the interface circuitry 24 are not even on the memory module and cannot possibly receive module control signals from any module control device on the memory module 22 or be controlled by at least one of the module control signals.

Applicant also disagrees with the Examiner's assertion that Manohararajah teaches "the plurality of buffer circuits" being distributed across a surface of the memory module 22. As shown in Manohararajah, FIG. 3, any buffer circuits that could be part of the interface circuitry 24 are not even on the memory module 22. They are separated from the memory module by the

Micron Technology Inc. et al.
Ex. 1010, p. 139

memory bus including the signal lines 34-1 to 34-N and 36. So, they cannot be distributed across a surface of the memory module 22.

Applicant also disagrees with the Examiner's assertion that "the buffer circuitry being on the memory module itself or on the interface would be an obvious variant as the utility function is the same." This assertion is baseless. For a claim to be rejected under pre-AIA 35 U.S.C. 102(e), a single piece of prior art must disclose each and every claimed limitation. Persons of ordinary skill in the art will understand that there is clear difference between data buffers on a memory module and data buffers on a memory controller. The Examiner cannot make up what is lacking in the prior art by a simple assertion that it is obvious. Furthermore, claim 1 recited not just data buffers on a memory module, but data buffers that receive module control signals from a module control device on the memory module and are controlled by the module control signals. The Examiner has not shown how the interface circuitry 24 in Manohararajah can be controlled by any module control device on the memory module 22, as claimed in claim 1.

Claim 1 as amended includes further limitations not disclosed in Manohararajah. For example, Manohararajah does not teach or disclose data buffers coupled between respective groups of memory devices and respective sets of the plurality of sets of data/strobe signal lines in the memory bus and distributed across a surface of the memory module in positions corresponding to respective sets of the plurality of sets of data/strobe signal lines. Further, Manohararajah also does not disclose that each respective buffer circuit is configured to determine a respective time interval based on signals received by the each respective buffer circuit during a memory write operation, and is further configured to time transmission of a respective set of read data signals received from a respective group of memory devices in accordance with the time interval and a read latency parameter of the memory system during a memory read operation.

The Examiner cites Col. 7, lines 42-53, Col. 8, lines 51-67, and Col. 9, lines 1-18 Manohararajah in the Office Action, but these passages in Manohararajah merely states that the DQ/DQS signals can be delayed so that the DQ/DQS signals are phase aligned for optimal setup and hold times. Nowhere in Manohararajah it is described the transmission of read data signal from each respective buffer is timed in accordance with a time interval determined based signals

Micron Technology Inc. et al.
Ex. 1010, p. 140

received by the each respective buffer. Thus, Manohararajah does not anticipate claim 1 as amended.

Therefore, claim 1 as amended is patentable over Manohararajah.

Claim 2 as amended, claim 3, claims 4-5 as amended depend from claim 1 as amended and include further limitations in addition to the limitations in claim 1. Therefore, claim 2 as amended, claim 3, claims 4-5 as amended are patentable for at least the same reasons claim 1 as amended is patentable.

Claim 6 as amended includes the limitations of:

> controlling a plurality of sets data paths in a plurality of data buffers using the module control signals such that each respective set of the plurality of sets of data paths receive and conduct a respective portion of the read data; and
>
> outputting read data from the plurality of data paths onto the memory bus, wherein a first set of the plurality of data paths output a first portion of the read data onto a first set of data/strobe signal lines in accordance with a first time interval, wherein a second set of the plurality of data paths output a second portion of the read data onto a second set of data/strobe signal lines in accordance with a second time interval different from the first time interval, the first time interval and the second time interval being determined during a prior write operation performed in the memory system.

These features in claim 6 as amended are not disclosed in Manohararajah because Manoahararajah does not teach or disclose a plurality of data paths for receiving and conducting read data from the memory devices, wherein the plurality of data paths are controlled by module control signals generated by a module control device on the memory module in accordance to a read command received via a set of control/address signal lines in the memory bus. As discussed above, the interface circuitry 24, wherein the Examiner speculated to include a plurality of data paths are not on the memory module 22 and could not possibly be controlled by module control signals generated by any module control device on the memory module 22 in accordance to a read command received via a set of control/address signal lines in the memory bus.

Manohararajah also does not teach or disclose the plurality of data paths outputting read data onto the memory bus. As shown in FIG. 3 of Manohararajah, the interface circuitry 24, being on the memory controller 10, cannot possibly output read data onto the memory bus. It can only receive read data from the memory bus. Furthermore, Manohararajah does not disclose the

Micron Technology Inc. et al.
Ex. 1010, p. 141

additional features of a first set of the plurality of data paths outputting a first portion of the read data onto a first set of data/strobe signal lines in accordance with a first time interval, a second set of the plurality of data paths outputting a second portion of the read data onto a second set of data/strobe signal lines in accordance with a second time interval different from the first time interval, the first time interval and the second time interval being determined during a prior write operation performed in the memory system.

Therefore, claim 6 as amended is patentable over Manohararajah.

Claim 7, claim 8 as amended, and claims 9-11 depend from claim 6 as amended and include further limitations in addition to the limitations in claim 6 as amended. Therefore, claim 7, claim 8 as amended, and claims 9-11 are patentable for at least the same reasons claim 6 as amended is patentable.

The arguments regarding claim 1 as amended apply to claim 12 as amended. Therefore, claim 12 as amended is also patentable over Manohararajah.

Claims 13-14 as amended, claims 15-17, and claims 18-20 as amended depend from claim 12 as amended and include further limitations in addition to the limitations in claim 12 as amended. Therefore, claims 13-14 as amended, claims 15-17, and claims 18-20 as amended are patentable for at least the same reasons claim 12 as amended is patentable.

Although the present communication may include alterations to the application or claims, or characterizations of claim scope or referenced art, Applicants are not conceding in this application that previously pending claims are not patentable over the cited references. Rather, any alterations or characterizations are being made to facilitate expeditious prosecution of this application. Applicants reserve the right to pursue at a later date any previously pending or other broader or narrower claims that capture any subject matter supported by the present disclosure, including subject matter found to be specifically disclaimed herein or by any prior prosecution. Accordingly, reviewers of this or any parent, child or related prosecution history shall not reasonably infer that Applicants have made any disclaimers or disavowals of any subject matter supported by the present application.

Micron Technology Inc. et al.
Ex. 1010, p. 142

The Examiner is invited to call the undersigned at (650) 273-6008 if such phone call would help resolve any remaining issues.

Respectfully submitted,

Date: December 31, 2014

By: /Jamie J. Zheng/
Jamie J. Zheng
Registration No. 51,167
Customer No. 79,141

Micron Technology Inc. et al.
Ex. 1010, p. 143

Electronic Patent Application Fee Transmittal

Application Number:	13952599
Filing Date:	27-Jul-2013
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Filer:	Jamie Jie Zheng
Attorney Docket Number:	NT003US

Filed as Small Entity

Filing Fees for Utility under 35 USC 111(a)

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Basic Filing:				
Pages:				
Claims:				
Miscellaneous-Filing:				
Petition:				
Patent-Appeals-and-Interference:				
Post-Allowance-and-Post-Issuance:				
Extension-of-Time:				

Description	Fee Code	Quantity	Amount	Sub-Total in USD($)
Extension - 2 months with $0 paid	2252	1	300	300
Miscellaneous:				
		Total in USD ($)		**300**

Electronic Acknowledgement Receipt	
EFS ID:	21104017
Application Number:	13952599
International Application Number:	
Confirmation Number:	6784
Title of Invention:	MEMORY MODULE WITH DISTRIBUTED DATA BUFFERS AND METHOD OF OPERATION
First Named Inventor/Applicant Name:	Hyun Lee
Customer Number:	79141
Filer:	Jamie Jie Zheng
Filer Authorized By:	
Attorney Docket Number:	NT003US
Receipt Date:	31-DEC-2014
Filing Date:	27-JUL-2013
Time Stamp:	23:52:35
Application Type:	Utility under 35 USC 111(a)

Payment information:

Submitted with Payment	yes
Payment Type	Credit Card
Payment was successfully received in RAM	$300
RAM confirmation Number	11075
Deposit Account	
Authorized User	
The Director of the USPTO is hereby authorized to charge indicated fees and credit any overpayment as follows:	

File Listing:

Document Number	Document Description	File Name	File Size(Bytes)/ Message Digest	Multi Part /.zip	Pages (if appl.)
1		2014-12-31_Amendment_NT003US.pdf	132651 cfad8065bafb9d79ac14decc757e99d22550e4f2	yes	11

Multipart Description/PDF files in .zip description		
Document Description	**Start**	**End**
Amendment/Req. Reconsideration-After Non-Final Reject	1	1
Claims	2	6
Applicant Arguments/Remarks Made in an Amendment	7	11

Warnings:

Information:

2	Fee Worksheet (SB06)	fee-info.pdf	30398 4543ef2223365250f1ae8fa1dd884471f17b71d5	no	2

Warnings:

Information:

Total Files Size (in bytes):	163049

This Acknowledgement Receipt evidences receipt on the noted date by the USPTO of the indicated documents, characterized by the applicant, and including page counts, where applicable. It serves as evidence of receipt similar to a Post Card, as described in MPEP 503.

New Applications Under 35 U.S.C. 111
If a new application is being filed and the application includes the necessary components for a filing date (see 37 CFR 1.53(b)-(d) and MPEP 506), a Filing Receipt (37 CFR 1.54) will be issued in due course and the date shown on this Acknowledgement Receipt will establish the filing date of the application.

National Stage of an International Application under 35 U.S.C. 371
If a timely submission to enter the national stage of an international application is compliant with the conditions of 35 U.S.C. 371 and other applicable requirements a Form PCT/DO/EO/903 indicating acceptance of the application as a national stage submission under 35 U.S.C. 371 will be issued in addition to the Filing Receipt, in due course.

New International Application Filed with the USPTO as a Receiving Office
If a new international application is being filed and the international application includes the necessary components for an international filing date (see PCT Article 11 and MPEP 1810), a Notification of the International Application Number and of the International Filing Date (Form PCT/RO/105) will be issued in due course, subject to prescriptions concerning national security, and the date shown on this Acknowledgement Receipt will establish the international filing date of the application.

PTO/SB/06 (09-11)
Approved for use through 1/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

PATENT APPLICATION FEE DETERMINATION RECORD Substitute for Form PTO-875	Application or Docket Number 13/952,599	Filing Date 07/27/2013	☐ To be Mailed

ENTITY: ☐ LARGE ☒ SMALL ☐ MICRO

APPLICATION AS FILED – PART I

FOR	(Column 1) NUMBER FILED	(Column 2) NUMBER EXTRA	RATE ($)	FEE ($)
☐ BASIC FEE (37 CFR 1.16(a), (b), or (c))	N/A	N/A	N/A	
☐ SEARCH FEE (37 CFR 1.16(k), (i), or (m))	N/A	N/A	N/A	
☐ EXAMINATION FEE (37 CFR 1.16(o), (p), or (q))	N/A	N/A	N/A	
TOTAL CLAIMS (37 CFR 1.16(i))	minus 20 = *		X $ =	
INDEPENDENT CLAIMS (37 CFR 1.16(h))	minus 3 = *		X $ =	
☐ APPLICATION SIZE FEE (37 CFR 1.16(s))	If the specification and drawings exceed 100 sheets of paper, the application size fee due is $310 ($155 for small entity) for each additional 50 sheets or fraction thereof. See 35 U.S.C. 41(a)(1)(G) and 37 CFR 1.16(s).			
☐ MULTIPLE DEPENDENT CLAIM PRESENT (37 CFR 1.16(j))				
* If the difference in column 1 is less than zero, enter "0" in column 2.			TOTAL	

APPLICATION AS AMENDED – PART II

AMENDMENT	12/31/2014	(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	* 20	Minus	** 20	= 0	X $40 =	0
	Independent (37 CFR 1.16(h))	* 3	Minus	***3	= 0	X $210 =	0
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	**0**

AMENDMENT		(Column 1) CLAIMS REMAINING AFTER AMENDMENT		(Column 2) HIGHEST NUMBER PREVIOUSLY PAID FOR	(Column 3) PRESENT EXTRA	RATE ($)	ADDITIONAL FEE ($)
	Total (37 CFR 1.16(i))	*	Minus	**	=	X $ =	
	Independent (37 CFR 1.16(h))	*	Minus	***	=	X $ =	
	☐ Application Size Fee (37 CFR 1.16(s))						
	☐ FIRST PRESENTATION OF MULTIPLE DEPENDENT CLAIM (37 CFR 1.16(j))						
						TOTAL ADD'L FEE	

* If the entry in column 1 is less than the entry in column 2, write "0" in column 3.
** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 20, enter "20".
*** If the "Highest Number Previously Paid For" IN THIS SPACE is less than 3, enter "3".
The "Highest Number Previously Paid For" (Total or Independent) is the highest number found in the appropriate box in column 1.

LIE
/RENEE HAWKINS/

This collection of information is required by 37 CFR 1.16. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 12 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 1 816 570 A2

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
08.08.2007 Bulletin 2007/32

(51) Int Cl.:
G06F 13/16 *(2006.01)*

(21) Application number: **06125954.5**

(22) Date of filing: **16.04.1991**

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR
Designated Extension States:
AL BA HR MK YU

(30) Priority: **18.04.1990 US 510898**

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
00100018.1 / 1 022 641
91908374.1 / 0 525 068

(71) Applicant: **RAMBUS INC.**
Los Altos,
California 94022 (US)

(72) Inventors:
- **Farmwald, Michael**
  Portola Valley, CA 94028 (US)
- **Horowitz, Mark**
  Palo Alto, CA 94306 (US)

(74) Representative: **Eisenführ, Speiser & Partner**
Patentanwälte Rechtsanwälte
Postfach 10 60 78
28060 Bremen (DE)

Remarks:
This application was filed on 12 - 12 - 2006 as a divisional application to the application mentioned under INID code 62.

(54) Integrated circuit I/O using a high performance bus interface

(57) The present invention includes a semiconductor device, the semiconductor device comprising connection means adapted to connect the semiconductor device to an external bus which is a part of a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to the external bus, wherein the external bus includes a plurality of bus lines for carrying substantially all address, data, control and device select information needed by the semiconductor device for communication with substantially every other semiconductor device connected to the external bus, and wherein the external bus has substantially fewer bus lines than the number of bits in a single address, wherein each said bus line is a terminated transmission line, and wherein the external bus carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to the semiconductor device, a bus line driver capable of producing a low-voltage-swing signal on one of said terminated transmission lines, wherein the semiconductor device is a semiconductor transceiver device to interface multiple units to the bus, wherein the multiple units are DRAMs.

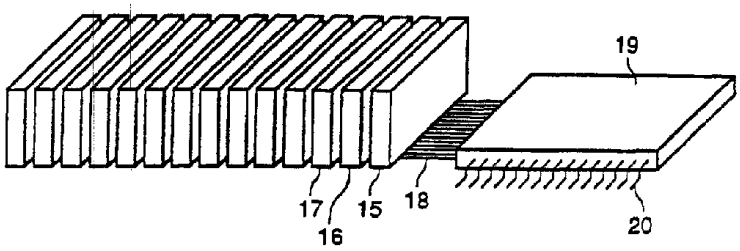


FIG. 3

EP 1 816 570 A2

Printed by Jouve, 75001 PARIS (FR)

Description

FIELD OF THE INVENTION

[0001] An integrated circuit bus interface for computer and video systems is described which allows high speed transfer of blocks of data, particularly to and from memory devices, with reduced power consumption and increased system reliability. A new method of physically implementing the bus architecture is also described.

BACKGROUND OF THE INVENTION

[0002] Semiconductor computer memories have traditionally been designed and structured to use one memory device for each bit, or small group of bits, of any individual computer word, where the word size is governed by the choice of computer. Typical word sizes range from 4 to 64 bits. Each memory device typically is connected in parallel to a series of address lines and connected to one of a series of data lines. When the computer seeks to read from or write to a specific memory location, an address is put on the address lines and some or all of the memory devices are activated using a separate device select line for each needed device. One or more devices may be connected to each data line but typically only a small number of data lines are connected to a single memory device. Thus data line 0 is connected to device(s) 0, data line 1 is connected to device(s) 1, and so on. Data is thus accessed or provided in parallel for each memory read or write operation. For the system to operate properly, every single memory bit in every memory device must operate dependably and correctly.

[0003] To understand the concept of the present invention, it is helpful to review the architecture of conventional memory devices. Internal to nearly all types of memory devices (including the most widely used Dynamic Random Access Memory (DRAM), Static RAM (SRAM) and Read Only Memory (ROM) devices), a large number of bits are accessed in parallel each time the system carries out a memory access cycle. However, only a small percentage of accessed bits which are available internally each time the memory device is cycled ever make it across the device boundary to the external world.

[0004] Referring to Fig. 1, all modern DRAM, SRAM and ROM designs have internal architectures with row (word) lines 5 and column (bit) lines 6 to allow the memory cells to tile a two dimensional area 1. One bit of data is stored at the intersection of each word and bit line. When a particular word line is enabled, all of the corresponding data bits are transferred onto the bit lines. Some prior art DRAMs take advantage of this organization to reduce the number of pins needed to transmit the address. The address of a given memory cell is split into two addresses, row and column, each of which can be multiplexed over a bus only half as wide as the memory cell address of the prior art would have required.

COMPARISON WITH PRIOR ART

[0005] Prior art memory systems have attempted to solve the problem of high speed access to memory with limited success. U.S. Patent No. 3,821,715 (Hoff et. al.), was issued to Intel Corporation for the earliest 4-bit micro-processor. That patent describes a bus connecting a single central processing unit (CPU) with multiple RAMs and ROMs. That bus multiplexes addresses and data over a 4-bit wide bus and uses point-to-point control signals to select particular RAMs or ROMs. The access time is fixed and only a single processing element is permitted. There is no block-mode type of operation, and most important, not all of the interface signals between the devices are bused (the ROM and RAM control lines and the RAM select lines are point-to-point).

[0006] In U.S. Patent No. 4,315,308 (Jackson), a bus connecting a single CPU to a bus interface unit is described. The invention uses multiplexed address, data, and control information over a single 16-bit wide bus. Block-mode operations are defined, with the length of the block sent as part of the control sequence. In addition, variable access-time operations using a "stretch" cycle signal are provided. There are no multiple processing elements and no capability for multiple outstanding requests, and again, not all of the interface signals are bused.

[0007] In U.S. Patent No. 4,449,207 (Kung, et. al.), a DRAM is described which multiplexes address and data on an internal bus. The external interface to this DRAM is conventional, with separate control, address and data connections.

[0008] In U.S. Patent Nos. 4,764,846 and 4,706,166 (Go), a 3-D package arrangement of stacked die with connections along a single edge is described. Such packages are difficult to use because of the point-to-point wiring required to interconnect conventional memory devices with processing elements. Both patents describe complex schemes for solving these problems. No attempt is made to solve the problem by changing the interface.

[0009] In U.S. Patent No. 3,969,706 (Proebsting, et. al.), the current state-of-the-art DRAM interface is described. The address is two-way multiplexed, and there are separate pins for data and control (RAS, CAS, WE, CS). The number of pins grows with the size of the DRAM, and many of the connections must be made point-to-point in a memory system using such DRAMs.

[0010] There are many backplane buses described in the prior art, but not in the combination described or having the

Micron Technology Inc. et al.
Ex. 1010, p. 150

features of this invention. Many backplane buses multiplex addresses and data on a single bus (e.g., the NU bus). ELXSI and others have implemented split-transaction buses (U.S. Patent No. 4,595,923 and 4,481,625 (Roberts)). ELXSI has also implemented a relatively low-voltage-swing current-mode ECL driver (approximately 1 V swing). Address-space registers are implemented on most backplane buses, as is some form of block mode operation.

[0011] Nearly all modern backplane buses implement some type of arbitration scheme, but the arbitration scheme used in this invention differs from each of these. U.S. Patent Nos. 4,837,682 (Culler), 4,818,985 (Ikeda), 4,779,089 (Theus) and 4,745,548 (Blahut) describe prior art schemes. All involve either log N extra signals, (Theus, Blahut), where N is the number of potential bus requestors, or additional delay to get control of the bus (Ikeda, Culler). None of the buses described in patents or other literature use only bused connections. All contain some point-to-point connections on the backplane. None of the other aspects of this invention such as power reduction by fetching each data block from a single device or compact and low-cost 3-D packaging even apply to backplane buses.

[0012] The clocking scheme used in this invention has not been used before and in fact would be difficult to implement in backplane buses due to the signal degradation caused by connector stubs. U.S. Patent No. 4,247,817 (Heller) describes a clocking scheme using two clock lines, but relies on ramp-shaped clock signals in contrast to the normal rise-time signals used in the present invention.

[0013] In U.S. Patent No. 4,646,279 (Voss), a video RAM is described which implements a parallel-load, serial-out shift register on the output of a DRAM. This generally allows greatly improved bandwidth (and has been extended to 2, 4 and greater width shift-out paths.) The rest of the interfaces to the DRAM (RAS, CAS, multiplexed address, etc.) remain the same as for conventional DRAMS.

[0014] One object of the present invention is to use a new bus interface built into semiconductor devices to support high-speed access to large blocks of data from a single memory device by an external user of the data, such as a microprocessor, in an efficient and cost-effective manner.

[0015] Another object of this invention is to provide a clocking scheme to permit high speed clock signals to be sent along the bus with minimal clock skew between devices.

[0016] Another object of this invention is to allow mapping out defective memory devices or portions of memory devices.

[0017] Another object of this invention is to provide a method for distinguishing otherwise identical devices by assigning a unique identifier to each device.

[0018] Yet another object of this invention is to provide a method for transferring address, data and control information over a relatively narrow bus and to provide a method of bus arbitration when multiple devices seek to use the bus simultaneously.

[0019] Another object of this invention is to provide a method of distributing a high-speed memory cache within the DRAM chips of a memory system which is much more effective than previous cache methods.

[0020] Another object of this invention is to provide devices, especially DRAMs, suitable for use with the bus architecture of the invention.

SUMMARY OF INVENTION

[0021] The present invention includes a memory subsystem comprising at least two semiconductor devices, including at least one memory device, connected in parallel to a bus, where the bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by said memory devices, where the control information includes device-select information and the bus has substantially fewer bus lines than the number of bits in a single address, and the bus carries device-select information without the need for separate device-select lines connected directly to individual devices.

[0022] Referring to Fig. 2, a standard DRAM 13, 14, ROM (or SRAM) 12, microprocessor CPU 11, I/O device, disk controller or other special purpose device such as a high speed switch is modified to use a wholly bus-based interface rather than the prior art combination of point-to-point and bus-based wiring used with conventional versions of these devices. The new bus includes clock signals, power and multiplexed address, data and control signals. In a preferred implementation, 8 bus data lines and an AddressValid bus line carry address, data and control information for memory addresses up to 40 bits wide. Persons skilled in the art will recognize that 16 bus data lines or other numbers of bus data lines can be used to implement the teaching of this invention. The new bus is used to connect elements such as memory, peripheral, switch and processing units.

[0023] In the system of this invention, DRAMs and other devices receive address and control information over the bus and transmit or receive requested data over the same bus. Each memory device contains only a single bus interface with no other signal pins. Other devices that may be included in the system can connect to the bus and other non-bus lines, such as input/output lines. The bus supports large data block transfers and split transactions to allow a user to achieve high bus utilization. This ability to rapidly read or write a large block of data to one single device at a time is an important advantage of this invention.

[0024] The DRAMs that connect to this bus differ from conventional DRAMs in a number of ways. Registers are

Micron Technology Inc. et al.
Ex. 1010, p. 151

provided which may store control information, device identification, device-type and other information appropriate for the chip such as the address range for each independent portion of the device. New bus interface circuits must be added and the internals of prior art DRAM devices need to be modified so they can provide and accept data to and from the bus at the peak data rate of the bus. This requires changes to the column access circuitry in the DRAM, with only a minimal increase in die size. A circuit is provided to generate a low skew internal device clock for devices on the bus, and other circuits provide for demultiplexing input and multiplexing output signals.

[0025] High bus bandwidth is achieved by running the bus at a very high clock rate (hundreds of MHz). This high clock rate is made possible by the constrained environment of the bus. The bus lines are controlled-impedance, doubly-terminated lines. For a data rate of 500 MHz, the maximum bus propagation time is less than 1 ns (the physical bus length is about 10 cm). In addition, because of the packaging used, the pitch of the pins can be very close to the pitch of the pads. The loading on the bus resulting from the individual devices is very small. In a preferred implementation, this generally allows stub capacitances of 1-2 pF and inductances of 0.5 - 2 nH. Each device 15, 16, 17, shown in Figure 3, only has pins on one side and these pins connect directly to the bus 18. A transceiver device 19 can be included to interface multiple units to a higher order bus through pins 20.

[0026] A primary result of the architecture of this invention is to increase the bandwidth of DRAM access. The invention also reduces manufacturing and production costs, power consumption, and increases packing density and system reliability.

BRIEF DESCRIPTION OF THE DRAWINGS

[0027]

Figure 1 is a diagram which illustrates the basic 2-D organization of memory devices.
Figure 2 is a schematic block diagram which illustrates the parallel connection of all bus lines and the serial Reset line to each device in the system.
Figure 3 is a perspective view of a system of the invention which illustrates the 3-D packaging of semiconductor devices on the primary bus.
Figure 4 shows the format of a request packet.
Figure 5 shows the format of a retry response from a slave.
Figure 6 shows the bus cycles after a request packet collision occurs on the bus and how arbitration is handled.
Figure 7 shows the timing whereby signals from two devices can overlap temporarily and drive the bus at the same time.
Figure 8 shows the connection and timing between bus clocks and devices on the bus.
Figure 9 is a perspective view showing how transceivers can be used to connect a number of bus units to a transceiver bus. Figure 10 is a block and schematic diagram of input/output circuitry used to connect devices to the bus.
Figure 11 is a schematic diagram of a clocked sense-amplifier used as a bus input receiver.
Figure 12 is a block diagram showing how the internal device clock is generated from two bus clock signals using a set of adjustable delay lines.
Figure 13 is a timing diagram showing the relationship of signals in the block diagram of Figure 12.
Figure 14 is timing diagram of a preferred means of implementing the reset procedure of this invention.
Figure 15 is a diagram illustrating the general organization of a 4 Mbit DRAM divided into 8 subarrays.

DETAILED DESCRIPTION

[0028] The present invention is designed to provide a high speed, multiplexed bus for communication between processing devices and memory devices and to provide devices adapted for use in the bus system. The invention can also be used to connect processing devices and other devices, such as I/O interfaces or disk controllers, with or without memory devices on the bus. The bus consists of a relatively small number of lines connected in parallel to each device on the bus. The'bus carries substantially all address, data and control information needed by devices for communication with other devices on the bus. In many systems using the present invention, the bus carries almost every signal between every device in the entire system. There is no need for separate device-select lines since device-select information for each device on the bus is carried over the bus. There is no need for separate address and data lines because address and data information can be sent over the same lines. Using the organization described herein, very large addresses (40 bits in the preferred implementation) and large data blocks (1024 bytes) can be sent over a small number of bus lines (8 plus one control line in the preferred implementation).

[0029] Virtually all of the signals needed by a computer system can be sent over the bus. Persons skilled in the art recognize that certain devices, such as CPUs, may be connected to other signal lines and possibly to independent buses, for example a bus to an independent cache memory, in addition to the bus of this invention. Certain devices, for

Micron Technology Inc. et al.
Ex. 1010, p. 152

example cross-point switches, could be connected to multiple, independent buses of this invention. In the preferred implementation, memory devices are provided that have no connections other than the bus connections described herein and CPUs are provided that use the bus of this invention as the principal, if not exclusive, connection to memory and to other devices on the bus.

[0030] All modern DRAM, SRAM and ROM designs have internal architectures with row (word) and column (bit) lines to efficiently tile a 2-D area. Referring to Fig. 1, one bit of data is stored at the intersection of each word line 5 and bit line 6. When a particular word line is enabled, all of the corresponding data bits are transferred onto the bit lines. This data, about 4000 bits at a time in a 4 MBit DRAM, is then loaded into column sense amplifiers 3 and held for use by the I/O circuits.

[0031] In the invention presented here, the data from the sense amplifiers is enabled 32 bits at a time onto an internal device bus running at approximately 125 MHz. This internal device bus moves the data to the periphery of the devices where the data is multiplexed into an 8-bit wide external bus interface, running at approximately 500 MHz.

[0032] The bus architecture of this invention connects master or bus controller devices, such as CPUs, Direct Memory Access devices (DMAs) or Floating Point Units (FPUs), and slave devices, such as DRAM, SRAM or ROM memory devices. A slave device responds to control signals; a master sends control signals. Persons skilled in the art realize that some devices may behave as both master and slave at various times, depending on the mode of operation and the state of the system. For example, a memory device will typically have only slave functions, while a DMA controller, disk controller or CPU may include both slave and master functions. Many other semiconductor devices, including I/O devices, disk controllers, or other special purpose devices such as high speed switches can be modified for use with the bus of this invention.

[0033] Each semiconductor device contains a set of internal registers, preferably including a device identification (device ID) register, a device-type descriptor register, control registers and other registers containing other information relevant to that type of device. In a preferred implementation, semiconductor devices connected to the bus contain registers which specify the memory addresses contained within that device and access-time registers which store a set of one or more delay times at which the device can or should be available to send or receive data.

[0034] Most of these registers can be modified and preferably are set as part of an initialization sequence that occurs when the system is powered up or reset. During the initialization sequence each device on the bus is assigned a unique device ID number, which is stored in the device ID register. A bus master can then use these device ID numbers to access and set appropriate registers in other devices, including access-time registers, control registers, and memory registers, to configure the system. Each slave may have one or several access-time registers (four in a preferred embodiment). In a preferred embodiment, one access-time register in each slave is permanently or semi-permanently programmed with a fixed value to facilitate certain control functions. A preferred implementation of an initialization sequence is described below in more detail.

[0035] All information sent between master devices and slave devices is sent over the external bus, which, for example, may be 8 bits wide. This is accomplished by defining a protocol whereby a master device, such as a microprocessor, seizes exclusive control of the external bus (i.e., becomes the bus master) and initiates a bus transaction by sending a request packet (a sequence of bytes comprising address and control information) to one or more slave devices on the bus. An address can consist of 16 to 40 or more bits according to the teachings of this invention. Each slave on the bus must decode the request packet to see if that slave needs to respond to the packet. The slave that the packet is directed to must then begin any internal processes needed to carry out the requested bus transaction at the requested time. The requesting master may also need to transact certain internal processes before the bus transaction begins. After a specified access time the slave(s) respond by returning one or more bytes (8 bits) of data or by storing information made available from the bus. More than one access time can be provided to allow different types of responses to occur at different times.

[0036] A request packet and the corresponding bus access are separated by a selected number of bus cycles, allowing the bus to be used in the intervening bus cycles by the same or other masters for additional requests or brief bus accesses. Thus multiple, independent accesses are permitted, allowing maximum utilization of the bus for transfer of short blocks of data. Transfers of long blocks of data use the bus efficiently even without overlap because the overhead due to bus address, control and access times is small compared to the total time to request and transfer the block.

Device Address Mapping

[0037] Another unique aspect of this invention is that each memory device is a complete, independent memory sub-system with all the functionality of a prior art memory board in a conventional backplane-bus computer system. Individual memory devices may contain a single memory section or may be subdivided into more than one discrete memory section. Memory devices preferably include memory address registers for each discrete memory section. A failed memory device (or even a subsection of a device) can be "mapped out" with only the loss of a small fraction of the memory, maintaining essentially full system capability. Mapping out bad devices can be accomplished in two ways, both compatible with this

Micron Technology Inc. et al.
Ex. 1010, p. 153

invention.

[0038] The preferred method uses address registers in each memory device (or independent discrete portion thereof) to store information which defines the range of bus addresses to which this memory device will respond. This is similar to prior art schemes used in memory boards in conventional backplane bus systems. The address registers can include a single pointer, usually pointing to a block of known size, a pointer and a fixed or variable block size value or two pointers, one pointing to the beginning and one to the end (or to the "top" and "bottom") of each memory block. By appropriate settings of the address registers, a series of functional memory devices or discrete memory sections can be made to respond to a contiguous range of addresses, giving the system access to a contiguous block of good memory, limited primarily by the number of good devices connected to the bus. A block of memory in a first memory device or memory section can be assigned a certain range of addresses, then a block of memory in a next memory device or memory section can be assigned addresses starting with an address one higher (or lower, depending on the memory structure) than the last address of the previous block.

[0039] Preferred devices for use in this invention include device-type register information specifying the type of chip, including how much memory is available in what configuration on that device. A master can perform an appropriate memory test, such as reading and writing each memory cell in one or more selected orders, to test proper functioning of each accessible discrete portion of memory (based in part on information like device ID number and device-type) and write address values (up to 40 bits in the preferred embodiment, 10^{12} bytes), preferably contiguous, into device address-space registers. Non-functional or impaired memory sections can be assigned a special address value which the system can interpret to avoid using that memory.

[0040] The second approach puts the burden of avoiding the bad devices on the system master or masters. CPUs and DMA controllers typically have some sort of translation look-aside buffers (TLBs) which map virtual to physical (bus) addresses. With relatively simple software, the TLBs can be programmed to use only working memory (data structures describing functional memories are easily generated). For masters which don't contain TLBs (for example, a video display generator), a small, simple RAM can be used to map a contiguous range of addresses onto the addresses of the functional memory devices.

[0041] Either scheme works and permits a system to have a significant percentage of non-functional devices and still continue to operate with the memory which remains. This means that systems built with this invention will have much improved reliability over existing systems, including the ability to build systems with almost no field failures.

Bus

[0042] The preferred bus architecture of this invention comprises 11 signals: BusData[0:7]; AddrValid; Clkl and Clk2; plus an input reference level and power and ground lines connected in parallel to each device. Signals are driven onto the bus during conventional bus cycles. The notation "Signal[i:j]" refers to a specific range of signals or lines, for example, BusData[0:7] means BusData0, BusData1, ..., BusData7. The bus lines for BusData[0:7] signals form a byte-wide, multiplexed data/address/control bus. AddrValid is used to indicate when the bus is holding a valid address request, and instructs a slave to decode the bus data as an address and, if the address is included on that slave, to handle the pending request. The two clocks together provide a synchronized, high speed clock for all the devices on the bus. In addition to the bused signals, there is one other line (ResetIn, ResetOut) connecting each device in series for use during initialization to assign every device in the system a unique device ID number (described below in detail).

[0043] To facilitate the extremely high data rate of this external bus relative to the gate delays of the internal logic, the bus cycles are grouped into pairs of even/odd cycles. Note that all devices connected to a bus should preferably use the same even/odd labeling of bus cycles and preferably should begin operations on even cycles. This is enforced by the clocking scheme.

Protocol and Bus Operation

[0044] The bus uses a relatively simple, synchronous, split-transaction, block-oriented protocol for bus transactions. One of the goals of the system is to keep the intelligence concentrated in the masters, thus keeping the slaves as simple as possible (since there are typically many more slaves than masters). To reduce the complexity of the slaves, a slave should preferably respond to a request in a specified time, sufficient to allow the slave to begin or possibly complete a device-internal phase including any internal actions that must precede the subsequent bus access phase. The time for this bus access phase is known to all devices on the bus - each master being responsible for making sure that the bus will be free when the bus access begins. Thus the slaves never worry about arbitrating for the bus. This approach eliminates arbitration in single master systems, and also makes the slave-bus interface simpler.

[0045] In a preferred implementation of the invention, to initiate a bus transfer over the bus, a master sends out a request packet, a contiguous series of bytes containing address and control information. It is preferable to use a request packet containing an even number of bytes and also preferable to start each packet on an even bus cycle.

Micron Technology Inc. et al.
Ex. 1010, p. 154

[0046] The device-select function is handled using the bus data lines. AddrValid is driven, which instructs all slaves to decode the request packet address, determine whether they contain the requested address, and if they do, provide the data back to the master (in the case of a read request) or accept data from the master (in the case of a write request) in a data block transfer. A master can also select a specific device by transmitting a device ID number in a request packet. In a preferred implementation, a special device ID number is chosen to indicate that the packet should be interpreted by all devices on the bus. This allows a master to broadcast a message, for example to set a selected control register of all devices with the same value.

[0047] The data block transfer occurs later at a time specified in the request packet control information, preferably beginning on an even cycle. A device begins a data block transfer almost immediately with a device-internal phase as the device initiates certain functions, such as setting up memory addressing, before the bus access phase begins. The time after which a data block is driven onto the bus lines is selected from values stored in slave access-time registers. The timing of data for reads and writes is preferably the same; the only difference is which device drives the bus. For reads, the slave drives the bus and the master latches the values from the bus. For writes the master drives the bus and the selected slave latches the values from the bus.

[0048] In a preferred implementation of this invention shown in Figure 4, a request packet 22 contains 6 bytes of data -- 4.5 address bytes and 1.5 control bytes. Each request packet uses all nine bits of the multiplexed data/address lines (AddrValid 23 + BusData[0:7] 24) for all six bytes of the request packet. Setting 23 AddrValid = 1 in an otherwise unused even cycle indicates the start of an request packet (control information). In a valid request packet, AddrValid 27 must be 0 in the last byte. Asserting this signal in the last byte invalidates the request packet. This is used for the collision detection and arbitration logic (described below). Bytes 25-26 contain the first 35 address bits, Address[0:35]. The last byte contains AddrValid 27 (the invalidation switch) and 28, the remaining address bits, Address[36:39], and BlockSize [0:3] (control information).

[0049] The first byte contains two 4 bit fields containing control information, AccessType[0:3], an op code (operation code) which, for example, specifies the type of access, and Master[0:3], a position reserved for the master sending the packet to include its master ID number. Only master numbers 1 through 15 are allowed - master number 0 is reserved for special system commands. Any packet with Master[0:3] = 0 is an invalid or special packet and is treated accordingly.

[0050] The AccessType field specifies whether the requested operation is a read or write and the type of access, for example, whether it is to the control registers or other parts of the device, such as memory. In a preferred implementation, AccessType[0] is a Read/Write switch: if it is a 1, then the operation calls for a read from the slave (the slave to read the requested memory block and drive the memory contents onto the bus); if it is a 0, the operation calls for a write into the slave (the slave to read data from the bus and write it to memory). AccessType[1:3] provides up to 8 different access types for a slave. AccessType[1:2] preferably indicates the timing of the response, which is stored in an access-time register, AccessRegN. The choice of access-time register can be selected directly by having a certain op code select that register, or indirectly by having a slave respond to selected op codes with pre-selected access times (see table below). The remaining bit, AccessType[3] may be used to send additional information about the request to the slaves.

[0051] One special type of access is control register access, which involves addressing a selected register in a selected slave. In the preferred implementation of this invention, AccessType[1:3] equal to zero indicates a control register request and the address field of the packet indicates the desired control register. For example, the most significant two bytes can be the device ID number (specifying which slave is being addressed) and the least significant three bytes can specify a register address and may also represent or include data to be loaded into that control register. Control register accesses are used to initialize the access-time registers, so it is preferable to use a fixed response time which can be preprogrammed or even hard wired, for example the value in AccessReg0, preferably 8 cycles. Control register access can also be used to initialize or modify other registers, including address registers.

[0052] The method of this invention provides for access mode control specifically for the DRAMs. One such access mode determines whether the access is page mode or normal RAS access. In normal mode (in conventional DRAMS and in this invention), the DRAM column sense amps or latches have been precharged to a value intermediate between logical 0 and 1. This precharging allows access to a row in the RAM to begin as soon as the access request for either inputs (writes) or outputs (reads) is received and allows the column sense amps to sense data quickly. In page mode (both conventional and in this invention), the DRAM holds the data in the column sense amps or latches from the previous read or write operation. If a subsequent request to access data is directed td the same row, the DRAM does not need to wait for the data to be sensed (it has been sensed already) and access time for this data is much shorter than the normal access time. Page mode generally allows much faster access to data but to a smaller block of data (equal to the number of sense amps). However, if the requested data is not in the selected row, the access time is longer than the normal access time, since the request must wait for the RAM to precharge before the normal mode access can start. Two access-time registers in each DRAM preferably contain the access times to be used for normal and for page-mode accesses, respectively.

[0053] The access mode also determines whether the DRAM should precharge the sense amplifiers or should save the contents of the sense amps for a subsequent page mode access. Typical settings are "precharge after normal

Micron Technology Inc. et al.
Ex. 1010, p. 155

access" and "save after page mode access" but "precharge after page mode access" or "save after normal access" are allowed, selectable modes of operation. The DRAM can also be set to precharge the sense amps if they are not accessed for a selected period of time.

[0054] In page mode, the data stored in the DRAM sense amplifiers may be accessed within much less time than it takes to read out data in normal mode (~10-20 nS vs. 40-100 nS). This data may be kept available for long periods. However, if these sense amps (and hence bit lines) are not precharged after an access, a subsequent access to a different memory word (row) will suffer a precharge time penalty of about 40-100 nS because the sense amps must precharge before latching in a new value.

[0055] The contents of the sense amps thus may be held and used as a cache, allowing faster, repetitive access to small blocks of data. DRAM-based page-mode caches have been attempted in the prior art using conventional DRAM organizations but they are not very effective because several chips are required per computer word. Such a conventional page-mode cache contains many bits (for example, 32 chips x 4Kbits) but has very few independent storage entries. In other words, at any given point in time the sense amps hold only a few different blocks or memory "locales" (a single block of 4K words, in the example above). Simulations have shown that upwards of 100 blocks are required to achieve high hit rates (>90% of requests find the requested data already in cache memory) regardless of the size of each block. *See, for* example, Anant Agarwal, et. al., "An Analytic Cache Model," ACM Transactions on Computer Systems, Vol. 7 (2), pp. 184-215 (May 1989).

[0056] The organization of memory in the present invention allows each DRAM to hold one or more (4 for 4MBit DRAMS) separately- addressed and independent blocks of data. A personal computer or workstation with 100 such DRAMs (i.e. 400 blocks or locales) can achieve extremely high, very repeatable hit rates (98-99% on average) as compared to the lower (50-80%), widely varying hit rates using DRAMS organized in the conventional fashion. Further, because of the time penalty associated with the deferred precharge on a "miss" of the page-mode cache, the conventional DRAM-based page-mode cache generally has been found to work less well than no cache at all.

[0057] For DRAM slave access, the access types are preferably used in the following way:

AccessType[1:3]	Use	AccessTime
0	Control Register Access	Fixed, 8[AccessReg0]
1	Unused	Fixed, 8[AccessReg0]
2-3	Unused	AccessReg1
4-5	Page Mode DRAM access	AccessReg2
6-7	Normal DRAM access	AccessReg3

Persons skilled in the art will recognize that a series of available bits could be designated as switches for controlling these access modes. For example:

AccessType[2] = page mode/normal switch

AccessType[3] = precharge/save-data switch

BlockSize[0:3] specifies the size of the data block transfer. If BlockSize[0] is 0, the remaining bits are the binary representation of the block size (0-7). If BlockSize[0] is 1, then the remaining bits give the block size as a binary power of 2, from 8 to 1024. A zero-length block can be interpreted as a special command, for example, to refresh a DRAM without returning any data, or to change the DRAM from page mode to normal access mode or vice-versa.

BlockSize[0:2]	Number of Bytes in Block
0-7	0-7 respectively
8	8
9	16
10	32
11	64
12	128
13	256
14	512
15	1024

Persons skilled in the art will recognize that other block size encoding schemes or values can be used.

[0058] In most cases, a slave will respond at the selected access time by reading or writing data from or to the bus over bus lines BusData[0:7] and AddrValid will be at logical 0. In a preferred embodiment, substantially each memory

Micron Technology Inc. et al.
Ex. 1010, p. 156

access will involve only a single memory device, that is, a single block will be read from or written to a single memory device.

Retry Format

[0059] In some cases, a slave may not be able to respond correctly to a request, e.g., for a read or write. In such a situation, the slave should return an error message, sometimes called a N(o)ACK(nowledge) or retry message. The retry message can include information about the condition requiring a retry, but this increases system requirements for circuitry in both slave and masters. A simple message indicating only that an error has occurred allows for a less complex slave, and the master can take whatever action is needed to understand and correct the cause of the error.

[0060] For example, under certain conditions a slave might not be able to supply the requested data. During a page-mode access, the DRAM selected must be in page mode and the requested address must match the address of the data held in the sense amps or latches. Each DRAM can check for this match during a page-mode access. If no match is found, the DRAM begins precharging and returns a retry message to the master during the first cycle of the data block (the rest of the returned block is ignored). The master then must wait for the precharge time (which is set to accommodate the type of slave in question, stored in a special register, PreChargeReg), and then resend the request as a normal DRAM access (AccessType = 6 or 7).

[0061] In the preferred form of the present invention, a slave signals a retry by driving AddrValid true at the time the slave was supposed to begin reading or writing data. A master which expected to write to that slave must monitor AddrValid during the write and take corrective action if it detects a retry message. Figure 5 illustrates the format of a retry message 28 which is useful for read requests, consisting of 23 AddrValid=1 with Master[0:3] = 0 in the first (even) cycle. Note that AddrValid is normally 0 for data block transfers and that there is no master 0 (only 1 through 15 are allowed). All DRAMs and masters can easily recognize such a packet as an invalid request packet, and therefore a retry message. In this type of bus transaction all of the fields except for Mastery[0:3] and AddrValid 23 may be used as information fields, although in the implementation described, the contents are undefined. Persons skilled in the art recognize that another method of signifying a retry message is to add a DataInvalid line and signal to the bus. This signal could be asserted in the case of a NACK.

Bus Arbitration

[0062] In the case of a single master, there are by definition no arbitration problems. The master sends request packets and keeps track of periods when the bus will be busy in response to that packet. The master can schedule multiple requests so that the corresponding data block transfers do not overlap.

[0063] The bus architecture of this invention is also useful in configurations with multiple masters. When two or more masters are on the same bus, each master must keep track of all the pending transactions, so each master knows when it can send a request packet and access the corresponding data block transfer. Situations will arise, however, where two or more masters send a request packet at about the same time and the multiple requests must be detected, then sorted out by some sort of bus arbitration.

[0064] There are many ways for each master to keep track of when the bus is and will be busy. A simple method is for each master to maintain a bus-busy data structure, for example by maintaining two pointers, one to indicate the earliest point in the future when the bus will be busy and the other to indicate the earliest point in the future when the bus will be free, that is, the end of the latest pending data block transfer. Using this information, each master can determine whether and when there is enough time to send a request packet (as described above under Protocol) before the bus becomes busy with another data block transfer and whether the corresponding data block transfer will interfere with pending bus transactions. Thus each master must read every request packet and update its bus-busy data structure to maintain information about when the bus is and will be free.

[0065] With two or more masters on the bus, masters will occasionally transmit independent request packets during the same bus cycle. Those multiple requests will collide as each such master drives the bus simultaneously with different information, resulting in scrambled request information and neither desired data block transfer. In a preferred form of the invention, each device on the bus seeking to write a logical 1 on a BusData or AddrValid line drives that line with a current sufficient to sustain a voltage greater than or equal to the high-logic value for the system. Devices do not drive lines that should have a logical 0; those lines are simply held at a voltage corresponding to a low-logic value. Each master tests the voltage on at least some, preferably all, bus data and the AddrValid lines so the master can detect a logical '1' where the expected level is '0' on a line that it does not drive during a given bus cycle but another master does drive.

[0066] Another way to detect collisions is to select one or more bus lines for collision signalling. Each master sending a request drives that line or lines and monitors the selected lines for more than the normal drive current (or a logical value of ">1"), indicating requests by more than one master. Persons skilled in the art will recognize that this can be implemented with a protocol involving BusData and AddrValid lines or could be implemented using an additional bus line.

Micron Technology Inc. et al.
Ex. 1010, p. 157

[0067] In the preferred form of this invention, each master detects collisions by monitoring lines which it does not drive to see if another master is driving those lines. Referring to Fig. 4, the first byte of the request packet includes the number of each master attempting to use the bus (Master[0:3]). If two masters send packet requests starting at the same point in time, the master numbers will be logical "or"ed together by at least those masters, and thus one or both of the masters, by monitoring the data on the bus and comparing what it sent, can detect a collision. For instance if requests by masters number 2 (0010) and 5 (0101) collide, the bus will be driven with the value Master[0:3]=7 (0010 + 0101 = 0111). Master number 5 will detect that the signal Master[2] = 1 and master 2 will detect that Master[1] and Master[3] = 1, telling both masters that a collision has occurred. Another example is masters 2 and 11, for which the bus will be driven with the value Master[0:3]=11 (0010 + 1011 = 1011), and although master 11 can't readily detect this collision, master 2 can. When any collision is detected, each master detecting a collision drives the value of AddrValid 27 in byte 5 of the request packet 22 to 1, which is detected by all masters, including master 11 in the second example above, and forces a bus arbitration cycle, described below.

[0068] Another collision condition may arise where master A sends a request packet in cycle 0 and master B tries to send a request packet starting in cycle 2 of the first request packet, thereby overlapping the first request packet. This will occur from time to time because the bus operates at high speeds, thus the logic in a second-initiating master may not be fast enough to detect a request initiated by a first master in cycle 0 and to react fast enough by delaying its own request. Master B eventually notices that it wasn't supposed to try to send a request packet (and consequently almost surely destroyed the address that master A was trying to send), and, as in the example above of a simultaneous collision, drives a 1 on AddrValid during byte 5 of the first request packet 27 forcing an arbitration. The logic in the preferred implementation is fast enough that a master should detect a request packet by another master by cycle 3 of the first request packet, so no master is likely to attempt to send a potentially colliding request packet later than cycle 2.

[0069] Slave devices not need to detect a collision directly, but they must wait to do anything irrecoverable until the last byte (byte 5) is read to ensure that the packet is valid. A request packet with Master[0:3] equal to 0 (a retry signal) is ignored and does not cause a collision. The subsequent bytes of such a packet are ignored.

[0070] To begin arbitration after a collision, the masters wait a preselected number of cycles after the aborted request packet (4 cycles in a preferred implementation), then use the next free cycle to arbitrate for the bus (the next available even cycle in the preferred implementation). Each colliding master signals to all other colliding masters that it seeks to send a request packet, a priority is assigned to each of the colliding masters, then each master is allowed to make its request in the order of that priority.

[0071] Figure 6 illustrates one preferred way of implementing this arbitration. Each colliding master signals its intent to send a request packet by driving a single BusData line during a single bus cycle corresponding to its assigned master number (1-15 in the present example). During two-byte arbitration cycle 29, byte 0 is allocated to requests 1-7 from masters 1-7, respectively, (bit 0 is not used) and byte 1 is allocated to requests 8-15 from masters 8-15, respectively. At least one device and preferably each colliding master reads the values on the bus during the arbitration cycles to determine and store which masters desire to use the bus. Persons skilled in the art will recognize that a single byte can be allocated for arbitration requests if the system includes more bus lines than masters. More than 15 masters can be accommodated by using additional bus cycles.

[0072] A fixed priority scheme (preferably using the master numbers, selecting lowest numbers first) is then used to prioritize, then sequence the requests in a bus arbitration queue which is maintained by at least one device. These requests are queued by each master in the bus-busy data structure and no further requests are allowed until the bus arbitration queue is cleared. Persons skilled in the art will recognize that other priority schemes can be used, including assigning priority according to the physical location of each master.

System Configuration/Reset

[0073] In the bus-based system of this invention, a mechanism is provided to give each device on the bus a unique device identifier (device ID) after power-up or under other conditions as desired or needed by the system. A master can then use this device ID to access a specific device, particularly to set or modify registers of the specified device, including the control and address registers. In the preferred embodiment, one master is assigned to carry out the entire system configuration process. The master provides a series of unique device ID numbers for each unique device connected to the bus system. In the preferred embodiment, each device connected to the bus contains a special device-type register which specifies the type of device, for instance CPU, 4 MBit memory, 64 MBit memory or disk controller. The configuration master should check each device, determine the device type and set appropriate control registers, including access-time registers. The configuration master should check each memory device and set all appropriate memory address registers.

[0074] One means to set up unique device ID numbers is to have each device to select a device ID in sequence and store the value in an internal device ID register. For example, a master can pass sequential device ID numbers through shift registers in each of a series of devices, or pass a token from device to device whereby the device with the token

Micron Technology Inc. et al.
Ex. 1010, p. 158

reads in device ID information from another line or lines. In a preferred embodiment, device ID numbers are assigned to devices according to their physical relationship, for instance, their order along the bus.

[0075] In a preferred embodiment of this invention, the device ID setting is accomplished using a pair of pins on each device, ResetIn and ResetOut. These pins handle normal logic signals and are used only during device ID configuration. On each rising edge of the clock, each device copies ResetIn (an input) into a four-stage reset shift register. The output of the reset shift register is connected to ResetOut, which in turn connects to ResetIn for the next sequentially connected device. Substantially all devices on the bus are thereby daisy-chained together. A first reset signal, for example, while ResetIn at a device is a logical 1, or when a selected bit of the reset shift register goes from zero to non-zero, causes the device to hard reset, for example by clearing all internal registers and resetting all state machines. A second reset signal, for example, the falling edge of ResetIn combined with changeable values on the external bus, causes that device to latch the contents of the external bus into the internal device ID register (Device[0:7]).

[0076] To reset all devices on a bus, a master sets the ResetIn line of the first device to a "1" for long enough to ensure that all devices on the bus have been reset (4 cycles times the number of devices -- note that the maximum number of devices on the preferred bus configuration is 256 (8 bits), so that 1024 cycles is always enough time to reset all devices.) Then ResetIn is dropped to "0" and the BusData lines are driven with the first followed by successive device ID numbers, changing after every 4 clock pulses. Successive devices set those device ID numbers into the corresponding device ID register as the falling edge of ResetIn propagates through the shift registers of the daisy-chained devices. Figure 14 shows ResetIn at a first device going low while a master drives a first device ID onto the bus data lines BusData[0:3]. The first device then latches in that first device ID. After four clock cycles, the master changes BusData[0:3] to the next device ID number and ResetOut at the first device goes low, which pulls ResetIn for the next daisy-chained device low, allowing the next device to latch in the next device ID number from BusData[0:3]. In the preferred embodiment, one master is assigned device ID 0 and it is the responsibility of that master to control the ResetIn line and to drive successive device ID numbers onto the bus at the appropriate times. In the preferred embodiment, each device waits two clock cycles after ResetIn goes low before latching in a device ID number from BusData[0:3].

[0077] Persons skilled in the art recognize that longer device ID numbers could be distributed to devices by having each device read in multiple bytes from the bus and latch the values into the device ID register. Persons skilled in the art also recognize that there are alternative ways of getting device ID numbers to unique devices. For instance, a series of sequential numbers could be clocked along the ResetIn line and at a certain time each device could be instructed to latch the current reset shift register value into the device ID register.

[0078] The configuration master should choose and set an access time in each access-time register in each slave to a period sufficiently long to allow the slave to perform an actual, desired memory access. For example, for a normal DRAM access, this time must be longer than the row address strobe (RAS) access time. If this condition is not met, the slave may not deliver the correct data. The value stored in a slave access-time register is preferably one-half the number of bus cycles for which the slave device should wait before using the bus in response to a request. Thus an access time value of '1' would indicate that the slave should not access the bus until at least two cycles after the last byte of the request packet has been received. The value of AccessReg0 is preferably fixed at 8 (cycles) to facilitate access to control registers.

[0079] The bus architecture of this invention can include more than one master device. The reset or initialization sequence should also include a determination of whether there are multiple masters on the bus, and if so to assign unique master ID numbers to each. Persons skilled in the art will recognize that there are many ways of doing this. For instance, the master could poll each device to determine what type of device it is, for example, by reading a special register then, for each master device, write the next available master ID number into a special register.

ECC

[0080] Error detection and correction ("ECC") methods well known in the art can be implemented in this system. ECC information typically is calculated for a block of data at the time that block of data is first written into memory. The data block usually has an integral binary size, e.g. 256 bits, and the ECC information uses significantly fewer bits. A potential problem arises in that each binary data block in prior art schemes typically is stored with the ECC bits appended, resulting in a block size that is not an integral binary power.

[0081] In a preferred embodiment of this invention, ECC information is stored separately from the corresponding data, which can then be stored in blocks having integral binary size. ECC information and corresponding data can be stored, for example, in separate DRAM devices. Data can be read without ECC using a single request packet, but to write or read error-corrected data requires two request packets, one for the data and a second for the corresponding ECC information. ECC information may not always be stored permanently and in some situations the ECC information may be available without sending a request packet or without a bus data block transfer.

[0082] In a preferred embodiment, a standard data block size can be selected for use with ECC, and the ECC method will determine the required number of bits of information in a corresponding ECC block. RAMs containing ECC information

Micron Technology Inc. et al.
Ex. 1010, p. 159

can be programmed to store an access time that is equal to: (1) the access time of the normal RAM (containing data) plus the time to access a standard data block (for corrected data) minus the time to send a request packet (6 bytes); or (2) the access time of a normal RAM minus the time to access a standard ECC block minus the time to send a request packet. To read a data block and the corresponding ECC block, the master simply issues a request for the data immediately followed by a request for the ECC block. The ECC RAM will wait for the selected access time then drive its data onto the bus right after (in case (1) above)) the data RAM has finished driving out the data block. Persons skilled in the art will recognize that the access time described in case (2) above can be used to drive ECC data before the data is driven onto the bus lines and will recognize that writing data can be done by analogy with the method described for a read. Persons skilled in the art will also recognize the adjustments that must be made in the bus-busy structure and the request packet arbitration methods of this invention in order to accommodate these paired ECC requests.

[0083] Since this system is quite flexible, the system designer can choose the size of the data blocks and the number of ECC bits using the memory devices of this invention. Note that the data stream on the bus can be interpreted in various ways. For instance the sequence can be 2^n data bytes followed by 2^m ECC bytes (or vice versa), or the sequence can be 2^k iterations of 8 data bytes plus 1 ECC byte. Other information, such as information used by a directory-based cache coherence scheme, can also be managed this way. See, for example, Anant Agarwal, et al., "Scaleable Directory Schemes for Cache Consistency," 15th International Symposium on Computer Architecture, June 1988, pp. 280-289. Those skilled in the art will recognize alternative methods of implementing ECC schemes that are within the teachings of this invention.

Low Power 3-D Packaging

[0084] Another major advantage of this invention is that it drastically reduces the memory system power consumption. Nearly all the power consumed by a prior art DRAM is dissipated in performing row access. By using a single row access in a single RAM to supply all the bits for a block request (compared to a row-access in each of multiple RAMs in conventional memory systems) the power per bit can be made very small. Since the power dissipated by memory devices using this invention is significantly reduced, the devices potentially can be placed much closer together than with conventional designs.

[0085] The bus architecture of this invention makes possible an innovative 3-D packaging technology. By using a narrow, multiplexed (time-shared) bus, the pin count for an arbitrarily large memory device can be kept quite small - on the order of 20 pins. Moreover, this pin count can be kept constant from one generation of DRAM density to the next. The low power dissipation allows each package to be smaller, with narrower pin pitches (spacing between the IC pins). With current surface mount technology supporting pin pitches as low as 20 mils, all off-device connections can be implemented on a single edge of the memory device. Semiconductor die useful in this invention preferably have connections or pads along one edge of the die which can then be wired or otherwise connected to the package pins with wires having similar lengths. This geometry also allows for very short leads, preferably with an effective lead length of less than 4 mm. Furthermore, this invention uses only bused interconnections, i.e., each pad on each device is connected by the bus to the corresponding pad of each other device.

[0086] The use of a low pin count and an edge-connected bus permits a simple 3-D package, whereby the devices are stacked and the bus is connected along a single edge of the stack. The fact that all of the signals are bused is important for the implementation of a simple 3-D structure. Without this, the complexity of the "backplane" would be too difficult to make cost effectively with current technology. The individual devices in a stack of the present invention can be packed quite tightly because of the low power dissipated by the entire memory system, permitting the devices to be stacked bumper-to-bumper or top to bottom. Conventional plastic-injection molded small outline (SO) packages can be used with a pitch of about 2.5 mm (100 mils), but the ultimate limit would be the device die thickness, which is about an order of magnitude smaller, 0.2-0.5 mm using current wafer technology.

Bus Electrical Description

[0087] By using devices with very low power dissipation and close physical packing, the bus can be made quite short, which in turn allows for short propagation times and high data rates. The bus of a preferred embodiment of the present invention consists of a set of resistor-terminated controlled impedance transmission lines which can operate up to a data rate of 500 MHz (2 ns cycles). The characteristics of the transmission lines are strongly affected by the loading caused by the DRAMs (or other slaves) mounted on the bus. These devices add lumped capacitance to the lines which both lowers the impedance of the lines and decreases the transmission speed. In the loaded environment, the bus impedance is likely to be on the order of 25 ohms and the propagation velocity about $c/4$ (c = the speed of light) or 7.5 cm/ns. To operate at a 2 ns data rate, the transit time on the bus should preferably be kept under 1 ns, to leave 1 ns for the setup and hold time of the input receivers (described below) plus clock skew. Thus the bus lines must be kept quite short, under about 8 cm for maximum performance. Lower performance systems may have much longer lines, e.g. a 4 ns bus

Micron Technology Inc. et al.
Ex. 1010, p. 160

may have 24 cm lines (3 ns transit time, 1 ns setup and hold time).

[0088] In the preferred embodiment, the bus uses current source drivers. Each output must be able to sink 50 mA, which provides an output swing of about 500 mV or more. In the preferred embodiment of this invention, the bus is active low. The unasserted state (the high value) is preferably considered a logical zero, and the asserted value (low state) is therefore a logical 1. Those skilled in the art understand that the method of this invention can also be implemented using the opposite logical relation to voltage. The value of the unasserted state is set by the voltage on the termination resistors, and should be high enough to allow the outputs to act as current sources, while being as low as possible to reduce power dissipation. These constraints may yield a termination voltage about 2V above ground in the preferred implementation. Current source drivers cause the output voltage to be proportional to the sum of the sources driving the bus.

[0089] Referring to Fig.7, although there is no stable condition where two devices drive the bus at the same time, conditions can arise because of propagation delay on the wires where one device, A 41, can start driving its part of the bus 44 while the bus is still being driven by another device, B 42 (already asserting a logical 1 on the bus). In a system using current drivers, when B 42 is driving the bus (before time 46), the value at points 44 and 45 is logical 1. If B 42 switches off at time 46 just when A 41 switches on, the additional drive by device A 41 causes the voltage at the output 44 of A 41 to drop briefly below the normal value. The voltage returns to its normal value at time 47 when the effect of device B 42 turning off is felt. The voltage at point 45 goes to logical 0 when device B 42 turns off, then drops at time 47 when the effect of device A 41 turning on is felt. Since the logical 1 driven by current from device A 41 is propagated irrespective of the previous value on the bus, the value on the bus is guaranteed to settle after one time of flight (t_f) delay, that is, the time it takes a signal to propagate from one end of the bus to the other. If a voltage drive was used (as in ECL wired-ORing), a logical 1 on the bus (from device B 42 being previously driven) would prevent the transition put out by device A 41 being felt at the most remote part of the system, e.g., device 43, until the turnoff waveform from device B 42 reached device A 41 plus one time of flight delay, giving a worst case settling time of twice the time of flight delay.

Clocking

[0090] Clocking a high speed bus accurately without introducing error due to propagation delays can be implemented by having each device monitor two bus clock signals and then derive internally a device clock, the true system clock. The bus clock information can be sent on one or two lines to provide a mechanism for each bused device to generate an internal device clock with zero skew relative to all the other device clocks. Referring to Figure 8, in the preferred implementation, a bus clock generator 50 at one end of the bus propagates an early bus clock signal in one direction along the bus, for example on line 53 from left to right, to the far end of the bus. The same clock signal then is passed through the direct connection shown to a second line 54, and returns as a late bus clock signal along the bus from the far end to the origin, propagating from right to left. A single bus clock line can be used if it is left unterminated at the far end of the bus, allowing the early bus clock signal to reflect back along the same line as a late bus clock signal.

[0091] Figure 8b illustrates how each device 51, 52 receives each of the two bus clock signals at a different time (because of propagation delay along the wires), with constant midpoint in time between the two bus clocks along the bus. At each device 51, 52, the rising edge 55 of Clock1 53 is followed by the rising edge 56 of Clock2 54. Similarly, the falling edge 57 of Clock1 53 is followed by the falling edge 58 of Clock2 54. This waveform relationship is observed at all other devices along the bus. Devices which are closer to the clock generator have a greater separation between Clock1 and Clock2 relative to devices farther from the generator because of the longer time required for each clock pulse to traverse the bus and return along line 54, but the midpoint in time 59, 60 between corresponding rising or falling edges is fixed because, for any given device, the length of each clock line between the far end of the bus and that device is equal. Each device must sample the two bus clocks and generate its own internal device clock at the midpoint of the two.

[0092] Clock distribution problems can be further reduced by using a bus clock and device clock rate equal to the bus cycle data rate divided by two, that is, the bus clock period is twice the bus cycle period. Thus a 500 MHz bus preferably uses a 250 MHz clock rate. This reduction in frequency provides two benefits. First it makes all signals on the bus have the same worst case data rates -- data on a 500 MHz bus can only change every 2 ns. Second, clocking at half the bus cycle data rate makes the labeling of the odd and even bus cycles trivial, for example, by defining even cycles to be those when the internal device clock is 0 and odd cycles when the internal device clock is 1.

Multiple Buses

[0093] The limitation on bus length described above restricts the total number of devices that can be placed on a single bus. Using 2.5 mm spacing between devices, a single 8 cm bus will hold about 32 devices. Persons skilled in the art will recognize certain applications of the present invention wherein the overall data rate on the bus is adequate but memory or processing requirements necessitate a much larger number of devices (many more than 32). Larger systems can easily be built using the teachings of this invention by using one or more memory subsystems, designated primary

Micron Technology Inc. et al.
Ex. 1010, p. 161

bus units, each of which consists of two or more devices, typically 32 or close to the maximum allowed by bus design requirements, connected to a transceiver device.

Referring to Figure 9, each primary bus unit can be mounted on a single circuit board 66, sometimes called a memory stick. Each transceiver device 19 in turn connects to a transceiver bus 65, similar or identical in electrical and other respects to the primary bus 18 described at length above. In a preferred implementation, all masters are situated on the transceiver bus so there are no transceiver delays between masters and all memory devices are on primary bus units so that all memory accesses experience an equivalent transceiver delay, but persons skilled in the art will recognize how to implement systems which have masters on more than one bus unit and memory devices on the transceiver bus as well as on primary bus units. In general, each teaching of this invention which refers to a memory device can be practiced using a transceiver device and one or more memory devices on an attached primay bus unit. Other devices, generically referred to as peripheral devices, including disk controllers, video controllers or I/O devices can also be attached to either the transceiver bus or a primary bus unit, as desired. Persons skilled in the art will recognize how to use a single primary bus unit or multiple primary bus units as needed with a transceiver bus in certain system designs.

[0094] The transceivers are quite simple in function. They detect request packets on the transceiver bus and transmit them to their primary bus unit. If the request packet calls for a write to a device on a transceiver's primary bus unit, that transceiver keeps track of the access time and block size and forwards all data from the transceiver bus to the primary bus unit during that time. The transceivers also watch their primary bus unit, forwarding any data that occurs there to the transceiver bus. The high speed of the buses means that the transceivers will need to be pipelined, and will require an additional one or two cycle delay for data to pass through the transceiver in either direction. Access times stored in masters on the transceiver bus must be increased to account for transceiver delay but access times stored in slaves on a primary bus unit should not be modified.

[0095] Persons skilled in the art will recognize that a more sophisticated transceiver can control transmissions to and from primary bus units. An additional control line, TrncvrRW can be bused to all devices on the transceiver bus, using that line in conjunction with the AddrValid line to indicate to all devices on the transceiver bus that the information on the data lines is: 1) a request packet, 2) valid data to a slave, 3) valid data from a slave, or 4) invalid data (or idle bus). Using this extra control line obviates the need for the transceivers to keep track of when data needs to be forwarded from its primary bus to the transceiver bus - all transceivers send all data from their primary bus to the transceiver bus whenever the control signal indicates condition 2) above. In a preferred implementation of this invention, if AddrValid and TrncvrRW are both low, there is no bus activity and the transceivers should remain in an idle state. A controller sending a request packet will drive AddrValid high, indicating to all devices on the transceiver bus that a request packet is being sent which each transceiver should forward to its primary bus unit. Each controller seeking to write to a slave should drive both AddrValid and TrncvrRW high, indicating valid data for a slave is present on the data lines. Each transceiver device will then transmit all data from the transceiver bus lines to each primary bus unit. Any controller expecting to receive information from a slave should also drive the TrncvrRW line high, but not drive AddrValid, thereby indicating to each transceiver to transmit any data coming from any slave on its primary local bus to the transceiver bus. A still more sophisticated transceiver would recognize signals addressed to or coming from its primary bus unit and transmit signals only at requested times.

[0096] An example of the physical mounting of the transceivers is shown in Figure 9. One important feature of this physical arrangement is to integrate the bus of each transceiver 19 with the original bus of DRAMs or other devices 15, 16, 17 on the primary bus unit 66. The transceivers 19 have pins on two sides, and are preferably mounted flat on the primary bus unit with a first set of pins connected to primary bus 18. A second set of transceiver pins 20, preferably orthogonal to the first set of pins, are oriented to allow the transceiver 19 to be attached to the transceiver bus 65 in much the same way as the DRAMs were attached to the primary bus unit. The transceiver bus can be generally planar and in a different plane, preferably orthogonal to the plane of each primary bus unit. The transceiver bus can also be generally circular with primary bus units mounted perpendicular and tangential to the transceiver bus.

[0097] Using this two level scheme allows one to easily build a system that contains over 500 slaves (16 buses of 32 DRAMs each). Persons skilled in the art can modify the device ID scheme described above to accommodate more than 256 devices, for example by using a longer device ID or by using additional registers to hold some of the device ID. This scheme can be extended in yet a third dimension to make a second-order transceiver bus, connecting multiple transceiver buses by aligning transceiver bus units parallel to and on top of each other and busing corresponding signal lines through a suitable transceiver. Using such a second-order transceiver bus, one could connect many thousands of slave devices into what is effectively a single bus.

Device Interface

[0098] The device interface to the high-speed bus can be divided into three main parts. The first part is the electrical interface. This part includes the input receivers, bus drivers and clock generation circuitry. The second part contains the address comparison circuitry and timing registers. This part takes the input request packet and determines if the request

Micron Technology Inc. et al.
Ex. 1010, p. 162

is for this device, and if it is, starts the internal access and delivers the data to the pins at the correct time. The final part, specifically for memory devices such as DRAMs, is the DRAM column access path. This part needs to provide bandwidth into and out of the DRAM sense amps greater than the bandwidth provided by conventional DRAMs. The implementation of the electrical interface and DRAM column access path are described in more detail in the following sections. Persons skilled in the art recognize how to modify prior-art address comparison circuitry and prior-art register circuitry in order to practice the present invention.

Electrical Interface - Input/Output Circuitry

[0099] A block diagram of the preferred input/output circuit for address/data/control lines is shown in Figure 10. This circuitry is particularly well-suited for use in DRAM devices but it can be used or modified by one skilled in the art for use in other devices connected to the bus of this invention. It consists of a set of input receivers 71, 72 and output driver 76 connected to input/output line 69 and pad 75 and circuitry to use the internal clock 73 and internal clock complement 74 to drive the input interface. The clocked input receivers take advantage of the synchronous nature of the bus. To further reduce the performance requirements for device input receivers, each device pin, and thus each bus line, is connected to two clocked receivers, one to sample the even cycle inputs, the other to sample the odd cycle inputs. By thus de-multiplexing the input 70 at the pin, each clocked amplifier is given a full 2 ns cycle to amplify the bus low-voltage-swing signal into a full value CMOS logic signal. Persons skilled in the art will recognize that additional clocked input receivers can be used within the teachings of this invention. For example, four input receivers could be connected to each device pin and clocked by a modified internal device clock to transfer sequential bits from the bus to internal device circuits, allowing still higher external bus speeds or still longer settling times to amplify the bus low-voltage-swing signal into a full value CMOS logic signal.

[0100] The output drivers are quite simple, and consist of a single NMOS pulldown transistor 76. This transistor is sized so that under worst case conditions it can still sink the 50 mA required by the bus. For 0.8 micron CMOS technology, the transistor will need to be about 200 microns long. Overall bus performance can be improved by using feedback techniques to control output transistor current so that the current through the device is roughly 50 mA under all operating conditions, although this is not absolutely necessary for proper bus operation. An example of one of many methods known to persons skilled in the art for using feedback techniques to control current is described in Hans Schumacher, et al., "CMOS Subnanosecond True-ECL Output Buffer," J. Solid State Circuits, Vol. 25 (1), pp. 150-154 (Feb. 1990). Controlling this current improves performance and reduces power dissipation. This output driver which can be operated at 500 MHz, can in turn be controlled by a suitable multiplexer with two or more (preferably four) inputs connected to other internal chip circuitry, all of which can be designed according to well known prior art.

[0101] The input receivers of every slave must be able to operate during every cycle to determine whether the signal on the bus is a valid request packet. This requirement leads to a number of constraints on the input circuitry. In addition to requiring small acquisition and resolution delays, the circuits must take little or no DC power, little AC power and inject very little current back into the input or reference lines. The standard clocked DRAM sense amp shown in Figure 11 satisfies all these requirements except the need for low input currents. When this sense amp goes from sense to sample, the capacitance of the internal nodes 83 and 84 in Figure11 is discharged through the reference line 68 and input 69, respectively. This particular current is small, but the sum of such currents from all the inputs into the reference lines summed over all devices can be reasonably large.

[0102] The fact that the sign of the current depends upon on the previous received data makes matters worse. One way to solve this problem is to divide the sample period into two phases. During the first phase, the inputs are shorted to a buffered version of the reference level (which may have an offset). During the second phase, the inputs are connected to the true inputs. This scheme does not remove the input current completely, since the input must still charge nodes 83 and 84 from the reference value to the current input value, but it does reduce the total charge required by about a factor of 10 (requiring only a 0.25V change rather than a 2.5V change). Persons skilled in the art will recognize that many other methods can be used to provide a clocked amplifier that will operate on very low input currents.

[0103] One important part of the input/output circuitry generates an internal device clock based on early and late bus clocks. Controlling clock skew (the difference in clock timing between devices) is important in a system running with 2 ns cycles, thus the internal device clock is generated so the input sampler and the output driver operate as close in time as possible to midway between the two bus clocks.

[0104] A block diagram of the internal device clock generating circuit is shown in Figure 12 and the corresponding timing diagram in Figure 13. The basic idea behind this circuit is relatively simple. A DC amplifier 102 is used to convert the small-swing bus clock into a full-swing CMOS signal. This signal is then fed into a variable delay line 103. The output of delay line 103 feeds three additional delay lines: 104 having a fixed delay; 105 having the same fixed delay plus a second variable delay; and 106 having the same fixed delay plus one half of the second variable delay. The outputs 107, 108 of the delay lines 104 and 105 drive clocked input receivers 101 and 111 connected to early and late bus clock inputs 100 and 110, respectively. These input receivers 101 and 111 have the same design as the receivers described

Micron Technology Inc. et al.
Ex. 1010, p. 163

above and shown in Fig. 11. Variable delay lines 103 and 105 are adjusted via feedback lines 116, 115 so that input receivers 101 and 111 sample the bus clocks just as they transition. Delay lines 103 and 105 are adjusted so that the falling edge 120 of output 107 precedes the falling edge 121 of the early bus clock, Clock1 53, by an amount of time 128 equal to the delay in input sampler 101. Delay line 108 is adjusted in the same way so that falling edge 122 precedes the falling edge 123 of late bus clock, Clock2 54, by the delay 128 in input sampler 111.

[0105] Since the outputs 107 and 108 are synchronized with the two bus clocks and the output 73 of the last delay line 106 is midway between outputs 107 and 108, that is, output 73 follows output 107 by the same amount of time 129 that output 73 precedes output 108, output 73 provides an internal device clock midway between the bus clocks. The falling edge 124 of internal device clock 73 precedes the time of actual input sampling 125 by one sampler delay. Note that this circuit organization automatically balances the delay in substantially all device input receivers 71 and 72 (Fig. 10), since outputs 107 and 108 are adjusted so the bus clocks are sampled by input receivers 101 and 111 just as the bus clocks transition.

[0106] In the preferred embodiment, two sets of these delay lines are used, one to generate the true value of the internal device clock 73, and the other to generate the complement 74 without adding any inverter delay. The dual circuit allows generation of truly complementary clocks, with extremely small skew. The complement internal device clock is used to clock the 'even' input receivers to sample at time 127, while the true internal device clock is used to clock the 'odd' input receivers to sample at time 125. The true and complement internal device clocks are also used to select which data is driven to the output drivers. The gate delay between the internal device clock and output circuits driving the bus is slightly greater than the corresponding delay for the input circuits, which means that the new data always will be driven on the bus slightly after the old data has been sampled.

DRAM Column Access Modification

[0107] A block diagram of a conventional 4 MBit DRAM 130 is shown in Figure 15. The DRAM memory array is divided into a number of subarrays 150-157, for example, 8. Each subarray is divided into arrays 148, 149 of memory cells. Row address selection is performed by decoders 146. A column decoder 147A, 147B, including column sense amps on either side of the decoder, runs through the core of each subarray. These column sense amps can be set to precharge or latch the most-recently stored value, as described in detail above. Internal I/O lines connect each set of sense-amps, as gated by corresponding column decoders, to input and output circuitry connected ultimately to the device pins. These internal I/O lines are used to drive the data from the selected bit lines to the data pins (some of pins 131-145), or to take the data from the pins and write the selected bit lines. Such a column access path organized by prior art constraints does not have sufficient bandwidth to interface with a high speed bus. The method of this invention does not require changing the overall method used for column access, but does change implementation details. Many of these details have been implemented selectively in certain fast memory devices, but never in conjunction with the bus architecture of this invention.

[0108] Running the internal I/O lines in the conventional way at high bus cycle rates is not possible. In the preferred method, several (preferably 4) bytes are read or written during each cycle and the column access path is modified to run at a lower rate (the inverse of the number of bytes accessed per cycle, preferably 1/4 of the bus cycle rate). Three different techniques are used to provide the additional internal I/O lines required and to supply data to memory cells at this rate. First, the number of I/O bit lines in each subarray running through the column decoder 147 is increased, for example, to 16, eight for each of the two columns of column sense amps and the column decoder selects one set of columns from the "top" half 148 of subarray 150 and one set of columns from the "bottom" half 149 during each cycle, where the column decoder selects one column sense amp per I/O bit line. Second, each column I/O line is divided into two halves, carrying data independently over separate internal I/O lines from the left half 147A and right half 147B of each subarray (dividing each subarray into quadrants) and the column decoder selects sense amps from each right and left half of the subarray, doubling the number of bits available at each cycle. Thus each column decode selection turns on *n* column sense amps, where n equals four (top left and right, bottom left and right quadrants) times the number of I/O lines in the bus to each subarray quadrant (8 lines each x 4=32 lines in the preferred implementation). Finally, during each RAS cycle, two different subarrays, e.g. 157 and 153, are accessed. This doubles again the available number of I/O lines containing data. Taken together, these changes increase the internal I/O bandwidth by at least a factor of 8. Four internal buses are used to route these internal I/O lines. Increasing the number of I/O lines and then splitting them in the middle greatly reduces the capacitance of each internal I/O line which in turn reduces the column access time, increasing the column access bandwidth even further.

[0109] The multiple, gated input receivers described above allow high speed input from the device pins onto the internal I/O lines and ultimately into memory. The multiplexed output driver described above is used to keep up with the data flow available using these techniques. Control means are provided to select whether information at the device pins should be treated as an address, and therefore to be decoded, or input or output data to be driven onto or read from the internal I/O lines.

[0110] Each subarray can access 32 bits per cycle, 16 bits from the left subarray and 16 from the right subarray. With

Micron Technology Inc. et al.
Ex. 1010, p. 164

8 I/O lines per sense-amplifier column and accessing two subarrays at a time, the DRAM can provide 64 bits per cycle. This extra I/O bandwidth is not needed for reads (and is probably not used), but may be needed for writes. Availability of write bandwidth is a more difficult problem than read bandwidth because over-writing a value in a sense-amplifier may be a slow operation, depending on how the sense amplifier is connected to the bit line. The extra set of internal I/O lines provides some bandwidth margin for write operations.

[0111] Persons skilled in the art will recognize that many variations of the teachings of this invention can be practiced that still fall within the claims of this invention which follow.

[0112] It follows a list of further embodiments of the invention:

Embodiment 1. A memory subsystem comprising
two memory devices connected in parallel to a bus,
said bus including a plurality of bus lines for carrying substantially all address, data and control information needed by said memory devices,
said control information including device-select information,
said bus containing substantially fewer bus lines than the number of bits in a single address, and
said bus carrying device-select information without the need for separate device-select lines connected directly to individual memory devices.

Embodiment 2. The memory subsystem with the features of embodiment 1 wherein said bus contains at least 8 bus lines adapted to carry at least 16 address bits and at least 8 data bits.

Embodiment 3. The memory subsystem with the features of embodiment 1 wherein said bus also includes parallel lines for clock and power.

Embodiment 4. A system comprising
a memory subsystem with the features of embodiment 1 wherein each bus of said memory subsystem is connected to its own transceiver device,
a transceiver bus connecting said transceiver devices, and
a means for transferring information between each of said buses of said memory subsystems and said transceiver bus, whereby memory subsystems may be integrated into a larger system having more memory than an individual memory subsystem.

Embodiment 5. The system with the features of embodiment 4 having a plurality of memory subsystems.

Embodiment 6. The system with the features of embodiment 4 further comprising a master device connected to said transceiver bus.

Embodiment 7. The system with the features of embodiment 6 wherein said master device is selected from the group consisting of a central processing unit, a floating point unit and a direct memory access unit.

Embodiment 8. The system with the features of embodiment 4 further comprising a peripheral device connected to the transceiver bus, said peripheral device adapted for connection to other devices not on the bus.

Embodiment 9. The system with the features of embodiment 8 wherein said peripheral device is selected from the group consisting of an I/O interface port, a video controller and a disk controller.

Embodiment 10. The system with the features of embodiment 5 wherein said transceiver bus is in a different plane than the plane of the bus of each of said memory subsystems.

Embodiment 11. The system with the features of embodiment 5 wherein the bus of each memory subsystem lies substantially in a subsystem bus plane and said transceiver bus lies substantially in a plane orthogonal to said subsystem bus plane.

Embodiment 12. The system with the features of embodiment 4 having at least two transceiver buses, each transceiver bus having a plurality of memory subsystem buses connected through a first transceiver to said transceiver bus,
each of said transceiver buses being further connected to a second transceiver adapted to interface to a second-order transceiver bus, whereby each transceiver bus is connected through said second transceiver to form a second-order transceiver bus unit.

Embodiment 13. A semiconductor subsystem bus for interconnecting semiconductor devices comprising
a plurality of semiconductor devices connected in parallel to a bus, at least one of said semiconductor devices being a memory device or a transceiver device which in turn is connected to a memory subsystem,
said bus including a plurality of bus lines for carrying substantially all address, data and control information needed by said semiconductor devices,
said control information including semiconductor device-select information,
said bus containing substantially fewer bus lines than the number of bits in a single address, and
said bus carrying device-select information without the need for separate device-select lines connected directly to individual semiconductor devices, and
at least one modifiable register in each of the semiconductor devices on said bus, said modifiable registers being accessible from said bus, whereby the subsystem can be configured using signals transmitted on said bus.

Embodiment 14. The semiconductor subsystem bus with the features of embodiment 13 wherein one type of modifiable

Micron Technology Inc. et al.
Ex. 1010, p. 165

register is an access-time register designed to store a time delay after which a device may take some specified action on said bus.

Embodiment 15. The semiconductor subsystem bus with the features of embodiment 13 further comprising a semiconductor device having at least two access-time registers and

one of said access-time registers is permanently programmed to contain a fixed value and at least one of said access-time registers can be modified by information carried on said bus.

Embodiment 16. The semiconductor subsystem bus with the features of embodiment 13 further comprising a memory device having at least one discrete memory section and also having a modifiable address register adapted to store memory address information which corresponds to each said discrete memory section.

Embodiment 17. The semiconductor subsystem bus with the features of embodiment 16 wherein said memory address information comprises a pointer to said discrete memory section.

Embodiment 18. The semiconductor subsystem bus with the features of embodiment 16 wherein said discrete memory section has a top and a bottom and said memory address information comprises pointers to said top and said bottom.

Embodiment 19. The semiconductor subsystem bus with the features of embodiment 16 wherein said memory address information comprises

a pointer to said discrete memory section and

a range value indicating the size of said discrete memory section.

Embodiment 20. The semiconductor subsystem bus with the features of embodiment 16 wherein said address registers of each of said discrete memory sections of each of said memory devices connected to said bus are set to contain memory address information that is different for each discrete memory section and such that the highest memory address in each discrete memory section is one less than the lowest memory address in another discrete memory section, whereby memory may be organized into one or a small number of contiguous memory blocks.

Embodiment 21. The semiconductor subsystem bus with the features of embodiment 16 further comprising a means for testing each of said discrete memory sections of each of said memory devices for proper function, and

for each non-functional discrete memory section, a means for setting at least one address register which corresponds to said discrete memory section to indicate that said discrete memory section is non-functional,

for each functional discrete memory section, a means for setting at least one address register which corresponds to said discrete memory section to contain such corresponding address information.

Embodiment 22. The semiconductor subsystem bus with the features of embodiment 21 wherein said address registers corresponding to said discrete memory sections are set to provide one contiguous memory block within the subsystem.

Embodiment 23. The semiconductor subsystem bus with the features of embodiment 13 wherein one of said modifiable registers is a device identification register which can be modified to contain a value unique to that semiconductor device.

Embodiment 24. The semiconductor subsystem bus with the features of embodiment 23 wherein said device identification register is set to contain a unique value which is a function of the physical position of that semiconductor device either along said bus or in relationship to other semiconductor devices or said bus.

Embodiment 25. A bus subsystem comprising

two semiconductor devices connected in parallel to a bus, wherein one of said semiconductor devices is a master device, said master device including a means for initiating bus transactions,

said bus including a plurality of bus lines for carrying substantially all address, data and control information needed by said devices,

said control information including device-select information,

said bus containing substantially fewer lines than the number of bits in a single address, and

said bus carrying device-select information without the need for separate device-select lines connected directly to individual devices on said bus, whereby said master device initiates bus transactions which transfer information between said semiconductor devices on said bus.

Embodiment 26. The bus subsystem with the features of embodiment 25 wherein one of said semiconductor devices is a memory device connected to said bus, said memory device having at least one discrete memory section and also having a modifiable address register adapted to store memory address information which corresponds to each said discrete memory section.

Embodiment 27. The bus subsystem with the features of embodiment 26 wherein one of said semiconductor devices comprises a transceiver device connected in parallel to said bus and connected in parallel to a memory device on a bus other than said bus.

28. The bus subsystem with the features of embodiment 26 further including a means Embodiment for said master device to request said memory device to prepare for a bus transaction by sending a request packet along said bus, said memory device and said master device each having a device-internal means to prepare to begin said bus transaction during a device-internal phase and further having a bus access means to effect said bus transaction during a bus access phase, said request packet including

a sequence of bytes containing address and control information,

Micron Technology Inc. et al.
Ex. 1010, p. 166

said control information including information about the requested bus transaction and about the access time, which corresponds to a number of bus cycles, which needs to intervene before beginning said bus-access phase, and
said address information pointing to at least one memory location within one of said discrete memory sections of said memory device.

Embodiment 29. The bus subsystem with the features of embodiment 28 wherein said memory device includes a means to read said control information and initiate said device-internal means at a time so as to complete said device-internal phase within said access time and begin said bus access phase after said number of bus cycles.

Embodiment 30. The bus subsystem with the features of embodiment 28 wherein said control information comprises an op code.

Embodiment 31. The bus subsystem with the features of embodiment 30 wherein said memory device includes sense amplifiers adapted to hold a bit of information or to precharge after a selected time and a means to transfer a data block during a data block transfer either reading data from said memory device or writing data into said memory device, and wherein said op code instructs said memory device to activate a response means, said response means including a means to

initiate a data block transfer,
select the size of said data block,
select the time to initiate said data block transfer,
access a control register, including reading from or writing to said control register,
precharge said sense amplifiers after each of said data block transfers is complete,
hold a bit of information in each of said sense amplifiers after each of said data block transfers is complete, or
select normal or page-mode access.

Embodiment 32. The bus subsystem with the features of embodiment 31 wherein said data block transfer comprises a read from or a write to memory within a single memory device.

Embodiment 33. The bus subsystem with the features of embodiment 28 further comprising a means for said master device to send control information to a specific one of said semiconductor devices on said bus by including in said request packet a device identification number unique to said semiconductor device.

Embodiment 34. The bus subsystem with the features of embodiment 28 further comprising a means for said master device to send control information to a selected one of said discrete memory portions by including in said request packet a specific memory address.

Embodiment 35. The bus subsystem with the features of embodiment 28 further comprising a means for said master device to send control information to substantially all semiconductor devices on said bus by including in said request packet a special device identification number which is recognized by said semiconductor devices.

Embodiment 15 36. The bus subsystem with the features of embodiment 28 wherein said control information specifies directly or indirectly the number of bus cycles for said master device and said memory device to wait before beginning said bus access phase.

Embodiment 37. The bus subsystem with the features of embodiment 36 wherein, for a data block transfer, said master device and said memory device use the same access time and same data block size regardless of whether said data block transfer is a read or write operation.

Embodiment 38. The bus subsystem with the features of embodiment 28 wherein said control information further includes a block-size value that encodes and specifies the size of the block of data to be transferred.

Embodiment 39. The bus subsystem with the features of embodiment 38 wherein said block-size value is encoded as a linear value for relatively small block sizes values and is encoded as a logarithmic value for relatively larger block sizes.

Embodiment 40. The bus subsystem with the features of embodiment 38 wherein said block-size value is encoded using four bits, and where the encoded value is

Encoded Value	Block Size (Bytes)
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	16

Micron Technology Inc. et al.
Ex. 1010, p. 167

(continued)

Encoded Value	Block Size (Bytes)
10	32
11	64
12	128
13	256
14	512
15	1024

Embodiment 41. The bus subsystem with the features of embodiment 26 wherein said memory device is a DRAM device containing
a plurality of sense amplifiers,
a means to hold said sense amplifiers in an unmodified state after a read or write operation, leaving the device in page mode,
a means to precharge said sense amplifiers and
a means for selecting whether to precharge said sense amplifiers or to hold said sense amplifiers in an unmodified state.
Embodiment 42. The bus subsystem with the features of embodiment 28 wherein said request packet comprises an even number of bytes.
Embodiment 43. The bus subsystem with the features of embodiment 28 further including a means for generating and controlling a plurality of bus cycles, during which said bus carries said address, data and control information, and wherein alternate said bus cycles are designated odd cycles and even cycles, respectively, and wherein said request packet begins only on an even cycle.
Embodiment 44. The bus subsystem with the features of embodiment 28 further including a means for generating ECC information corresponding to a block of data and a means for using said ECC information to correct errors in storing or reading said block of data, wherein said ECC information may be stored separately from said block of data.
Embodiment 45. The bus subsystem with the features of embodiment 44 further comprising at least two of said memory devices wherein said ECC information and said corresponding block of data are stored in a first and a second said memory device, respectively, and said master device includes a means to write or read said block of data with error correction by sending separate ones of said request packets for said ECC information and for said corresponding block of data.
Embodiment 46. A bus subsystem comprising
a memory device and a master device connected in parallel on a bus,
a means for said master device to send a request packet and initiate a bus transaction and
a means for said master device to keep track of current and pending bus transactions,
said bus including a plurality of bus lines for carrying substantially all address, data and control information needed by said memory devices,
said bus containing substantially fewer lines than the number of bits in a single address, and
said bus carrying device-select information without the need for separate device-select lines connected directly to individual devices on said bus, whereby said master device initiates bus transactions which transfer information between devices on said bus and collisions on said bus are avoided because said master device avoids initiating bus transactions which would conflict with current or pending bus transactions.
Embodiment 47. The bus subsystem with the features of embodiment 46 having at least two of said master devices and including
a collision detecting means whereby a first said master device sending a first said request packet can detect a second said master device sending one of said colliding request packets, where one of said said colliding request packet may be sent simultaneous with the initial sending of or overlapping the sending of said first request packet, and
an arbitration means whereby said first and said second master devices select a priority order in which each of said master devices will be allowed to access said bus sequentially .
Embodiment 48. The bus subsystem with the features of embodiment 47 wherein each of said master devices has a master ID number and each of said request packets includes a master ID position which is a predetermined number of bits in a predetermined position in said request packet, and wherein said collision detection means comprises
a means included in each master device for sending a request packet including said master ID number of said master device in said master ID position of said request packet and
a means to detect a collision and invoke said arbitration means if any master device detects any other master ID number in said master ID position.
Embodiment 49. The bus subsystem with the features of embodiment 47 wherein each of said master devices includes

Micron Technology Inc. et al.
Ex. 1010, p. 168

a means for sending a request packet,
a means for driving a selected bus line or lines during at least one selected bus cycle while said request packet is being sent,
a means for monitoring said selected bus line or lines to see if a said master device is sending a colliding request packet and
a means for informing all other master devices that a collision has occurred and for invoking said arbitration means.

Embodiment 50. The bus subsystem with the features of embodiment 47 wherein each of said master devices includes
a means, when sending a request packet, to drive a selected bus line or lines with a certain current during at least one selected bus cycle,
a means for monitoring said selected bus line or lines for a greater than normal current to see if another master device is driving that line or lines,
a means for detecting said greater than normal current, and
a means for informing all said master devices that a collision has occurred and for invoking said arbitration means.

Embodiment 51. The bus subsystem with the features of embodiment 47 wherein said arbitration means comprises
a means for initiating an arbitration cycle,
a means for allocating a single bus line to each master device during at least one selected bus cycle relative to the start of said arbitration cycle,
a means for allocating each master device to a single bus line during one of said selected bus cycles if there are more master devices than available bus lines,
a means for each of said master devices which sent a colliding request packet to drive said bus line allocated to said master device during said selected bus cycle, and
a means in at least one of said master devices for storing information about which master devices sent a colliding request packet,
whereby said master devices can monitor selected bus lines during said arbitration cycle and identify each said master device which sent a colliding request packet.

Embodiment 52. The bus subsystem with the features of embodiment 47 wherein said arbitration means comprises
a means included in a first one of said master devices which sent colliding request packets for identifying each of said master devices which sent colliding request packets,
a means for assigning a priority to each said master device which sent a colliding request packet, and
a means for allowing each said master device which sent a colliding request packet to access the bus sequentially according to that priority.

Embodiment 53. The bus subsystem with the features of embodiment 52 wherein said priority is based on the physical location of each of said master devices.

Embodiment 54. The bus subsystem with the features of embodiment 52 wherein said priority is based on said master ID number of said master devices.

Embodiment 55. The bus subsystem with the features of embodiment 52 wherein each of said master devices includes
a means, when sending a colliding request packet, for deciding which master device can send the next request packet in what order or at what time, whereby no master device may send a new request packet until responses to each pending request packet have been completed or scheduled.

Embodiment 56. A bus subsystem comprising
a plurality of semiconductor devices connected in parallel to a bus,
said bus including a plurality of bus lines for carrying substantially all address, data and control information needed by said semiconductor devices,
said control information including device-select information,
said bus containing substantially fewer lines than the number of bits in a single address,
said bus carrying said device-select information without the need for separate device-select lines connected directly to individual semiconductor devices,
said semiconductor devices including a reset means having an input and an output, the output of the reset means of one semiconductor device being connected to the input of the reset means of the next semiconductor device in series.

Embodiment 57. The bus subsystem with the features of embodiment 56 further including system reset means comprising
a means for generating a first and a second reset signal,
a means for passing said first reset signal to a first of said semiconductor devices and then to subsequent ones of said semiconductor devices in series and
a means for passing a second reset signal to said first semiconductor device and then to said subsequent semiconductor devices in series,
said bus subsystem including one of said semiconductor devices containing
a device identification register adapted to contain a number unique to said semiconductor device within said bus subsystem,
a device identification register setting means, and
a device reset means for resetting said semiconductor device to some desired, known reset state in response to said

Micron Technology Inc. et al.
Ex. 1010, p. 169

first reset signal and for setting said device identification register in response to said second reset signal,
whereby said bus subsystem can be reset to a known reset state with a unique device identification value in said device identification register of each of said semiconductor devices.
Embodiment 58. The bus subsystem with the features of embodiment 57 wherein said desired, known reset state is where all registers in the semiconductor device are cleared and the state machines are reset.
Embodiment 59. The bus subsystem with the features of embodiment 57 wherein said device identification register setting means comprises
a means for detecting said second reset signal,
a means for reading a device identification number from said bus lines at a specific time relative to said second reset signal and
a means for storing said device identification number in said device identification register of said semiconductor device.
Embodiment 60. The bus subsystem with the features of embodiment 57 wherein said second reset signal comprises multiple pulse sequences and wherein said device identification setting means includes
a means for interpreting said pulse sequences as a device identification number and
a means for storing said device identification number in said device identification register of said semiconductor device.
Embodiment 61. The bus subsystem with the features of embodiment 57 wherein said device reset means comprises an n-stage shift register capable of storing n-bit values, wherein said device reset means interprets a specific value in said shift register as said first reset signal and interprets a specific value in said shift register as said second reset signal.
Embodiment 62. The bus subsystem with the features of embodiment 57 wherein one of said semiconductor devices is a master device, said master device including a means for generating said first and said second reset signals.
Embodiment 63. The bus subsystem with the features of embodiment 57 wherein one of said semiconductor devices is a master device, said master device including
a master ID register,
a means for assigning a master ID number to said master device and
a means for storing said master ID number in said master ID register.
Embodiment 64. The bus subsystem with the features of embodiment 63 further comprising a second one of said master devices, and a means for a first one of said master devices to assign a master ID number to substantially all other said master devices, whereby said first master device assigns one of said master ID numbers to each of said master devices on said bus subsystem and each said master device stores said assigned master ID number in said master ID register.
Embodiment 65. The bus subsystem with the features of embodiment 57 wherein one of said semiconductor devices includes a device-type register adapted to contain an identifier characteristic of that type of semiconductor device, and one or more modifiable registers, at least one of which is an access-time register adapted for storing access times.
Embodiment 66. The bus subsystem with the features of embodiment 65 wherein one of said semiconductor devices is a master device having
a means for selecting a semiconductor device,
a means for reading said device-type register of said selected semiconductor device,
a means for determining the device type of said selected semiconductor device,
a means for determining access-time values appropriate for said selected semiconductor device and for storing said access-time values in said access-time registers of said selected semiconductor device, and
a means for selecting and storing other values appropriate for said selected semiconductor device in corresponding registers of said selected semiconductor device,
whereby said master device can select a semiconductor device, determine what type it is, and set said access-time and other registers to contain appropriate values.
Embodiment 67. The bus subsystem with the features of embodiment 66 further comprising a memory device having at least one discrete memory section and at least one modifiable address register adapted to store memory address information which corresponds to each of said discrete memory sections, and
said master device further comprising a means for selecting and testing each of said discrete memory sections and a means for storing address information in said address registers corresponding to each of said discrete memory sections,
whereby said master device can test all said discrete memory sections and assign unique address values thereto.
Embodiment 68. A bus subsystem comprising
two semiconductor devices connected in parallel to a bus, one of said semiconductor devices being a master device,
said bus including a plurality of bus data lines for carrying substantially all address, data and control information needed by said semiconductor devices,
said control information including device-select information,
said bus containing substantially fewer of said bus data lines than the number of bits in a single address, and
said bus carrying device-select information without the need for separate device-select lines connected directly to individual semiconductor devices,
wherein all of said bus data lines are terminated transmission lines and all of said address, data and control information

Micron Technology Inc. et al.
Ex. 1010, p. 170

is carried on said bus data lines as a sequential series of bits in the form of low-voltage-swing signals.

Embodiment 69. The bus subsystem with the features of embodiment 68 further comprising a semiconductor device including a current-mode driver connected to drive one of said bus data lines.

Embodiment 70. The bus subsystem with the features of embodiment 69 further comprising a semiconductor device having a means to measure the voltage of said low-voltage-swing signals on a selected one of said bus data lines, whereby said semiconductor device can determine whether zero, one, or more than one of said current-mode drivers are driving said selected bus data line.

Embodiment 71. The bus subsystem with the features of embodiment 70 further comprising a semiconductor device having

a plurality of input receivers connected to one of said bus data lines, and

a selection means for selecting said input receivers one by one to sense and store, one at a time, the bits of said sequential series of bits.

Embodiment 72. The bus subsystem with the features of embodiment 70 further comprising a semiconductor device having two input receivers connected to one of said bus data lines.

Embodiment 73. A bus subsystem comprising

two semiconductor devices connected in parallel to a bus having a first and a second end, said bus including a bus clock line, said bus clock line having first and second ends corresponding to said first and second ends of said bus, respectively,

a clock generator connected to said first end of said bus clock line to generate early bus clock signals with a normal rise time, and

signal return means at said second end of said bus clock line to return said early bus clock signals to said first end of said bus as corresponding late bus clock signals,

whereby each of said early bus clock signals will propagate from said clock generator along said clock line starting from said first end to said second end of said bus and then return at a later time to said first end of said bus as a corresponding late bus clock signal, whereby each semiconductor device on said bus can detect said early bus clock signals and said corresponding late bus clock signals.

Embodiment 74. The bus subsystem with the features of embodiment 73 further comprising a first and a second said bus clock line having first and second ends at said first and said second ends of said bus, respectively, wherein said signal return means directly connects said second ends of said first and said second bus clock lines whereby each of said early bus clock signals will propagate from said clock generator at said first end of said bus along said first bus clock line to said second end of said bus and then return on said second bus clock line to said first end of said bus as one of said corresponding late bus clock signals.

Embodiment 75. The bus subsystem with the features of embodiment 73 wherein said signal return means comprises said first bus clock line without a line terminator at said second end thereof whereby each of said early bus clock signals reaching said second end of said first bus clock line will be reflected back along said first bus clock line as said corresponding late bus clock signals.

Embodiment 76. The bus subsystem with the features of embodiment 73 further comprising

a means for operating said bus in bus cycles timed to have a certain bus cycle frequency and a corresponding bus cycle period and

a means for operating said clock generator with a period of twice the bus cycle period.

Embodiment 77. The bus subsystem with the features of embodiment 76 wherein said bus cycle frequency is greater than approximately 50 MHz and less than or equal to approximately 500 MHz.

Embodiment 78. The bus subsystem with the features of embodiment 73 further including a semiconductor device having an internal device clock generating means to derive the midpoint time between said early and corresponding late bus clock signals and to generate an internal device clock synchronized to said midpoint time.

Embodiment 79. The bus subsystem with the features of embodiment 73 further including a semiconductor device having a low-skew clock generator circuit comprising

a first delay line having an input, an output and a basic delay and means for synchronizing the output of said first delay line with said early bus clock signal,

a second delay line having said basic delay plus a variable delay, said second delay line having an output and a means for synchronizing the output of said second delay line with said late bus clock signal, and

a third delay line having a third delay and a means to set said third delay midway between the delays of said first and second delay lines, said third delay line having an output which provides an internal device clock signal synchronized to a time halfway between said early and said late bus clock signals.

Embodiment 80. The bus subsystem with the features of embodiment 73 wherein said early and said late bus clock signals are low-voltage-swing signals that transition cyclically between low and high logical values, and further including a semiconductor device having a low-skew clock generator circuit comprising

a DC amplifier to convert said early and said late bus clock signals into full-swing logic signals,

a first variable delay line having a first variable delay and an input and an output, the input of said first variable delay

Micron Technology Inc. et al.
Ex. 1010, p. 171

line being connected to said DC amplifier
a first, a second and a third additional delay line, each having an input and an output, the input of each of said additional delay lines being connected to the output of said first delay line,
said first additional delay line having a fixed delay,
said second additional delay line having said fixed delay plus a second variable delay, and
said third additional delay line having said fixed delay plus one half of said second variable delay,
a first clocked input receiver connected to sample said early bus clock signal and gated by said output of said first additional delay line,
a means for adjusting said first variable delay so said first clocked input receiver samples said early bus clock signal just as said early bus clock signal transitions,
a second clocked input receiver connected to sample said late bus clock signal and gated by said output of said second additional delay line,
a means for adjusting said second variable delay so said second clocked input receiver samples said late bus clock signal just as said late bus clock signal transitions,
whereby said output of said third additional delay line is synchronized to a time halfway between said outputs of said first and said second additional delay lines, and said output of said third additional delay line provides an internal device clock signal.

Embodiment 81. The bus subsystem with the features of embodiment 80 further comprising a semiconductor device having
a first one of said low-skew clock generator circuits which generates a "true" internal device clock signal and
a second one of said low-skew clock generator circuits connected to generate a "complement" internal device clock signal synchronized with but opposite in logical value to said "true" internal device clock signal.

Embodiment 82. A DRAM device designed to be connected to an external bus having a plurality of bus lines for carrying substantially all address, data and control information needed by said DRAM device as a sequential series of bits, said control information including device-select information, said external bus containing substantially fewer said bus lines than the number of bits in a single address, and said bus carrying device-select information without the need for separate device-select lines connected directly to said DRAM device, said DRAM device comprising
an array of memory cells connected in rows and columns, each of said memory cells adapted to store one of said bits,
a row address selection means for selecting one of said rows,
a column sense amp connected to each of said columns, each of said column sense amps adapted to latch one of said bits as a binary logical value or to precharge to a selected state,
a column decoding means connected to each of said column sense amps for selecting a plurality of said column sense amps for inputting one of said bits to or outputting one of said bits from said memory cells,
an internal I/O bus having a plurality of internal I/O lines wherein each of said internal I/O lines is connected to a plurality of said column sense amps, and
a plurality of bus connection means designed to connect said internal I/O lines to said external bus,
whereby a selected bit of said sequential series of bits can be transferred from said external bus to a selected one of said memory cells or said bit contained in a selected one of said memory cells can be transferred to said external bus.

Embodiment 83. The DRAM device with the features of embodiment 82 further comprising
an output driver connected to one said bus connection means,
an output multiplexer having an output connected to said output driver and a plurality of inputs, each of said inputs being connected to one of said internal I/O lines, and
a control means to select whether said output driver can drive said external bus,
whereby a plurality of memory cells are selected using said row address selection means and said column decoding means and a plurality of bits contained in said plurality of memory cells are output through said column sense amps to said internal I/O bus to said output multiplexer to said output driver to said external bus.

Embodiment 84. The DRAM device with the features of embodiment 82 further comprising
a plurality of input receivers connected to one of said bus data lines and to said internal I/O bus,
a selection means for selecting said input receivers one by one to sense and store, one at a time, the bits of said sequential series of bits, and
a control means to select whether an input receiver can drive said internal I/O bus, whereby a bit of said sequential series of bits is input from said external bus through one of said input receivers to one of said internal I/O lines to one of said column sense amps to one of said memory cells.

Embodiment 85. The DRAM device with the features of embodiment 82 further comprising
a first and a second half-array of said memory cells wherein each said row of said array of said memory cells is subdivided into two parts,
a first and a second one of said internal I/O buses connected to said column sense amps in said first and said second half-arrays, respectively, and

Micron Technology Inc. et al.
Ex. 1010, p. 172

a column decoder means to gate selected ones of said column sense amps connected to said memory cells in a selected row of said first and said second half-arrays simultaneously.

Embodiment 86. The DRAM device with the features of embodiment 85 wherein said column decoder means selects sixteen column sense amps at a time.

Embodiment 87. The DRAM device with the features of embodiment 82 wherein said external bus operates at a certain speed and wherein said DRAM device includes four of said internal I/O buses, each of which operates at one-fourth the speed of said external bus.

Embodiment 88. The DRAM device with the features of embodiment 82 further comprising

a means for precharging one of said column sense amps to a precharged state from which a binary logical value can quickly be loaded into said column sense amp,

if said column sense amp contains a binary logical value, a means for latching the logical value currently contained in said column sense amp and

a means for instructing said DRAM device to precharge said column sense amp or latch said binary logical value in said column sense amp.

Embodiment 89. The DRAM device with the features of embodiment 88 further comprising a means for instructing said DRAM device to precharge said column sense amp without further instruction whenever said row address selection means selects a different one of said rows.

Embodiment 90. The DRAM device with the features of embodiment 88 further comprising a means for instructing said DRAM device to precharge said column sense amp without further instruction at a first or a second preselected time after latching the latest said binary logical value, said first preselected time being long enough for said DRAM to latch said binary logical value into said column sense amp and transfer said binary logical value into memory or onto one of said internal I/O lines, and said second preselected time being a variable which can be stored in said DRAM device whereby said DRAM can latch a binary logical value into said column sense amp for transferring said binary logical value into or out of a selected said memory cell, then precharge to allow a faster subsequent read or write.

Embodiment 91. A package containing

a semiconductor die having a side, circuitry and a plurality of connecting areas positioned along or near said side, spaced at a selected pitch and connected to said circuitry,

said package comprising a plurality of bus connecting means for connecting to a plurality of external bus lines, each of said external bus lines corresponding to one of said connecting areas, each of said bus connecting means being positioned on a first side of said package,

connected to one said external bus line and to said corresponding connecting area on said semiconductor die, and spaced at a pitch substantially identical to said selected pitch of said connecting areas,

whereby each of said external bus lines can be connected to said corresponding connecting area on said semiconductor die by bus connection means positioned along a single side of said package.

Embodiment 92. The package with the features of embodiment 91 further comprising a plurality of said bus connecting means wherein each of said bus connecting means includes

a pin adapted for connection to one of said external bus lines and

a wire connecting said pin to one of said connecting areas on said semiconductor die,

said wire having an effective lead length less than about 4 millimeters and wherein the effective lead length of said wire of each of said bus connection means for said package is approximately equal.

Embodiment 93. A plurality of packages with the features of embodiment 91 wherein at least two of said semiconductor die are memory devices, each of said packages being generally flat, having a top and a bottom, and wherein

said packages are physically secured adjacent and parallel to each other in a stack,

where a first one of said packages is adjacent to a second one of said packages in said stack, said top of said first package is substantially aligned with said bottom of said second package, and

said bus connecting means of each of said packages are substantially aligned and are lying substantially in a plane.

Embodiment 94. The plurality of packages with the features of embodiment 93 further comprising a plurality of stacks wherein each of said bus connecting means can be electrically connected to corresponding said bus connecting means in each of said stacks.

Embodiment 95. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus, and

at least one modifiable identification register accessible to said bus through said connection means, whereby data may

Micron Technology Inc. et al.
Ex. 1010, p. 173

be transmitted to said register via said bus and enable said device thereafter to be uniquely identified.

Embodiment 96. The semiconductor device with the features of embodiment 95 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 97. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus, and at least one modifiable register to hold device address information, said modifiable register accessible to said bus through said connection means, whereby data may be transmitted to said register via said bus which enables said device thereafter to respond to a predetermined range of addresses.

Embodiment 98. The semiconductor device with the features of embodiment 97 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 99. The semiconductor device with the features of embodiment 98 wherein said memory device has at least one discrete memory section and also has at least one modifiable address register adapted to store memory address information which corresponds to each said discrete memory section.

Embodiment 100. The semiconductor device with the features of embodiment 99 wherein said memory address information comprises a pointer to said discrete memory section.

Embodiment 101. The semiconductor device with the features of embodiment 100 wherein said discrete memory section has a top and a bottom and said memory address information comprises pointers to said top and said bottom.

Embodiment 102. The semiconductor device with the features of embodiment 100 wherein said memory address information comprises

a pointer to said discrete memory section and

a range value indicating the size of said discrete memory section.

Embodiment 103. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus, and

at least one modifiable access-time register accessible to said bus through said connection means, whereby data may be transmitted to said register via said bus which establishes a predetermined amount of time that said semiconductor device thereafter must wait before using said bus in response to a request.

Embodiment 104. The semiconductor device with the features of embodiment 103 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 105. The semiconductor device with the features of embodiment 103 further comprising at least two access-time registers and one of said access-time registers is permanently programmed to contain a fixed value and at least one of said access-time registers can be modified by information carried on said bus.

Embodiment 106. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, and wherein each said bus line is a terminated transmission line, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus, and

a bus line driver capable of producing a low-voltage-swing signal on one of said terminated transmission lines.

Embodiment 107. The semiconductor device with the features of embodiment 106 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 108. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying

Micron Technology Inc. et al.
Ex. 1010, p. 174

substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said bus further including at least one bus clock line for carrying early and late bus clock signals, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus, and

an internal device clock generating means which generates an internal device clock synchronized to a time halfway between said early and said late bus clock signals.

embodiment 109. The semiconductor device with the features of embodiment 108 wherein said bus further includes a first and a second one of said bus clock lines, said first bus clock line carries said early bus clock signal and said second bus clock line carries said late bus clock signal, said semiconductor device further comprising a means to detect said early bus clock signal on said first bus clock line and a means to detect said late bus clock signal on said second bus clock line.

Embodiment 110. The semiconductor device with the features of embodiment 109 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 111. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying as a sequential series of bits substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus,

a plurality of input receivers connected to one of said bus data lines and

a selection means for selecting said input receivers one by one to sense and store, one at a time, the bits of said sequential series of bits.

Embodiment 112. The semiconductor device with the features of embodiment 111 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 113. The semiconductor device with the features of embodiment 112 wherein two input receivers are connected to one of said bus lines.

Embodiment 114. A semiconductor device capable of use in an architecture for a semiconductor system bus including a plurality of semiconductor devices connected in parallel to a bus wherein said bus system includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said system bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said system bus,

an internal input/output bus within said semiconductor device having more lines than said system bus, and

a means for multiplexing the lines of said internal bus to the lines of said system bus, whereby said system bus can run at a higher speed than said internal bus.

Embodiment 115. The semiconductor device with the features of embodiment 114 wherein said semiconductor device is a memory device which connects substantially only to said system bus and sends and receives substantially all address, data and control information over said system bus.

Embodiment 116. A semiconductor device capable of use in an architecture for a semiconductor system bus including a plurality of semiconductor devices connected in parallel to a bus wherein said system bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said system bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said system bus,

an internal input/output bus within said semiconductor device having more lines than said system bus,

a means for multiplexing the lines of said internal bus to the lines of said system bus, whereby said system bus can run at a higher speed than said internal bus, and

Micron Technology Inc. et al.
Ex. 1010, p. 175

at least one modifiable identification register accessible to said system bus through said connection means, whereby data may be transmitted to said register via said system bus and which enables said device thereafter to be uniquely identified.

Embodiment 117. The semiconductor device with the features of embodiment 116 wherein said semiconductor device is a memory device which connects substantially only to said system bus and sends and receives substantially all address, data and control information over said system bus.

Embodiment 118. A semiconductor device capable of use in an architecture for a semiconductor system bus including a plurality of semiconductor devices connected in parallel to a bus wherein said system bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said system bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said system bus,

an internal input/output bus within said semiconductor device having more lines than said system bus,

a means for multiplexing the lines of said internal bus to the lines of said system bus, whereby said system bus can run at a higher speed than said internal bus, and

at least one modifiable register to hold device address information, said modifiable register accessible to said system bus through said connection means, whereby data may be transmitted to said register via said system bus which enables said device thereafter to respond to a predetermined range of addresses.

Embodiment 119. The semiconductor device with the features of embodiment 118 wherein said semiconductor device is a memory device which connects substantially only to said system bus and sends and receives substantially all address, data and control information over said system bus.

Embodiment 120. The semiconductor device with the features of embodiment 119 wherein said memory device has at least one discrete memory section and also has at least one modifiable address register adapted to store memory address information which corresponds to each said discrete memory section.

Embodiment 121. A semiconductor device capable of use in an architecture for a semiconductor system bus including a plurality of semiconductor devices connected in parallel to a bus wherein said system bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said system bus, and has substantially fewer bus lines than the number of bits in a single address, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said system bus,

an internal input/output bus within said semiconductor device having more lines than said system bus,

a means for multiplexing the lines of said internal bus to the lines of said system bus, whereby said system bus can run at a higher speed than said internal bus, and

at least one modifiable access-time register accessible to said system bus through said connection means, whereby data may be transmitted to said register via said system bus which establishes a predetermined amount of time that said semiconductor device thereafter must wait before using said system bus in response to a request.

Embodiment 122. The semiconductor device with the features of embodiment 121 wherein said semiconductor device is a memory device which connects substantially only to said system bus and sends and receives substantially all address, data and control information over said system bus.

Embodiment 123. The semiconductor device with the features of embodiment 121 further comprising at least two access-time registers and one of said access-time registers is permanently programmed to contain a fixed value and at least one of said access-time registers can be modified by information carried on said system bus.

Embodiment 124. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, wherein said address, data, control and device-select information is carried over said bus in the form of request packets and bus transactions, said semiconductor device comprising

connection means adapted to connect said semiconductor device to said bus,

a means to receive said request packets over said bus,

a means to decode information in said request packets, and

a means to respond to said information in said request packets.

Embodiment 125. The semiconductor device with the features of embodiment 124 wherein said means to decode information in said request packet further comprises

Micron Technology Inc. et al.
Ex. 1010, p. 176

a means to identify and decode said control information in said request packet,
a means to identify and decode said device-select information in said request packet,
a means to identify and decode said address information in said request packet and
a means to determine whether said control information or said address information instructs said semiconductor device to begin a response.

Embodiment 126. The semiconductor device with the features of embodiment 124 wherein each of said bus transactions is carried out in response to said address and said control information in one of said request packets, and wherein said means to identify and decode information in said request packets includes a means to identify a sequence of bytes on said bus as one of said request packets containing said address and said control information, said control information including information about the type of said bus transaction being requested and the access time which needs to intervene before beginning said bus transaction over said bus and said address and said control information includes device-select information instructing one or more said semiconductor devices to respond to said address and said control information.

Embodiment 127. The semiconductor device with the features of embodiment 124 further comprising
a plurality of sense amplifiers adapted to precharge to a selected state or to latch a bit of information,
a means to hold said sense amplifiers in an unmodified state after latching one of said bits of information,
a means to precharge said sense amplifiers and
a means for selecting whether said semiconductor device should precharge said sense amplifiers or should hold said sense amplifiers in an unmodified state.

Embodiment with the features of embodiment 128. The semiconductor device 124 wherein said means to respond to said information, where said information is control information, further comprises a means to
transfer a data block during a data block transfer, further including a means to
read data from said semiconductor device and
write data into said semiconductor device, and
initiate a data block transfer,
transfer a data block of a selected size,
transfer a data block at a selected time,
access a control register, including a means to read from or write to said control register, or
select normal or page-mode access.

Embodiment 129. The semiconductor device with the features of embodiment 124 further comprising a means to respond to said information in said request packet if said information includes a device identification number unique to said semiconductor device.

Embodiment 130. The semiconductor device with the features of embodiment 124 further comprising a means to respond to said information in said request packet if said information includes a special device identification number which calls for said semiconductor device to respond.

Embodiment 131. The semiconductor device with the features of embodiment 124 further comprising a means to respond to said information in said request packet if said information includes an address unique to said semiconductor device.

Embodiment 132. The semiconductor device with the features of embodiment 124 further comprising a means to interpret said control information and decode the time to wait before beginning said bus transaction over said bus.

Embodiment 133. The semiconductor device with the features of embodiment 124 further comprising a means to interpret said control information and decode the size of a data block to transfer during one of said bus transactions.

Embodiment 134. The semiconductor device with the features of embodiment 124, 125, 126, 127, 128, 129, 130, 131, 132 or 133 wherein said semiconductor device is a memory device which connects substantially only to said bus and sends and receives substantially all address, data and control information over said bus.

Embodiment 135. A semiconductor device capable of use in a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to a bus wherein said bus includes a plurality of bus lines for carrying substantially all address, data, control and device-select information needed by said semiconductor device for communication with substantially every other semiconductor device connected to said bus, and has substantially fewer bus lines than the number of bits in a single address, and carries device-select information for said semiconductor device without the need for a separate device-select line connected directly to said individual semiconductor device, wherein said address, data, control and device-select information is carried over said bus in the form of request packets and bus transactions, said semiconductor device comprising
connection means adapted to connect said semiconductor device to said bus,
a means to encode address and control information in said request packets and
a means to send said request packets over said bus.

Embodiment 136. The semiconductor device 135 further comprising a means to request a bus transaction wherein each of said bus transactions is carried out in response to said address and said control information in one of said request packets, and wherein said means to encode information in said request packets includes a means to mark a sequence

Micron Technology Inc. et al.
Ex. 1010, p. 177

of bytes on said bus as one of said request packets, said control information including information about the type of said bus transaction being requested and the access time which needs to intervene before beginning said bus transaction over said bus and said address and said control information includes device-select information instructing one or more said semiconductor devices to respond to said address and said control information.

Embodiment 137. The semiconductor device with the features of embodiment 135 wherein one or more of said plurality of semiconductor devices has a unique device identification number, said semiconductor device further comprising a means to send control information to a specific one of said plurality of semiconductor devices by including in said request packet a selected said device identification number.

Embodiment 138. The semiconductor device with the features of embodiment 135 wherein each of said plurality of semiconductor devices is adapted to respond to a special device identification number, said semiconductor device further comprising a means to send control information to each of said plurality of semiconductor devices by including in said request packet said special device identification number.

Embodiment 139. The semiconductor device with the features of embodiment 135 wherein one or more of said plurality of semiconductor devices is a memory device having a plurality of addresses, said semiconductor device further comprising a means to send control information to a specific address or range of addresses in one of said plurality of semiconductor devices by including said specific address or range of addresses in said request packet.

Embodiment 140. The semiconductor device with the features of embodiment 135 wherein at least one of said request packets is a request packet requesting a bus transaction which is followed by a corresponding one of said bus transactions, said semiconductor device further comprising a means to encode said control information to specify directly or indirectly the time between the end of said request packet requesting a bus transaction and said corresponding bus transaction over said bus.

Embodiment 141. The semiconductor device with the features of embodiment 140 wherein one type of said bus transactions is a transfer of a data block, said semiconductor device further comprising a means to encode said control information to specify the size of said data block to transfer.

Embodiment 142. The semiconductor device with the features of embodiment 140 further comprising a means to keep track of current and pending bus transactions, whereby collisions on said bus are avoided because said semiconductor device avoids initiating bus transactions which would conflict with current or pending bus transactions.

Embodiment 143. The semiconductor device with the features of embodiment 135 wherein said semiconductor device is a first master device and one of said plurality of semiconductor devices is a second master device, further comprising a collision detecting means whereby said first master device when sending a first one of said request packets can detect said second master device sending a colliding one of said request packets, where said colliding request packet may be sent simultaneous with the initial sending of or overlapping the sending of said first request packet, and
an arbitration means whereby said first and said second master devices select a priority order in which each of said master devices will be allowed to access said bus sequentially.

Embodiment 144. The semiconductor device with the features of embodiment 143 wherein said semiconductor device is a master device and at least one of said plurality of semiconductor devices is a master device, each of said master devices has a master ID number and each of said request packets includes a master ID position which is a predetermined number of bits in a predetermined position in said request packet, and wherein said collision detection means comprises a means for said semiconductor device to send its master ID number in said request packet and
a means to detect a collision and invoke said arbitration means if said semiconductor device detects any other master ID number in said master ID position.

Embodiment 145. The semiconductor device with the features of embodiment 144 wherein said system bus architecture includes a means for carrying information on said bus during bus cycles, said semiconductor device further comprising a means for driving a selected bus line or lines during at least one selected bus cycle while sending each said request packet,
a means for monitoring said selected bus line or lines to see if another said master device is sending one of said colliding request packets and
a means for informing all said master devices that a collision has occurred and for invoking said arbitration means.

Embodiment 146. The semiconductor device with the features of embodiment 145 further comprising
a means, when sending a request packet, for driving a selected bus line or lines with a certain current during at least one selected bus cycle,
a means for monitoring said selected bus line or lines for a greater than normal current to see if another said master device is driving that line or lines,
a means for detecting said greater than normal current, and
a means for informing all said master devices that a collision has occurred and for invoking said arbitration means.

Embodiment 147. The semiconductor device with the features of embodiment 143 wherein said arbitration means comprises
a means for initiating an arbitration cycle,

Micron Technology Inc. et al.
Ex. 1010, p. 178

a means for allocating a single bus line to each said master device during at least one selected bus cycle relative to the start of said arbitration cycle,
a means for allocating each said master device to a single bus line during one of said selected bus cycles if there are more master devices than available bus lines,
a means for each of said master devices which sent one of said colliding request packets to drive said bus line allocated to said master device during said selected bus cycle, and
a means in at least one of said master devices for storing information about which master devices sent one of said colliding request packets,
whereby said master devices can monitor selected bus lines during said arbitration cycle and identify each said master device which sent one of said colliding request packets.
Embodiment 148. The semiconductor device with the features of embodiment 143 wherein said arbitration means comprises
a means for identifying each of said master devices which sent one of said colliding request packets,
a means for assigning a priority to each said master device which sent one of said colliding request packets, and
a means for allowing each said master device which sent one of said colliding request packets to access the bus sequentially according to that priority.
Embodiment 149. The semiconductor device with the features of embodiment 143 wherein said priority is based on the physical location of each of said master devices.
Embodiment 150. The semiconductor device with the features of embodiment 143 wherein said priority is based on said master ID number of said master devices.

Claims

1. A semiconductor device, the semiconductor device comprising:

   connection means adapted to connect the semiconductor device to an external bus which is a part of a semiconductor bus architecture including a plurality of semiconductor devices connected in parallel to the external bus, wherein the external bus includes a plurality of bus lines for carrying substantially all address, data, control and device select information needed by the semiconductor device for communication with substantially every other semiconductor device connected to the external bus, and wherein the external bus has substantially fewer bus lines than the number of bits in a single address, wherein each bus line of the plurality of bus lines is a terminated transmission line, and wherein the external bus carries device-select information for the semiconductor device without the need for a separate device-select line connected directly to the semiconductor device,
   a bus line driver capable of producing a low-voltage-swing signal on one of the terminated transmission lines,

   wherein the semiconductor device is a semiconductor transceiver device to interface multiple units to the bus, wherein the multiple units are DRAMs.

2. The semiconductor device of claim 1,
   wherein the bus line driver is capable of producing a low-voltage-swing signal having a voltage swing of not more than 1 Volt on one of said terminated transmission lines.

3. The semiconductor device of claim 1 or 2,
   wherein the bus line driver is capable of producing a low-voltage-swing signal having a voltage swing of at least 500 mV on one of said terminated transmission lines.

Micron Technology Inc. et al.
Ex. 1010, p. 179

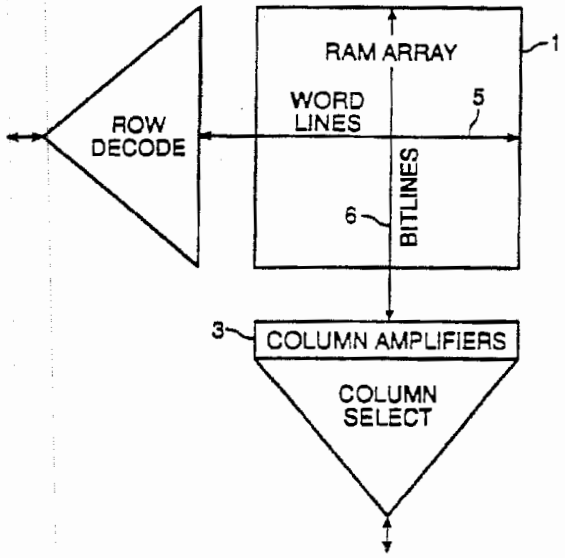


FIG. 1

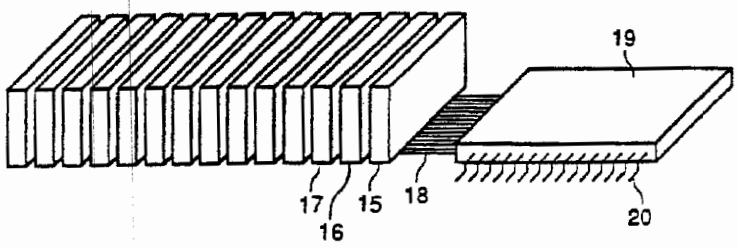


FIG. 3

Micron Technology Inc. et al.
Ex. 1010, p. 180

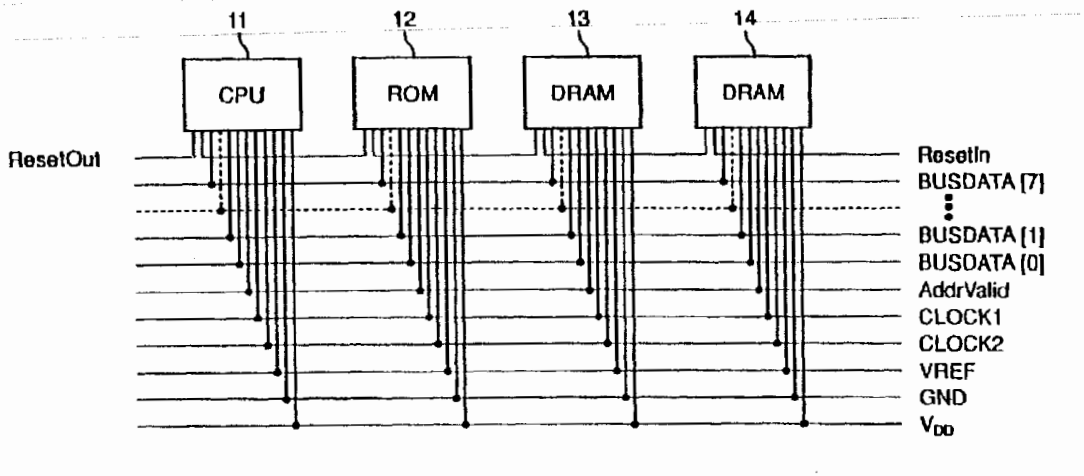


FIG. 2

Micron Technology Inc. et al.
Ex. 1010, p. 181

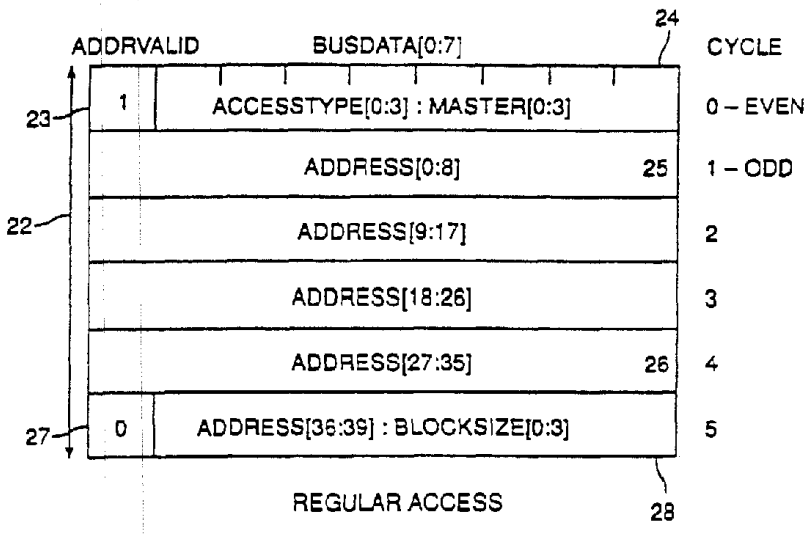


FIG. 4

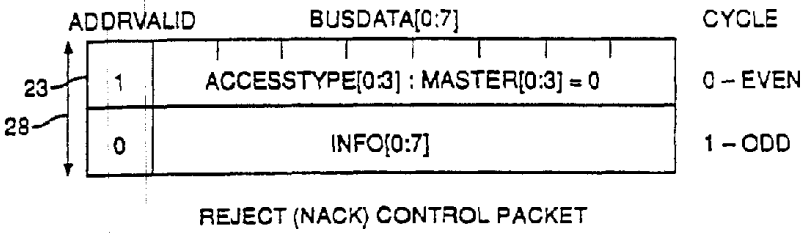


FIG. 5

Micron Technology Inc. et al.
Ex. 1010, p. 182

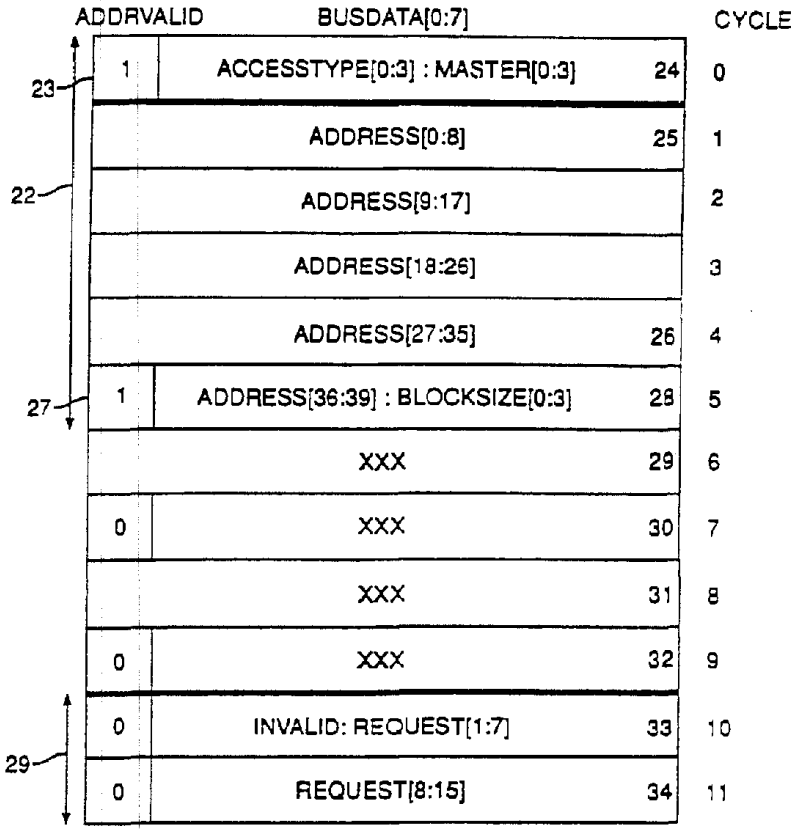


FIG. 6

Micron Technology Inc. et al.
Ex. 1010, p. 183



FIG. 7a

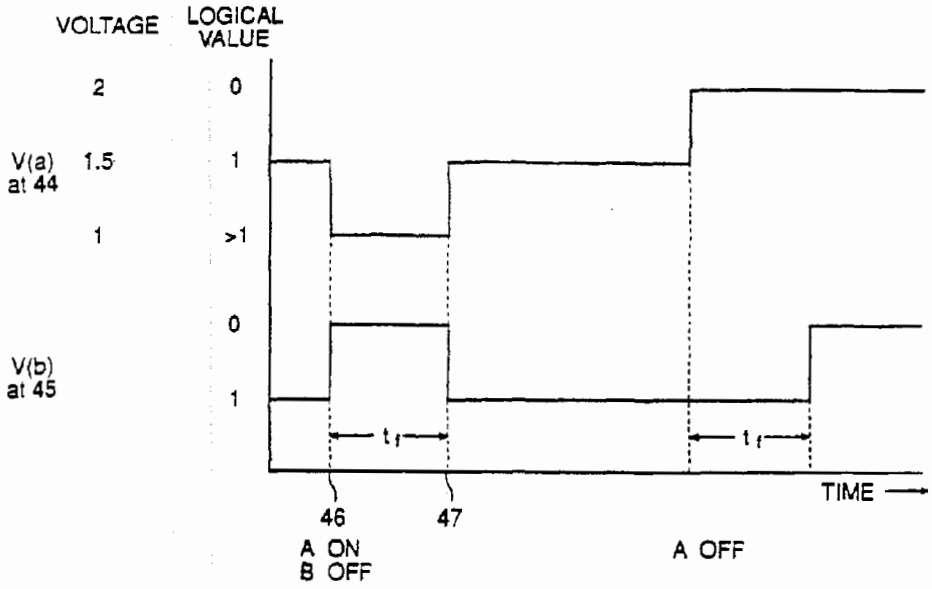


FIG. 7b

Micron Technology Inc. et al.
Ex. 1010, p. 184

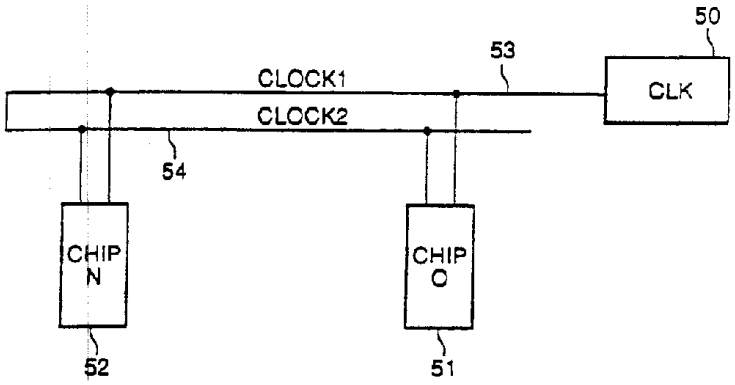


FIG. 8a

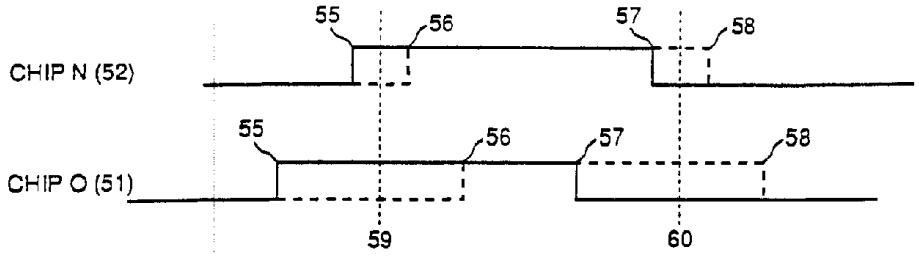


FIG. 8b

Micron Technology Inc. et al.
Ex. 1010, p. 185

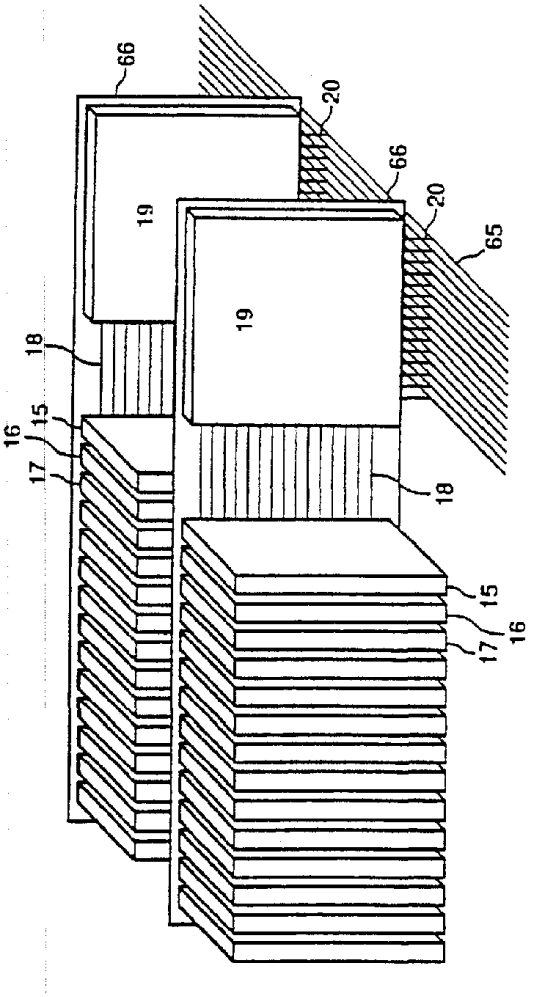


FIG. 9

Micron Technology Inc. et al.
Ex. 1010, p. 186

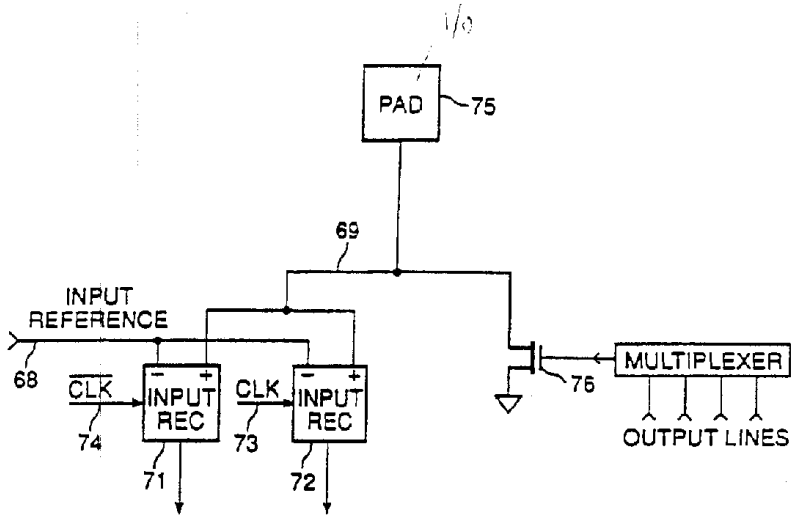


FIG. 10

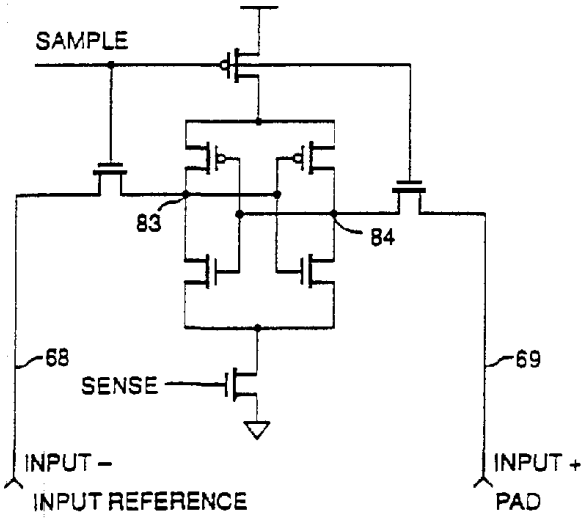


FIG. 11

Micron Technology Inc. et al.
Ex. 1010, p. 187

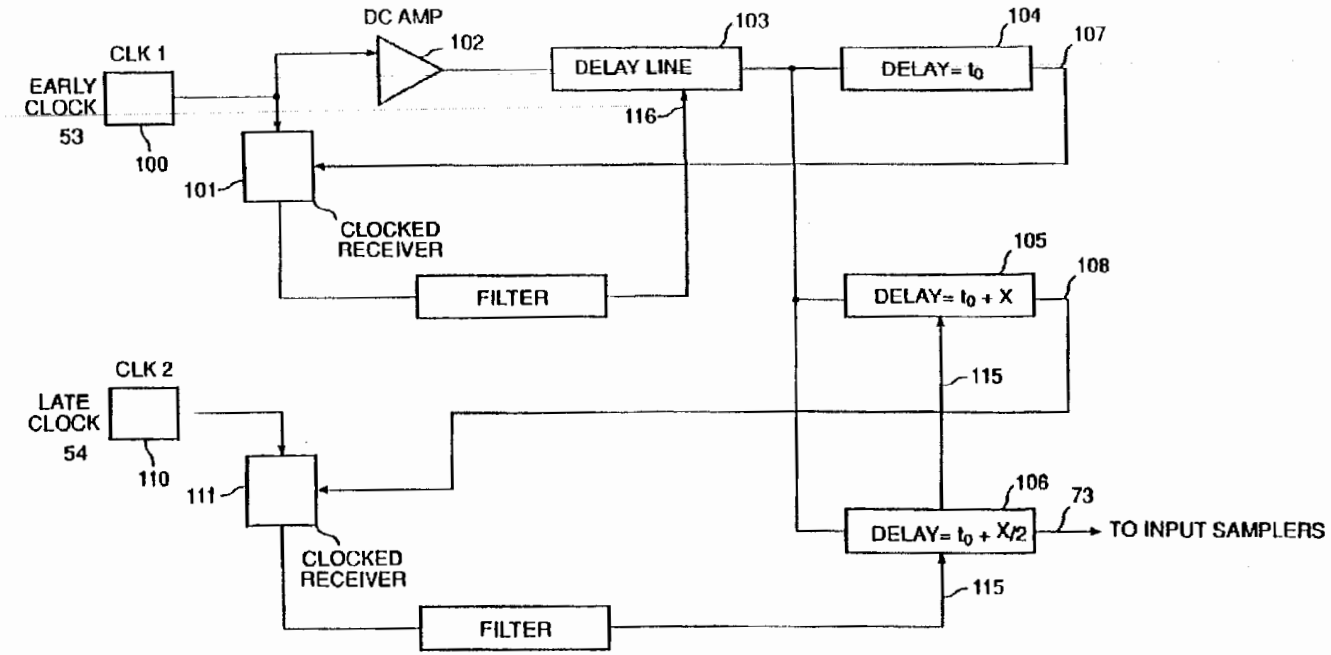

DC AMP
102
CLK 1
EARLY CLOCK
53
100
101
CLOCKED RECEIVER
FILTER
DELAY LINE
103
116
104
DELAY= t_0
107
105
DELAY= t_0 + X
108
115
CLK 2
LATE CLOCK
54
110
111
CLOCKED RECEIVER
FILTER
106
DELAY= t_0 + X/2
73
TO INPUT SAMPLERS
115

FIG. 12

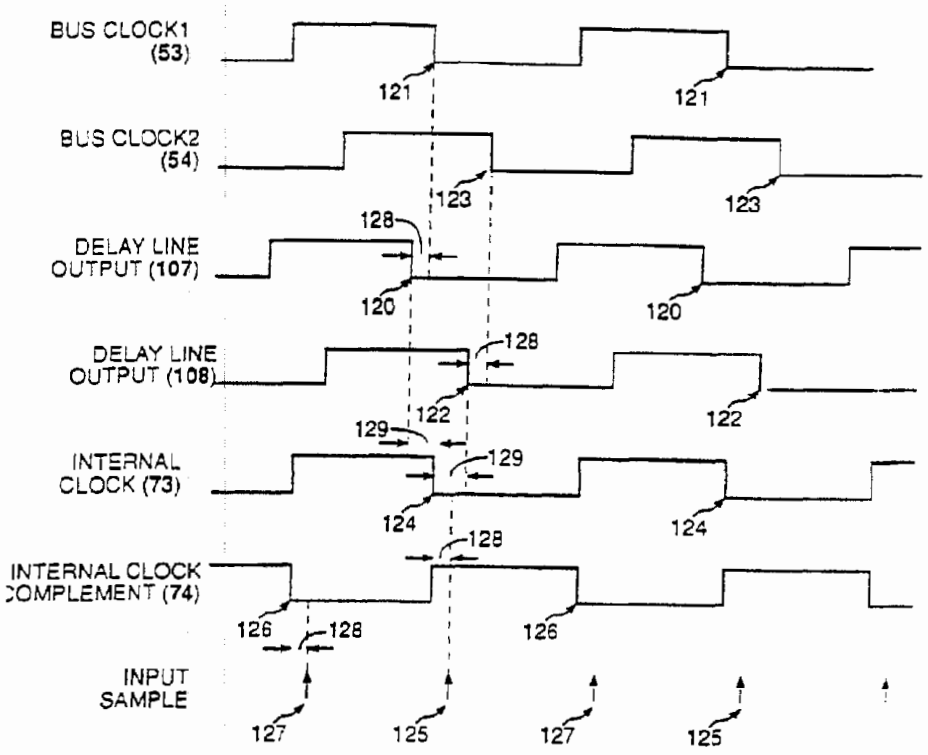


FIG. 13

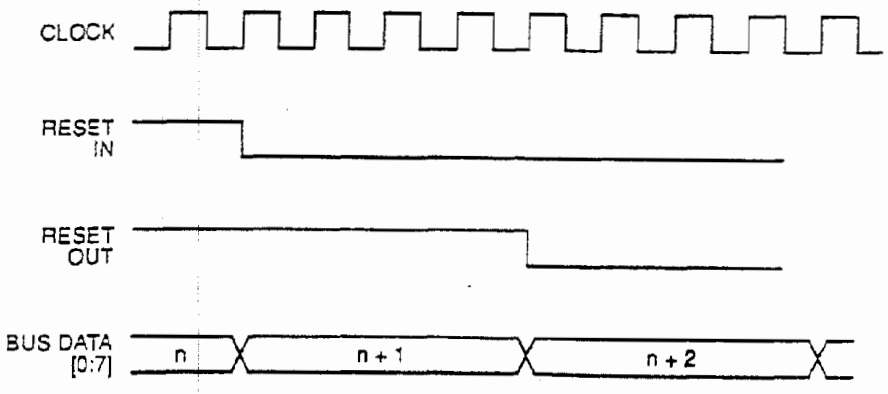


FIG. 14

Micron Technology Inc. et al.
Ex. 1010, p. 189

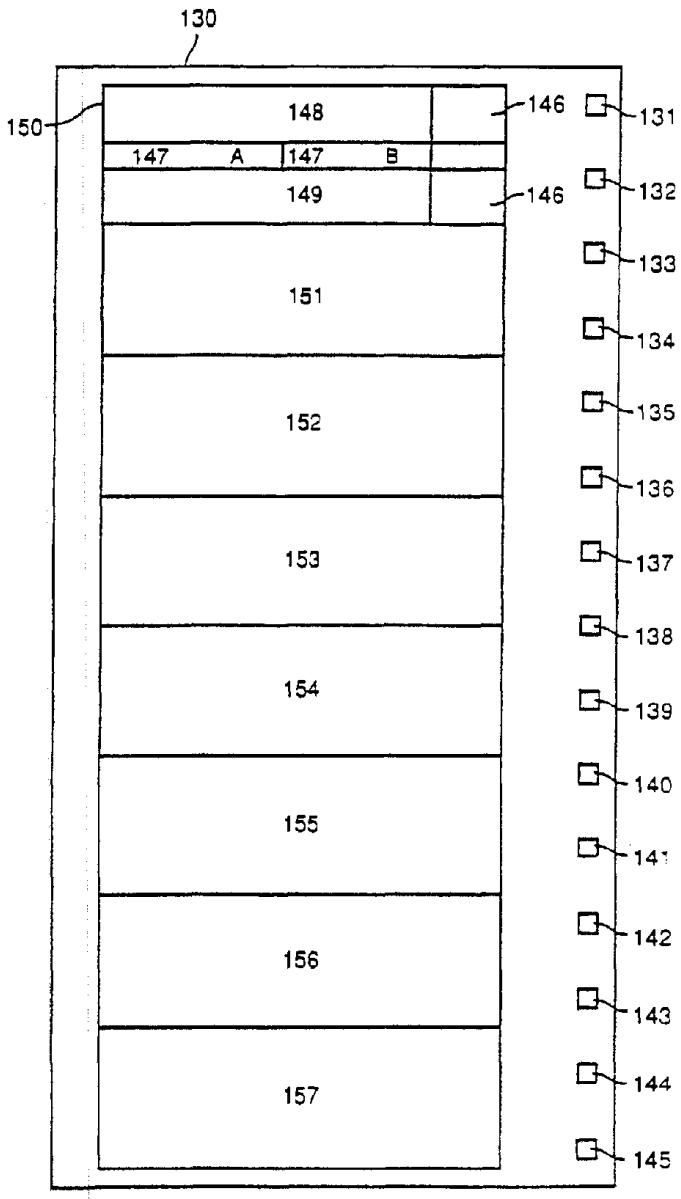


FIG. 15

Micron Technology Inc. et al.
Ex. 1010, p. 190

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 3821715 A, Hoff **[0005]**
- US 4315308 A, Jackson **[0006]**
- US 4449207 A, Kung **[0007]**
- US 4764846 A **[0008]**
- US 4706166 A, Go **[0008]**
- US 3969706 A, Proebsting **[0009]**
- US 4595923 A **[0010]**
- US 4481625 A, Roberts **[0010]**
- US 4837682 A, Culler **[0011]**
- US 4818985 A, Ikeda **[0011]**
- US 4779089 A, Theus **[0011]**
- US 4745548 A, Blahut **[0011]**
- US 4247817 A, Heller **[0012]**
- US 4646279 A, Voss **[0013]**

Non-patent literature cited in the description

- **ANANT AGARWAL.** An Analytic Cache Model. *ACM Transactions on Computer Systems,* May 1989, vol. 7 (2), 184-215 **[0055]**
- **ANANT AGARWAL et al.** Scaleable Directory Schemes for Cache Consistency. *15th International Symposium on Computer Architecture,* June 1988, 280-289 **[0083]**
- **HANS SCHUMACHER et al.** CMOS Subnanosecond True-ECL Output Buffer. *J. Solid State Circuits,* February 1990, vol. 25 (1), 150-154 **[0100]**

Micron Technology Inc. et al.
Ex. 1010, p. 191

Espacenet

Bibliographic data: JP2000311485 (A) — 2000-11-07

MEMORY MODULE SYSTEM CONNECTING ONLY SELECTED MEMORY MODULE TO DATA LINE AND DATA INPUT/OUTPUT METHOD UTILIZING THE SYSTEM

No documents available for this priority number.

Inventor(s): KIN CHIKYOKU ± (KIN CHIKYOKU)

Applicant(s): SAMSUNG ELECTRONICS CO LTD ± (SAMSUNG ELECTRONICS CO LTD)

Classification: - international: ***G06F12/00; G06F12/06; G06F13/16; G11C11/401; G11C11/407; G11C5/00; G11C7/10;*** (IPC1-7): G06F13/16; G11C11/401; G11C11/407
- cooperative: **G11C5/04**; **G11C7/1006**

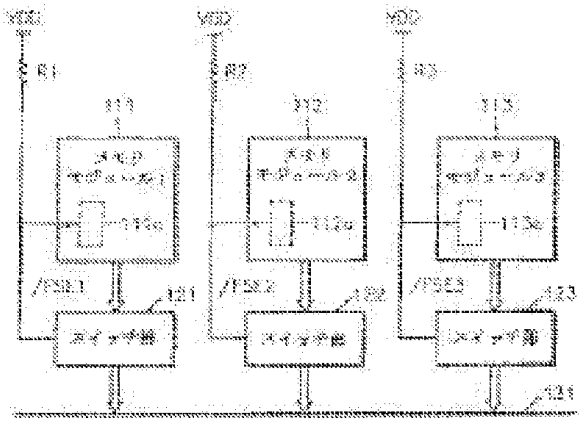
Application number: JP20000072737 20000315

Priority number(s): KR19990012027 19990407

Also published as: JP4319320 (B2) GB2348722 (A) GB2348722 (B) US6526473 (B1) TW475110 (B) KR20000065601 (A) FR2792105 (A1) FR2792105 (B1) DE19960005 (A1) less

Abstract of JP2000311485 (A)

PROBLEM TO BE SOLVED: To improve data write and read rates by generating a connection control signal by activation while data having burst length specified by a write mode is inputted. SOLUTION: As resistors R1-R3 arranged at the outside of memory modules 1-3, 111-113 are made high resistance value, voltage levels of first to third connection control signals /FSE1-/FSE3 are decided by an internal signal of the memory module 1-3, 111-113 connected to the resistors, but when an internal signal is in a floating state, voltage levels of the first to third connection control signal /FSE1-/FSE3 are made external power source voltage VDD through the resistors R1-R3, that is, made a high state.; Therefore, when the module 111 is selected, second, third connection control signals /FSE2, /FSE3 are made non-activated and second, third switches 122, 124 are turned off, loads of data pins of modules 2, 3, 112, 113 are separated from a data base line 131.



US006526473B1

(12) **United States Patent**
Kim

(10) **Patent No.: US 6,526,473 B1**
(45) **Date of Patent: Feb. 25, 2003**

(54) **MEMORY MODULE SYSTEM FOR CONTROLLING DATA INPUT AND OUTPUT BY CONNECTING SELECTED MEMORY MODULES TO A DATA LINE**

(75) Inventor: **Chi-wook Kim**, Suwon (KR)

(73) Assignee: **Samsung Electronics Co., Ltd.**, Suwon (KR)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/440,728**

(22) Filed: **Nov. 16, 1999**

(30) **Foreign Application Priority Data**

Apr. 7, 1999 (KR) 99-12027

(51) **Int. Cl.**[7] **G06F 12/02**

(52) **U.S. Cl.** **711/105**; 711/5; 711/111; 711/209; 365/230.03; 710/131

(58) **Field of Search** 711/5, 105, 111, 711/209; 365/230.03, 230.08; 710/131

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,881,203	A	11/1989	Watanabe et al.	365/203
5,513,135	A *	4/1996	Dell et al.	365/52
5,522,064	A *	5/1996	Aldereguia et al.	395/550
5,555,215	A	9/1996	Nakagome et al.	365/226
5,586,300	A *	12/1996	Wilcox et al.	395/477
5,715,211	A	2/1998	Toda	365/233
5,802,395	A	9/1998	Connolly et al.	395/834
5,850,368	A *	12/1998	Ong et al.	365/238.5
5,940,328	A *	8/1999	Iwamoto et al.	365/189.01
5,953,215	A *	9/1999	Karabatsos	361/767
5,987,579	A *	11/1999	Nishtala et al.	711/169
6,070,217	A *	5/2000	Connolly et al.	710/131
6,098,132	A *	8/2000	Olarig et al.	710/103
6,161,165	A *	12/2000	Solomon et al.	711/114
6,209,074	B1 *	3/2001	Dell et al.	711/170
6,338,113	B1 *	1/2002	Kubo et al.	711/105

FOREIGN PATENT DOCUMENTS

EP	0090591	10/1983	
EP	0335715	10/1989	
JP	61145629	3/1986	
WO	WO9941667	8/1999	711/5

OTHER PUBLICATIONS

Micron technology Inc., Synchronous DRAM (64Mb).*
Micron technology Inc., Synchronous DRAM (256Mb).*

* cited by examiner

Primary Examiner—Do Hyun Yoo
Assistant Examiner—Jasmine Song
(74) *Attorney, Agent, or Firm*—F. Chau & Associates, LLP

(57) **ABSTRACT**

A memory module system for connecting only selected memory modules to a data line to control data input and output is disclosed. The memory module system has a multiplicity of memory modules for outputting data to a data bus line, and more particularly, only the memory modules outputting data is electrically connected to the data bus line in response to activation of a predetermined connection control signal. The connection control signal has an activation width corresponding to a burst length of the output data. Only selected memory modules are connected to the data line during the data burst length, so that load per data pin is minimized, to thereby improve speed of writing and reading data.

27 Claims, 13 Drawing Sheets

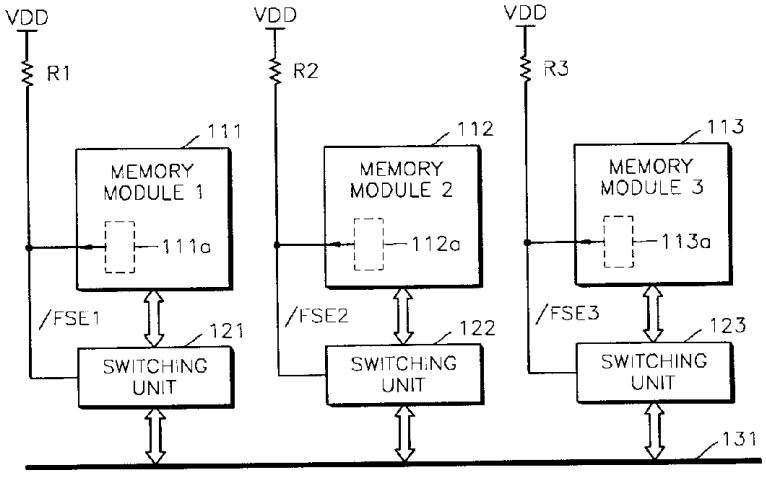
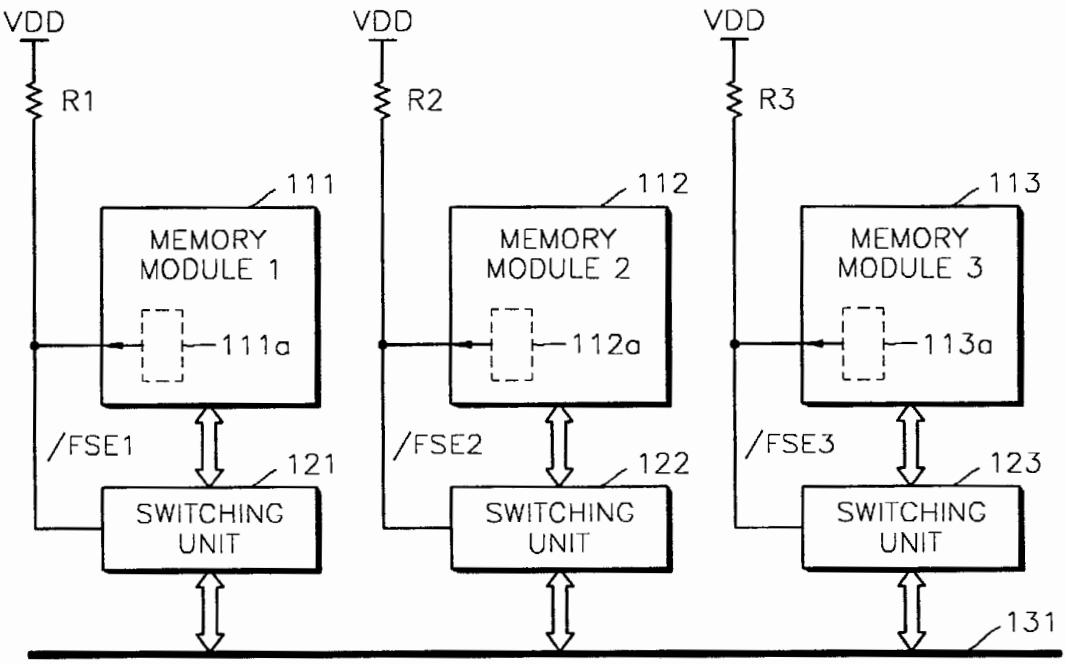


FIG. 1


VDD
VDD
VDD
R1
R2
R3
111
112
113
MEMORY MODULE 1
MEMORY MODULE 2
MEMORY MODULE 3
111a
112a
113a
/FSE1
/FSE2
/FSE3
121
122
123
SWITCHING UNIT
SWITCHING UNIT
SWITCHING UNIT
131

Micron Technology Inc. et al.
Ex. 1010, p. 196

FIG. 2

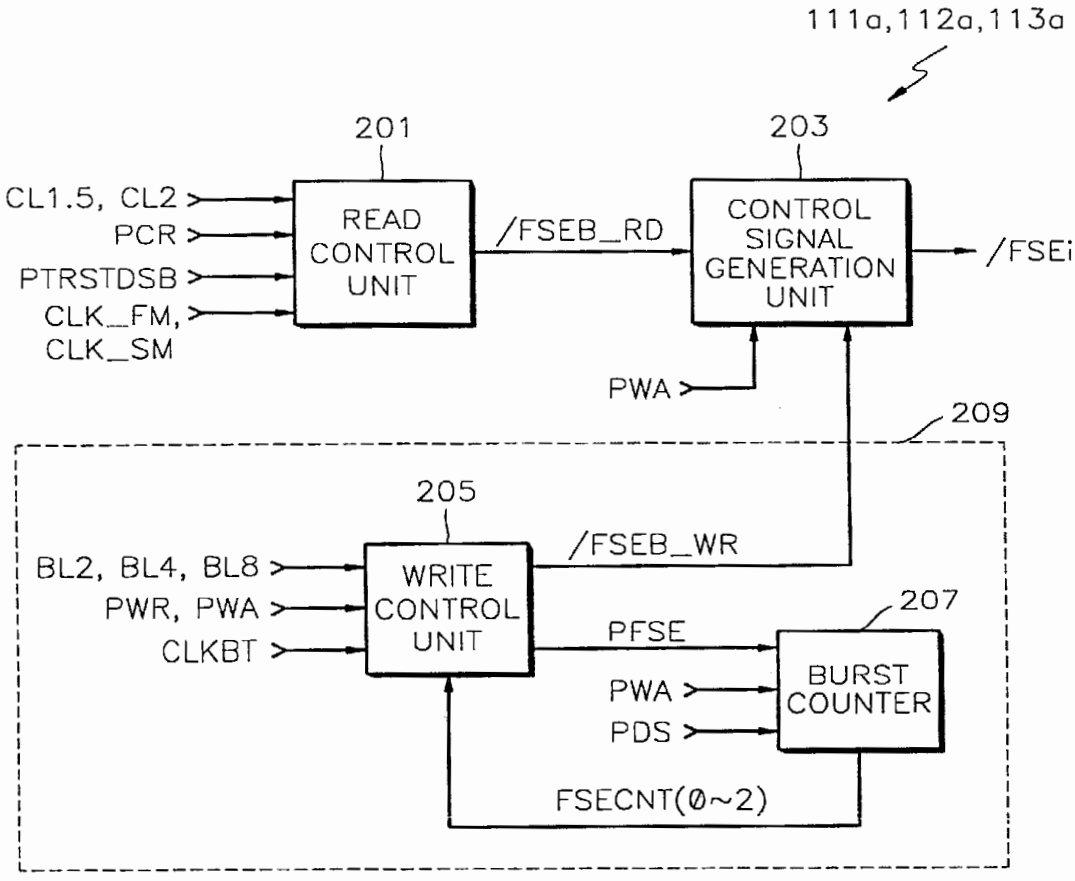


FIG. 3

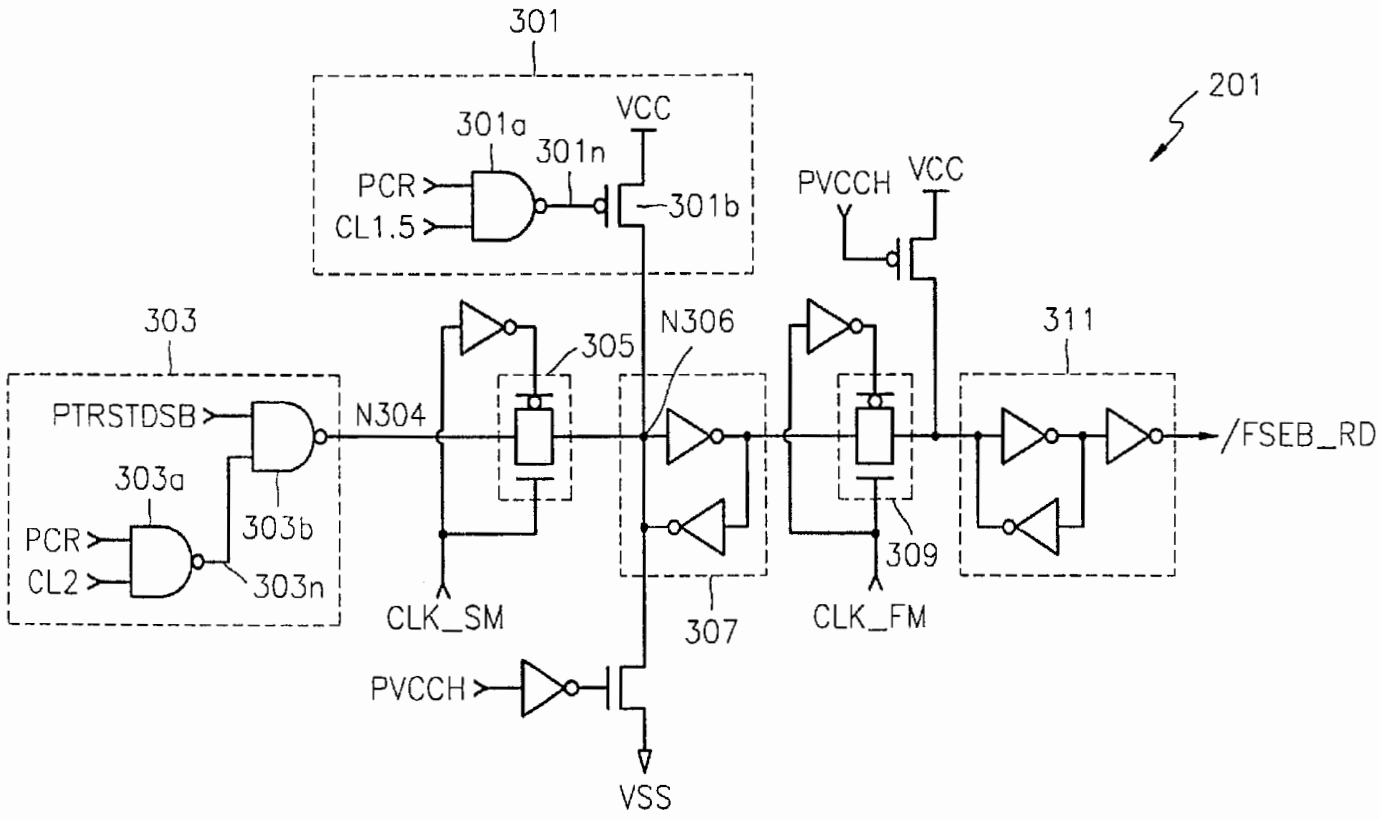

301
201
VCC
301a
301n
PCR
CL1.5
301b
PVCCH
VCC
303
305
N306
311
PTRSTDSB
N304
/FSEB_RD
303a
303b
PCR
CL2
303n
309
CLK_SM
307
CLK_FM
PVCCH
VSS

FIG. 4

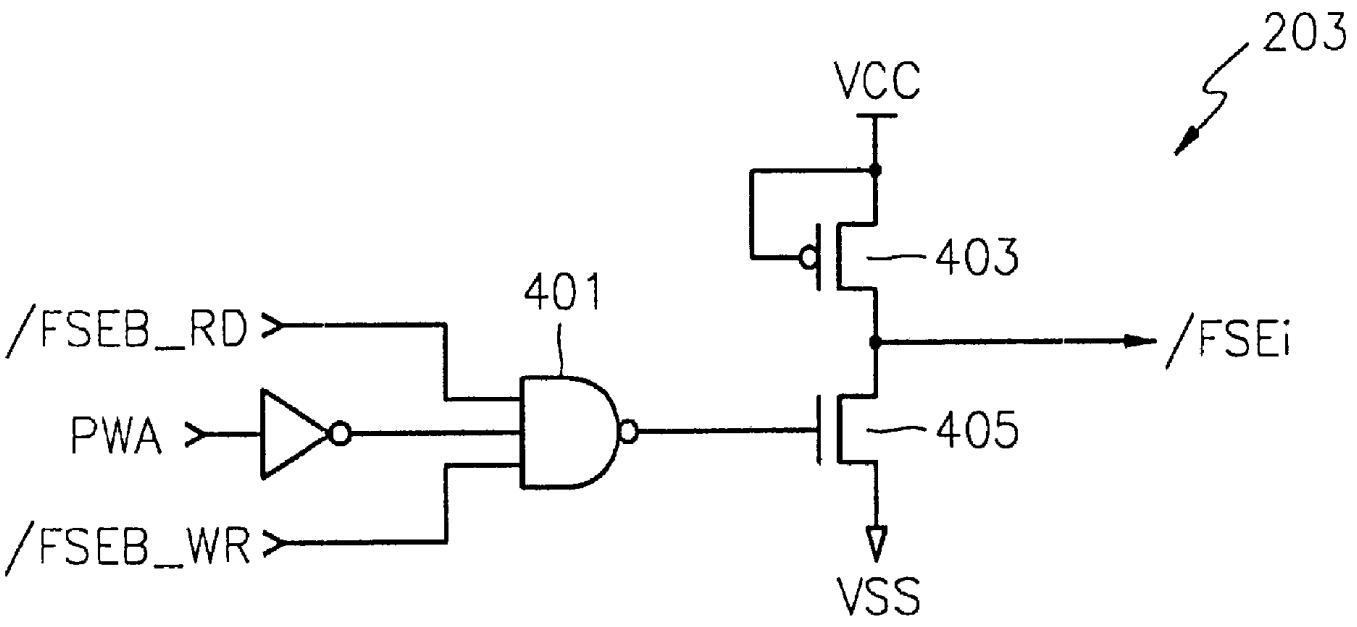


FIG. 5

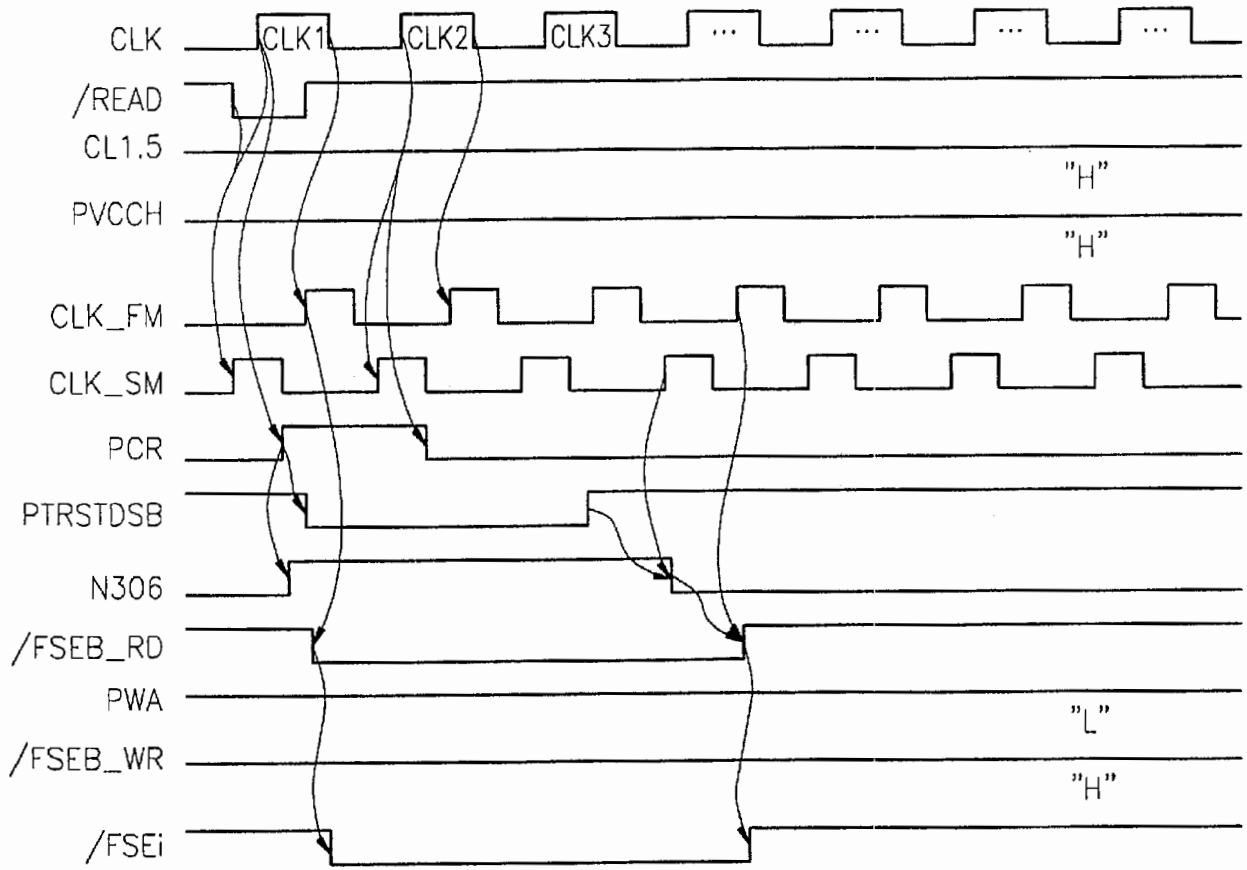

CLK
CLK1
CLK2
CLK3
/READ
CL1.5
"H"
PVCCH
"H"
CLK_FM
CLK_SM
PCR
PTRSTDSB
N306
/FSEB_RD
PWA
"L"
/FSEB_WR
"H"
/FSEi

FIG. 6

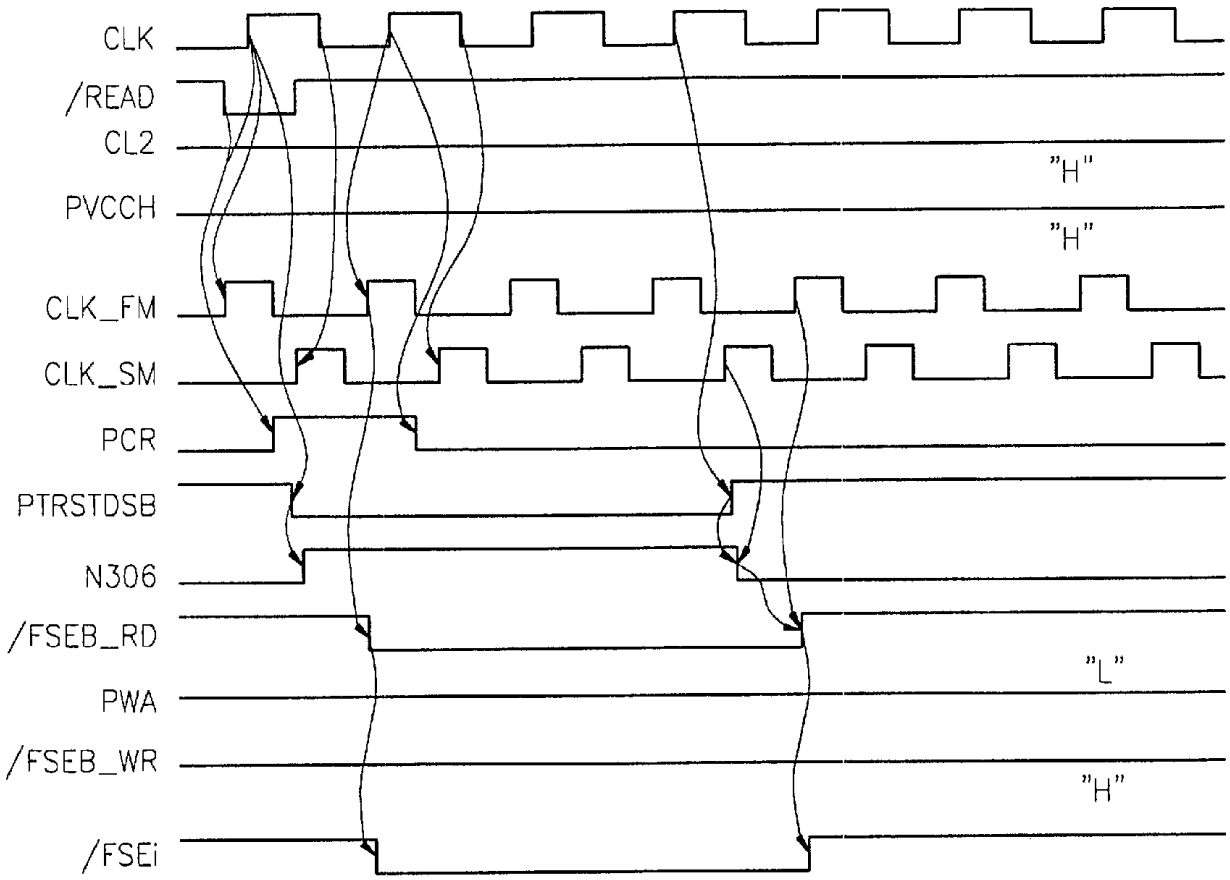

CLK
/READ
CL2
PVCCH
CLK_FM
CLK_SM
PCR
PTRSTDSB
N306
/FSEB_RD
PWA
/FSEB_WR
/FSEi
"H"
"H"
"L"
"H"

FIG. 7

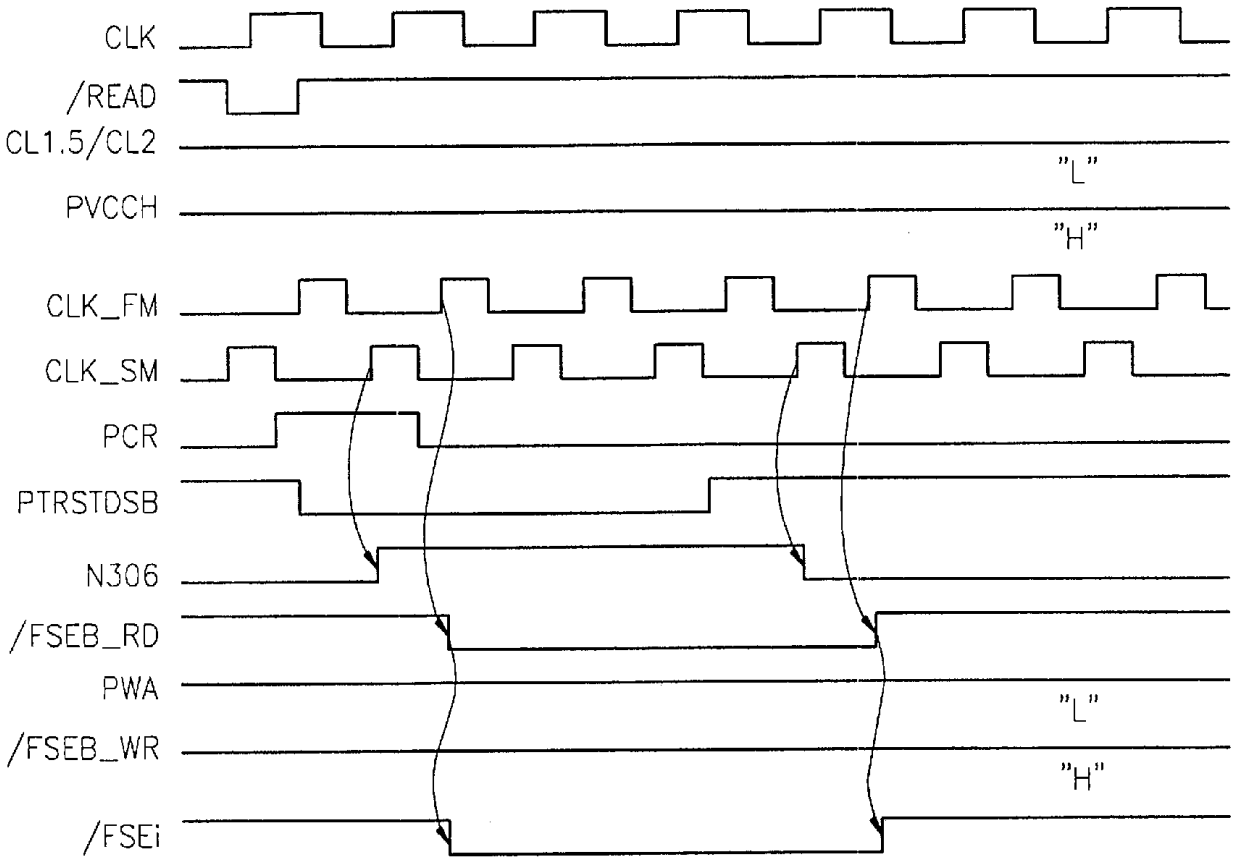

CLK
/READ
CL1.5/CL2
"L"
PVCCH
"H"
CLK_FM
CLK_SM
PCR
PTRSTDSB
N306
/FSEB_RD
PWA
"L"
/FSEB_WR
"H"
/FSEi

FIG. 8

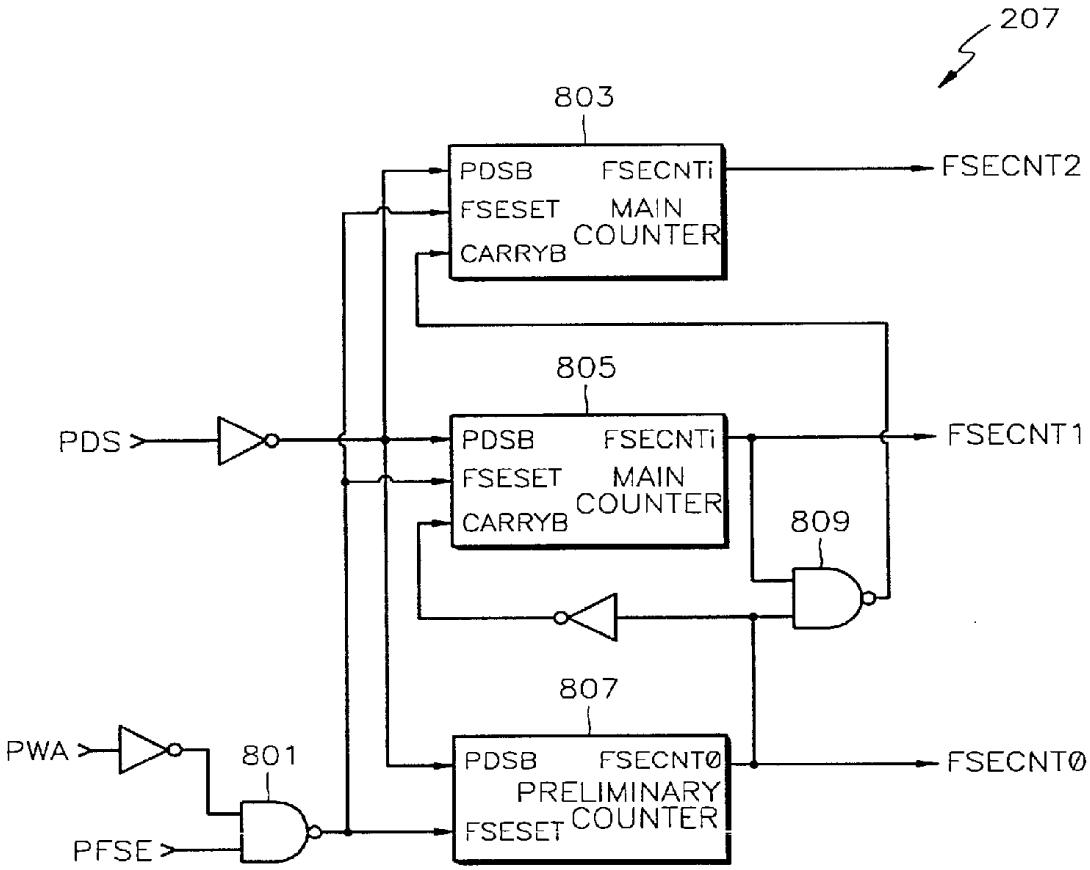

207
803
PDSB
FSECNTi
FSESET
MAIN
COUNTER
CARRYB
FSECNT2
805
PDS
PDSB
FSECNTi
FSESET
MAIN
COUNTER
CARRYB
FSECNT1
809
807
PWA
801
PDSB
FSECNT0
PRELIMINARY
COUNTER
FSESET
FSECNT0
PFSE

FIG. 9

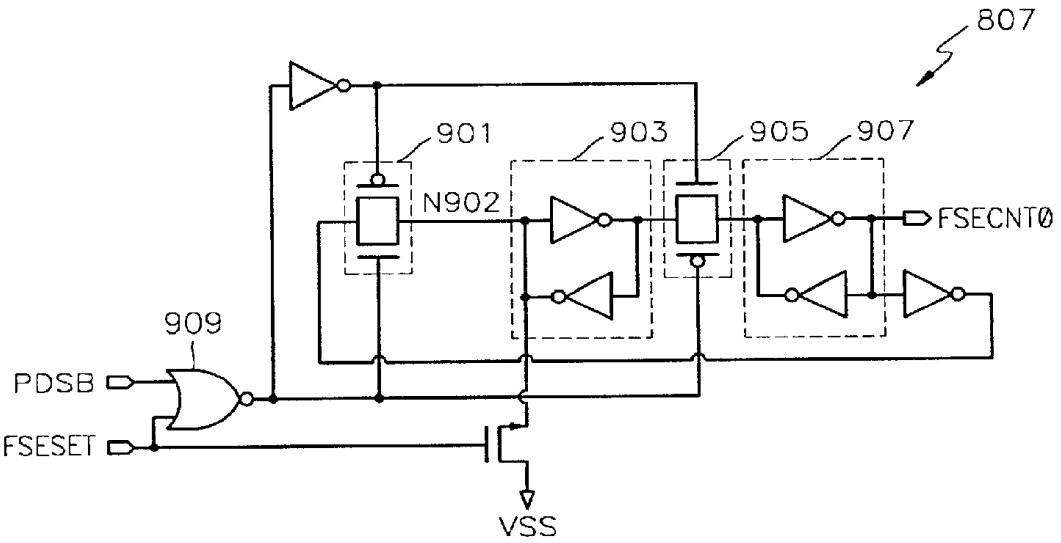

807
901
903
905
907
N902
FSECNT0
909
PDSB
FSESET
VSS

FIG. 10

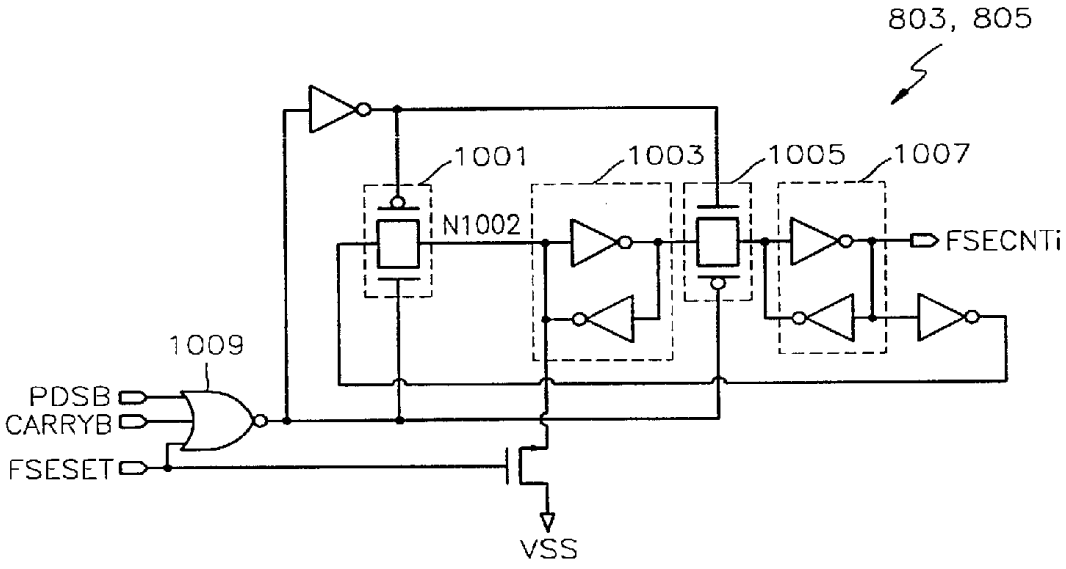

803, 805
1001
1003
1005
1007
N1002
FSECNTi
1009
PDSB
CARRYB
FSESET
VSS

FIG. 11

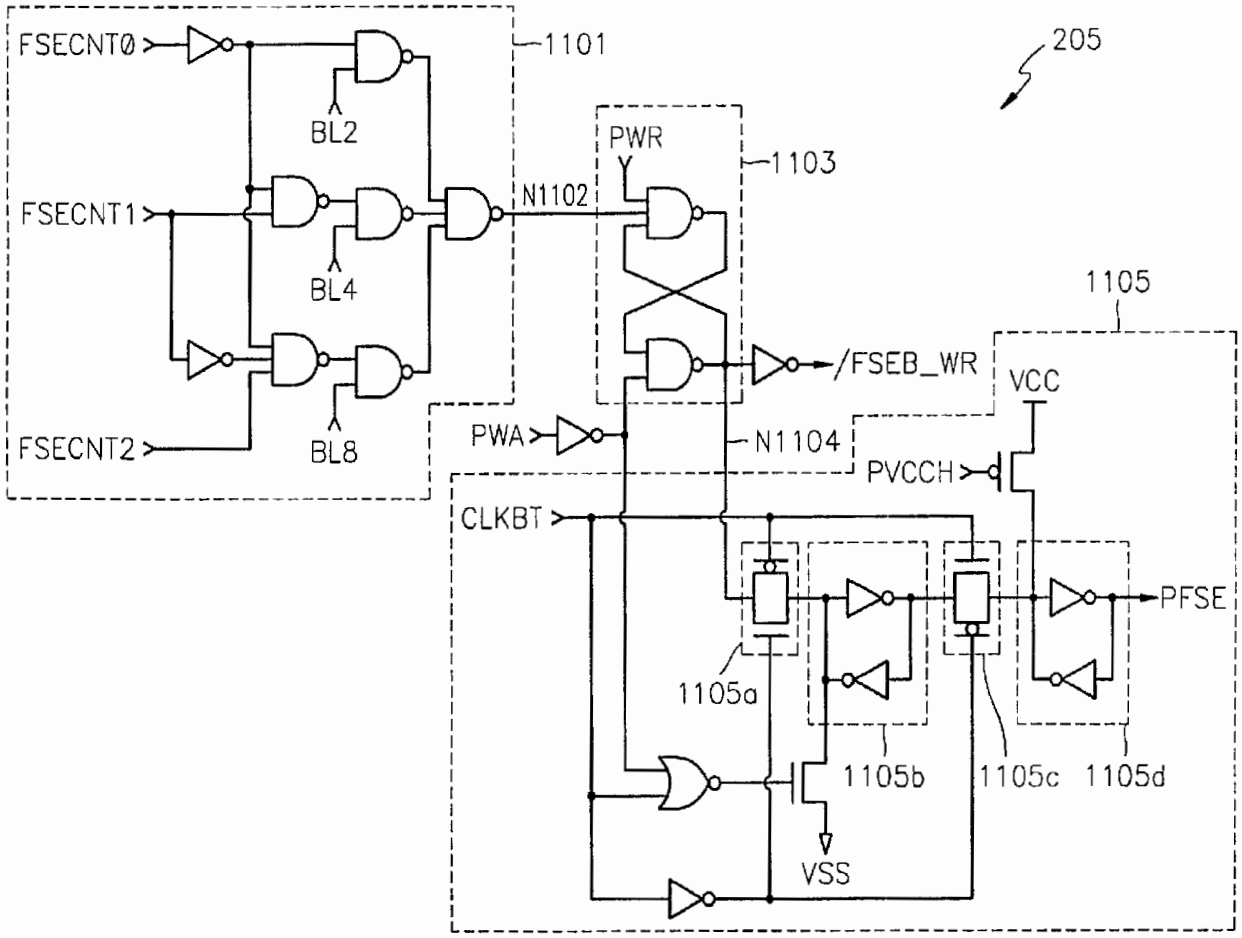

FSECNT0
FSECNT1
FSECNT2
BL2
BL4
BL8
1101
N1102
PWR
1103
205
1105
/FSEB_WR
VCC
PWA
N1104
PVCCH
CLKBT
PFSE
1105a
1105b
1105c
1105d
VSS

FIG. 12

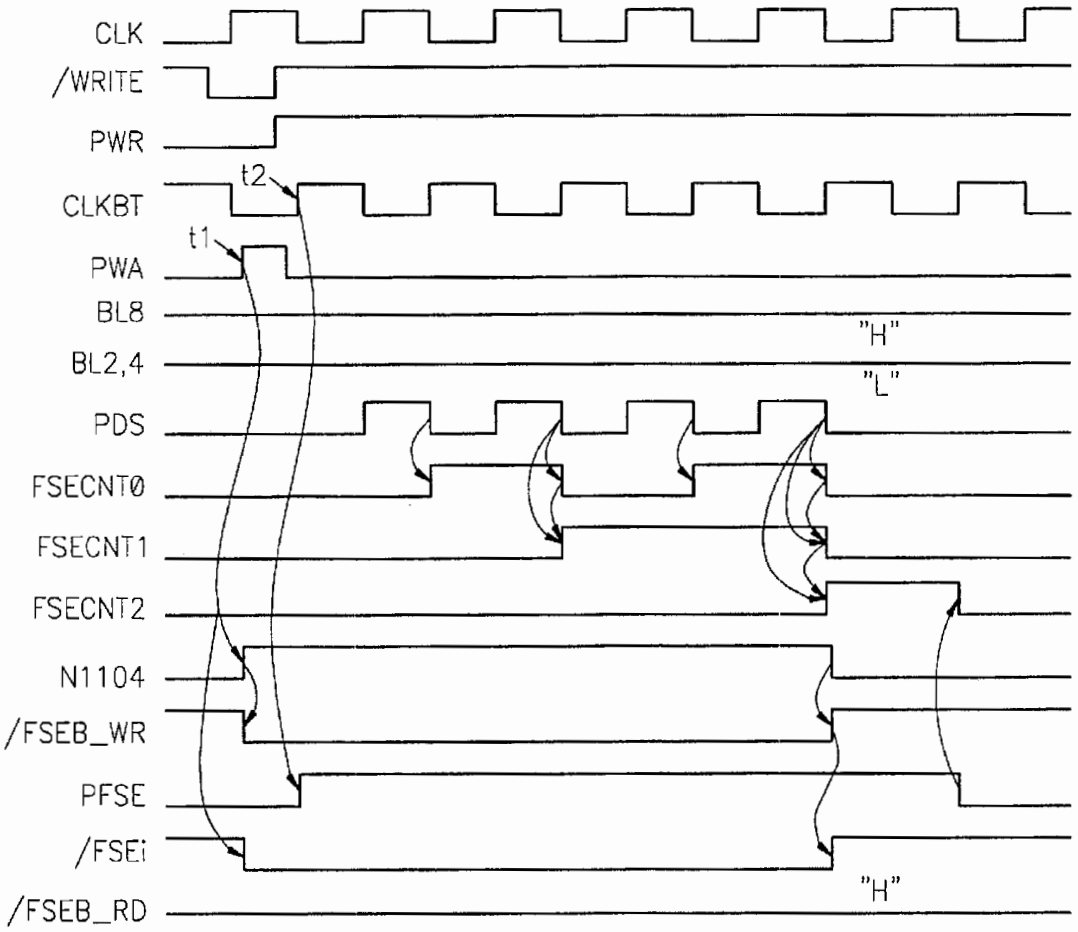

CLK
/WRITE
PWR
t2
CLKBT
t1
PWA
BL8
"H"
BL2,4
"L"
PDS
FSECNT0
FSECNT1
FSECNT2
N1104
/FSEB_WR
PFSE
/FSEi
"H"
/FSEB_RD

FIG. 13

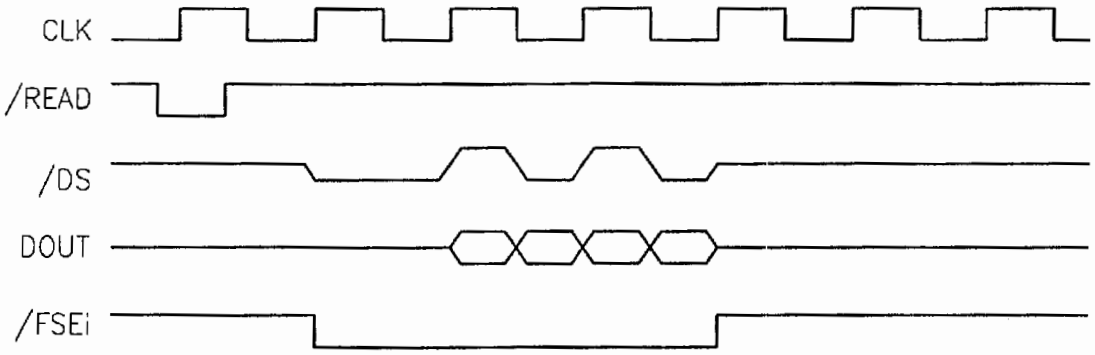


FIG. 14

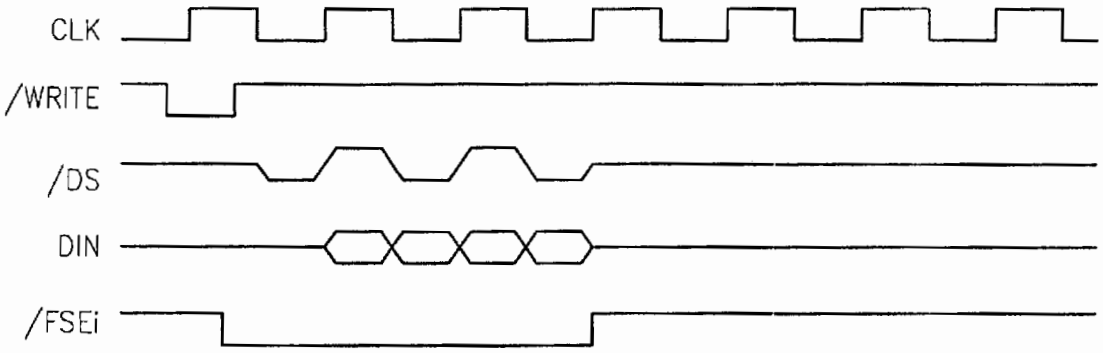
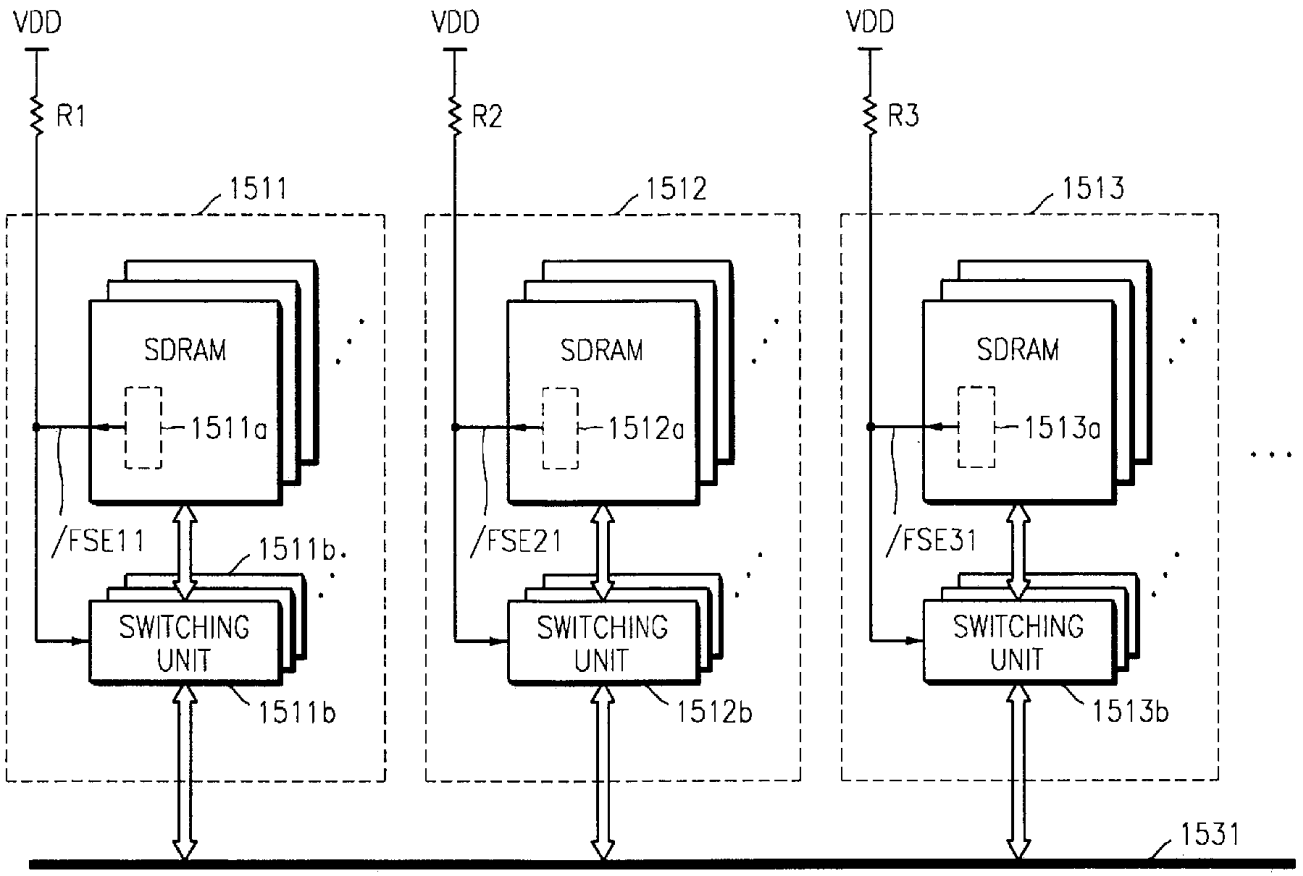


FIG. 15


VDD
R1
1511
SDRAM
1511a
/FSE11
1511b
SWITCHING UNIT
1511b
VDD
R2
1512
SDRAM
1512a
/FSE21
SWITCHING UNIT
1512b
VDD
R3
1513
SDRAM
1513a
/FSE31
SWITCHING UNIT
1513b
1531

MEMORY MODULE SYSTEM FOR CONTROLLING DATA INPUT AND OUTPUT BY CONNECTING SELECTED MEMORY MODULES TO A DATA LINE

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to a system with memory modules for controlling data input and output, and more particularly, to a system having a multiplicity of memory modules for controlling data input and output from and to a common data line by connecting selected memory modules to the common data line.

2. Description of the Related Art

A semiconductor memory device such as a dynamic random access memory (DRAM) can be used as a main memory in a system such as a computer. A memory using semiconductor memory devices generally has memory modules, such as dual in-line memory modules (DIMM) and single in-line memory modules (SIMM). A memory may be required to include a plurality of memory modules for operations of a system. In such a system, each of the memory modules uses a common data line. For instance, assuming that a main memory has a capacity of 32M bytes and is made with DIMMs each of which has a capacity of 8M bytes and made of four 16M (1M×16) DRAMs, the main memory of 32M bytes requires four slots, i.e., four DIMMs. Since the number of data bus lines is 64 and the number of data pins per DIMM is 64, the main memory of 32M bytes requires a multiplicity of data pins. Thus, four DIMMs constituting the main memory of 32M bytes use a common data bus line of a system. As the capacity of the main memory is increased, the necessity of using a common data line is increased.

The capacitance per data pin of a DIMM module is 20~25 pF. Thus, when the 32M byte memory is composed of four DIMMs of 1M×64, the capacitance per data pin becomes approximately 80~100 pF. Thus, the capacitance per data pin is so great that an increase in the load of an output driver increases power consumption, and that speed for writing and reading data is remarkably reduced.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide a memory module with minimized load capacitance per data pin to improve speed of writing and reading data and of the memory module.

To achieve the above and other objects, according to an aspect of the present invention, there is provided a memory module system for controlling data output to a data bus line, the system comprising a multiplicity of memory modules each of which includes a control signal generation circuit for generating a connection control signal having a width of activation corresponding to a burst length of output data, wherein the control signal generation circuit has a read control unit for activating input data during output of data of a predetermined burst length in a read mode to generate a read control signal, the read control unit receiving CAS latency information, a read command signal and burst length information of read data, and a multiplicity of switching units for electrically connecting output terminals of the memory modules to the data bus line in response to activation of the connection control signal. In the memory module system according to the present invention, the memory modules for receiving outputting data are electrically connected to the data bus line in response to activation of the connection control signals. The control signal generation circuit further comprises a write information generation unit. The write information generation unit receives a write command signal and burst length information of write data, and it is activated in a write mode during inputting data of a predetermined burst length to generate a write control signal.

In a memory module system for controlling data input and output according to the present invention, selected memory modules are connected to a data line during a data burst length in which data is input and output. Thus, load is minimized per data pin, so that speed of writing and reading data is improved.

BRIEF DESCRIPTION OF THE DRAWINGS

The above and other objects and advantages of the present invention will become more apparent by describing in detail a preferred embodiment thereof with reference to the attached drawings in which:

FIG. **1** is a block diagram schematically showing a system having a multiplicity of memory modules according to a first embodiment of the present invention;

FIG. **2** is a block diagram of one of the control signal generation circuits of FIG. **1**;

FIG. **3** is a circuit diagram of the read control unit of FIG. **2**;

FIG. **4** is a circuit diagram of the control signal generation unit of FIG. **2**;

FIGS. **5** through **7** are timing diagrams of control and data signals in FIGS. **3** and **4** in a read mode having CAS latencies of 1.5, 2 and 2.5, respectively;

FIG. **8** is a circuit diagram of the burst counter of FIG. **2**;

FIG. **9** is a circuit diagram of the preliminary counter of FIG. **8**;

FIG. **10** is a detailed circuit diagram of one of the main counters of FIG. **8**;

FIG. **11** is a circuit diagram of the write control unit of FIG. **2**;

FIG. **12** is a timing diagram of control and data signals related to the burst counter of FIG. **2**;

FIG. **13** is a timing diagram of external signals of memory modules in a read mode of a DDR mode in a system of the present invention;

FIG. **14** is a timing diagram of external signals of memory modules in a write mode of a DDR mode in a system of the present invention; and

FIG. **15** is a block diagram schematically showing a system having a multiplicity of memory modules according to a second embodiment of the present invention.

DESCRIPTION OF PREFERRED EMBODIMENTS

Referring to FIG. **1**, a first embodiment of a memory module system according to the present invention has a multiplicity of memory modules. For convenience, three memory modules, i.e., first through third memory modules **111**, **112** and **113**, are shown. The first through third memory modules **111**, **112** and **113** are selected, so that output data is transferred to a common data bus line **131** through switching units **121**, **122** and **123**.

Each of the switching units includes switches, through which data is output from a corresponding memory module.

Micron Technology Inc. et al.
Ex. 1010, p. 208

Also, the data output through the switches is transferred to the data line bus **131**. Each of the switching units **121**, **122** and **123** is turned on when data is input to and output from the corresponding memory module. However, when data is not input to or output from the corresponding memory module, each of the switching units **121**, **122** and **123** is turned off. For instance, assuming that the first memory module **111** is selected, to input or output data to or from the first memory module **111**, a first connection control signal /FSE1 is activated to logic "low", so that the first switching unit **121** is turned on. Thus, output data of the first memory module **111** is transferred to the data bus line **131** through the first switching unit **121**. At this time, second and third connection control signals /FSE2 and /FSE3 are deactivated to logic "high" by an external power supply voltage VDD applied through resistors R2 and R3, respectively. The resistors R2 and R3 may be external to the memory modules **112** and **113**.

In detail, the resistors R1 to R3 have relatively high resistances. Thus, voltage levels of the first through third connection control signals /FSE1, /FSE2 and /FSE3 are determined by the first through third memory modules **111**, **112** and **113**, respectively. However, when signals to determine the voltage levels of the connection control signals /FSE1, /FSE2 and /FSE3 in the respective memory modules are in floating state, the voltage levels of the first through third connection control signals /FSE1, /FSE2 and /FSE3 become an external power supply voltage VDD, i.e., logic "high", applied through the resistors R1 to R3 connected to the respective memory modules. That is, when the first memory module **111** is selected, the second and the third connection control signals /FSE2 and /FSE3 become logic "high". The second and the third connection control signals /FSE2 and /FSE3 deactivated to the "high" then turn off the second and the third switching units **122**, **123**. Thus, load of data pins of the second and third memory modules **112** and **113** is separated from the data bus line **131**. In other words, memory modules other than the selected memory module are separated from the data line **131**, to thereby reduce load of the data bus line **131**.

Preferably, each of the first through third connection control signals /FSE1, /FSE2 and /FSE3 is generated by the corresponding memory modules **111**, **112** and **113**. Control signal generation circuits **111***a*, **112***a* and **113***a* which are included in the corresponding memory modules **111**, **112** and **113** generate the first through third connection control signals /FSE1, /FSE2 and /FSE3, respectively. Also, activation intervals of the first through third connection control signals /FSE1, /FSE2 and /FSE3 are preferably determined by burst lengths of input or output data. Also, the first through third connection control signals /FSE1, /FSE2 and /FSE3 may be provided from a control unit (not shown) which transfers the input data from the outside of the memory modules.

FIG. **2** is a block diagram illustrating a preferred embodiment of a control signal generation circuit according to the present invention. As described above, the control signal generation circuits **111***a*, **112***a* and **113***a* each having the same configuration as in FIG. **2** are included in the memory modules **111**, **112** and **113**, respectively. For convenience, the first through third connection control signals /FSE1, /FSE2 and /FSE3 are indicated by /FSEi.

Referring to FIG. **2**, the control signal generation circuit includes a read control unit **201**, a control signal generation unit **203**, a write control unit **205** and a burst counter **207**. The read control unit **201** receives column address strobe (CAS) latency signals CL1.5 and CL2, a read command signal PCR and a read strobe signal PTRSTDSB, and generates a read control signal /FSEB_RD. The CAS latency signals CL1.5 and CL2 include /CAS latency information, and the read command signal PCR introduces a read command for a read mode. The read strobe signal PTRSTDSB includes information on a burst length of read data, for example, information on burst signals BL2, BL4 and BL8 input to the write control unit **205**. Also, the read control signal /FSEB_RD is activated in response to the burst length of the read data in the read mode. The read control unit **201** will be described in detail with reference to FIG. **3**.

The control signal generation unit **203** supplies the connection control signal /FSEi in response to the read control signal /FSEB_RD and a write control signal /FSEB_WR. The write control signal /FSEB_WR supplied by the write control unit **205**, is activated for an interval corresponding to a burst length of data input in a write mode. Thus, the connection control signal /FSEi is activated in response to the read control signal /FSEB_RD in the read mode, and also activated in response to the write control signal /FSEB_WR in the write mode. The control signal generation unit **203** will be described in detail, with reference to FIG. **4**.

The write control unit **205** receives burst signals BL2, BL4 and BL8 and first and second write command signals PWA and PWR, and generates a write control signal /FSEB_WR and a counter control signal PFSE. Here, the burst signals BL2, BL4 and BL8 include information on burst lengths of data in a write mode. The first write command signal PWA is a pulse indicating that a memory module enters into a write mode, in response to a write command. The second write command signal PWR indicates that a memory module is in a write mode.

The write control signal /FSEB_WR which is activated for an interval corresponding to the data burst length in which data is input in the write mode, is supplied to the control signal generation unit **203**. Also, the counter control signal PFSE which is supplied to the burst counter **207**, is activated in response to an inverted clock signal CLKBT which is firstly generated after generating the write command signal, and deactivated after receiving a predetermined number of data. The inverted clock signal CLKBT is an inverted signal of an external clock signal CLK (not shown). The write control unit **205** will be described in detail with reference to FIG. **11**.

The burst counter **207** is enabled in response to the activation of the counter control signal PFSE. The burst counter **207** receives a first write command signal PWA and a data strobe signal PDS, and counts the number of input data. Also, a counting signal FSECNTi (where, i=0 to 2) indicating the number of input data is supplied to the write control unit **205**. The data strobe signal PDS indicating an input of data, is formed by an external data strobe signal of a memory module. The data strobe signal PDS indicates an input or output of one of the data by one transition in a double data rate (DDR). The memory module system of the present invention will be described in the DDR mode. The number of input data can be indicated by a combination of the counting signals FSECNTi, where i=0 to 2. The burst counter **207** will be described in detail with reference to FIG. **8**.

A control signal generation circuit which generates the connection control signal /FSEi can be realized with the read control unit **201** and the control signal generation unit **203** of FIG. **2** in a read mode. A write information generation unit

209 for activating the write control signal /FSEB_WR for an interval corresponding to the burst length of data input may be realized in a write mode with the write control unit 205 and the burst counter 207. Thus, the control signal generation circuit for generating the connection control signal /FSEi may be realized in the write mode with the write information generation unit 209 and the control signal generation unit 203 of FIG. 2.

FIG. 3 is a circuit diagram of a preferred embodiment of the read control unit 201 in FIG. 2. Referring to FIG. 3, the read control unit 201 includes a first latency response unit 301, a second latency response unit 303, a first transfer unit 309, a first latch unit 307, a second transfer unit 305 and a second latch unit 311.

The first latency response unit 301 is enabled in response to a CAS latency shorter than a predetermined latency. According to a preferred embodiment of the present invention, the first latency response unit 301 is enabled, when the CAS latency is 1.5. The first latency response unit 301 includes a logic gate 301*a* and a drive transistor 301*b*. The logic gate 301*a* is enabled by activation of CL1.5 indicating the CAS latency of 1.5, and responds to a read command signal PCR. The read command signal PCR is activated in response to a first clock signal from generation of a read command, and. is maintained in the activated state during a period of one clock. The drive transistor 301*b* is gated by an output signal 301*n* of the logic gate 301*a*. Thus, the transistor 301*b* outputs an output signal to the first latch unit 307 in response to activation of the read command signal PCR. Preferably, the logic gate 301*a* is a NAND gate, and the transistor 301*b* is a PMOS transistor.

When the CAS latency is two or more, the second latency response unit 303 is enabled. An output signal N304 of the second latency response unit 303 maintains a logic "high" state by activation of a read strobe signal PTRSTDSB. The second latency response unit 303 includes logic gates 303*a* and 303*b*. The logic gate 303*a* is enabled by activation of CL2 indicating the CAS latency of 2, and responds to the read command signal PCR. Also, the logic gate 303*b* is enabled by the output signal 303*n* of the logic gate 303*a*, and responds to the read strobe signal PTRSTDSB. Preferably, the logic gates 303*a* and 303*b* are NAND gates. The second transfer unit 305 transfers the output signal N304 of the second latency response unit 303 in response to a first clock signal CLK_SM.

The first latch unit 307 latches either the output signal of the first latency response unit 301 or the output signal N304 of the second latency response unit 303 transferred by the second transfer unit 305. The first transfer unit 309 transfers signals latched by the first latch unit 307 in response to a second clock signal CLK_FM. The second latch unit 311 latches a signal N306 transferred by the first transfer unit 309, and outputs the latched signal as a read control signal /FSEB_RD.

Here, the first and the second clock signals CLK_SM and CLK_FM are pulse signals generated in response to rising and falling transitions of the external clock signal CLK. The first and the second clock signals CLK_SM and CLK_FM which control data output timing, have timing with respect to the external clock signal CLK as shown in FIGS. 5 through 7.

FIG. 4 is a circuit diagram illustrating a preferred embodiment of the control signal generation unit 203 of FIG. 2. Referring to FIG. 4, the control signal generation unit 203 includes a NAND gate 401, a PMOS transistor 403 and an NMOS transistor 405. The NAND gate 401 receives an inverted signal of the read control signal /FSEB_RD, an inverted signal of the write control signal /FSEB_WR and the first write command signal PWA. That is, the NAND gate 401 is enabled by the first write command signal PWA, and acts as a logic gate responding to the read control signal /FSEB_RD or the write control signal /FSEB_WR.

The NMOS transistor 405 is gated by an output signal of the NAND gate 401. Thus, the NMOS transistor 405 acts for driving connection control signal in response to the read control signal /FSEB_RD or the write control signal /FSEB_WR.

The PMOS transistor 403 has source and gate terminals which are in common connected. A drain terminal of the PMOS transistor 403 and a drain terminal of the NMOS transistor 405 are in common connected to generate the connection control signal /FSEi.

In a precharge state, the read control signal /FSEB_RD and the write control signal /FSEB_WR are both logic "high" and the first write command signal PWA is logic low. Thus, the output signal of the NAND gate 401 becomes low and the NMOS transistor 405 is turned off. The connection control signal /FSEi becomes logic "high" in response to an external power supply voltage VDD applied through a resistor outside a memory module.

An embodiment of a control signal generation circuit for generating a connection control signal can be realized with the read control unit 201 in FIG. 3 and the control signal generation unit 203 in FIG. 4 in a read mode.

FIGS. 5 through 7 are timing diagrams of control signals and data signals in FIGS. 3 and 4 in a read mode in which CAS latencies are 1.5, 2, and 2.5, respectively. Assuming that a data burst length is 4, and the memory modules operate in a DDR mode, the operation of the read control unit 201 when the CAS latency is 1.5, will FIG. 5 be described with reference to FIGS. 3, 4 and 5.

Referring to FIG. 5, a logic state of the CAS latency signal CL1.5 is high, and a logic state of the CAS latency signal CL2 is low. Activation of /READ signal to low state indicates a generation of a read command. The read command signal PCR is activated for one clock period in response to a rising edge of the first clock pulse CLK1 after the generation of the read command. Thus, the output signal N306 of the first latency response unit 301 becomes logic "high" in response to the read command signal PCR, and is latched to high by the first latch unit 307. That is, when the read command is generated, the output signal N306 of the first latency response unit 301 is latched to high by the first latch unit 307, regardless of a level of the clock signal.

The read strobe signal PTRSTDSB is activated to logic "low" in response to activation of the read command signal, and is activated to logic "high" after a period of time corresponding to the data burst length. The output signal N304 of the second latency response unit 303 is transferred to the first latch unit 307 by the second transfer unit 305 turned on in response to the first clock signal CLK_SM. Thus, when the read strobe signal PTRSTDSB is activated to logic "low", the input signal at the terminal N306 of the first latch unit 307 is maintained at logic "high" even when the read command signal PCR is deactivated to logic "low".

Data latched by the first latch unit 307 is transferred to the second latch unit 311 by the first transfer unit 309 turned on in response to the second clock signal CLK_FM. The output signal of the second latch unit 311 is connected to an output terminal of the read control signal /FSEB_RD. Thus, the read control signal /FSEB_RD is activated to logic "low" in response to the generation of the read command, and then

deactivated to high in a period of time corresponding to the burst length. That is, a time period of the activation of the read control signal /FSEB_RD corresponds to that of the burst length.

In a read mode, the first write command signal PWA is maintained at logic "low", and the write control signal /FSEB_WR is maintained at logic "high". Thus, in the read mode, activation of the connection control signal /FSEi depends on activation of the read control signal /FSEB_RD.

In a preferred embodiment of the present invention, when the CAS latency is 1.5, the input signal **N306** of the first latch unit **307** is activated regardless of the clock signal, and the activation is maintained by the read strobe signal PTRSTDSB. When the read control signal /FSEB_RD is activated in response to the first clock signal CLK_SM, the connection control signal /FSEi is not activated until a first data burst is read, despite of the CAS latency of 1.5.

The operation of the read control unit **201** with the CAS latency of 2 will now be described, with reference to FIGS. **3**, **4** and **6**. A logic state of the CAS latency signal CL**2** is high, and a logic state of the CAS latency signal CL**1.5** is low. Thus, the output signal **N306** of the first latency response unit **301** does not respond to the read command signal PCR.

The output signal **N304** of the second latency response unit **303** is activated to high in response to an activation of the read command signal PCR, and is transferred to the first latency unit **307** in response to an activation of the first clock signal CLK_SM. The logic state of the input signal **N306** of the first latch unit **307** is maintained at high in response to activations of both the read strobe signal PTRSTDSB and the first clock signal CLK_SM.

Thus, the read control signal /FSEB_RD, similar to the case of the CAS latency of 1.5, is activated to low in response to a generation of the read command, and deactivated to high after a period corresponding to the burst length. That is, the read control signal /FSEB_RD is activated for an interval corresponding to the burst length.

The operation of the read control unit **201** when the CAS latency is 2 or more, i.e., the logic states of the CAS latency signals CL**1.5** and CL**2** are low, will now be described with reference to FIGS. **3**, **4** and **7**. The output signal **N304** of the second latency response unit **303** becomes high after the read strobe signal PTRSTDSB is activated to low. The output signal **N304** of the second latency response unit **303** is transferred to the first latch unit **307** in response to activation of the first clock signal CLK_SM. When the read strobe signal PTRSTDSB is deactivated to high, the input signal **N306** of the first latch unit **307** becomes low. Thus, the read control signal /FSEB_RD is activated or deactivated by the read strobe signal PTRSTDSB, unlike the case in which the CAS latency is 1.5 or 2.

As described above, the control signal generation circuit according to a preferred embodiment of the present invention generates a connection control signal activated in response to the data burst length in which data is read from a memory module in a read mode. Also, only the memory module selected by the connection control signal is connected to a data line during the data burst length in which data is output.

FIG. **8** is a circuit diagram illustrating an embodiment of the burst counter **207** of FIG. **2**. In FIG. **8**, a burst counter capable of counting eight data bursts is shown. The burst counter **207** shown in FIG. **8** includes one preliminary counter **807** and two main counters **803** and **805**. The preliminary counter **807** and the main counters **803** and **805** are reset by activation of a counter control signal PFSE. That is, when a write command is generated so that the counter control signal PFSE is activated to high, the preliminary counter **807** and the main counters **803** and **805** are reset. Output signals of the preliminary counter **807** and the main counters **803** and **805**, i.e., first through third counting signals FSECNT**0**~**2** are reset to low.

As described above, a data strobe signal PDS indicates inputs of external data during rising and falling transitions in a DDR mode. Structure and operation of the burst counter **207** in a DDR mode will now be described.

In an embodiment of the present invention, data is input during rising and falling edges of the data strobe signal PDS, and the data is continuously input when a phase transition of the data strobe signal PDS is repeated. The first counting signal FSECNT**0** performs the phase transition at every falling edge of the input data strobe signal. Also, when carries CARRYBs of the main counters **803** and **805** are activated to low, the second and the third counting signals FSECNT**1**~**2** repeat the phase transition every two inputs of data.

Referring to FIG. **8**, the carry CARRYB of the main counter **805** is activated to low, when the first counting signal FSECNT**0** is activated to high. The carry CARRYB of the main counter **803** is activated to low, when the first and the second counting signals FSECNT**0**~**1** are high.

The preliminary counter **807** and the main counters **803** and **805** of FIG. **8** will now be described with reference to FIGS. **9** and **10**, respectively. FIG. **9** is a detailed circuit diagram illustrating an embodiment of the preliminary counter **807** of FIG. **8**. Referring to FIG. **9**, the preliminary counter **807** includes first and second transfer units **901** and **905**, first and second latch units **903** and **907**, and a NOR gate **909**. The NOR gate **909** performs an OR operation on a reset signal FSESET and an inverted signal PDSB of the data strobe signal PDS and inverts the OR-operated signals. The reset signal FSESET is connected to the counter control signal PFSE. Thus, when the reset signal FSESET is disabled to low, the preliminary counter **807** is enabled to respond to the phase transition of the inverted signal PDSB. That is, when a phase of the inverted signal PDSB is transited from high to low, the first transfer unit **901** is turned on. When the phase of the inverted signal PDSB is transited from low to high, the second transfer unit **905** is turned on. Whenever the inverted signal PDSB is deactivated, the first counting signal FSECNT**0** repeats the phase transition.

FIG. **10** is a detailed circuit diagram illustrating an embodiment of the main counters **803** and **805** of FIG. **8**. The main counters **803** and **805** may have the same structure as shown in FIG. **10**. For convenience, output signals of the main counters **803** and **805**, i.e., second and third counting signals FSECNT**1**~**2**, are indicated by a counting signal FSECNTi.

A main counter has a structure and operation similar to those of the preliminary counter **807** of FIG. **9**. However, a NOR gate **1009** with three input terminals is used, to which a reset signal FSESET, an inverted signal PDSB and a carry CARRYB are provided. Thus, when the reset signal FSESET and the carry CARRYB are low, the main counter is enabled in response to a phase transition of the inverted signal PDSB. That is, when the reset signal FSESET and the carry CARRYB are low, the counting signal FSECNTi repeats the phase transition whenever the inverted signal PDSB is deactivated.

FIG. **11** is a detailed circuit diagram illustrating an embodiment of the write control unit **205** of FIG. **2**. Refer-

ring to FIG. **11**, the write control unit **205** includes a counting response unit **1101**, a latch unit **1103** and a counter control signal generation unit **1105**. The counting response unit **1101** indicates a data burst of input data. That is, a phase of an output signal **N1102** is transited when a predetermined number of data is input.

When the first write command signal PWA is activated to high, the latch unit **1103** is enabled to respond to the output signal **N1102** of the counting response unit **1101**. When the second write command signal PWR becomes high, an output signal **N1104** of the latch unit **1103** is latched to high. At this time, a write control signal /FSEB_WR connected to an output terminal inverting the output signal **N1104** of the latch unit **1103** is activated to low. When the first write command signal PWA becomes low, the output signal **N1104** of the latch unit **1103** is maintained at high. When the second write command signal PWR is high and the predetermined number of data is input so that the output signal **N1102** of the counting response unit **1101** is transited to low, the output signal **N1104** of the latch unit **1103** becomes low, and the write control signal /FSEB_WR is deactivated to high.

The counter control signal generation unit **1105** includes first and second transfer units **1105***a* and **1105** and first and second latches **1105***b* and **1105***d*. The first transfer unit **1105***a* transfers the output signal **N1104** of the latch unit **1103** in response to a falling edge of an inverted clock signal CLKBT. The first latch unit **1105***b* latches a signal transferred by the first transfer unit **1105***a*. The second transfer unit **1105***c* transfers an output signal of the first latch **1105***b* in response to a rising edge of the inverted clock signal CLKBT. Also, the second latch **1105***d* latches a signal transferred by the second transfer unit **1105***c*, and generates a counter control signal PFSE. Here, the inverted clock signal CLKBT is an inverted signal of an external clock signal CLK.

FIG. **12** is a timing diagram of signals related to the burst counter illustrated in FIGS. **2** and **8** through **11** in a write mode. FIG. **12** shows a case of a data burst of 8 in a DDR mode. Operation of the write control unit **205**, the burst counter **207**, the preliminary counter **807** and the main counters **803** and **805** of the burst counter **207** will be described with reference to FIG. **12**.

The data burst is 8, that is, burst signal BL**8** is high and burst signals BL**2** and BL**4** are low. The DDR mode indicates that data is input in every phase transition of the data strobe signal PDS. When the first write command signal PWA becomes high (t1), the output signal **N104** of the latch unit **1103** is latched to high, and the write control signal /FSEB_WR is activated to low. When the second write command signal PWR is activated to high and activation and deactivation of the data strobe signal PDS are performed four times, i.e., 8 data are input, both the first and second counting signals FSECNT**0**~**1** become low and the third counting signal FSECNT**2** becomes high. At this time, the phase of the output signal **N1102** of the counting response unit **1101** is transited to low from high, and the output signal **N1104** of the latch unit **1103** becomes low. The write control signal /FSEB_WR is deactivated to high. The counter control signal PFSE is deactivated to low in response to a rising edge of the inverted clock signal CLKBT. The preliminary counter **807** and the main counters **803** and **805** are reset, so that the first through third counting signals FSECN**10**~**2** become low.

Referring to FIG. **4** again, a connection control signal /FSEi is activated and deactivated in response to the write control signal /FSEB_WR in a write mode. During the write mode the read control signal /FSEB_RD is deactivated to high.

FIG. **13** is a timing diagram of external signals of memory modules when reading data in a DDR mode in a system of the present invention. Referring to FIG. **13**, the CAS latency is 2, and the burst length is 4. When a predetermined preamble interval, i.e., 1 clock cycle, passes after a read command is input, i.e., a signal /READ is activated, an external data strobe signal /DS is activated to low. The connection control signal /FSEi is also enabled to low, so that switching units **121**, **122** and **123** (see FIG. **1**) connected to the activated memory modules are turned on. First, data are output from the memory modules, and then the output data are transferred to the data line **131** (see FIG. **1**). When a postamble interval, i.e., a half clock, passes after the last data of a data burst is output, the external data strobe signal /DS becomes high impedance state. At this time, the connection control signal /FSEi of low is disabled to high, to thereby turn off the switching units **121**, **122** and **123**.

FIG. **14** is a timing diagram of external signals of memory modules in a write mode of a DDR mode in a system of the present invention. In FIG. **14**, the burst length is 4. In a ¼ clock cycle after a write command is input, i.e., a signal /WRITE is activated, the connection control signal /FSEi is enabled to low. Activated memory modules receive data corresponding to the burst length from a control unit (not shown) through the data bus line **131** (see FIG. **1**). When data holding time passes after last data is received, the connection control signal /FSEi of low is disabled to high to turn off the corresponding switching units **121**, **122** and **123**.

As described above, a control signal generation circuit according to a preferred embodiment of the present invention generates a connection control signal activated during an interval corresponding to the burst length in which data is written in a write mode. Only the memory module selected by the connection control signal is connected to a data line during the data burst length in which data is input.

FIG. **15** is a block diagram illustrating a second embodiment of the memory module system according to the present invention. Structure and operation of the memory module system according to the second embodiment are similar to those of the first embodiment. Switching units **1511***b*, **1512***b* and **1513***b* corresponding to the switching units **121**, **122** and **123** of FIG. **1** are installed in memory modules **1511**, **1512** and **1513**, respectively.

As shown in FIG. **15**, a multiplicity of SDRAMs (synchronous DRAMs) are installed in each of the memory modules. Also, a multiplicity of switching units each corresponding to each of the SDRAMs are installed in each of the memory modules. Connection control signals /FSE**11**, /FSE**21** and /FSE**31** controlling the switching units are generated by the SDRAMs included in the memory modules. For convenience, the connection control signals /FSE**11**, /FSE**21** and /FSE**31** are shown for controlling the switching units **1511***b*, **1512***b* and **1513***b*, respectively, each of which connects with a first SDRAM in a memory module. In FIG. **15**, although not shown, switching units connected to other SDRAMs exist, and connection control signals for controlling the switching units are generated from the corresponding SDRAMs. A switching unit may preferably have as many switches as output terminals for outputting data from an SDRAM corresponding to the switching unit. The data output through the multiple switches are transferred to a data bus line **1531**. Structure and operation of control signal generation circuits **1511***a*, **1512***a* and **1513***a* of the second embodiment in FIG. **15** are the same as those of the control signal generation circuits **111***a*, **112***a* and **113***a* of the first embodiment in FIGS. **2** through **14**.

Having described preferred embodiments of a memory module system of the present invention, it is noted that

modification and variations can be made by persons skilled in the art in light of the above teachings. For example, the switches which are arranged outside the SDRAM may be arranged in the SDRAM. Also, other integrated memory circuit devices such as DRAMs or SRAMs may be used instead of SDRAMs. In a DDR mode, the counter circuit of the first embodiment of the present invention counts eight data by using one preliminary counter and two main counters. However, the number of main counters may be increased to increase the number of counted data.

Also, in a memory module system of the present invention, only selected memory modules are connected to a data line during a time period of data burst length in which data is input and output, so that load per data pin is minimized and data writing and reading speed is improved.

What is claimed is:

1. A memory module system for controlling data output to a data bus, the system comprising:

a multiplicity of memory modules each of which includes a control signal generation circuit for generating a connection control signal having a width of activation corresponding to a burst length of output data, wherein the control signal generation circuit has a read control unit for activating input data during output of data of a predetermined burst length in a read mode to generate a read control signal, the read control unit receiving CAS latency information, a read command signal and burst length information of read data; and

a multiplicity of switching units for electrically connecting output terminals of the memory modules to the data bus in response to activation of the connection control signal;

wherein the read control unit comprises:

a first latency response unit for being enabled in a first CAS latency;

a second latency response unit for being enabled in a second CAS latency;

a first latch unit for latching one of an output signal of the first latency response unit and an output signal of the second latency response unit; and

a first transfer unit for transferring a latched signal by the first latch unit to generate the read control signal in response to a first internal clock signal that is synchronized with a first external clock signal.

2. The memory module system of claim **1**, wherein the first CAS latency is 1.5, and the second CAS latency is 2.

3. The memory module system of claim **2**, wherein the connection control signal is activated in response to activation of a read strobe signal including the burst length information in case of a third CAS latency that is greater than the first and the second CAS latencies.

4. The memory module system of claim **1**, wherein the first latency response unit comprises:

a logic gate responding to activation of the first CAS latency and a read command signal; and

a drive transistor for driving an output signal to the first latch unit in response to an output signal of the logic gate.

5. The memory module system of claim **1**, wherein the second latency response unit comprises a logic gate responding to a predetermined read strobe signal which is activated during a period corresponding to the burst length of data.

6. The memory module system of claim **1**, wherein the memory modules are synchronized with a second external clock signal, and the read control unit further comprises:

a second transfer unit for transferring the output signal of the second latency response unit in response to a second internal clock signal synchronized with the second external clock signal; and

a second latch unit for latching a signal transferred by the first transfer unit to generate the read control signal.

7. The memory module system of claim **1**, wherein the connection control signal is activated by a read strobe signal including the first CAS latency and the burst length information.

8. The memory module system of claim **1**, wherein the control single generation circuit further comprising a control signal generation unit for generating the connection control signal in response to activation of the read control signal output from the read control unit.

9. The memory module system of claim **8**, wherein the control signal generation unit comprises:

a logic gate responding to the read control signal; and

a drive transistor for activating the connection control signal in accordance with the activation of the read control signal, the drive transistor being gated by an output signal of the logic gate.

10. A memory module system for controlling data output to a data bus, comprising:

a multiplicity of memory modules; and

a multiplicity of switching units for electrically connecting only memory module outputting data to the data bus in response to activation of a connection control signal, wherein the connection control signal is generated in the memory module outputting data and has a width of activation corresponding to a burst length of the output data;

wherein the connection control signal, with respect to a first and a second CAS latencies, is activated in response to a CAS latency and a read command signal, and the activation is maintained by a read strobe signal including the burst length information, and the connection control signal, with respect to a third CAS latency longer than the first and the second CAS latencies, is activated in response to activation of the read strobe signal.

11. A memory module system for controlling data input from a data bus, comprising:

a multiplicity of memory modules each of which includes a control signal generation circuit for generating a connection control signal having an activation width corresponding to a burst length of input data, wherein the control signal generation circuit includes a write information generation unit for receiving a write command signal and burst length information of write data, the write information generation unit being activated in a write mode during inputting data of a predetermined burst length to generate a write control signal; and

a multiplicity of switches for electrically connecting input terminals of the memory modules to the data bus in response to activation of the connection control signal;

wherein the write information generation unit comprises:

a write control unit for receiving the write command signal and for supplying the write control signal and a counter control signal; and

a burst counter for counting the number of external input data to supply the counted data to the write control unit, wherein the burst counter is enabled in response to activation of the counter control signal;

wherein the write control unit comprises:

a counting response unit for generating an output signal in response to input of data of a predetermined burst length in response to an output signal of the burst counter;

a latch unit for being enabled in response to the write command signal, for latching the output signal of the counting response unit, and for supplying the write control signal; and

a counting control signal generation unit for being enabled in response to the write command signal and for inputting the write control signal to supply the counter control signal.

12. The memory module system of claim **11**, wherein the memory modules are synchronized with an external clock signal, and the latch unit comprises:

a first logic gate for being enabled by a first write command signal to generate the write control signal; and

a second logic gate for being enabled by a second write command signal and for responding to an output signal of the counting response unit;

wherein the data of the burst length is input, and then the write control signal is deactivated by the output signal of the counting response unit.

13. The memory module system of claim **11**, wherein the counting control signal generation unit comprises:

a first transfer unit for transferring an output signal of the latch unit in response to a first transition of the external clock signal;

a first latch for latching a signal transferred by the first transfer unit;

a second transfer unit for transferring the signal latched by the first latch in response to a second transition of the external clock signal; and

a second latch for latching the signal transferred by the first transfer unit to generate the counter control signal;

wherein a direction of the first transition is opposite to that of the second transition.

14. The memory module system of claim **11**, wherein the control signal generation circuit further comprises a control signal generation unit for being activated in response to the write command signal, and the activation is maintained by the write control signal.

15. The memory module system of claim **14**, wherein the control signal generation unit comprises:

a logic gate responding to the write information signal; and

a drive transistor for activating the connection control signal in response to the activation of the write control signal, the drive transistor being gated by an output signal of the logic gate.

16. A semiconductor memory device comprising:

a multiplicity of integrated memory circuit devices including a control signal generation circuit for generating a connection control signal having an activation width corresponding to a burst length of output data, wherein the control signal generation circuit comprises a read control unit for activating input data during output of data of a predetermined burst length in a read mode to generate a read control signal, the read control unit receiving CAS latency information, a read command signal and burst length information of data to be read; and

at least one switching unit for connecting output terminals of the integrated memory circuit devices where data is output and a data bus in response to activation of the connection control signal;

wherein the read control unit comprises:

a first latency response unit for being enabled in a first CAS latency;

a second latency response unit for being enabled in a second CAS latency;

a first latch unit for latching one of an output signal of the first latency response unit and an output signal of the second latency response unit; and

a first transfer unit for transferring a latched signal by the first latch unit to generate the read control signal in response to a first internal clock signal that is synchronized with a first external clock signal.

17. The semiconductor memory device of claim **16**, wherein the first CAS latency is 1.5, and the second CAS latency is 2.

18. The semiconductor memory device of claim **16**, wherein the first latency response unit comprises:

a logic gate responding to the activation of the first CAS latency and a predetermined read command signal; and

a drive transistor for driving an output signal to the first latch unit in response to an output signal of the logic gate.

19. The semiconductor memory device of claim **16**, wherein the second latency response unit comprises a logic gate responding to a predetermined read strobe signal, wherein activation of the read strobe signal is maintained during a period corresponding to a burst length of minimum data.

20. The semiconductor memory device of claim **16**, wherein the semiconductor memory device is synchronized with a second external clock signal, and the read control further comprises:

a second transfer unit for transferring the output signal of the second latency response unit to the first latch unit in response to a second internal clock signal synchronized with the second external clock signal; and

a second latch unit for latching a signal transferred by the first transfer unit to generate the read control signal.

21. The semiconductor memory device of claim **16**, wherein activation of the read control signal is maintained by a read strobe signal including the first CAS latency and the burst length information.

22. The semiconductor memory device of claim **21**, wherein the read control signal, with respect to a third CAS latency longer than the first and the second CAS latencies, is activated in response to activation of the read strobe signal including the burst length information.

23. A semiconductor memory device comprising:

a multiplicity of the integrated memory circuits including at least one control signal generation circuit for generating a connection control signal having an activation width corresponding to a burst length of input data, wherein the control signal generation circuit includes a write information generation unit for receiving burst length information of write command and write data, the control signal generation circuit being activated during input of data of a predetermined burst length in a write mode to generate a write control signal; and

at least one switching unit for connecting input terminals of the integrated memory circuits where data is input and a data bus in response to activation of the connection control signal;

wherein the write information generation unit comprises:

a write control unit for receiving the write command signal to supply the write control signal and a counter control signal; and

a burst counter for counting the number of external input data to supply the counted data to the write control unit, wherein the burst counter is enabled in response to activation of the counter control signal;

wherein the write control unit comprises:

a counting response unit for generating an output signal responding to input of data having a predetermined burst length in response to an output signal of the burst counter;

a latch unit for being enabled in response to the write command signal to latch the output signal of the counting response unit, and for supplying the write control signal; and

a counting control signal generation unit for being enabled in response to the write command signal and for inputting the write control signal to supply the counter control signal.

24. The semiconductor memory device of claim **23**, wherein the semiconductor memory device is synchronized with an external clock signal, and the latch unit comprises:

a first logic gate for being enabled by a first write command signal to generate the write control signal; and

a second logic gate for being enabled by a second write command signal to respond to an output signal of the counting response unit;

wherein the write control signal is deactivated by the output signal of the counting response unit after input of data having a predetermined burst length.

25. The semiconductor memory device of claim **23**, wherein the counting control signal generation unit comprises:

a first transfer unit for transferring an output signal of the latch unit in response to a first transition of the external clock signal;

a first latch for latching a signal transferred by the first transfer unit;

a second transfer unit for transferring a signal latched by the first latch in response to a second transition of the external clock signal; and

a second latch for latching a signal transferred by the first transfer unit, to generate the counter control signal;

wherein a direction of the first transition is opposite to that of the second transition.

26. The semiconductor memory device of claim **23**, wherein the control signal generation circuit further comprises a control signal generation unit for generating the connection control signal activated in response to the write command signal and for maintaining the activation in accordance with the write control signal.

27. The semiconductor memory device of claim **26**, wherein the control signal generation unit comprises:

a logic gate responding to the write information signal; and

a drive transistor for activating the connection control signal in accordance with activation of the write control signal, the drive transistor being gated by an output signal by the logic gate.

* * * * *

PATENT ABSTRACTS OF JAPAN

(11)Publication number : 09-237492
(43)Date of publication of application : 09.09.1997

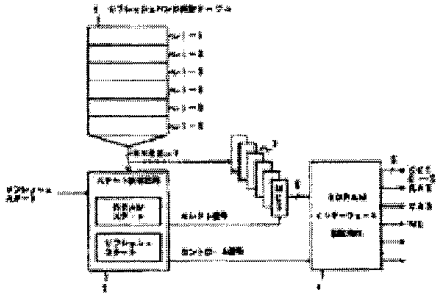
(51)Int.Cl. G11C 11/406
G06F 12/00

(21)Application number : 08-045057
(22)Date of filing : 01.03.1996

(71)Applicant : TOSHIBA CORP
TOSHIBA COMPUT ENG CORP
(72)Inventor : MATOBA TSUKASA
TOYODA YASUHIRO

(54) MEMORY CONTROL DEVICE

(57)Abstract:
PROBLEM TO BE SOLVED: To always suppress the amount of the current flowing at the time of refreshing synchronous DRAMs (dynamic random access memory 0) bellow a prescribed value regardless of kinds and the number of synchronous DRAMs being used.
SOLUTION: The content of refresh control flags is successively checked from the register 1-1 of the first stage in a reflesh bank specifying table 1 and refleshings are executed only to shychronous banks whose executions of refleshings are specified by the reflesh control flags. Thus, the number of synchronous DRAM banks whose refleshings are executed simultaneously and the combination of the banks at that time become to be arbitrarily stipulated by setting the pattern of reflesh control flags to the refresh bank specifying table 1 and it is made possible to execute the refleshing of plural synchronous DRAM banks while dividing the refleshing into one cycle or more.



Micron Technology Inc. et al.
Ex. 1010, p. 216

(19)日本国特許庁（ＪＰ）　(12) **公開特許公報（Ａ）**　(11)特許出願公開番号

特開平9－237492

(43)公開日　平成９年(1997)９月９日

(51)Int.Cl.[6]	識別記号	庁内整理番号	ＦＩ		技術表示箇所
Ｇ１１Ｃ　11/406			Ｇ１１Ｃ　11/34	３６３Ｋ	
Ｇ０６Ｆ　12/00	５５０		Ｇ０６Ｆ　12/00	５５０Ｂ	

審査請求　未請求　請求項の数５　ＯＬ　（全 ８ 頁）

(21)出願番号　特願平8－45057

(22)出願日　平成８年(1996)３月１日

(71)出願人　000003078
株式会社東芝
神奈川県川崎市幸区堀川町72番地

(71)出願人　000221052
東芝コンピュータエンジニアリング株式会社
東京都青梅市新町1381番地１

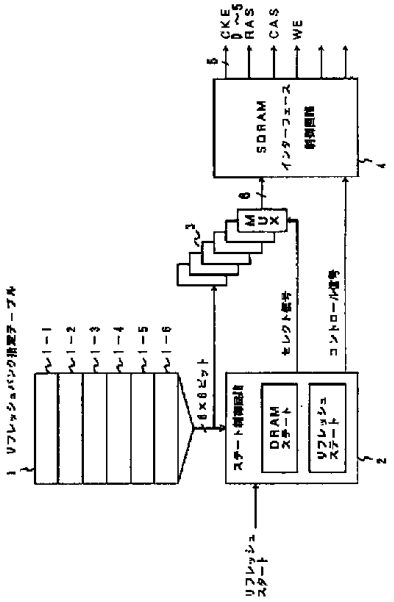
(72)発明者　的場　可
東京都青梅市末広町２丁目９番地　株式会社東芝青梅工場内

(72)発明者　豊田　康裕
東京都青梅市新町1381番地１　東芝コンピュータエンジニアリング株式会社内

(74)代理人　弁理士　鈴江　武彦

(54)【発明の名称】　メモリ制御装置

(57)【要約】
【課題】使用されているシンクロナスＤＲＡＭの種類や数によらずに、シンクロナスＤＲＡＭのリフレッシュ時に流れる電流量を常に所定値以下に抑える。
【解決手段】リフレッシュバンク指定テーブル１の中の１段目のレジスタ１－１から順にリフレッシュ制御フラグの内容が調べられ、そのリフレッシュ制御フラグによってリフレッシュの実行が指定されているシンクロナスＤＲＡＭバンクに対してのみのリフレッシュが実行される。このため、リフレッシュバンク指定テーブル１に対してどのようなリフレッシュ制御フラグのパターンを設定するかによって、同時にリフレッシュが実行されるシンクロナスＤＲＡＭバンクの数およびその時のバンクの組み合わせを任意に規定できるようになり、複数のシンクロナスＤＲＡＭバンクのリフレッシュを１以上のサイクルに分けて実行することが可能となる。



Micron Technology Inc. et al.
Ex. 1010, p. 217

【特許請求の範囲】

【請求項１】　１以上のシンクロナスＤＲＡＭを各々が有する複数のシンクロナスＤＲＡＭバンクをアクセス制御するメモリ制御装置において、

前記複数のシンクロナスＤＲＡＭバンクの数に対応する複数のレジスタから構成され、各レジスタには、前記シンクロナスＤＲＡＭバンク毎にリフレッシュを行うか否かを指定する複数のリフレッシュ制御情報が設定されているリフレッシュバンク指定テーブルと、

リフレッシュサイクルの要求に応答して、前記リフレッシュバンク指定テーブルの複数段のレジスタを１段ずつ順番に選択し、その選択したレジスタに設定されているリフレッシュ制御情報によってリフレッシュの実行が指定されているシンクロナスＤＲＡＭバンクに対してのみリフレッシュを実行させるリフレッシュサイクル制御手段とを具備し、

同時にリフレッシュが実行されるシンクロナスＤＲＡＭバンクを任意に組み合わせることにより、前記複数のシンクロナスＤＲＡＭバンクのリフレッシュを１以上のサイクルに分けて実行できるようにしたことを特徴とするメモリ制御装置。

【請求項２】　前記リフレッシュサイクル制御手段は、前記リフレッシュ制御情報によってリフレッシュの実行が指定されているシンクロナスＤＲＡＭバンクそれぞれに対応するクロックイネーブル信号をアクティブにして、リフレッシュコマンドを発行することを特徴とする請求項１記載のメモリ制御装置。

【請求項３】　前記リフレッシュサイクル制御手段は、前記選択したレジスタに設定されているリフレッシュ制御情報が全てリフレッシュの非実行を示すとき、後段のレジスタを参照せずにリフレッシュサイクルの実行を終了することを特徴とする請求項１記載のメモリ制御装置。

【請求項４】　コンピュータシステムの主記憶として使用される複数のシンクロナスＤＲＡＭバンクをアクセス制御するメモリ制御装置において、

前記コンピュータシステムに実装可能なシンクロナスＤＲＡＭバンクの最大数に対応する複数のレジスタから構成され、各レジスタには、前記シンクロナスＤＲＡＭバンク毎にリフレッシュを行うか否かを指定する複数のリフレッシュ制御情報が設定されているリフレッシュバンク指定テーブルと、

リフレッシュサイクルの要求に応答して、前記リフレッシュバンク指定テーブルの複数段のレジスタを１段ずつ順番に選択し、その選択したレジスタに設定されているリフレッシュ制御情報によってリフレッシュの実行が指定されているシンクロナスＤＲＡＭバンクに対してのみリフレッシュを実行させるリフレッシュサイクル制御手段とを具備し、

同時にリフレッシュが実行されるシンクロナスＤＲＡＭバンクを任意に組み合わせることにより、前記複数のシンクロナスＤＲＡＭバンクのリフレッシュを１以上のサイクルに分けて実行できるようにしたことを特徴とするメモリ制御装置。

【請求項５】　１以上のＤＲＡＭを各々が有する複数のＤＲＡＭバンクをアクセス制御するメモリ制御装置において、

前記複数のＤＲＡＭバンクの数に対応する複数のレジスタから構成され、各レジスタには、前記ＤＲＡＭバンク毎にリフレッシュを行うか否かを指定する複数のリフレッシュ制御情報が設定されているリフレッシュバンク指定テーブルと、

リフレッシュサイクルの要求に応答して、前記リフレッシュバンク指定テーブルの複数段のレジスタを１段ずつ順番に選択し、その選択したレジスタに設定されているリフレッシュ制御情報によってリフレッシュの実行が指定されているＤＲＡＭバンクに対してのみリフレッシュを実行させるリフレッシュサイクル制御手段とを具備し、

同時にリフレッシュが実行されるＤＲＡＭバンクを任意に組み合わせることにより、前記複数のＤＲＡＭバンクのリフレッシュを１以上のサイクルに分けて実行できるようにしたことを特徴とするメモリ制御装置。

【発明の詳細な説明】

【０００１】

【発明の属する技術分野】この発明はシンクロナスＤＲＡＭをアクセス制御するメモリ制御装置に関し、特に１以上のシンクロナスＤＲＡＭを各々が含む複数のシンクロナスＤＲＡＭバンクをアクセス制御するメモリ制御装置に関する。

【０００２】

【従来の技術】近年、非同期ＤＲＡＭに代わる高速ＤＲＡＭとして、シンクロナスＤＲＡＭが注目されている。シンクロナスＤＲＡＭは、外部クロックに同期した入出力回路構成、コマンド形式のアクセス、バースト転送による連続アクセス、２バンク構成などの特徴を持つ。

【０００３】このシンクロナスＤＲＡＭをコンピュータシステムの主記憶として利用すれば、例えばＣＰＵとシンクロナスＤＲＡＭのクロックの共通化によりアクセス時のロスを少なくでき、またバースト転送を利用することによってＣＰＵをノーウエイトで動作させること等を実現できる。

【０００４】従来の非同期ＤＲＡＭへのアクセスを行う場合、ＲＡＳ＃信号をアサートしローアドレスをＤＲＡＭに与えた後、ＣＡＳ＃信号をアサートしカラムアドレスをＤＲＡＭに与える事でアクセスを行った。

【０００５】これに対し、シンクロナスＤＲＡＭへのアクセスを行う場合は、クロックイネーブルとなるＣＫＥ信号をアクティブにした状態で、ＡＣＴ（バンク・アクティブ）コマンドにより、ローアドレスをシンクロナス

Micron Technology Inc. et al.
Ex. 1010, p. 218

DRAMに与えた後、リード／ライトコマンドによりカラムアドレスをシンクロナスDRAMに与える事でアクセスを行なう。シンクロナスSDRAMの上記コマンド受け取りは、そのシンクロナスDRAMに入力されたクロックCLKに同期して行われる。

【0006】ACT（バンク・アクティブ）コマンドは非同期DRAMにおけるRAS#信号の立ち下がりに相当し、リード／ライトコマンドは非同期DRAMにおけるCAS#信号の立ち下がりに相当する。シンクロナスDRAMは、ACT（バンク・アクティブ）コマンドの入力に応答してアクティブ状態となり、以降、プリチャージコマンドが入力されるまでそのアクティブ状態を維持する。

【0007】また、シンクロナスDRAMのリフレッシュを行う場合は、対象となるシンクロナスDRAMがアクティブ状態の時は、PRC（プリチャージ）コマンドにより、非アクテイブ状態にし、tRP時間（PRE to ACTIVE Command Period）後、CKE信号をアクティブにした状態で、CBR（CAS Before RAS Refresh）コマンドにより、シンクロナスDRAM内部リフレッシュ動作を起動した後、tRC時間（REF to REF／ACTIVE Command Period）時間待つ事で可能となる。

【0008】しかしながら、このようなシンクロナスDRAMは、その高速動作のための構造上、非同期DRAMに比べてその消費電流は大きい。このため、シンクロナスDRAMを主記憶として利用した場合には、そのリフレッシュ時に多大な電流が流れることになり、システム上の電源回路の容量を大きくする必要がある。

【0009】リフレッシュ時に大電流が流れるのは、次の理由による。すなわち、主記憶は、複数のシンクロナスDRAMチップを各々が含む複数のシンクロナスDRAMバンクから構成されるが、リフレッシュ時にはこれら全てのバンクのチップそれぞれに対して同時にリフレッシュ動作が行われる。したがって、主記憶全体では、シンクロナスDRAM1チップ当たりのリフレッシュ時のピーク電流の全チップ数倍もの電流が流れることになる。

【0010】

【発明が解決しようとする課題】上述したように、従来では、全てのシンクロナスDRAMバンクのチップそれぞれに対して同時にリフレッシュ動作が行われるので、リフレッシュ時に多大な電流が流れるという問題があった。

【0011】この発明はこのような点に鑑みてなされたものであり、使用されているシンクロナスDRAMの種類や数によらずに、シンクロナスDRAMのリフレッシュ時に流れる電流量を常に所定値以下に抑えることができるメモリ制御装置を提供することを目的とする。

【0012】

【課題を解決するための手段】この発明は、1以上のシンクロナスDRAMを各々が有する複数のシンクロナスDRAMバンクをアクセス制御するメモリ制御装置において、前記複数のシンクロナスDRAMバンクの数に対応する複数のレジスタから構成され、各レジスタには、前記シンクロナスDRAMバンク毎にリフレッシュを行うか否かを指定する複数のリフレッシュ制御情報が設定されているリフレッシュバンク指定テーブルと、リフレッシュサイクルの要求に応答して、前記リフレッシュバンク指定テーブルの複数段のレジスタを1段ずつ順番に選択し、その選択したレジスタに設定されているリフレッシュ制御情報によってリフレッシュの実行が指定されているシンクロナスDRAMバンクに対してのみリフレッシュを実行させるリフレッシュサイクル制御手段とを具備し、同時にリフレッシュが実行されるシンクロナスDRAMバンクを任意に組み合わせることにより、前記複数のシンクロナスDRAMバンクのリフレッシュを1以上のサイクルに分けて実行できるようにしたことを特徴とする。

【0013】このメモリ制御装置においては、リフレッシュサイクルの実行が要求されると、まず、リフレッシュバンク指定テーブルの中の1段目のレジスタが選択される。そして、その1段目のレジスタに設定されているリフレッシュ制御情報によってリフレッシュの実行が指定されているシンクロナスDRAMバンクに対してリフレッシュが実行される。この後、2段目以降のレジスタについても、リフレッシュの実行を指定するリフレッシュ制御情報が含まれている場合には、そのリフレッシュ制御情報に対応するバンクに対してのリフレッシュが実行される。このため、リフレッシュバンク指定テーブルに対してどのようなリフレッシュ制御情報のパターンを設定するかによって、同時にリフレッシュが実行されるシンクロナスDRAMバンクの数およびその時のバンクの組み合わせを任意に規定できるようになり、複数のシンクロナスDRAMバンクのリフレッシュを1以上のサイクルに分けて実行することが可能となる。

【0014】したがって、使用されているシンクロナスDRAMの種類や数によらずに、シンクロナスDRAMのリフレッシュ時に流れるピーク電流値を常に所定値以下に抑えることができる。

【0015】

【発明の実施の形態】以下、図面を参照してこの発明の実施形態を説明する。図1には、この発明の一実施形態に係るシンクロナスDRAMコントローラの構成が示されている。このシンクロナスDRAMコントローラは、コンピュータシステムに主記憶として設けられた複数のシンクロナスDRAM（SDRAM）バンクをアクセス制御するためのものであり、図示のように、リフレッシュバンク指定テーブル1、ステート制御回路2、セレク

Micron Technology Inc. et al.
Ex. 1010, p. 219

タ3、SDRAMインターフェース制御回路4を備えている。

【0016】ここでは、コンピュータシステムに実装可能な最大シンクロナスDRAMバンク数が6個であり、システムボード上には図2に示されているように4個のシンクロナスDRAMバンク0～3が実装されており、残りの2個のシンクロナスDRAMバンクは拡張用として使用される場合を例示して説明する。各シンクロナスDRAMバンクにはクロックイネーブル信号CKEが個別に割り当て割れるので、SDRAMインターフェース制御回路4からは6本のクロックイネーブル信号CKEが発生されることになる。

【0017】リフレッシュバンク指定テーブル1は、システムのCPUによってリードライト可能な6段のレジスタ1－1～1－6から構成されており、レジスタ1－1～1－6それぞれのビット0～5にはシンクロナスDRAMバンク0～5に対応するリフレッシュ制御フラグが設定される。ここで、リフレッシュ制御フラグの値が“1”時は、そのフラグのビット位置に対応するシンクロナスDRAMバンクがリフレッシュの対象であることが指定される。

【0018】ステート制御回路2は、各シンクロナスDRAMバンクの状態の遷移を管理するDRAMステート回路と、リフレッシュサイクルの状態遷移を管理するリフレッシュステート回路とを含んでおり、リフレッシュバンク指定テーブル1に設定されたフレッシュ制御フラグの読み取り、セレクタ3の選択動作の制御、SDRAMインターフェース制御回路4によるリフレッシュサイクルの実行制御を行う。

【0019】このステート制御回路2は、リフレッシュタイマなどから定期的に発生されるリフレッシュスタート要求に応答して、リフレッシュバンク指定テーブル1の1段目のレジスタ1－1から順番に1段ずつ選択し、その選択したレジスタに設定されているリフレッシュ制御フラグの内容を調べる。6ビットのリフレッシュ制御フラグの中に“1”のリフレッシュ制御フラグが含まれている場合には、セレクト信号を発生してそのフラグをセレクタ3に選択させる。また、SDRAMインターフェース制御回路4にコントロール信号を供給して、リフレッシュサイクルの実行を制御する。

【0020】セレクタ3は6ビットto1ビットのセレクタであり、6本のCKE信号それぞれに対応して設けられている。すなわち、各セレクタ3は、6個のレジスタ1－1～1－6それぞれの同一ビット位置に設定されている6ビットのリフレッシュ制御フラグを入力し、そのうちの1ビットを選択する。したがって6ビットのリフレッシュ制御フラグの中に“1”のリフレッシュ制御フラグが含まれている場合には、6個のセレクタ3からは、その6ビットのリフレッシュ制御フラグの値がそのまま出力されることになる。

【0021】SDRAMインターフェイス制御回路4は、ステート制御回路2からのコントロール信号に従って、各SDRAMバンクへの制御信号を生成する。この図1の構成においては、リフレッシュスタート要求が入力されると、まず、リフレッシュバンク指定テーブル1の中の1段目のレジスタ1－1の内容がステート制御回路2によって参照される。その1段目のレジスタ1－1に設定されている6ビットのリフレッシュ制御フラグの中に“1”のリフレッシュ制御フラグが含まれている場合には、ステート制御回路2からのセレクト信号により、6個のセレクタ3それぞれが1段目のレジスタ1－1に対応するビット位置の入力を選択する。そして、ステート制御回路2の制御の下、“1”のリフレッシュ制御フラグに対応するシンクロナスDRAMバンクのリフレッシュがSDRAMインターフェイス制御回路4によって同時実行される。この後、2段目以降のレジスタについても順次同様の処理が行われ、リフレッシュの実行を指定する“1”のリフレッシュ制御フラグが含まれている場合には、そのリフレッシュ制御フラグに対応するシンクロナスDRAMバンクに対してのリフレッシュが実行される。

【0022】このようにして、1段目から6番目までのレジスタ1－1～1－6の内容が順次読み込まれて、“1”のリフレッシュ制御フラグに対応するシンクロナスDRAMバンクに対してリフレッシュが実行される。もし、読み込んだ6ビットのリフレッシュ制御フラグが全て“0”であった場合には、読み込んでないレジスタが残っていても、その時点でリフレッシュサイクルは終了される。

【0023】このため、リフレッシュバンク指定テーブル1に対してどのようなリフレッシュ制御フラグのパターンを設定するかによって、同時にリフレッシュが実行されるシンクロナスDRAMバンクの数およびその時のバンクの組み合わせを任意に規定できるようになり、複数のシンクロナスDRAMバンクのリフレッシュを1以上のサイクルに分けて（最大で6サイクルに分割される）実行することが可能となる。

【0024】次に、図3乃至図6を参照して、リフレッシュバンク指定テーブル1に設定される6X6ビットのリフレッシュ制御フラグの具体的なパターンについて説明する。

【0025】図3のパターンでは、1段のレジスタ毎に“1”のリフレッシュ制御フラグが互いに異なるビット位置に1つずつ設定されている。この場合、シンクロナスDRAMバンクのリフレッシュは1バンク単位で6回に分けて行われる。したがって、リフレッシュ時のピーク電流は1バンクに属するチップの合計分だけで済み、全てのバンクを同時にリフレッシュする場合の1／6にピーク電流を低減できる。

【0026】図4のパターンでは、1段目のレジスタ1

Micron Technology Inc. et al.
Ex. 1010, p. 220

－1の全てのビット位置に“1”のリフレッシュ制御フラグが設定されている。この場合、全てのバンクが同時にリフレッシュされ、それが終了した時点でリフレッシュサイクルが終了される。このパターンは、システムに用意された電源回路の容量が十分大きい場合に好適なものであり、リフレッシュサイクルに要する時間を低減することができる。

【００２７】図5のパターンでは、1段目から3段目までのレジスタ1－1～1－3それぞれの互いに異なるビット位置に“1”のリフレッシュ制御フラグが2つずつ設定されている。この場合、シンクロナスDRAMバンクのリフレッシュは2バンク単位で3回に分けて行われ、3回目のリフレッシュが終了した時点でリフレッシュサイクルが終了される。したがって、リフレッシュ時のピーク電流は2バンクに属するチップの合計分だけで済み、全てのバンクを同時にリフレッシュする場合の1／3にピーク電流を低減できる。

【００２８】図6のパターンは、標準実装されているバンク0～3と、拡張スロットを介して増設されたバンク4，5のバンク構成やチップの種類が異なっており、バンク0～3に比べ、バンク4，5に大きなピーク電流が流れる場合に好適なパターンである。すなわち、このパターンにおいては、バンク0～3については同時にリフレッシュされ、その後に、バンク4，5のリフレッシュが2回に分けて行われる。

【００２９】つぎに、図7および図8を参照して、リフレッシュ制御動作の手順を具体的に説明する。図7はステート制御回路2によって管理されているリフレッシュサイクルの状態遷移を示す状態遷移図であり、また図8はリフレッシュサイクルの動作を示すタイミングチャートである。

【００３０】まず、システム起動時に、リフレッシュバンク指定テーブル1のレジスタに対してリフレッシュ制御フラグの設定がCPUによって行われる。ここでは、図3のリフレッシュ制御パターンが設定される場合を想定する。

【００３１】リフレッシュスタート要求が発行されると、ステート制御回路2は、まず、SDRAM状態ステート回路からアクティブなバンクの有無を調べる。いずれかのバンクがアクテイブな時は、IDLEサイクルからPRCサイクルへ進み、そのアクティブなバンク（図8ではバンク1）に対してプリチャージコマンド（RAS#＝0，CAS#＝1，WE#＝0）を送り、tRP時間確保後、CHECKサイクルへ進む。またアクテイブなバンクが無い時は、PRCサイクルに移行せずに、CHECKサイクルに直接進む。

【００３２】CHECKサイクルでは、レジスタ1－1の内容がチェックされる。もし、6ビット全てが“0”ならば、リフレッシュを実行せずにIDLEサイクルに戻る。

【００３３】この例では、バンク0に対応するリフレッシュ制御フラグが“1”であるので、それがセレクタ3を介してSDRAMインターフェース制御回路4に送られる。SDRAMインターフェース制御回路4は、CBRコマンドを発行する。このコマンドは最後のバンクについてのリフレッシュが完了するまで保持される。

【００３４】ステートはCBR00サイクルへ進み、“1”のリフレッシュ制御フラグに対応するバンク（ここでは、バンク0）のCKE0信号がアクティブにされる。次にCBR01サイクルへ進み、チップセレクト信号CSがアクテイブにされる。これにより、バンク0に対してリフレッシュコマンドが送り出される。

【００３５】また、このCBR01サイクルでは、レジスタ1－2の内容のチェックも同時に行われる。もし、6ビット全てが“0”であれば、tRCサイクルに進みtRC時間確保後、IDLEサイクルに戻りリフレッシュサイクルを終了する。

【００３６】“1”にセットされたビットがあれば（この例では、バンク1に対応するリフレッシュ制御フラグが“1”）、CBR10サイクルへ進む。このサイクルでは、バンク1に対応するCKE1信号がアクティブにされる。またこの時、CBR00においてアクティブにしたCKE0信号をインアクティブにする。次にCBR11サイクルへ進み、チップセレクト信号CSがアクテイブにされる。これにより、バンク1に対してリフレッシュコマンドが送り出される。また、このCBR11サイクルでは、レジスタ1－3の内容のチェックも同時に行われる。以降、同様にして、バンク2～5についてのリフレッシュが行われる。

【００３７】以上説明したように、この実施形態においては、リフレッシュバンク指定テーブル1の中の1段目のレジスタ1－1から順にリフレッシュ制御フラグの内容が調べられ、そのリフレッシュ制御フラグによってリフレッシュの実行が指定されているシンクロナスDRAMバンクに対してのみのリフレッシュが実行される。このため、リフレッシュバンク指定テーブル1に対してどのようなリフレッシュ制御フラグのパターンを設定するかによって、同時にリフレッシュが実行されるシンクロナスDRAMバンクの数およびその時のバンクの組み合わせを任意に規定できるようになり、複数のシンクロナスDRAMバンクのリフレッシュを1以上のサイクルに分けて実行することが可能となる。なお、この様なリフレッシュ制御は非同期のDRAMバンクに対して適用しても良い。これにより、DRAMのさらなる低消費電力化を図ることができる。

【００３８】

【発明の効果】以上説明したように、この発明のメモリ制御装置によれば、同時にリフレッシュが実行されるシンクロナスDRAMバンクの数およびその時のバンクの組み合わせを任意に規定できるので、使用されているシ

Micron Technology Inc. et al.
Ex. 1010, p. 221

ンクロナスＤＲＡＭの種類や数によらずに、シンクロナスＤＲＡＭのリフレッシュ時に流れる電流量を常に所定値以下に抑えることが可能となる。

【図面の簡単な説明】

【図１】この発明の一実施形態に係るＳＤＲＡＭコントローラの構成を示すブロック図。

【図２】同実施形態のＳＤＲＡＭコントローラによって制御されるＳＤＲＡＭバンクの構成を示す図。

【図３】同実施形態のＳＤＲＡＭコントローラのリフレッシュバンク指定テーブルに設定されるリフレッシュ制御フラグの第１パターンを説明するための図。

【図４】同実施形態のＳＤＲＡＭコントローラのリフレッシュバンク指定テーブルに設定されるリフレッシュ制御フラグの第２パターンを説明するための図。

【図５】同実施形態のＳＤＲＡＭコントローラのリフレッシュバンク指定テーブルに設定されるリフレッシュ制御フラグの第３パターンを説明するための図。

【図６】同実施形態のＳＤＲＡＭコントローラのリフレッシュバンク指定テーブルに設定されるリフレッシュ制御フラグの第４パターンを説明するための図。

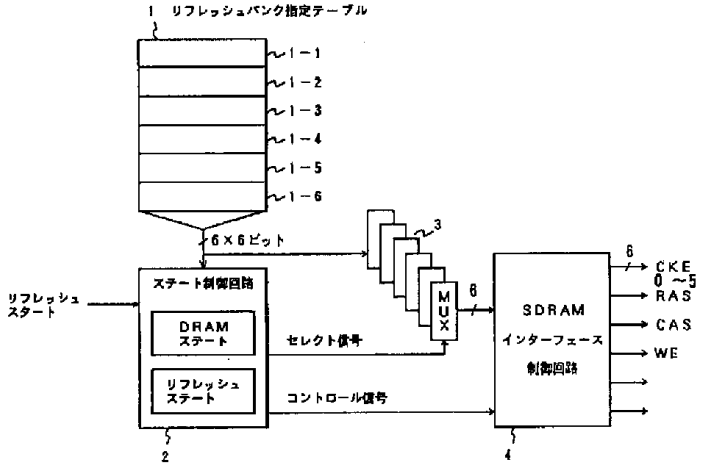
【図７】同実施形態のＳＤＲＡＭコントローラによって管理されているリフレッシュサイクルの状態遷移を示す図。

【図８】同実施形態のＳＤＲＡＭコントローラによって実行されるリフレッシュ制御動作を示すタイミングチャート。

【符号の説明】

1…リフレッシュバンク指定テーブル、１－１～１－６…レジスタ、２…ステート制御回路、４…ＳＤＲＡＭインターフェース制御回路。

【図1】



【図3】

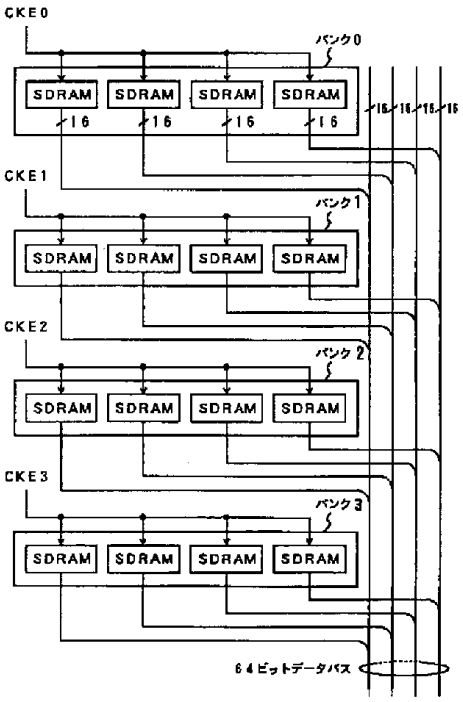
バンク0	1	2	3	4	5	
1	0	0	0	0	0	1-1
0	1	0	0	0	0	1-2
0	0	1	0	0	0	1-3
0	0	0	1	0	0	1-4
0	0	0	0	1	0	1-5
0	0	0	0	0	1	1-6

【図4】

バンク0	1	2	3	4	5	
1	1	1	1	1	1	1-1
0	0	0	0	0	0	1-2
0	0	0	0	0	0	1-3
0	0	0	0	0	0	1-4
0	0	0	0	0	0	1-5
0	0	0	0	0	0	1-6

Micron Technology Inc. et al.
Ex. 1010, p. 222

【図2】



【図5】

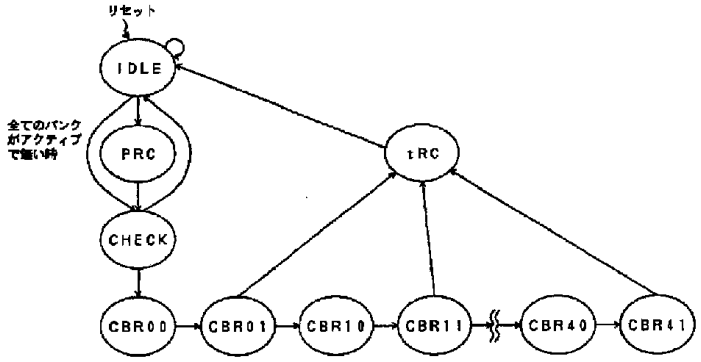
バンク0	1	2	3	4	5	
1	0	0	1	0	0	1-1
0	1	0	0	1	0	1-2
0	0	1	0	0	1	1-3
0	0	0	0	0	0	1-4
0	0	0	0	0	0	1-5
0	0	0	0	0	0	1-6

【図6】

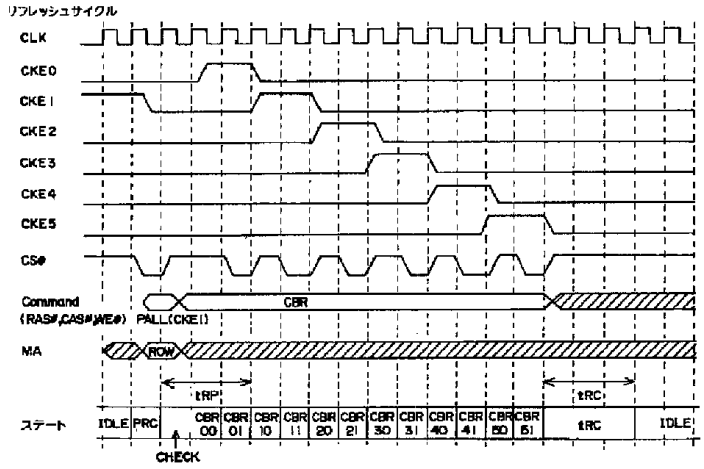
				拡張メモリ		
バンク0	1	2	3	4	5	
1	1	1	1	0	0	1-1
0	0	0	0	1	0	1-2
0	0	0	0	0	1	1-3
0	0	0	0	0	0	1-4
0	0	0	0	0	0	1-5
0	0	0	0	0	0	1-6

Micron Technology Inc. et al.
Ex. 1010, p. 223

【図7】


リセット
IDLE
全てのバンクがアクティブで無い時
PRC
tRC
CHECK
CBR00
CBR01
CBR10
CBR11
CBR40
CBR41

【図8】


リフレッシュサイクル
CLK
CKE0
CKE1
CKE2
CKE3
CKE4
CKE5
CS#
Command (RAS#,CAS#,WE#)
PALL(CKE1)
CBR
MA
ROW
tRP
tRC
ステート
IDLE
PRC
CBR 00
CBR 01
CBR 10
CBR 11
CBR 20
CBR 21
CBR 30
CBR 31
CBR 40
CBR 41
CBR 50
CBR 51
tRC
IDLE
CHECK

* NOTICES *

JPO and INPIT are not responsible for any
damages caused by the use of this translation.

1.This document has been translated by computer. So the translation may not reflect the original precisely.
2.**** shows the word which can not be translated.
3.In the drawings, any words are not translated.

CLAIMS

[Claim(s)]
[Claim 1]In a memory controller which carries out access control of two or more synchronous DRAM banks where each has one or more Synchronous DRAMs,
A memory controller characterized by dividing refreshment of two or more aforementioned synchronous DRAM banks into one or more cycles, and enabling it to perform it by having the following and combining arbitrarily synchronous DRAM banks where refreshment is performed simultaneously.
The refreshment bank appointed table where it comprises two or more registers corresponding to the number of two or more aforementioned synchronous DRAM banks, and two or more refreshment control information which specifies whether it refreshes for said every synchronous DRAM banks is set to each register.
Answer a demand of a refresh cycle and it chooses two or more steps of registers [one step of] of the aforementioned refreshment bank appointed table at a time in order, A refresh cycle control means which performs refreshment only to synchronous DRAM banks where execution of refreshment is specified by refreshment control information set as the selected register.

[Claim 2]The aforementioned refresh cycle control means activates a clock enable signal corresponding to each synchronous DRAM banks where execution of refreshment is specified by the aforementioned refreshment control information, The memory controller according to claim 1 publishing a refresh command.
[Claim 3]The memory controller according to claim 1 when the aforementioned refresh cycle control means shows [all refreshment control information set as a register which chose / aforementioned] un-performing of refreshment, wherein it ends execution of a refresh cycle, without referring to a latter register.
[Claim 4]In a memory controller which carries out access control of two or more synchronous DRAM banks used as main memory of a computer system,
A memory controller characterized by dividing refreshment of two or more aforementioned synchronous DRAM banks into one or more cycles, and enabling it to perform it by having the following and combining arbitrarily synchronous DRAM banks where refreshment is performed simultaneously.
Comprise two or more registers corresponding to the maximum number of synchronous DRAM banks which can be mounted in the aforementioned computer system, and to each register, The

Micron Technology Inc. et al.
Ex. 1010, p. 225

refreshment bank appointed table where two or more refreshment control information which specifies whether it refreshes for said every synchronous DRAM banks is set up.
Answer a demand of a refresh cycle and it chooses two or more steps of registers [one step of] of the aforementioned refreshment bank appointed table at a time in order, A refresh cycle control means which performs refreshment only to synchronous DRAM banks where execution of refreshment is specified by refreshment control information set as the selected register.

[Claim 5]In a memory controller which carries out access control of two or more DRAM banks where each has one or more DRAMs,
A memory controller characterized by dividing refreshment of two or more aforementioned DRAM banks into one or more cycles, and enabling it to perform it by having the following and combining arbitrarily DRAM banks where refreshment is performed simultaneously.
The refreshment bank appointed table where it comprises two or more registers corresponding to the number of two or more aforementioned DRAM banks, and two or more refreshment control information which specifies whether it refreshes for every aforementioned DRAM banks is set to each register.
Answer a demand of a refresh cycle and it chooses two or more steps of registers [one step of] of the aforementioned refreshment bank appointed table at a time in order, A refresh cycle control means which performs refreshment only to DRAM banks where execution of refreshment is specified by refreshment control information set as the selected register.

DETAILED DESCRIPTION

[Detailed Description of the Invention]
[0001]
[Field of the Invention]This invention relates to the memory controller which carries out access control of two or more synchronous DRAM banks where each contains especially one or more Synchronous DRAMs about the memory controller which carries out access control of the Synchronous DRAM.
[0002]
[Description of the Prior Art]In recent years, the Synchronous DRAM attracts attention as high-speed DRAM replaced with asynchronous DRAM. A Synchronous DRAM has the characteristics, such as input output circuit composition in sync with an external clock, access of command format, continuous access by burst transmission, and 2 bank configurations.
[0003]If this Synchronous DRAM is used as main memory of a computer system, For example, it is realizable by being able to lessen the loss at the time of access by communalization of the clock of CPU and a Synchronous DRAM, and using burst transmission to operate CPU by nowait etc.
[0004]When access to the conventional asynchronous DRAM is performed, after asserting the RAS# signal and giving a row address to DRAM, it accessed by asserting a CAS# signal and giving a column address to DRAM.
[0005]On the other hand, when performing access to a Synchronous DRAM, where the CKE signal used as clock enabling is activated, after giving a row address to a Synchronous DRAM by the ACT (bank - active) command, it accesses by giving a column address to a Synchronous DRAM with Read/write command. The above-mentioned command receipt of synchronous SDRAM is performed synchronizing with the clock CLK input into the Synchronous DRAM.

Micron Technology Inc. et al.
Ex. 1010, p. 226

[0006]the ACT (bank - active) command is equivalent to falling of the RAS# signal in asynchronous DRAM, and Read/write command is equivalent to falling of the CAS# signal in asynchronous DRAM. a Synchronous DRAM will answer the input of the ACT (bank - active) command, will be in an active state, and henceforth, it maintains the active state until a precharge command is input.
[0007]When refreshing a Synchronous DRAM, When the target Synchronous DRAM is bitter taste tape shape voice, Where it used non-bitter taste tape shape voice and a CKE signal is made into AKUTIPU after tRP time (PRE to ACTIVE Command Period) by the PRC (precharge) command, By the CBR (CAS Before RASRefresh) command, after starting Synchronous DRAM internal refreshment operation, it becomes possible by waiting tRC time (REF to REF/ACTIVE Command Period) time.
[0008]However, compared with asynchronous DRAM, the consumed electric current of such a Synchronous DRAM is large on the structure for the high-speed operation. For this reason, when a Synchronous DRAM is used as main memory, current great at the time of that refreshment will flow, and it is necessary to carry out capacity of the power supply circuit on a system largely.
[0009]It is based on the following Reason that a high current flows at the time of refreshment. That is, although main memory comprises two or more synchronous DRAM banks where each contains two or more synchronous DRAM chips, at the time of refreshment, refresh operation is simultaneously performed to each chip of all of a bank of these. Therefore, in the whole main memory, all the chip number double thing current of the peak current at the time of refreshment per Synchronous DRAM 1 chip will flow.
[0010]
[Problem to be solved by the invention]As mentioned above, since refresh operation was simultaneously performed to each chip of all the synchronous DRAM banks, by the former, there was a problem that current great at the time of refreshment flowed.
[0011]The present invention is made in view of such a point, and is a thing.
The purpose is to provide a memory controller which can always hold down a current amount which flows at the time of refreshment of a Synchronous DRAM to below a specified value, without being based on a kind and number of Synchronous DRAMs.

[0012]
[Means for solving problem]In the memory controller with which this invention carries out access control of two or more synchronous DRAM banks where each has one or more Synchronous DRAMs, Comprise two or more registers corresponding to the number of two or more aforementioned synchronous DRAM banks, and to each register, The refreshment bank appointed table where two or more refreshment control information which specifies whether it refreshes for every above-mentioned synchronous DRAM banks is set up, Answer the demand of a refresh cycle and it chooses two or more steps of registers [one step of] of the aforementioned refreshment bank appointed table at a time in order, The refresh cycle control means which performs refreshment only to the synchronous DRAM banks where execution of refreshment is specified by the refreshment control information set as the selected register is provided, Refreshment of two or more aforementioned synchronous DRAM banks is divided into one or more cycles, and it enabled it to perform it by combining arbitrarily the synchronous DRAM banks where refreshment is performed simultaneously.
[0013]In this memory controller, a demand of execution of a refresh cycle will choose the 1st step of register in the refreshment bank appointed table first. And refreshment is performed to the

Micron Technology Inc. et al.
Ex. 1010, p. 227

synchronous DRAM banks where execution of refreshment is specified by the refreshment control information set as the 1st step of the register. Then, when the refreshment control information which specifies execution of refreshment is included also about the register after the 2nd step, refreshment to the bank corresponding to that refreshment control information is performed. For this reason, by of what kind of refreshment control information a pattern is set up to the refreshment bank appointed table. The number of the synchronous DRAM banks where refreshment is performed simultaneously, and the combination of the bank at that time can be arbitrarily specified now, and it becomes possible to divide refreshment of two or more synchronous DRAM banks into one or more cycles, and to perform it.

[0014]Therefore, the peak current value which flows at the time of refreshment of a Synchronous DRAM can always be held down below to a specified value, without being based on the kind and number of Synchronous DRAMs which are used.

[0015]

[Mode for carrying out the invention]Hereafter, the embodiment of this invention is described with reference to Drawings. The composition of the synchronous DRAM controller concerning one embodiment of this invention is shown in Fig.1. Like [it is for this synchronous DRAM controller carrying out access control of two or more Synchronous DRAM (SDRAM) banks provided by the computer system as main memory, and] a graphic display, It has the refreshment bank appointed table 1, the State control circuit 2, the selector 3, and the SDRAM interface control circuit 4.

[0016]Here, the number of the maximum synchronous DRAM banks which can be mounted in a computer system is six pieces, On the system board, four synchronous DRAM banks 0-3 are mounted as shown in Fig.2, and the two remaining synchronous DRAM banks illustrate and describe the case where it is used as an object for extension. Since the clock enable signal CKE assigns individually each synchronous DRAM banks and it is divided, the six clock enable signals CKE will be generated from the SDRAM interface control circuit 4.

[0017]The refreshment bank appointed table 1 is constituted from six steps of registers 1-1 to 1-6 in which read/write is possible by CPU of the system, the register 1-1 to 1-6 -- the refreshment control flag corresponding to synchronous DRAM banks 0-5 is set to each of pits 0-5. Here, it is specified that the value of a refreshment control flag is an object of refreshment of the synchronous DRAM banks corresponding to the Flagg's bit position at the time of "1."

[0018]The DRAM State circuit where the State control circuit 2 manages transition of the state of each synchronous DRAM banks, The refreshment State circuit which manages the change state of a refresh cycle is included, Reading of the fresh control flag set as the refreshment bank appointed table 1, control of the selection operation of the selector 3, and execution control of the refresh cycle by the SDRAM interface control circuit 4 are performed.

[0019]This State control circuit 2 answers the refreshment start demand periodically generated from refresh timers etc., It chooses one step at a time in an order from the 1st step of register 1-1 of the refreshment bank appointed table 1, and the contents of the refreshment control flag set as the selected register are investigated. When the refreshment control flag of "1" is contained in a 6-bit refreshment control flag, a select signal is generated and the Flagg is made to choose it as the selector 3. A control signal is supplied to the SDRAM interface control circuit 4, and execution of a refresh cycle is controlled.

[0020]The selector 3 is 6-bit a selector of 1 bit of to(es), and is provided corresponding to each six CKE signals. namely, each selector 3 -- the six registers 1-1 to 1-6 -- the 6-bit refreshment control flag set as each identical bit position is input, and 1 bit of them is chosen. Therefore,

Micron Technology Inc. et al.
Ex. 1010, p. 228

when the refreshment control flag of "1" is contained in a 6-bit refreshment control flag, from the six selectors 3, the value of the 6-bit refreshment control flag will be outputted as it is.
[0021]The SDRAM interface control circuit 4 generates the control signal to each SDRAM bank according to the control signal from the State control circuit 2. In the composition of this Fig.1, if a refreshment start demand is input, refer to the contents of the 1st step of register 1-1 in the refreshment bank appointed table 1 for the State control circuit 2 first. When the refreshment control flag of "1" is contained in the 6-bit refreshment control flag set as the 1st step of the register 1-1, the select signal from the State control circuit 2 -- the six selectors 3 -- each chooses the input of the bit position corresponding to the 1st step of register 1-1. And the concurrency of the refreshment of the synchronous DRAM banks corresponding to the refreshment control flag of "1" is carried out by the SDRAM interface control circuit 4 under control of the State control circuit 2. Then, same processing is sequentially performed also about the register after the 2nd step, and when the refreshment control flag of "1" which specifies execution of refreshment is contained, refreshment to the synchronous DRAM banks corresponding to that refreshment control flag is performed.
[0022]Thus, the contents of the register 1-1 to 1-6 from the 1st step to the 6th are read sequentially, and refreshment is performed to the synchronous DRAM banks corresponding to the refreshment control flag of "1." When all of the read 6-bit refreshment control flag are "0", even if the register which has not been read remains, a refresh cycle is ended at the time.
[0023]For this reason, by of what kind of refreshment control flag a pattern is set up to the refreshment bank appointed table 1. The number of the synchronous DRAM banks where refreshment is performed simultaneously, and the combination of the bank at that time can be arbitrarily specified now, It becomes possible to divide refreshment of two or more synchronous DRAM banks into one or more cycles, and to perform it (divided into six cycles at the maximum).
[0024]Next, with reference to Fig.3 thru/or Fig.6, it describes about the specific pattern of the 6X6-bit refreshment control flag set as the refreshment bank appointed table 1.
[0025]In the pattern of Fig.3, the refreshment control flag of "1" is set as every one mutually different bit position for one step of every register. In this case, refreshment of synchronous DRAM banks is performed in 6 steps by 1 bank unit. Therefore, the peak current at the time of refreshment requires only a totaled part of the chip belonging to one bank, and can reduce peak current to one sixth in the case of refreshing all the banks simultaneously.
[0026]In the pattern of Fig.4, the refreshment control flag of "1" is set to all the bit positions of the 1st step of register 1-1. In this case, all the banks are refreshed simultaneously, and a refresh cycle is ended when it is completed. This pattern is preferable when the capacity of the power supply circuit prepared for the system is sufficiently large, and it can reduce the time which a refresh cycle takes.
[0027]the pattern of Fig.5 -- the register 1-1 to 1-3 from the 1st step to the 3rd step -- every two refreshment control flags of "1" are set to the bit position where each differs mutually. In this case, refreshment of synchronous DRAM banks is performed in 3 steps by 2 bank units, and a refresh cycle is ended when the 3rd refreshment is completed. Therefore, the peak current at the time of refreshment requires only a totaled part of the chip belonging to two banks, and can reduce peak current to one third in the case of refreshing all the banks simultaneously.
[0028]The pattern of Fig.6 is a preferable pattern, when the bank configuration of the banks 0-3 by which standard mounting is carried out, and the banks 4 and 5 extended via the expansion slot differs from the kind of chip and big peak current flows into the banks 4 and 5 compared with

Micron Technology Inc. et al.
Ex. 1010, p. 229

the banks 0-3. That is, in this pattern, about the banks 0-3, it refreshes simultaneously, and refreshment of the banks 4 and 5 is performed in 2 steps after that.
[0029]Next, with reference to Fig.7 and Fig.8, the procedure of a refreshment control action is described specifically.Fig.7 is a state transition diagram showing the change state of the refresh cycle managed by the State control circuit 2, and Fig.8 is a timing chart which shows operation of a refresh cycle.
[0030]First, setting out of a refreshment control flag is performed by CPU to the register of the refreshment bank appointed table 1 at the time of a system startup. Here, the case where the refreshment control pattern of Fig.3 is set up is assumed.
[0031]If a refreshment start demand is published, the State control circuit 2 will investigate the existence of an AKUTIPU bank from a SDRAM state State circuit first. When AKUTEIPU [banks / one of], a precharge command (RAS#=0, CAS#=1, WE#=0) is sent from an IDLE cycle to PRC SAIKURUHE progress and its active bank (Fig.8 bank 1), and they are after tRP time reservation and CHECK SAIKURUHE ****. When there is no AKUTEIPU bank, it proceeds to a CHECK cycle directly, without shifting to a PRC cycle.
[0032]The contents of the register 1-1 are checked in a CHECK cycle. if -- all six pits -- "0" -- if it becomes, it will return to an IDLE cycle, without performing refreshment.
[0033]In this example, since the refreshment control flag corresponding to the bank 0 is "1", it is sent to the SDRAM interface control circuit 4 via the selector 3. The SDRAM interface control circuit 4 publishes the CBR command. This command is held until the refreshment about the last bank is completed.
[0034]CKE0 signal of the bank (here bank 0) corresponding to CBR00 SAIKURUHE progress and the refreshment control flag of "1" in the State is made into AKUTIPU. Next, CBR01 SAIKURUHE progress and chip select signal CS are made into AKUTEIPU. Thereby, a refresh command is sent out to the bank 0.
[0035]In this 01 cycle CBR, the check of the contents of the register 1-2 is also performed simultaneously. If 6 bits of all are "0", it proceeds to a tRC cycle, and after tRC time reservation, it will return to an IDLE cycle and a refresh cycle will be ended.
[0036]If there is a pit set to "1" (the refreshment control flag corresponding to the bank 1 is "1" at this example), it will be CBR10 SAIKURUHE ****. In this cycle, CKE1 signal corresponding to the bank 1 is made into AKUTIPU. At this time, CKE0 signal activated in CBR00 is made inactive. Next, CBR11 SAIKURUHE progress and chip select signal CS are made into AKUTEIPU. Thereby, a refresh command is sent out to the bank 1. In this 11 cycle CBR, the check of the contents of the register 1-3 is also performed simultaneously. Henceforth, refreshment about the banks 2-5 is performed similarly.
[0037]In [as described above] this embodiment, The contents of the refreshment control flag are investigated sequentially from the register 1-1 of one-step Sakka in the refreshment bank appointed table 1, and refreshment only to the synchronous DRAM banks where execution of refreshment is specified by the refreshment control flag is performed. For this reason, by of what kind of refreshment control flag a pattern is set up to the refreshment bank appointed table 1. The number of the synchronous DRAM banks where refreshment is performed simultaneously, and the combination of the bank at that time can be arbitrarily specified now, and it becomes possible to divide refreshment of two or more synchronous DRAM banks into one or more cycles, and to perform it. Such refreshment control may be applied to asynchronous DRAM banks. Thereby, further power consumption reduction of DRAM can be attained.
[0038]

Micron Technology Inc. et al.
Ex. 1010, p. 230

[Effect of the Invention]Since the number of the synchronous DRAM banks where refreshment is performed simultaneously, and the combination of the bank at that time can be arbitrarily specified according to the memory controller of this invention as described above, It becomes possible to always hold down the current amount which flows at the time of refreshment of a Synchronous DRAM to below a specified value, without being based on the kind and number of Synchronous DRAMs which are used.

[Translation done.]

Micron Technology Inc. et al.
Ex. 1010, p. 231

Espacenet

Bibliographic data: JP2000285674 (A) — 2000-10-13

SYNCHRONIZING TYPE SEMICONDUCTOR MEMORY

Inventor(s): UCHIDA TOSHIYA ± (UCHIDA TOSHIYA)

Applicant(s): FUJITSU LTD ± (FUJITSU LTD)

Classification: - international: ***G11C11/407; G11C11/4076; G11C11/409; G11C11/4093; G11C7/10; G11C7/22;*** (IPC1-7): G11C11/407

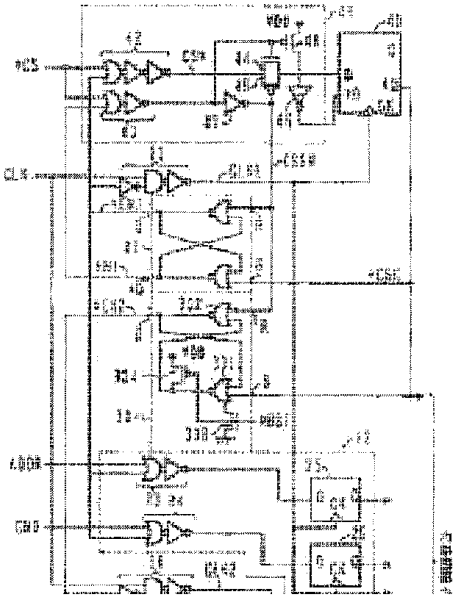
- European: G11C11/4076; G11C11/4093; G11C7/10S; G11C7/10W; G11C7/10W7; G11C7/22

Application number: JP19990092246 19990331

Priority number (s): JP19990092246 19990331

Also published as: JP4187346 (B2) US6188641 (B1)

Abstract of JP2000285674 (A)



Micron Technology Inc. et al.
Ex. 1010, p. 232



PROBLEM TO BE SOLVED: To reduce needless power consumption in an input circuit without increasing the number of external wirings. SOLUTION: This device is provided with command input circuits 24, 26 latching a command CMD synchronizing with an internal clock CLK1, a command decoder decoding the latched command, chip-select signal input circuits 41, 40, 20 activating an enable-signal EN1 responding to activation of a chipselect signal *CS, inactivating the enable-signal EN1 after arrival of a pulse being next to an external clock CLK, responding to inactivation of a chip-select signal *CS, and generating an internal chip-select signal *CSC by synchronizing a chip-select signal *CS with external clock CLK, a clock input circuit 21 operating an internal clock CLK1 responding to the external clock CLK during activation of the enable-signal EN1,; and D flip-flop generating a command decoder activating signal by synchronizing an internal chip-select signal *CSC by the internal clock CLK1.

Last updated: 24.09.2012 Worldwide Database 5.7.42; 93p

Micron Technology Inc. et al.
Ex. 1010, p. 233

(19)日本国特許庁（ＪＰ）　(12) **公開特許公報（Ａ）**　(11)特許出願公開番号
特開2000－285674
（P2000－285674A）
(43)公開日　平成12年10月13日(2000.10.13)

(51)Int.Cl.⁷	識別記号	ＦＩ	ﾃｰﾏｺｰﾄﾞ(参考)
Ｇ１１Ｃ　11/407		Ｃ１１Ｃ　11/34	３６２Ｓ　５Ｂ０２４ ３５４Ｃ

審査請求　未請求　請求項の数11　ＯＬ　（全 15 頁）

(21)出願番号　特願平11－92246

(22)出願日　平成11年３月31日(1999.3.31)

(71)出願人　000005223
富士通株式会社
神奈川県川崎市中原区上小田中４丁目１番１号

(72)発明者　内田　敏也
神奈川県川崎市中原区上小田中４丁目１番１号　富士通株式会社内

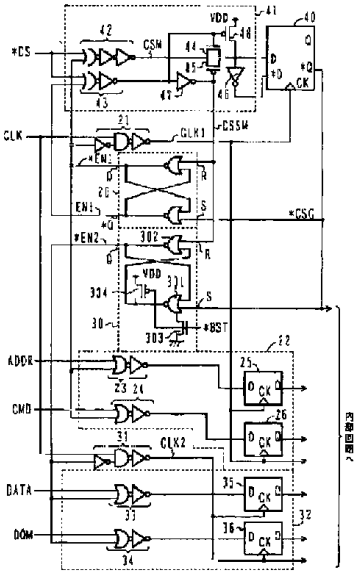
(74)代理人　100092587
弁理士　松本　眞吉

Ｆターム(参考)　5B024 AA01 BA21 BA23 CA07 CA16 DA18

(54)【発明の名称】　同期型半導体記憶装置

(57)【要約】
【課題】外部配線数を増加することなく入力回路での無駄な消費電力を低減する。
【解決手段】コマンドＣＭＤを内部クロックＣＬＫ１に同期してラッチするコマンド入力回路２４及び２６と、ラッチされたコマンドをデコードするコマンドデコーダと、チップセレクト信号＊ＣＳの活性化に応答してイネーブル信号ＥＮ１を活性化させ、チップセレクト信号＊ＣＳの不活性化に応答して、外部クロックＣＬＫの次のパルスが到来した後にイネーブル信号ＥＮ１を不活性にし、チップセレクト信号＊ＣＳを外部クロックＣＬＫで同期化して内部チップセレクト信号＊ＣＳＣを生成するチップセレクト信号入力回路４１、４０及び２０と、イネーブル信号ＥＮ１が活性である間、外部クロックＣＬＫに応答して内部クロックＣＬＫ１を動作させるクロック入力回路２１と、内部チップセレクト信号＊ＣＳＣを内部クロックＣＬＫ１で同期化してコマンドデコーダ活性化信号を生成するＤフリップフロップとを備えている。

図１中の１つのＳＤＲＡＭ内の入力回路を示す図



Micron Technology Inc. et al.
Ex. 1010, p. 234

【特許請求の範囲】

【請求項１】　外部コマンドを内部クロックに同期してラッチするコマンド入力回路と、ラッチされた該コマンドをデコードするコマンドデコーダとを備え、デコード結果に応じた動作を行う同期型半導体記憶装置において、さらに、

外部チップセレクト信号の活性化に応答してイネーブル信号を活性化させ、該外部チップセレクト信号の不活性化に応答して、外部クロックの次のパルスが到来した後に該イネーブル信号を不活性にし、該外部チップセレクト信号を該外部クロックで同期化して内部チップセレクト信号を生成するチップセレクト信号入力回路と、

該イネーブル信号が活性である間、該外部クロックに応答して該内部クロックを動作させるクロック入力回路と、

該内部チップセレクト信号を該内部クロックで同期化して、該コマンドデコーダを活性化させる信号を生成するデコードイネーブル回路と、

を有することを特徴とする同期型半導体記憶装置。

【請求項２】　上記チップセレクト信号入力回路は、

上記外部チップセレクト信号の不活性状態を上記内部クロックに同期して保持し第１状態になり、その状態を上記内部チップセレクト信号として出力する第１フリップフロップと、

該第１フリップフロップの出力を遅延させて上記イネーブル信号を生成する遅延回路と、

を有することを特徴とする請求項１記載の同期型半導体記憶装置。

【請求項３】　上記チップセレクト信号入力回路はさらに、上記外部チップセレクト信号の活性化を検出するタイミング回路を有し、

上記遅延回路は、上記第１フリップフロップの上記第１状態に応答して第１状態になり、該活性化の検出に応答して第２状態になる第２フリップフロップである、

ことを特徴とする請求項２記載の同期型半導体記憶装置。

【請求項４】　上記タイミング回路は、

上記イネーブル信号が活性であるときに上記外部チップセレクト信号を通す第１論理ゲートと、

該イネーブル信号が不活性であるときに該外部チップセレクト信号を通すことにより該外部チップセレクト信号の活性化を検出する第２論理ゲートと、

該第１論理ゲートの出力端と上記第１フリップフロップのデータ入力端との間に接続され、該第２論理ゲートの出力でオン／オフ制御される転送ゲートと、

該第１フリップフロップのデータ入力端と電源電位との間に接続され、該転送ゲートのオン／オフ制御と連動してオフ／オン制御されるスイッチング素子と、

を有することを特徴とする請求項３記載の同期型半導体記憶装置。

【請求項５】　上記チップセレクト信号入力回路は、

上記外部チップセレクト信号を上記外部クロックに同期して保持しその状態を上記内部チップセレクト信号として出力する第１フリップフロップと、

該第１フリップフロップの出力を遅延させて上記イネーブル信号を生成する遅延回路と、

を有することを特徴とする請求項１記載の同期型半導体記憶装置。

【請求項６】　上記第１フリップフロップは、

第１及び第２データ入力端、クロック入力端並びに第１及び第２ノードを備え、該クロック入力端が第１レベルのとき該第１及び第２ノードがプリチャージされ、該クロック入力端が第２レベルのとき該第１及び第２データ入力端の信号レベルに応じてそれぞれ該第１及び第２ノードがフローティング状態になり又はディスチャージされるダイナミック型フリップフロップと、

該第１及び第２ノードの一方がディスチャージされているときその状態が書き込まれ、該第１及び第２ノードがプリチャージされているとき該書き込まれた状態を保持するラッチ回路と、

を有することを特徴とする請求項５記載の同期型半導体記憶装置。

【請求項７】　上記チップセレクト信号入力回路はさらにタイミング回路を有し、該タイミング回路は、

上記外部チップセレクト信号の入力端子と上記第１データ入力端との間に接続され、上記外部チップセレクト信号が活性状態のときオンになるように上記第１フリップフロップの出力でオン／オフ制御される転送ゲートと、

該第１データ入力端と電源電位との間に接続され、該電源電位はこれを該第１データ入力端に供給すると上記第１ノードがフローティング状態になる電位であり、該転送ゲートのオン／オフ制御と連動してオフ／オン制御されるスイッチング素子と、

を有することを特徴とする請求項６記載の同期型半導体記憶装置。

【請求項８】　上記タイミング回路はさらに、上記外部チップセレクト信号が活性化した時、上記外部クロックが上記第２レベルの場合に上記第１レベルになるのを待って上記第２データ入力端のレベルを反転する同期化回路を有することを特徴とする請求項７記載の同期型半導体記憶装置。

【請求項９】　上記遅延回路は、上記第１フリップフロップの状態の変化に応答して状態が変化する第２フリップフロップであることを特徴とする請求項５乃至８のいずれか１つに記載の同期型半導体記憶装置。

【請求項１０】　入力端が上記第２フリップフロップと並列接続され、第２イネーブル信号を出力し、バースト信号が活性であるときに該第２イネーブル信号が強制的に活性状態になる第３フリップフロップと、

該第２イネーブル信号が活性である間、上記外部クロッ

Micron Technology Inc. et al.
Ex. 1010, p. 235

クに応答して第2内部クロックを動作させる第2クロック入力回路と、
該第2イネーブル信号が活性である期間、該第2クロックに同期してデータを入力するデータ入力回路と、
をさらに有することを特徴とする請求項3又は9記載の同期型半導体記憶装置。
【請求項11】 タイマーからの定期的なタイムアップ信号に応答して自動リフレッシュ動作を行う自動リフレッシュ回路をさらに有することを特徴とする請求項1乃至10のいずれか1つに記載の同期型半導体記憶装置。
【発明の詳細な説明】
【0001】
【発明の属する技術分野】本発明は、クロックに同期して信号を入出力する同期型半導体記憶装置に関する。
【0002】
【従来の技術】DRAMは、SRAMに比し記憶密度が高いが消費電力が大きいので、その低減化が要求されている。
【0003】図9は、SDRAMコントローラ10と複数のSDRAM11～14との間の配線接続図である。
【0004】SDRAMコントローラ10は、供給されるアドレスの上位2ビットをデコードしてチップセレクト信号*CS1～*CS4を生成し、該上位2ビットを除いたアドレスを行アドレスと列アドレスに分けて時分割多重化するとともに、行アドレス及び列アドレスに対応してそれぞれロウアドレスストローブ信号RAS及びコラムアドレスストローブ信号CASを生成する。これらRAS及びCASは、ライトイネーブル信号WEとともに多ビットのコマンドCMDを構成している。SDRAMコントローラ10は、SDRAM11～14のいずれかを選択する前に、外部クロックCLKとの論理積で内部クロックを生成するためのクロックイネーブル信号CKEを高レベルにする。
【0005】SDRAM11～14は、SDRAMコントローラ10を介して不図示のマイクロプロセッサに接続されている。SDRAMコントローラ10のチップセレクト信号出力端子*CS1～*CS4はそれぞれ、SDRAM11～14のチップセレクト信号入力端子*CSに接続され、SDRAMコントローラ10の外部クロックCLK、コマンドCMD、データDATA、アドレスADDR及びクロックイネーブル信号CKEの出力端子はSDRAM11～14の対応する端子に共通に接続されている。
【0006】チップセレクト信号*CS1～*CS4が全て高レベルでSDRAM11～14が非選択のときには、クロックイネーブル信号CKEが低レベルにされてSDRAM11～14内の内部クロックが停止し、入力を内部クロックに同期化させる入力回路が不活性になる。これにより、SDRAM11～14の低消費電力化が図られている。
【0007】
【発明が解決しようとする課題】しかし、図10に示す如く、例えばチップセレクト信号*CS2～*CS4を高レベルに維持してチップセレクト信号*CS1のみ低レベルにすることによりSDRAM11～14のうちSDRAM11のみを選択する場合には、その前にクロックイネーブル信号CKEが高レベルにされて、SDRAM11～14内の内部クロックが動作し入力回路が活性化されるので、非選択のSDRAM12～14において無駄な電力が消費される。
【0008】これを避けるために、SDRAM11～14に異なるクロックイネーブル信号を供給すれば、配線数が増加して複雑になる。
【0009】本発明の目的は、このような問題点に鑑み、外部配線数を増加することなく入力回路での無駄な消費電力を低減することが可能な同期型半導体記憶装置を提供することにある。
【0010】
【課題を解決するための手段及びその作用効果】請求項1では、外部コマンドを内部クロックに同期してラッチするコマンド入力回路と、ラッチされた該コマンドをデコードするコマンドデコーダとを備え、デコード結果に応じた動作を行う同期型半導体記憶装置において、さらに、例えば図8に示す如く、外部チップセレクト信号の活性化に応答してイネーブル信号を活性化させ、該外部チップセレクト信号の不活性化に応答して、外部クロックの次のパルスが到来した後に該イネーブル信号を不活性にし、該外部チップセレクト信号を該外部クロックで同期化して内部チップセレクト信号を生成するチップセレクト信号入力回路と、該イネーブル信号が活性である間、該外部クロックに応答して該内部クロックを動作させるクロック入力回路と、該内部チップセレクト信号を該内部クロックで同期化して、該コマンドデコーダを活性化させる信号を生成するデコードイネーブル回路（40N）とを有する。
【0011】この同期型半導体記憶装置によれば、該内部チップセレクト信号が活性である時に、該内部クロックに同期してコマンドデコーダを活性化させる信号が生成され、コマンドデコーダのデコード結果に応じた動作が行われるので、従来のクロックイネーブル信号を用いる必要がない。また、外部チップセレクト信号が不活性であればイネーブル信号が不活性であるので、非選択の同期型半導体記憶装置のクロック入力回路及びコマンド入力回路が不活性となって、その消費電力が低減される。
【0012】請求項2の同期型半導体記憶装置では、請求項1において例えば図3に示す如く、上記チップセレクト信号入力回路は、上記外部チップセレクト信号の不活性状態を上記内部クロック（CK1）に同期して保持し第1状態になり、その状態を上記内部チップセレクト

Micron Technology Inc. et al.
Ex. 1010, p. 236

信号として出力する第１フリップフロップ（４０）と、該第１フリップフロップの出力を遅延させて上記イネーブル信号を生成する遅延回路（２０）とを有する。

【００１３】この同期型半導体記憶装置によれば、外部チップセレクト信号の不活性化に応答して、外部クロックの次のパルスが到来した後に該イネーブル信号を不活性にすることを確実に行うことができる。

【００１４】請求項３の同期型半導体記憶装置では、請求項２において、上記チップセレクト信号入力回路はさらに、上記外部チップセレクト信号の活性化を検出するタイミング回路を有し、上記遅延回路は、上記第１フリップフロップの上記第１状態に応答して第１状態になり、該活性化の検出に応答して第２状態になる第２フリップフロップである。

【００１５】この同期型半導体記憶装置によれば、外部チップセレクト信号が不活性になるとイネーブル信号が不活性になって、クロック入力回路及びコマンド入力回路が不活性となるので、選択された同期型半導体記憶装置の消費電力も低減される。

【００１６】請求項４の同期型半導体記憶装置では、請求項３において例えば図３に示す如く、上記タイミング回路（４１）は、上記イネーブル信号（ＥＮ１）が活性であるときに上記外部チップセレクト信号（＊ＣＳ）を通す第１論理ゲート（４２）と、該イネーブル信号が不活性であるときに該外部チップセレクト信号を通すことにより該外部チップセレクト信号の活性化を検出する第２論理ゲート（４３）と、該第１論理ゲートの出力端と上記第１フリップフロップのデータ入力端との間に接続され、該第２論理ゲートの出力でオン／オフ制御される転送ゲート（４４及び４５）と、該第１フリップフロップのデータ入力端と電源電位との間に接続され、該転送ゲートのオン／オフ制御と連動してオフ／オン制御されるスイッチング素子（４１）とを有する。

【００１７】この同期型半導体記憶装置によれば、外部チップセレクト信号が不活性になったとき、イネーブル信号はまだ活性であるので第１論理ゲートを通り、第１フリップフロップにこの不活性状態が保持されて、イネーブル信号が不活性になる。これにより転送ゲートがオフ、スイッチング素子がオンになって、転送ゲートの入力側と出力側の信号の衝突が避けられ、かつ、スイッチング素子のオンにより第１フリップフロップの状態が維持される。イネーブル信号の不活性により第１論理ゲートが不活性になっているので、外部チップセレクト信号が活性化しても第１フリップフロップの状態は変化しない。しかし、イネーブル信号が不活性であるとき外部チップセレクト信号は第２論理ゲートを通るので、イネーブル信号が活性化すると第２論理ゲートの出力に応答して第２フリップフロップが第２状態に遷移し、これによりイネーブル信号が活性になる。

【００１８】請求項５の同期型半導体記憶装置では、請求項１において、上記チップセレクト信号入力回路は、例えば図６に示す如く、上記外部チップセレクト信号を上記外部クロックに同期して保持しその状態を上記内部チップセレクト信号として出力する第１フリップフロップ（４０）と、該第１フリップフロップの出力を遅延させて上記イネーブル信号を生成する遅延回路（２０）とを有する。

【００１９】請求項６の同期型半導体記憶装置では、請求項５において例えば図４に示す如く、上記第１フリップフロップは、第１及び第２データ入力端（Ｄ及び＊Ｄ）、クロック入力端（ＣＫ）並びに第１及び第２ノード（ＮＯ１及びＮＯ２）を備え、該クロック入力端が第１レベルのとき該第１及び第２ノードがプリチャージされ、該クロック入力端が第２レベルのとき該第１及び第２データ入力端の信号レベルに応じてそれぞれ該第１及び第２ノードがフローティング状態になり又はディスチャージされるダイナミック型フリップフロップ（５０）と、該第１及び第２ノードの一方がディスチャージされているときその状態が書き込まれ、該第１及び第２ノードがプリチャージされているとき該書き込まれた状態を保持するラッチ回路（４０）とを有する。

【００２０】この同期型半導体記憶装置によれば、ダイナミック型フリップフロップが用いられているので、スタティック型の場合よりも消費電力が低減される。

【００２１】請求項７の同期型半導体記憶装置では、請求項６において例えば図６に示す如く、上記チップセレクト信号入力回路はさらにタイミング回路（４１Ａ）を有し、該タイミング回路は、上記外部チップセレクト信号の入力端子と上記第１データ入力端との間に接続され、上記外部チップセレクト信号が活性状態のときオンになるように上記第１フリップフロップの出力でオン／オフ制御される転送ゲート（４４及び４５）と、該第１データ入力端と電源電位との間に接続され、該電源電位はこれを該第１データ入力端に供給すると上記第１ノードがフローティング状態になる電位であり、該転送ゲートのオン／オフ制御と連動してオフ／オン制御されるスイッチング素子（８１）とを有する。

【００２２】この同期型半導体記憶装置によれば、外部チップセレクト信号が不活性になると転送ゲートがオフ、スイッチング素子がオンになって、第１フリップフロップ内で第１及び第２ノードがフローティング状態になり、クロック入力端に外部クロックパルス列が供給されても第１及び第２ノード上の電荷がディスチャージされないので、消費電力が低減され、また、転送ゲートの入力側と出力側の信号の衝突が避けられる。

【００２３】請求項８の同期型半導体記憶装置では、請求項７において、上記タイミング回路はさらに、上記外部チップセレクト信号が活性化した時、上記外部クロックが上記第２レベルの場合に上記第１レベルになるのを待って上記第２データ入力端のレベルを反転する同期化

Micron Technology Inc. et al.
Ex. 1010, p. 237

回路（８２）を有する。

【００２４】この同期型半導体記憶装置によれば、外部チップセレクト信号の活性化により外部クロックに同期せずにコマンド入力回路が活性化されて誤動作するのが防止される。

【００２５】請求項９の同期型半導体記憶装置では、請求項５乃至８のいずれか１つにおいて、上記遅延回路は、上記第１フリップフロップの状態の変化に応答して状態が変化する第２フリップフロップ（２０）である。

【００２６】請求項１０の同期型半導体記憶装置では、請求項３又は９において、入力端が上記第２フリップフロップと並列接続され、第２イネーブル信号を出力し、バースト信号が活性であるときに該第２イネーブル信号が強制的に活性状態になる第３フリップフロップ（３０）と、該第２イネーブル信号が活性である間、上記外部クロックに応答して第２内部クロックを動作させる第２クロック入力回路（３１）と、該第２イネーブル信号が活性である期間、該第２クロックに同期してデータを入力するデータ入力回路（３２）とをさらに有する。

【００２７】この同期型半導体記憶装置によれば、バースト動作が終了すると第２イネーブル信号が不活性になるので、バーストモードにおいても消費電力低減効果が得られる。

【００２８】請求項１１の同期型半導体記憶装置では、請求項１乃至１０のいずれか１つにおいて例えば図８に示す如く、タイマーからの定期的なタイムアップ信号に応答して自動リフレッシュ動作を行う自動リフレッシュ回路（９６、９２、９１及び９４）をさらに有する。

【００２９】この同期型半導体記憶装置によれば、外部チップセレクト信号により選択されていないチップは自動的にパワーダウン状態になるとともに自動的にリフレッシュが行われるので、非選択チップに対しては一切の注意を払う必要がなくなり、これにより、コマンド供給はアクセスしたいチップに対してのみ行えばよく、このような同期型半導体記憶装置を用いたシステム設計が容易になるとともに、アクセスしたいチップに対するコマンド供給可能な時間が増えて、システムの高速化が可能となる。

【００３０】

【発明の実施の形態】以下、図面を参照して本発明の実施形態を説明する。

【００３１】［第１実施形態］図１は、本発明の第１実施形態に係る、ＳＤＲＡＭコントローラ１０ＡとＳＤＲＡＭ１１Ａ～１４Ａとの間の配線接続図であり、図９に対応している。図２は、図１の回路の動作を示す概略タイミングチャートである。

【００３２】ＳＤＲＡＭコントローラ１０Ａは、クロックイネーブル信号ＣＫＥを出力せず、また、図２に示す如く時点ｔ１及びｔ３でのコマンドＣＭＤ１及びＣＭＤ２の発行の１クロックサイクル前のクロックＣＬＫ立ち上がり時点ｔ０及びｔ２でチップセレクト信号＊ＣＳ１が低レベル（活性）になっている点で図９のＳＤＲＡＭコントローラ１０と異なり、他はＳＤＲＡＭコントローラ１０と同一である。

【００３３】ＳＤＲＡＭ１１Ａ～１４Ａは互いに同一構成であり、クロックイネーブル信号の外部端子がないことと、チップセレクト信号入力端子＊ＣＳの活性化（低レベル）に応答して図２に示すようなイネーブル信号を内部で生成し、これに応答して、内部クロックを動作させ入力回路を活性化する点で、図９のＳＤＲＡＭ１１～１４と異なる。ＳＤＲＡＭ１１Ａ～１４Ａの他の点は、図９のＳＤＲＡＭ１１～１４と同一である。

【００３４】１クロックサイクル前にチップセレクト信号＊ＣＳを活性化しておくのは、次のクロックサイクルの内部クロック立ち上がりでコマンドＣＭＤ等をラッチするのを確実に行うためである。図２では、ＳＤＲＡＭ１１Ａのみ選択され、ＳＤＲＡＭ１２Ａ～１４Ａ内のイネーブル信号ＥＮが低レベル（不活性）でこれらの内部クロックが停止し、これらの入力回路が不活性になってコマンドＣＭＤ等が該入力回路に取り込まれない。

【００３５】図３は、図１中の１つのＳＤＲＡＭ内の入力回路を示す。図３では簡単化のために、アドレスＡＤＤＲ、コマンドＣＭＤ、データＤＡＴＡ及びデータマスクＤＱＭがいずれも１ビットである場合を示している。多ビットの場合には、各ビットに対応した入力回路が並列して存在する。

【００３６】チップセレクト信号＊ＣＳを外部クロックＣＬＫの立ち上がりに同期化したイネーブル信号＊ＥＮ１は、ＲＳフリップフロップ２０に保持され、その非反転出力端Ｑから取り出されてクロック入力回路２１並びにアドレス・コマンド入力回路２２のオアゲート２３及び２４の一方の入力端に供給される。クロック入力回路２１、オアゲート２３及び２４の他方の入力端にはそれぞれ外部クロックＣＬＫ、アドレスＡＤＤＲ及びコマンドＣＭＤが供給される。オアゲート２３及び２４の出力はアドレス・コマンド入力回路２２のＤフリップフロップ２５及び２６のデータ入力端Ｄに供給され、Ｄフリップフロップ２５及び２６のクロック入力端ＣＫにはクロック入力回路２１の出力が内部クロックＣＬＫ１として供給される。

【００３７】イネーブル信号＊ＥＮ１が高レベルのとき、クロック入力回路２１の出力は外部クロックＣＬＫの変化によらず低レベルとなり、オアゲート２３及び２４の出力はそれぞれアドレスＡＤＤＲ及びコマンドＣＭＤの変化によらず高レベルとなるので、アドレス・コマンド入力回路２２は不活性状態になる。イネーブル信号＊ＥＮ１が低レベルのとき、外部クロックＣＬＫ、アドレスＡＤＤＲ及びコマンドＣＭＤはそれぞれクロック入力回路２１、オアゲート２３及び２４を通るので、アドレス・コマンド入力回路２２は活性状態になる。Ｄフリ

Micron Technology Inc. et al.
Ex. 1010, p. 238

ップフロップ２５及び２６の出力は内部回路に供給される。

【００３８】ここで、ＳＤＲＡＭ１１がバーストモードのとき、同一行アドレスに対し、列アドレスが不図示のカウンタに保持された後、そのカウントが内部クロックでインクリメントされてデータＤＡＴＡ及びデータマスクＤＱＭが読み込まれる。このとき、アドレス・コマンド入力回路２２は不活性である。そこで、ＲＳフリップフロップ２０とは別に、イネーブル信号＊ＥＮ２を保持するＲＳフリップフロップ３０が備えられている。

【００３９】ＲＳフリップフロップ３０は、ノアゲート３０１とノアゲート３０２とがクロス結合されてＲＳフリップフロップ２０と同一になっており、さらに、ノアゲート３０１のグランド電位供給端がｎＭＯＳトランジスタ３０３を介してグランド線に接続され、ノアゲート３０１の出力端がｐＭＯＳトランジスタ３０４を介して電源電位ＶＤＤに接続されている。ｎＭＯＳトランジスタ３０３及びｐＭＯＳトランジスタ３０４のゲートにはバースト信号＊ＢＳＴが供給されている。ノアゲート３０１は、ｎＭＯＳトランジスタ３０３がオフでセット入力端Ｓが高レベルのとき、出力がハイインピーダンス状態になる。バーストモードのとき、バースト信号＊ＢＳＴが低レベルになって、ｎＭＯＳトランジスタ３０３及びｐＭＯＳトランジスタ３０４がそれぞれオフ及びオンになり、セット入力端Ｓが高レベルであっても非反転出力端Ｑが低レベルに固定される。バースト信号＊ＢＳＴが高レベルのとき、ｎＭＯＳトランジスタ３０３及びｐＭＯＳトランジスタ３０４がそれぞれオン及びオフになって、ノアゲート３０１が活性になる。

【００４０】ＲＳフリップフロップ３０の非反転出力端Ｑから取り出されたイネーブル信号＊ＥＮ２は、クロック入力回路３１並びにデータ入力回路３２のオアゲート３３及び３４の一方の入力端に供給される。クロック入力回路３１、オアゲート３３及び３４の他方の入力端にはそれぞれ外部クロックＣＬＫ、データＤＡＴＡ及びデータマスクＤＱＭが供給される。オアゲート３３及び３４の出力はデータ入力回路３２のＤフリップフロップ３５及び３６のデータ入力端Ｄに供給され、Ｄフリップフロップ３５及び３６のクロック入力端ＣＫにはクロック入力回路３１の出力が内部クロックＣＬＫ２として供給される。

【００４１】イネーブル信号＊ＥＮ２が高レベルのとき、クロック入力回路３１の出力は外部クロックＣＬＫの変化によらず低レベルとなり、オアゲート３３及び３４の出力はそれぞれデータＤＡＴＡ及びデータマスクＤＱＭの変化によらず高レベルとなるので、データ入力回路３２は不活性状態になる。イネーブル信号＊ＥＮ２が低レベルのとき、外部クロックＣＬＫ、データＤＡＴＡ及びデータマスクＤＱＭはそれぞれクロック入力回路３１、オアゲート３３及び３４を通るので、データ入力回路３２は活性状態になる。Ｄフリップフロップ３５及び３６の出力は内部回路に供給される。

【００４２】ＲＳフリップフロップ２０及び３０のセット入力端Ｓには、Ｄフリップフロップ４０の反転出力端＊Ｑからチップセレクト信号＊ＣＳＣが供給される。Ｄフリップフロップ４０の相補データ入力端Ｄ及び＊Ｄには、チップセレクト信号＊ＣＳがタイミング回路４１を介して供給される。Ｄフリップフロップ４０は、チップセレクト信号＊ＣＳが高レベルになった後、外部クロックＣＬＫの最初の立ち上がりに同期してＲＳフリップフロップ２０及び３０をセットするためのものである。

【００４３】チップセレクト信号＊ＣＳの高レベルをタイミング回路４１で反転したものが内部クロックＣＬＫ１の立ち上がりでＤフリップフロップ４０に保持されると、チップセレクト信号＊ＣＳＣが高レベルとなってＲＳフリップフロップ２０がセットされ、イネーブル信号＊ＥＮ１が高レベルになり、クロック入力回路２１が閉じて内部クロックＣＬＫ１が停止する。このため、チップセレクト信号＊ＣＳが低レベルに遷移してもこれがＤフリップフロップ４０に保持されないので、Ｄフリップフロップ４０の非反転出力端ＱでＲＳフリップフロップ２０をリセットすることができない。また、外部クロックＣＬＫが高レベルのときに、チップセレクト信号＊ＣＳの立ち下がりでＲＳフリップフロップ２０をリセットすれば、クロック入力回路２１、オアゲート２３及び２４が活性化されると同時に、この高レベルがクロック入力回路２１を通って内部クロックＣＬＫ１が立ち上がるので、アドレス・コマンド入力回路２２の動作が１クロックサイクル早くなってしまう。

【００４４】そこで、構成要素４２～４８を備えたタイミング回路４１でこのような問題を解決している。

【００４５】タイミング回路４１では、ノアゲート４２及びオアゲート４３の一方の入力端にチップセレクト信号＊ＣＳが供給され、他方の入力端にそれぞれイネーブル信号＊ＥＮ１及びＲＳフリップフロップ２０の反転出力端＊Ｑからのイネーブル信号ＥＮ１が供給される。ノアゲート４２及びオアゲート４３はそれぞれ、チップセレクト信号＊ＣＳの不活性化（高レベルへの遷移）及び活性化（低レベルへの遷移）を検出するためのものである。ノアゲート４２の出力は、ｎＭＯＳトランジスタ４４とｐＭＯＳトランジスタ４５の並列接続を介し、一方ではＤフリップフロップ４０のデータ入力端Ｄに供給され、他方ではさらにインバータ４６を介してＤフリップフロップ４０のデータ入力端Ｄに供給される。ｎＭＯＳトランジスタ４４のゲートにはオアゲート４３の出力が供給され、ｐＭＯＳトランジスタ４５のゲートにはオアゲート４３の出力がインバータ４７を介して供給される。オアゲート４３の出力はさらに、ｐＭＯＳトランジスタ４８のゲートに供給される。ｐＭＯＳトランジスタ４８は、電源電位ＶＤＤとＤフリップフロップ４０のデ

Micron Technology Inc. et al.
Ex. 1010, p. 239

ータ入力端Dとの間に接続されている。インバータ４７の出力は、チップセレクト信号ＣＳＳＭとしてＲＳフリップフロップ２０及び３０のリセット入力端に供給される。

【００４６】上記構成において、ノアゲート４２が不活性（ＣＳＭが高レベル）でオアゲート４３が活性の状態においてチップセレクト信号＊ＣＳが低レベルに遷移すると、オアゲート４３の出力が低レベルに遷移し、ｐＭＯＳトランジスタ４８がオンになってＤフリップフロップ４０のデータ入力端D及び＊Dがそれぞれ高レベル及び低レベルに遷移し、また、インバータ４７の出力が高レベルに遷移してＲＳフリップフロップ２０及び３０がリセットされる。ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５は、ｐＭＯＳトランジスタがオンのときチップセレクト信号ＣＳＭが低レベルであるので、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５をオフにして衝突を避けるためのものである。

【００４７】チップセレクト信号入力回路は、タイミング回路４１、Ｄフリップフロップ４０及びＲＳフリップフロップ２０、３０で構成されている。

【００４８】次に、上記の如く構成された入力回路の動作を、図５を参照して説明する。

【００４９】図５において、図２と異なりチップセレクト信号＊ＣＳが複数クロックサイクル低レベルを維持しているのは、次の理由による。すなわち、ＳＤＲＡＭでは、あるバンクにコマンドを供給した後、次のクロックサイクルで他のバンクにコマンドを供給したり、また、同一バンクに対して１ビット読み出しコマンドを連続して供給する場合があるからである。図８について後述するように、クロックＣＬＫの立ち上がりの１クロックサイクル前にチップセレクト信号＊ＣＳが活性化していれば、すなわち内部チップセレクト信号＊ＣＳＣが活性のとき、内部クロックの立ち上がりに同期して、コマンドＣＭＤのデコードが行われ、その結果に応じた動作が行われる。

【００５０】（ｔ１）イネーブル信号＊ＥＮ１が低レベルであるので、チップセレクト信号＊ＣＳに対しノアゲート４２がインバータとして機能する。チップセレクト信号＊ＣＳが低レベルであるので、ノアゲート４２の出力ＣＳＭは高レベルになっている。イネーブル信号＊ＥＮ１及び＊ＥＮ２が低レベルであるので、クロック入力回路２１、オアゲート２３、２４、クロック入力回路３１、オアゲート３３及び３４が活性状態になっていて、Ｄフリップフロップ４０、アドレス・コマンド入力回路２２及びデータ入力回路３２も活性状態になっている。また、イネーブル信号ＥＮ１が高レベルであるので、オアゲート４３が不活性、すなわちその出力がチップセレクト信号＊ＣＳの変化によらず高レベルになっていて、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオン、ｐＭＯＳトランジスタ４８がオフとなって、Ｄフリップフロップ４０のデータ入力端D及び＊Dにそれぞれ高レベル及び低レベルが供給される。インバータ４７の出力ＣＳＳＭ及びＤフリップフロップ４０の反転出力端＊Ｑは低レベルであり、ＲＳフリップフロップ２０及び３０のリセット入力端Ｒ及びセット入力端Ｓは低レベルになっている。バースト信号＊ＢＳＴが低レベル（バーストモード）であり、ｎＭＯＳトランジスタ３０３及びｐＭＯＳトランジスタ３０４はそれぞれオフ及びオンになっている。

【００５１】（ｔ２）この状態からチップセレクト信号＊ＣＳが立ち上がって、ノアゲート４２の出力ＣＳＭが低レベルに遷移し、Ｄフリップフロップ４０のデータ入力端D及び＊Dがそれぞれ低レベル及び高レベルになる。

【００５２】（ｔ３）次に外部クロックＣＬＫが立ち上がり、内部クロックＣＬＫ１が立ち上がって、Ｄフリップフロップ４０の入力が保持され、その反転出力端＊Ｑからのチップセレクト信号＊ＣＳＣが高レベルに遷移する。チップセレクト信号＊ＣＳＣの高レベルは、ＲＳフリップフロップ２０内のノアゲート２０１の出力ＥＮ１を低レベルに反転させ、次いでノアゲート２０２の出力＊ＥＮ１を高レベルに反転させるので、ＲＳフリップフロップ２０は遅延回路としても機能している。この信号伝搬遅延により、ＣＬＫ１のパルスがクロック入力回路２１を通った後に、イネーブル信号＊ＥＮ１が高レベルに遷移してクロック入力回路２１が不活性になる。

【００５３】ＣＬＫ１の立ち上がりに同期して、アドレスＡＤＤＲ及びコマンドＣＭＤがアドレス・コマンド入力回路２２のＤフリップフロップに保持される。＊ＢＳＴが低レベルであるので、このアドレスＡＤＤＲは列アドレスであり、コマンドＣＭＤはリードコマンド又はライトコマンドである。

【００５４】イネーブル信号＊ＥＮ１が高レベルに遷移すると、クロック入力回路２１、オアゲート２３及び２４が不活性になり、内部クロックＣＬＫ１の動作が停止してＤフリップフロップ４０及びアドレス・コマンド入力回路２２も不活性になる。また、ノアゲート４２が不活性になると同時にオアゲート４３が活性になり、後に生ずるチップセレクト信号＊ＣＳの立ち下がりは、ノアゲート４２を通らずにオアゲート４３を通る。オアゲート４３の出力は高レベルを維持している。

【００５５】バースト信号＊ＢＳＴが低レベルであるので、チップセレクト信号＊ＣＳＣが高レベルに遷移してもイネーブル信号＊ＥＮ２は低レベルを維持、すなわちクロック入力回路３１及びデータ入力回路３２は活性状態を維持する。これによりバーストモードでの動作が行われ、データＤＡＴＡ及びデータマスクＤＱＭの内容に応じたデータが内部に書き込まれる。Ｄフリップフロップ３６の出力が高レベルのとき、Ｄフリップフロップ３５の出力は内部回路でマスクされ、指定アドレスのメモ

Micron Technology Inc. et al.
Ex. 1010, p. 240

リセルへの書き込みは行われない。Dフリップフロップ３６の出力が低レベルのとき、このマスクが解除されてDフリップフロップ３５の出力が指定アドレスのメモリセルへ書き込まれる。

【００５６】（ｔ４）バーストモードが終了してバースト信号＊ＢＳＴが高レベルに遷移し、これによりＲＳフリップフロップ３０のセット入力端Ｓに供給されている高レベルが有効になり、イネーブル信号＊ＥＮ２が高レベルとなって、クロック入力回路３１、オアゲート３３及び３４が不活性化され、内部クロックＣＬＫ２の動作が停止する。その結果、不図示のコラムアドレス系回路が不活性になる。

【００５７】（ｔ５）チップセレクト信号＊ＣＳが低レベルに遷移して、オアゲート４３及びインバータ４７の出力がそれぞれ低レベル及び高レベルに遷移し、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオフ、ｐＭＯＳトランジスタ４８がオンになる。これにより、チップセレクト信号ＣＳＳＭが高レベルに遷移して、ＲＳフリップフロップ２０及び３０がリセットされ、イネーブル信号＊ＥＮ１及び＊ＥＮ２が低レベルに遷移してノアゲート４２、クロック入力回路２１、２３、２４、クロック入力回路３１、オアゲート３３及び３４が活性化され、内部クロックＣＬＫ１及びＣＬＫ２が動作を開始し、ノアゲート４２、Ｄフリップフロップ４０、アドレス・コマンド入力回路２２及びデータ入力回路３２が活性状態になる。ノアゲート４２の活性化によりチップセレクト信号ＣＳＭが高レベルに遷移する。

【００５８】このような動作により、内部クロックＣＬＫ１及びＣＬＫ２が停止していてもチップセレクト信号＊ＣＳの活性化（立ち下がり）に応答してイネーブル信号＊ＥＮ１及び＊ＥＮ２が活性化されて、内部クロックＣＬＫ１及びＣＬＫ２の動作が再開される。

【００５９】（ｔ６）Ｄフリップフロップ４０のデータ入力端Ｄ及＊Ｄがそれぞれ高レベル及び低レベルであるので、内部クロックＣＬＫ１の立ち上がりに同期してＤフリップフロップ４０がセットされ、チップセレクト信号＊ＣＳＣが低レベルに遷移する。これにより、イネーブル信号ＥＮ１が高レベルに遷移してオアゲート４３が不活性になり、その出力が高レベルに遷移してｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオン、ｐＭＯＳトランジスタ４８がオフになり、またチップセレクト信号ＣＳＳＭが低レベルに遷移する。

【００６０】内部クロックＣＬＫ１の立ち上がり時点では内部クロック＊ＣＳＣが高レベルであるので、コマンドＣＭＤのデコードは行われない。

【００６１】（ｔ７）バースト信号＊ＢＳＴが低レベルに遷移して、時点ｔ１と同じ状態になる。

【００６２】本第１実施形態によれば、クロックＣＬＫの立ち上がりの１クロックサイクル前にチップセレクト信号＊ＣＳが活性（低レベル）であれば内部クロックＣＬＫ１が動作しており、内部チップセレクト信号＊ＣＳＣが活性のとき、内部クロックＣＬＫ１の立ち上がりでコマンドＣＭＤ等が内部回路に取り込まれてこれに応じた動作が行われる。また、チップセレクト信号＊ＣＳが不活性（高レベル）になってから、１クロックパルスの後に内部クロックＣＬＫ１が停止して、アドレス・コマンド入力回路２２が不活性になり、さらに、バーストモードによるデータ取り込み終了後に内部クロックＣＬＫ２が停止してデータ入力回路３２が不活性になるので、選択されたＳＤＲＡＭ１１においても入力回路の消費電力が低減される。しかも、非選択のＳＤＲＡＭ１１ではチップセレクト信号＊ＣＳが高レベルであるので、内部クロックが停止して入力回路の消費電力が低減される。クロックイネーブル信号ＣＫＥを用いずにこのような動作が自動的に行われるので、低消費電力化が達成されるとともに、複数のＳＤＲＡＭ１１を用いたシステムの制御及び設計が容易になる。

【００６３】図４は、図３中のＤフリップフロップ４０の構成例を示す。この回路自体は公知であるが、クロックのパルス列による無駄な電力消費を示すためにこの回路の構成をその動作とともに説明する。

【００６４】Ｄフリップフロップ４０は、入力段のダイナミックフリップフロップ５０と、その一対の出力の各々を相補信号化してダイナミックフリップフロップ５０をフィドバック制御する相補化回路６０と、回路６０の出力により状態が制御されるラッチ回路７０とからなる。

【００６５】ダイナミックフリップフロップ５０では、クロック入力端ＣＫが低レベルのとき、ｐＭＯＳトランジスタ５１Ａ及び５１Ｂがオン、ｎＭＯＳトランジスタ５２がオフとなって、ノードＮＯ１及びＮＯ２へのプリチャージが行われる。

【００６６】増幅回路５３は、フリップフロップ型であり、ｐＭＯＳトランジスタ５１Ａ及び５１Ｂにそれぞれ並列接続されたｐＭＯＳトランジスタ５４Ａ及び５４Ｂと、ｐＭＯＳトランジスタ５４Ａ及び５４Ｂにそれぞれ直列接続されたｎＭＯＳトランジスタ５５Ａ及び５５Ｂとからなり、これらがクロス結合されている。上記ノードＮＯ１及びＮＯ２へのプリチャージによりｐＭＯＳトランジスタ５４Ａ及び５４Ｂがオフ、ｎＭＯＳトランジスタ５５Ａ及び５５Ｂがオンで、増幅回路５３が不活性状態になっている。

【００６７】ｎＭＯＳトランジスタ５５Ａのソースとｎ ＭＯＳトランジスタ５２のドレインとの間にはｎＭＯＳトランジスタ５６Ａと５７Ａとが並列接続され、ｎＭＯＳトランジスタ５５ＢのソースとｎＭＯＳトランジスタ５２のドレインとの間にはｎＭＯＳトランジスタ５６Ｂと５７Ｂとが並列接続されている。ノードＮＯ１は相補化回路６０のインバータ６１Ａを介してｎＭＯＳトランジスタ５７Ａのゲートに接続され、ノードＮＯ２は相補

Micron Technology Inc. et al.
Ex. 1010, p. 241

化回路６０のインバータ６１Ｂを介してｎＭＯＳトランジスタ５７Ｂのゲートに接続されている。したがって、ノードＮＯ１及びＮＯ２がプリチャージされている状態では、インバータ６１Ａ及び６１Ｂの出力が低レベルとなってｎＭＯＳトランジスタ５７Ａ及び５７Ｂがオフになっている。ｎＭＯＳトランジスタ５６Ａ及び５６Ｂのゲートにはそれぞれ、データ入力端Ｄ及び＊Ｄからの信号が供給される。データ入力端Ｄ及び＊Ｄの信号レベルが互いに逆の時、ｎＭＯＳトランジスタ５６Ａ及び５６Ｂは一方がオン、他方がオフであり、データ入力端Ｄが高レベルのときにはｎＭＯＳトランジスタ５６Ａがオン、ｎＭＯＳトランジスタ５６Ｂがオフになっている。

【００６８】上記プリチャージ状態において、ラッチ回路７０では、ｎＭＯＳトランジスタ７１Ａ、ｐＭＯＳトランジスタ７１Ｂ、７２Ａ及びｎＭＯＳトランジスタ７２Ｂが全てオフになっており、その前の状態が、インバータ７３と７４とからなる保持回路に保持される。

【００６９】この状態から外部クロックＣＬＫが立ち上がると、ｐＭＯＳトランジスタ５１Ａ及び５１Ｂがオフ、ｎＭＯＳトランジスタ５２がオンになり、ノードＮＯ１上の電荷がｎＭＯＳトランジスタ５５Ａ、５６Ａ及び５２を通ってグランド線へ引き抜かれ、その電位が低下し、ノードＮＯ２とＮＯ１の電位差が増幅回路５３で増幅され、インバータ６１Ａの出力が高レベルに遷移してｎＭＯＳトランジスタ５７Ａがオンになり、ノードＮＯ１からの電荷引き抜きが助長される。また、ｎＭＯＳトランジスタ７１Ａ及びｐＭＯＳトランジスタ７２Ａがオンになって、出力端Ｑ及び＊Ｑがそれぞれ高レベル及び低レベルになる。

【００７０】逆にデータ入力端Ｄが低レベルの場合には、ｎＭＯＳトランジスタ５６Ａ及び５７Ａがオフ、ｎＭＯＳトランジスタ５６Ｂ及び５７Ｂがオン、ｐＭＯＳトランジスタ７１Ｂ及びｎＭＯＳトランジスタ７２Ｂがオンになって、出力端Ｑ及び＊Ｑがそれぞれ低レベル及び高レベルになる。

【００７１】次に、クロック入力端ＣＫが低レベルになると、上記プリチャージが行われ、ｐＭＯＳトランジスタ７１Ｂ、ｎＭＯＳトランジスタ７１Ａ、ｐＭＯＳトランジスタ７２Ａ及びｎＭＯＳトランジスタ７２Ｂがオフになって出力端Ｑ及び＊Ｑの電位が保持される。

【００７２】データ入力端Ｄが高レベルの状態でクロック入力端ＣＫにパルス列が供給されると、出力端Ｑ及び＊Ｑはそれぞれ高レベル及び低レベルで一定であるが、ノードＮＯ１のプリチャージ及びディスチャージが繰り返されて、無駄な電力が消費される。図５では内部クロックＣＫ１及びＣＫ２のパルスが不要になるとそのクロックが停止するので、消費電力が低減される。

【００７３】クロック入力端ＣＫが高レベルの状態でデータ入力端Ｄが変化しても、増幅回路５３の増幅動作によりその出力が保持されるので、Ｄフリップフロップ４０の出力も一定に保持される。

【００７４】［第２実施形態］図６は、本発明の第２実施形態のＳＤＲＡＭ内の入力回路を示し、図３に対応している。

【００７５】この回路では、図３のタイミング回路４１の替わりにタイミング回路４１Ａを用い、外部クロックＣＬＫをＤフリップフロップ４０のクロック入力端ＣＫに直接供給し、Ｄフリップフロップ４０の非反転出力端ＱをＲＳフリップフロップ２０及び３０のセット入力端Ｓに接続し、Ｄフリップフロップ４０の反転出力端＊ＱをＲＳフリップフロップ２０及び３０のリセット入力端Ｒに接続している。

【００７６】タイミング回路４１Ａでは、チップセレクト信号＊ＣＳの駆動能力がツルーバッファゲート８０で増幅され、その出力＊ＣＳＭが、ｎＭＯＳトランジスタ４４とｐＭＯＳトランジスタ４５との並列接続を介してＤフリップフロップ４０のデータ入力端Ｄに供給される。データ入力端Ｄとグランド線との間にはｎＭＯＳトランジスタ８１が接続されている。ｎＭＯＳトランジスタ４４のゲートはＤフリップフロップ４０の非反転出力端Ｑに接続され、ｐＭＯＳトランジスタ４５及びｎＭＯＳトランジスタ８１のゲートはＤフリップフロップ４０の反転出力端＊Ｑに接続されている。これにより、出力端Ｑ及び＊Ｑがそれぞれ低レベル及び高レベルの状態では、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオン、ｎＭＯＳトランジスタ８１がオフになり、低レベルのチップセレクト信号＊ＣＳがツルーバッファゲート８０、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５を介してＤフリップフロップ４０のデータ入力端Ｄに供給される。

【００７７】この状態からチップセレクト信号＊ＣＳが高レベルに遷移すると、外部クロックＣＬＫの立ち上がりに同期してこれがＤフリップフロップ４０に保持され、出力端Ｑ及び＊Ｑがそれぞれ高レベル及び低レベルに遷移して、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオフ、ｎＭＯＳトランジスタ８１がオンになり、Ｄフリップフロップ４０のデータ入力端Ｄが低レベルになる。図４において、クロック入力端ＣＫが低レベルのとき、ｐＭＯＳトランジスタ５１Ａ及び５１Ｂがオン、ｎＭＯＳトランジスタ５２がオフになってノードＮＯ１及びＮＯ２がプリチャージされ、これによりｎＭＯＳトランジスタ５５Ａ及び５５Ｂがオン、ｎＭＯＳトランジスタ５７Ａ及び５７Ｂがオフになっている。データ入力端Ｄ及び＊Ｄが共に低レベルであれば、ｎＭＯＳトランジスタ５６Ａ及び５６Ｂがオフであるので、外部クロックＣＬＫが高レベルになってもディスチャージが行われず、外部クロックＣＬＫのパルス列がＤフリップフロップ４０のクロック入力端ＣＫＤに供給されても、無駄な電力消費が防止される。

【００７８】チップセレクト信号＊ＣＳが低レベルに遷

Micron Technology Inc. et al.
Ex. 1010, p. 242

移する場合には、上記ｎＭＯＳトランジスタ８１のオンによりＤフリップフロップ４０のデータ入力端Ｄが既に低レベルになっているので、データ入力端＊Ｄのみ高レベルに遷移させればよい。この時、外部クロックＣＬＫが高レベルであれば、図４においてノードＮＯ２上の電荷がｎＭＯＳトランジスタ５５Ｂ、５６Ｂ及び５２を通って引き抜かれるので、直ちにＤフリップフロップ４０の出力が反転し、誤動作する。

【００７９】この問題を解決するために、ツルーバッファゲート８０の出力＊ＣＳＭは、立ち下がり同期化インバータ８２を介してＤフリップフロップ４０のデータ入力端＊Ｄに接続されている。立ち下がり同期化インバータ８２は、インバータ４６の入力を外部クロックＣＬＫの立ち下がりに同期して反転させるためのものであり、ツルーバッファゲート８０の出力端とＤフリップフロップ４０のデータ入力端＊Ｄとの間にインバータ４６が接続され、その電源入力端がｐＭＯＳトランジスタ８３を介して電源電位ＶＤＤに接続され、ｐＭＯＳトランジスタ８３のゲートが外部クロックＣＬＫに接続されている。外部クロックＣＬＫが低レベルのときには、ｐＭＯＳトランジスタ８３がオンであるので、立ち下がり同期化インバータ８２は通常のインバータとして機能する。外部クロックＣＬＫが高レベルのときには、ｐＭＯＳトランジスタ８３がオフであり、インバータ４６の出力が入力の変化に応答しないので、外部クロックＣＬＫが高レベルになる前のインバータ４６の出力レベルを保持するために、インバータ４６の出力端にインバータ８４と８５とが環状結合された保持回路が接続されている。

【００８０】図６の回路の他の構成は、図３のそれと同一である。

【００８１】図７は、図６の回路の動作を示すタイミングチャートである。

【００８２】図７では、時点ｔ１の前にチップセレクト信号＊ＣＳが低レベルで、これがＤフリップフロップ４０に保持され、その出力端Ｑ及び＊Ｑからのチップセレクト信号＊ＣＳＣ及びＣＳＣがそれぞれ低レベル及び高レベルになっていて、ＲＳフリップフロップ２０及び３０がリセット状態である。したがって、イネーブル信号＊ＥＮ１及び＊ＥＮ２が低レベルで、クロック入力回路２１、オアゲート２３、２４、クロック入力回路３１、オアゲート３３及び３４が活性状態であり、内部クロックＣＬＫ１及びＣＬＫ２が動作し、アドレス・コマンド入力回路２２及びデータ入力回路３２が活性状態になっている。また、Ｄフリップフロップ４０の出力端Ｑ及び＊Ｑがそれぞれ低レベル及び高レベルであるので、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５はオン、ｎＭＯＳトランジスタ８１はオフになっている。Ｄフリップフロップ４０の出データ入力端Ｄ及び＊Ｄがそれぞれ低レベル及び高レベルであるので図４のｎＭＯＳトランジスタ５６Ａ及び５６Ｂがそれぞれオフ及びオンになっている。

【００８３】（ｔ１）外部クロックＣＬＫの立ち上がりにより、図４においてクロック入力端ＣＫが高レベルに遷移し、ノードＮＯ２の電荷がｎＭＯＳトランジスタ５５Ｂ、５６Ｂ及び５２を通ってグランド線へ引き抜かれる。これに対し、ｎＭＯＳトランジスタ５６Ａ及び５７Ａがオフであるので、ノードＮＯ１は高レベルを維持している。

【００８４】（ｔ２）チップセレクト信号＊ＣＳが高レベルに遷移し、ツルーバッファゲート８０の出力＊ＣＳＭが高レベルに遷移し、図４中のノードＮＯ３が高レベルに遷移して、ｎＭＯＳトランジスタ５６Ａがオンになる。ｎＭＯＳトランジスタ５５Ａがオフであるので、ノードＮＯ１は高レベルを維持している。

【００８５】（ｔ３）外部クロックＣＬＫが低レベルに遷移し、ｐＭＯＳトランジスタ８３がオンになって、チップセレクト信号＊ＣＳＭがインバータ４６で反転され、図４中のノードＮＯ４が低レベルになってｎＭＯＳトランジスタ５６Ｂがオフになる。

【００８６】（ｔ４）内部クロックＣＬＫ１の立ち上がりに同期して、アドレスＡＤＤＲ及びコマンドＣＭＤがアドレス・コマンド入力回路２２のＤフリップフロップに保持される。

【００８７】外部クロックＣＬＫの立ち上がりによりチップセレクト信号＊ＣＳの高レベルがＤフリップフロップ４０に保持され、出力端Ｑ及び＊Ｑがそれぞれ高レベル及び低レベルに遷移する。これにより、ＲＳフリップフロップ２０がセットされてイネーブル信号＊ＥＮ１が高レベルに遷移し、クロック入力回路２１、オアゲート２３及び２４が不活性になる。ＲＳフリップフロップ３０については、バースト信号＊ＢＳＴが低レベルであるので、イネーブル信号＊ＥＮ２の低レベルが維持され、データ入力回路３２は活性状態を維持する。また、ｎＭＯＳトランジスタ４４及びｐＭＯＳトランジスタ４５がオフ、ｎＭＯＳトランジスタ８１がオンになり、図４中のノードＮＯ３が低レベルとなってｎＭＯＳトランジスタ５６Ａがオフになる。

【００８８】（ｔ５）外部クロックＣＬＫの立ち下がりにより、図４においてｐＭＯＳトランジスタ５１Ａ及び５１Ｂがオン、ｎＭＯＳトランジスタ５２がオフになり、ノードＮＯ１がプリチャージされ、増幅回路５３が不活性になる。この状態では、ｎＭＯＳトランジスタ５６Ａ、５７Ａ、５６Ｂ及び５７Ｂがオフであるので、外部クロックＣＬＫが動作してもノードＮＯ１及びＮＯ２のディスチャージが行われず、Ｄフリップフロップ４０での無駄な電力消費が防止される。

【００８９】（ｔ６）バーストモードが終了してバースト信号＊ＢＳＴが高レベルに遷移し、これによりＲＳフリップフロップ３０のセット入力端Ｓに供給されている高レベルが有効になり、イネーブル信号＊ＥＮ２が高レ

Micron Technology Inc. et al.
Ex. 1010, p. 243

ベルとなって、クロック入力回路３１、オアゲート３３及び３４が不活性化され、内部クロックＣＬＫ２の動作が停止する。その結果、不図示のコラムアドレス系回路が不活性になる。

【００９０】（ｔ７）チップセレクト信号＊ＣＳが立ち下がり、ツルーバッファゲート８０の出力＊ＣＳＭも立ち下がる。このとき、外部クロックＣＬＫが高レベルでｎＭＯＳトランジスタ５２がオンであっても、ｐＭＯＳトランジスタ８３がオフであるのでインバータ４６の出力がハイインピーダンス状態になってノードＮＯ４は低レベルを維持し、ｎＭＯＳトランジスタ５６Ｂはオフを維持し、Ｄフリップフロップ４０はセット状態を維持する。したがって、チップセレクト信号＊ＣＳの立ち下がりにより外部クロックＣＬＫに同期せずにアドレス・コマンド入力回路２２が活性化されて誤動作するのが防止される。

【００９１】（ｔ８）外部クロックＣＬＫの立ち下がりによりｐＭＯＳトランジスタ８３がオンになってインバータ４６が活性化され、ノードＮＯ４が高レベルに遷移し、ｎＭＯＳトランジスタ５６Ｂがオンになる。これと同時にｎＭＯＳトランジスタ５２がオフになるので、ノードＮＯ２は高レベルを維持する。

【００９２】（ｔ９）外部クロックＣＬＫの立ち上がりによりｐＭＯＳトランジスタ５１Ａ及び５１Ｂがオフ、ｎＭＯＳトランジスタ５２がオンになって、ノードＮＯ２の電荷がｎＭＯＳトランジスタ５５Ｂ、５６Ｂ及び５２を通ってグランド線に引き抜かれ、ノードＮＯ１とＮＯ２の電位差が増幅回路５３で増幅されてノードＮＯ２が低レベルに遷移する。これにより、出力端Ｑ及び＊Ｑがそれぞれ低レベル及び高レベルに遷移し、ＲＳフリップフロップ２０及び３０がリセットされて、イネーブル信号＊ＥＮ１及び＊ＥＮ２が低レベルとなり、クロック入力回路２１、２３、２４、クロック入力回路３１、オアゲート３３及び３４が活性化される。

【００９３】時点ｔ９は図５中の時点ｔ６に相当し、内部クロックＣＬＫ１が停止している点で図５の場合と異なる。しかし、次のクロックＣＬＫの立ち上がりでは図５及び図７のいずれについても、内部チップセレクト信号＊ＣＳＣが活性で内部クロックＣＬＫ１が立ち上がるので、図８について後述するように、コマンドＣＭＤのデコード開始時点は同じになる。

【００９４】（ｔ１０）バースト信号＊ＢＳＴが低レベルに遷移して、時点ｔ１の前と同じ状態になる。

【００９５】本第２実施形態についても、上記第１実施形態と同様の効果が得られる。

【００９６】［第３実施形態］図８は、本発明の第３実施形態のＳＤＲＡＭ１１Ｂの概略構成を示すブロック図である。

【００９７】入力回路９０は、図３又は図６に示す回路を含んでいる。図３及び図６では簡単化のためにアドレスＡＤＤＲ、コマンドＣＭＤ、データＤＡＴＡ及びデータマスクＤＱＭが１ビットである場合を示しているが、入力回路９０ではこれらがいずれも多ビットである。また、図３及び図６ではデータ出力回路が省略されているが、入力回路９０にはデータ出力回路が含まれている。

【００９８】ロウアドレス切換回路９１は、アービター９２からの制御信号に応答してアドレス入力回路９３又はリフレッシュアドレス発生回路９４からのロウアドレスの一方を選択し、ロウデコーダ９５へ供給する。アービター９２は、先に活性化された入力を受け付け、同時の場合には、リフレッシュタイマ９６から定期的に供給されるタイムアップ信号の活性化を優先的に受け付ける。アービター９２は、タイムアップ信号の活性化を受け付けた場合、ロウアドレス切換回路９１に対しリフレッシュアドレス発生回路９４の出力を選択させる。

【００９９】デコードイネーブル回路としてのＤフリップフロップ４０Ｎのデータ入力端Ｄ及びクロック入力端にはそれぞれ、入力回路９０の＊ＣＳ入力回路及びクロック入力回路から内部チップセレクト信号＊ＣＳＣ及び内部クロックＣＬＫ１が供給され、Ｄフリップフロップ４０Ｎの反転出力端Ｑからイネーブル信号ＥＮ３がコマンドデコーダに供給される。コマンド入力回路９７からの信号がコマンドデコーダ９８に供給され、イネーブル信号ＥＮ３が活性（高レベル）のとき、アクティベイトコマンド、リードコマンド、ライトコマンド又はプリチャージコマンド等のコマンドのデコードが行われる。すなわち、上述のように、内部チップセレクト信号＊ＣＳＣが活性のとき内部クロックＣＬＫ１の立ち上がりでコマンドＣＭＤのデコードが行われる（コマンドデコーダ９８が活性化される）。、アクティベイトコマンドはアービター９２に供給され、アービター９２はこれを受け付けると、ロウアドレス切換回路９１に対しアドレス入力回路９３の出力を選択させる。ロウデコーダ９５の出力によりメモリセル１００内のワード線が選択されてそのセル行の内容がセンスアンプに保持される。

【０１００】コラムアドレス切換回路１０１は、コマンドデコーダ９８の出力がリードコマンド又はライトコマンドであることを示している場合には、アドレス入力回路９３からのアドレスを列アドレスとして受け取り、バーストモードのときにはこれをバーストアドレス発生回路１０３のカウンタにロードさせる。このカウンタは、クロック入力回路からの内部クロックＣＬＫ２でインクリメントされる。コラムアドレス切換回路１０１は、バーストモードでない場合にはアドレス入力回路９３の出力を選択し、バーストモードの場合にはバーストアドレス発生回路１０３の出力を選択して、コラムデコーダ１０２に供給する。コラムデコーダ１０２の出力によりメモリセル１００内のコラムが選択される。

【０１０１】書き込みの場合には、入力回路９０のデータ入力回路からのデータがリード・ライトアンプ１０４

Micron Technology Inc. et al.
Ex. 1010, p. 244

を介してメモリセル100内の選択された行及び列のセルに書き込まれ、読み出しの場合には、この選択されたメモリセルの内容がリード・ライトアンプ104及び入力回路90のデータ出力回路を介して外部に取り出される。

【0102】アービター92を用いてアドレス入力回路93とリフレッシュアドレス発生回路94の出力の一方をロウアドレス切換回路91で選択するのは、非同期型RAMでは公知であるがSDRAMでは用いられていない。

【0103】本第3実施形態では、上記第1実施形態又は第2実施形態のような入力回路90を用い、かつ、リフレッシュタイマ96及アービター92を用いており、チップセレクト信号*CSにより選択されていないチップは自動的にパワーダウン状態になるとともに自動的にリフレッシュが行われるので、非選択チップに対しては一切の注意を払う必要がなくなる。これにより、コマンド供給はアクセスしたいチップに対してのみ行えばよく、このようなSDRAMを用いたシステム設計が容易になるとともに、アクセスしたいチップに対するコマンド供給可能な時間が増えて、システムの高速化が可能となる。

【図面の簡単な説明】

【図1】本発明の第1実施形態に係る、SDRAMコントローラと複数のSDRAMとの間の配線接続図である。

【図2】図1の回路の動作を示す概略タイミングチャートである。

【図3】図1中の1つのSDRAM内の入力回路を示す図である。

【図4】図3中のDフリップフロップの構成例を示す図である。

【図5】図3の回路の動作を示すタイミングチャートである。

【図6】本発明の第2実施形態のSDRAM内の入力回路を示す図である。

【図7】図6及び図4の回路の動作を示すタイミングチャートである。

【図8】本発明の第3実施形態のSDRAMを示す概略ブロック図である。

【図9】従来の、SDRAMコントローラと複数のSDRAMとの間の配線接続図である。

【図10】図9の回路の動作を示す概略タイミングチャートである。

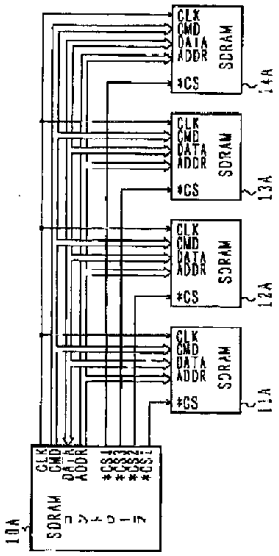
【符号の説明】

10、10A SDRAMコントローラ
11～14、11A～14A SDRAM
20、30 RSフリップフロップ
21、23、24、31、33、34、43 オアゲート
22 アドレス・コマンド入力回路
25、26、35、36、40 Dフリップフロップ
301、302、42 ノアゲート
304、45、48、51A、51B、54A、54B、71B、72A、45、83 pMOSトランジスタ
303、44、52、55A、55B、56A、56B、57A、57B、71A、72B、44、81 nMOSトランジスタ
32 データ入力回路
41、41A タイミング回路
46、47、61A、61B、73、74、84、85 インバータ
50、50A ダイナミックフリップフロップ
53 増幅回路
60 相補化回路
70 ラッチ回路
80 ツルーバッファゲート
82 立ち下がり同期化インバータ
90 入力回路
*CS、CSSM、CSM、*CSC チップセレクト信号
CLK 外部クロック
CLK1、CLK2 内部クロック
EN1、*EN1、*EN2 イネーブル信号
*BST バースト信号
ADDR アドレス
CMD コマンド
DATA データ
DQM データマスク
NO1～NO4 ノード

Micron Technology Inc. et al.
Ex. 1010, p. 245

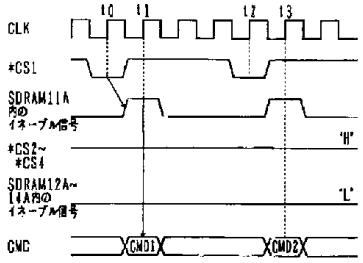
【図1】

本発明の第1実施形態に係る、SDRAMコントローラと複数のSDRAMとの間の配線接続図



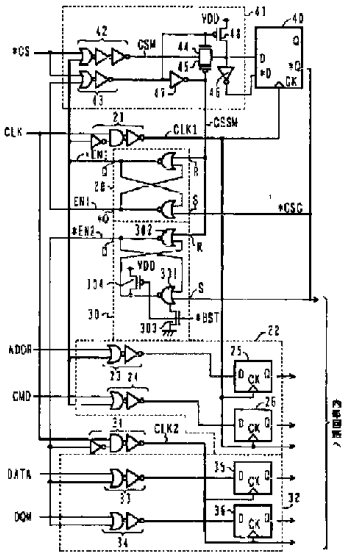
【図2】

図1の回路の動作を示す概略タイミングチャート



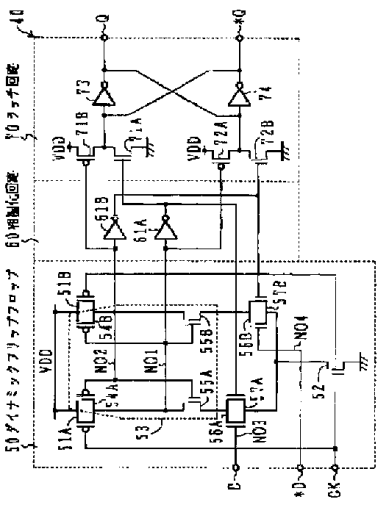
【図3】

図1中の1つのSDRAM内の入力回路を示す図



【図4】

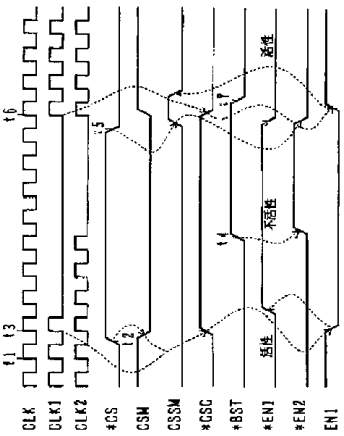
図3中のDフリップフロップの構成例を示す図



Micron Technology Inc. et al.
Ex. 1010, p. 246

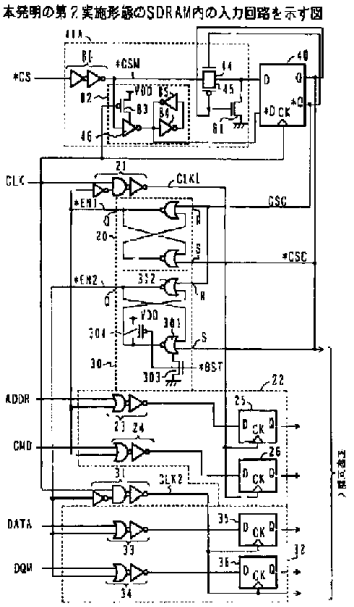
【図5】

図3の回路の動作を示すタイミングチャート



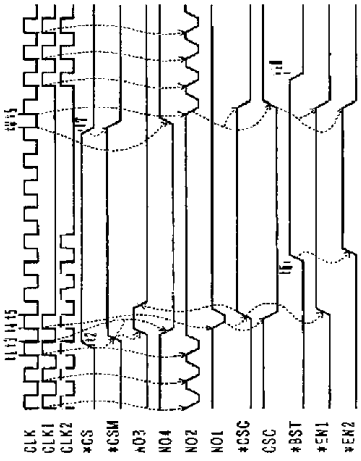
【図6】

本発明の第2実施形態のSDRAM内の入力回路を示す図



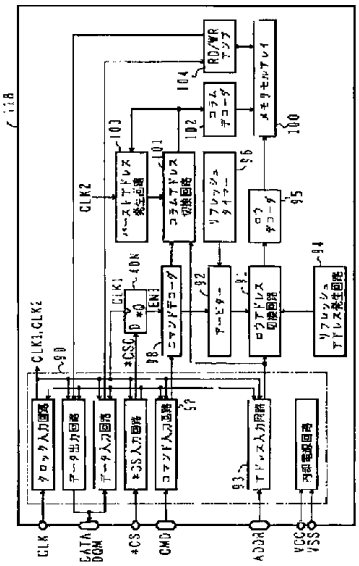
【図7】

図5及び図4の回路の動作を示すタイミングチャート



【図8】

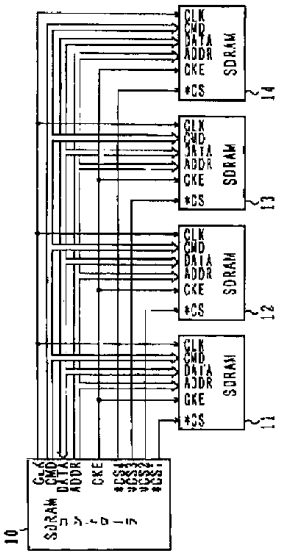
本発明の第3実施形態のSDRAMを示す概略ブロック図



Micron Technology Inc. et al.
Ex. 1010, p. 247

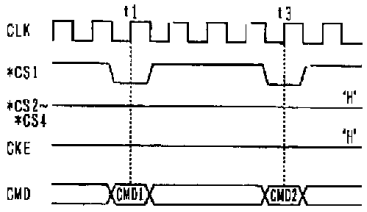
【図9】

従来の、SDRAMコントローラと複数のSDRAMとの間の配線接続図



【図10】

図9の回路の動作を示す概略タイミングチャート



Espacenet

Bibliographic data: JP2002184176 (A) — 2002-06-28

SDRAM CONTROL CIRCUIT AND SDRAM MODULE

Inventor(s): KIKUCHI MASASHI ± (KIKUCHI MASASHI)

Applicant(s): NEC TOHOKU LTD ± (NEC TOHOKU LTD)

Classification: - **international:** ***G11C11/401; G11C11/407; G11C5/00;*** (IPC1-7): G11C11/401; G11C11/407

- **European:**

Application number: JP20000378746 20001213

Priority number (s): JP20000378746 20001213

Abstract of JP2002184176 (A)

PROBLEM TO BE SOLVED: To provide a memory module being inexpensive, practical, and large capacity by constituting a memory module (DIMM) using an inexpensive SDRAM put to practical use now. SOLUTION: A control circuit outputting plural new chip selecting signal from an address signal, a clock enable-signal, and a chip selecting signal is formed to a SDRAM module, and the number of SDRAM being controllable is increased. Consequently, four inexpensive 16 Mbit× 4 bit.SDRAM can be used instead of expensive 32 Mbit× 8 bit.SDRAM.

Last updated: 24.09.2012 Worldwide Database 5.7.42; 93p

Micron Technology Inc. et al.
Ex. 1010, p. 249

(19)日本国特許庁（ＪＰ）　(12)公開特許公報（Ａ）　(11)特許出願公開番号
特開2002－184176
（Ｐ2002－184176Ａ）
(43)公開日　平成14年６月28日(2002.6.28)

(51)Int.Cl.[7]	識別記号	ＦＩ	ﾃｰﾏｺｰﾄﾞ(参考)
Ｇ１１Ｃ　11/407		Ｃ１１Ｃ　11/34	３６２Ｓ　５Ｂ０２４
11/401			３６２Ｈ

審査請求　有　　請求項の数８　ＯＬ　（全　７　頁）

(21)出願番号　特願2000－378746(Ｐ2000－378746)

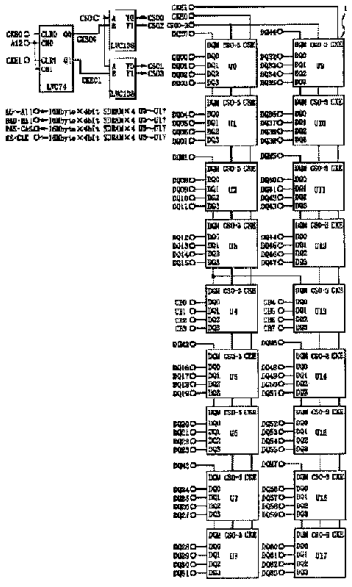
(22)出願日　平成12年12月13日(2000.12.13)

(71)出願人　000222060
東北日本電気株式会社
岩手県一関市柄貝１番地
(72)発明者　菊池　昌司
岩手県一関市柄貝１番地　東北日本電気株式会社内
(74)代理人　100082935
弁理士　京本　直樹　（外２名）

Ｆターム(参考)　5B024 AA15 BA21 CA16

(54)【発明の名称】　ＳＤＲＡＭ制御回路およびＳＤＲＡＭモジュール

(57)【要約】
【課題】年々ＳＤＲＡＭの容量増が要求されているが、これに応えるためには、より大容量のＳＤＲＡＭを開発し、モジュール化して使用することが望ましい。しかし、大容量のＳＤＲＡＭの開発には膨大な開発費、長期の開発期間などの問題がある。
【解決手段】アドレス信号とクロックイネーブル信号とチップセレクト信号から複数の新たなチップセレクト信号を出力する制御回路をＳＤＲＡＭモジュールに形成し、制御可能なＳＤＲＡＭの数を増やす。この結果、高価な３２Ｍｂｉｔ×８ｂｉｔ・ＳＤＲＡＭに代えて４個の安価な１６Ｍｂｉｔ×４ｂｉｔ・ＳＤＲＡＭを使用できる。



Micron Technology Inc. et al.
Ex. 1010, p. 250

【特許請求の範囲】

【請求項１】　ＳＤＲＡＭモジュールの制御信号を生成するＳＤＲＡＭ制御回路であって、アドレス信号とクロックイネーブル信号を入力させ新たな複数のクロックイネーブル信号を出力する第１の回路と、該第１の回路から出力されたクロックイネーブル信号とチップセレクト信号を入力させ複数の新たなチップセレクト信号を出力する第２の回路とを備えることを特徴とするＳＤＲＡＭ制御回路。

【請求項２】　ＳＤＲＡＭモジュールは複数の第１のＳＤＲＡＭ部を備え、各第１のＳＤＲＡＭ部は該第１のＳＤＲＡＭ部より小さいメモリ容量の第２のＳＤＲＡＭの複数個によって構成される請求項１記載のＳＤＲＡＭ制御回路。

【請求項３】　前記第２のＳＤＲＡＭには第２の回路が出力するチップセレクト信号がそれぞれ入力される請求項２記載のＳＤＲＡＭ制御回路。

【請求項４】　アドレス信号はバンクアドレス信号である請求項１、２、または３記載のＳＤＲＡＭ制御回路。

【請求項５】　複数の第１のＳＤＲＡＭ部を備えるＳＤＲＡＭモジュールであって、前記第１のＳＤＲＡＭ部は該ＳＤＲＡＭより小さいメモリ容量の第２のＳＤＲＡＭの複数個によって構成され、アドレス信号を入力し前記第２のＳＤＲＡＭを制御する信号を出力するＳＤＲＡＭ制御回路を備えることを特徴とするＳＤＲＡＭモジュール。

【請求項６】　ＳＤＲＡＭ制御回路の出力する信号は、前記第２のＳＤＲＡＭを制御するチップセレクト信号である請求項５記載のＳＤＲＡＭモジュール。

【請求項７】　ＳＤＲＡＭ制御回路は、アドレス信号とクロックイネーブル信号を入力させ新たな複数のクロックイネーブル信号を出力する第１の回路と、該第１の回路から出力されたクロックイネーブル信号とチップセレクト信号を入力させ複数の新たなチップセレクト信号を出力する第２の回路とを備える請求項５記載のＳＤＲＡＭモジュール。

【請求項８】　前記ＳＤＲＡＭモジュールはデュアル・インライン・メモリー・モジュールである請求項５、６または７記載のＳＤＲＡＭモジュール。

【発明の詳細な説明】

【０００１】

【発明の属する技術分野】本発明はシンクロナス・ダイナミック・ランダム・アクセス・メモリ（Ｓｙｎｃｈｒｏｎｏｕｓ　Ｄｙｎａｍｉｃ　Ｒａｎｄｏｍ　Ａｃｃｅｓｓ　Ｍｅｍｏｒｙ、以下ＳＤＲＡＭと記載する）制御回路、およびＳＤＲＡＭを使用したメモリー・モジュールに関する。

【０００２】

【従来の技術】ＳＤＲＡＭは、ＣＰＵやプロセッサモジュールなどのコンピュータあるいは通信ユニットシステムのメモリ部に使用されている。ＳＤＲＡＭは一般的にシステムに合ったメモリ容量が必要であるが、年々、アクセスデータや通信速度は増大する傾向にあり、ＳＤＲＡＭの容量増が要求されている。

【０００３】図５に、ＳＤＲＡＭを複数組み合わせ、容量の増加を図ったＳＤＲＡＭモジュールの例として、容量５１２ＭＢｙｔｅのデュアル・インライン・メモリー・モジュール（Ｄｕａｌ　Ｉｎ－Ｌｉｎｅ　Ｍｅｍｏｒｙ　Ｍｏｄｕｌｅ、以下ＤＩＭＭ）のブロック図を示す。このＤＩＭＭにおいては、１８個の３２Ｍｂｉｔ×８ｂｉｔのＳＤＲＡＭ（図中のＵ０～Ｕ１７、ただしこのうち２個はチェックビット用）に、それぞれ、１３信号のアドレス信号（Ａ０～Ａ１２）、およびチップセレクト信号（ＣＳ０）とクロックイネーブル信号（ＣＫＥ０～ＣＫＥ１）が１信号づつ供給され制御されている。さらにクロック信号（ＣＬＫ）、バンクアドレス信号（ＢＡ０～ＢＡ１）、ロウアドレスストローブ信号（ＲＡＳ）、カラムアドレスストローブ信号（ＣＡＳ）、ライトイネーブル信号（ＷＥ）が入力する。また入出力信号として、データマスク信号（ＤＱＭ０～ＤＱＭ７）、データ信号（ＤＱ００～ＤＱ６３）、チェックビット信号（ＣＢ０～ＣＢ７）が供給される。

【０００４】

【発明が解決しようとする課題】上記ＳＤＲＡＭ容量増の要求に応えるためには、より大容量のＳＤＲＡＭを開発し、モジュール化して使用することが望ましい。しかし、大容量のＳＤＲＡＭの開発には膨大な開発費、長期の開発期間などの問題がある。さらに、できあがったシステムは高価で実用的ではなくなるおそれもある。また、メモリだけでなく、これを制御するＬＳＩ等の開発も必要であり、ますます開発期間が長期化し、かつ高価なシステムになるという問題もある。

【０００５】本発明の目的は、現段階で実用化されている安価なＳＤＲＡＭを使用してメモリーモジュール（ＤＩＭＭ）を構成することにより、安価で実用的な大容量のメモリーモジュールを提供することにある。具体的には、ＤＩＭＭを構成する各ＳＤＲＡＭについて、このＳＤＲＡＭの容量より小容量のＳＤＲＡＭを複数使用して、同じ容量のＤＩＭＭを形成することを目的としている。

【０００６】

【課題を解決するための手段】本発明のＳＤＲＡＭ制御回路は、ＳＤＲＡＭモジュールを制御するＳＤＲＡＭ制御回路であって、アドレス信号とクロックイネーブル信号を入力させ新たな複数のクロックイネーブル信号を出力する第１の回路と、該第１の回路から出力されたクロックイネーブル信号とチップセレクト信号を入力させ複数の新たなチップセレクト信号を出力する第２の回路とを備える。

【０００７】また、本発明のＳＤＲＡＭモジュールは、

Micron Technology Inc. et al.
Ex. 1010, p. 251

複数の第1のSDRAM部を備え、それぞれの第1のSDRAMは該SDRAMより小さいメモリ容量の複数のSDRAMによって構成され、アドレス信号が入力され上記小さいメモリ容量のSDRAMを制御する信号を出力するSDRAM制御回路を備える。

【0008】上記発明により、現在入手可能な安価な小容量SDRAMを用いて大容量のSDRAMモジュールが構成できる。

【0009】

【発明の実施の形態】図1は、本発明のSDRAM制御回路とSDRAMモジュール（具体的にはDIMM）の例を示すブロック図である。図1の本発明のDIMMにおいては、右側の回路部分が18個の32Mbit×8bit・SDRAM（U0～U17）1からなる部分であり、左側の回路部分が右側のSDRAMを制御するSDRAM制御回路である。本SDRAMモジュールも、制御するため入力信号としては上述の従来例と同様の信号を用いる。しかし本発明では、SDRAM（U0～U17）のそれぞれが、図2で示されるような4個の16Mbit×4bit・SDRAM（U0－0～U0－3）2によって形成されている。1個の16Mbit×4bit・SDRAM2は、チップセレクト信号（CS）が1信号、アドレス信号（A0～A11）が12信号必要である。

【0010】すなわち、従来の32Mbit×8bit・DRAM1を使用したDIMMにおいて用いられたA12のアドレス信号は、カラムアドレス番号12を指定、制御する信号であり、本発明の16Mbit×4bit・SDRAM2を制御する場合には不要となる。このため本発明のSDRAMモジュールでは、従来のDIMMよりもアドレス信号が1信号余分となる。この余分なA12信号を使用して、16Mbit×4bitの4個のSDRAMを制御するチップセレクト信号（CSO0～CSO3）がSDRAM制御回路によって作られる。

【0011】まず、A12信号がクロックイネーブル信号（CKE0、CKE1）と共にD－フリップフロップ回路（LVC74）に入力され、処理（クリア・プリセット）されて、新たなクロックイネーブル信号（CKEO0、CKEO1）が出力される。このクロックイネーブル信号（CKEO0、CKEO1）は、チップセレクト信号（CS0）と共に2個のデコーダ回路（LVC138）にそれぞれ入力され、該回路にて処理（デコード）され、新しいチップセレクト信号（CSO0～CSO3）の4信号が出力される。かくして得られたチップセレクト信号（CSO0～CSO3）は、該モジュールに入力するチップセレクト信号（CS0）、アドレス信号（A12）、クロックイネーブル信号（CKE）と同期して動作する信号であり、図2の16Mbit×4bit・SDRAM（U0－0～U0－3）2のチップセレクト信号となって、これらSDRAMのリード／ライト、リフレッシュ、プリチャージなどの制御を行う。図2の他のSDRAM（U1～U17）でも同様である。なお、上述のチップセレクト信号（CSO0～CSO3）を出力するD－フリップフロップ回路とデコード回路は、高速で出力遅延の少ないゲートアレイでも構成できる。

【0012】従って、これらの信号によって72個の16Mbit×4bit・SDRAM2が制御できるようになり、図5のモジュールと同じ容量である512MByteのDIMMが構成可能となる。

【0013】次に、本SDRAMモジュールの動作について、図3（a）～（f）のタイミングチャートを用いて説明する。各信号はCLK信号に同期して制御される。

【0014】図3（a）では、D－フリップフロップ回路（LVC74）において、CLR端子に入力されるクロックイネーブル信号（CKE）がハイレベルで、CLK端子に入力されるアドレス信号（A12）がローレベルとなる場合、該回路の出力端子からは、ローレベルのクロックイネーブル信号（CKEO）が出力される。このローレベルのクロックイネーブル信号（CKEO）とチップセレクト信号（CS0）がデコーダ回路（LVC138）に入力され、チップセレクト信号（CS0）がローレベルの場合、出力端子にはCS0信号と同期したローレベルのチップセレクト信号（CSO0）が出力され、SDRAM（U0－0）はアクティブとなり、またハイレベルのチップセレクト信号（CSO2）が出力され続け、SDRAM（U0－2）はスタンバイ状態となる。

【0015】図3（b）では、D－フリップフロップ回路（LVC74）のCLR端子に入力されるクロックイネーブル信号（CKE）がハイレベルで、CLK端子に入力されるアドレス信号（A12）がハイレベルになると、それ以降CKE信号がローレベルになるまで、該回路から出力される新たなクロックイネーブル信号（CKEO）はハイレベルで出力される。このハイレベルのクロックイネーブル信号（CKEO）とチップセレクト信号（CS0）がデコーダ回路（LVC138）に入力され、チップセレクト信号（CS0）がローレベルになると、該デコード回路から出力される新たなチップセレクト信号（CSO0）はハイレベルで出力され続け、SDRAM（U0－0）はスタンバイ状態となり、チップセレクト信号（CSO2）はCS0信号と同期したローレベルの信号となり、SDRAM（U0－2）はアクティブとなる。

【0016】図3（c）では、D－フリップフロップ回路（LVC74）の入力信号と出力信号は図3（a）と同じである。このとき、デコード回路（LVC138）に入力されるチップセレクト信号（CS0）がハイレベ

Micron Technology Inc. et al.
Ex. 1010, p. 252

ルのままであるとき、出力されるＣＳＯ０信号およびＣＳＯ２信号はハイレベルのままである。

【００１７】図３（ｄ）では、Ｄ－フリップフロップ回路（ＬＶＣ７４）の入力信号と出力信号は図３（ｂ）と同じである。このとき、デコード回路（ＬＶＣ１３８）に入力されるチップセレクト信号（ＣＳ０）がハイレベルのままであるとき、出力されるＣＳＯ０信号およびＣＳＯ２信号はハイレベルのままである。

【００１８】図３（ｅ）、図３（ｆ）では、Ｄ－フリップフロップ回路（ＬＶＣ７４）に入力するクロックイネーブル信号（ＣＫＥ）がローレベルならば、アドレス信号（Ａ１２）がローレベルであり、またはハイレベルになっても、出力される新たなクロックイネーブル信号（ＣＫＥＯ）はローレベルである。さらに同様にデコード回路（ＬＶＣ１３８）に入力されるチップセレクト信号（ＣＳ０）がハイレベルのままであるとき、出力されるＣＳＯ０信号およびＣＳＯ２信号はハイレベルのままである。もう一方のデコード回路（ＬＶＣ１３８）から出力されるＣＳＯ１、ＣＳＯ３についても上述のタイムチャートと同様に動作する。

【００１９】このＤ－フリップフロップ回路（ＬＶＣ７４）とデコード回路（ＬＶＣ１３８）を備えた制御回路によって作り出されたチップセレクト信号（ＣＳＯ０～ＣＳＯ３）を使用し、通常ＤＩＭＭで使用している制御信号と組み合わせることで、ＳＤＲＡＭのリード／ライト、その他の制御を行うようにしている。

【００２０】なお、上述のＤ－フリップフロップ回路（ＬＶＣ７４）に入力させる信号として、アドレス信号（Ａ１２）に代えて、バンクアドレス信号（ＢＡ０）を使用することができる。図４にこの場合の制御回路部の構成を示す。ＢＡ０信号を用いても新しいＣＳＯ０～ＣＳＯ３信号によってＳＤＲＡＭを制御することができるが、ＢＡ信号を持たないＳＤＲＡＭ部品も使え、ＤＩＭＭにおけるＳＤＲＡＭ部品において使用できる部品の範囲が広がる。

【００２１】上記１６Ｍｂｉｔ×４ｂｉｔのＳＤＲＡＭ部品はＣＳＰ（Ｃｈｉｐ　ＳｉｚｅＰａｃｋａｇｅ）構造となっており、４段のスタック品となって１つのＳＤＲＡＭを形成している。

【００２２】

【発明の効果】このように、本発明では、チップセレクト信号（ＣＳ）、クロックイネーブル信号（ＣＫＥ）、アドレス信号（Ａ１２）を使用し、新たなチップセレクト信号（ＣＳＯ０～ＣＳＯ３）という４出力の制御信号を作り出すことによって、制御できるＳＤＲＡＭを増加させ、小容量のＳＤＲＡＭを使用し、大容量のＤＩＭＭを構成することが可能となる。また３２Ｍｂｉｔ×８ｂｉｔのＳＤＲＡＭ部品に比べ、１６Ｍｂｉｔ×４ｂｉｔのＳＤＲＡＭ部品の４個分はかなり安価であるので、大容量のＤＩＭＭを安価に構成できる。

【図面の簡単な説明】

【図１】本発明が適用されたＤＩＭＭの一例を示す構成図。

【図２】ＤＩＭＭを構成する３２Ｍｂｉｔ×８ｂｉｔ・ＳＤＲＡＭの構成図。

【図３】ＤＩＭＭの制御信号のタイムチャート。

【図４】バンクアドレス信号が入力するＤＩＭＭ制御部の構成図。

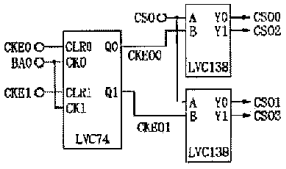
【図５】従来のＤＩＭＭの構成図。

【符号の説明】

１　３２Ｍｂｉｔ×８ｂｉｔ・ＳＤＲＡＭ

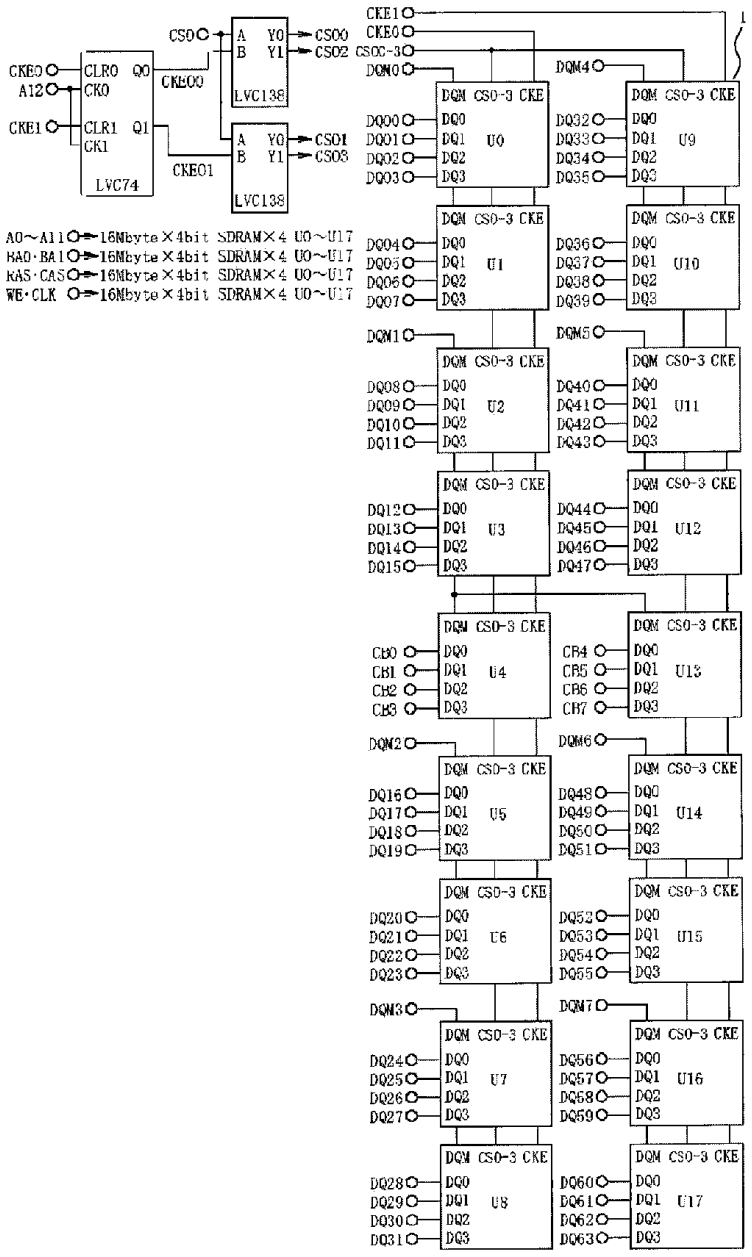
２　１６Ｍｂｉｔ×４ｂｉｔ・ＳＤＲＡＭ

【図４】

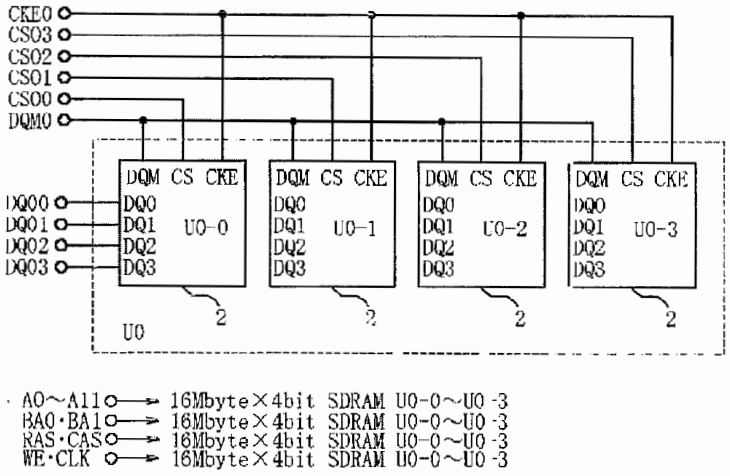


Micron Technology Inc. et al.
Ex. 1010, p. 253

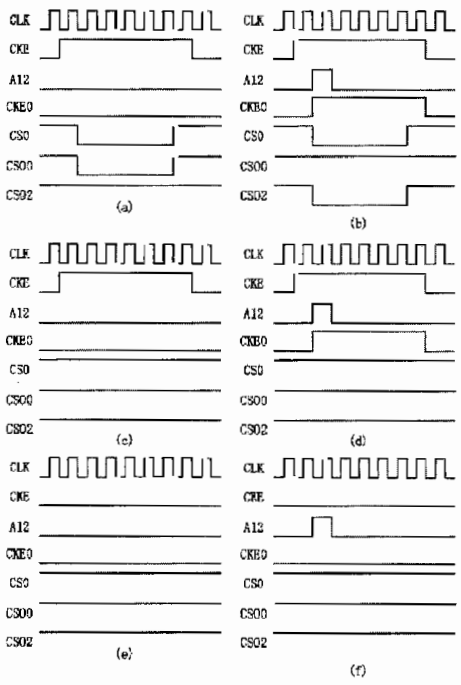
【図1】


CS0
CKE0
A12
CKE1
CLR0
CK0
CLR1
CK1
LVC74
CKE00
CKE01
A
B
Y0
Y1
LVC138
CS00
CS02
CS01
CS03
A0～A11 16Mbyte×4bit SDRAM×4 U0～U17
BA0·BA1 16Mbyte×4bit SDRAM×4 U0～U17
RAS·CAS 16Mbyte×4bit SDRAM×4 U0～U17
WE·CLK 16Mbyte×4bit SDRAM×4 U0～U17
CKE1
CKE0
CS0-3
DQM0
DQM4
DQM1
DQM5
DQM2
DQM6
DQM3
DQM7
DQM CS0-3 CKE
DQ0
DQ1
DQ2
DQ3
U0
U1
U2
U3
U4
U5
U6
U7
U8
U9
U10
U11
U12
U13
U14
U15
U16
U17
1

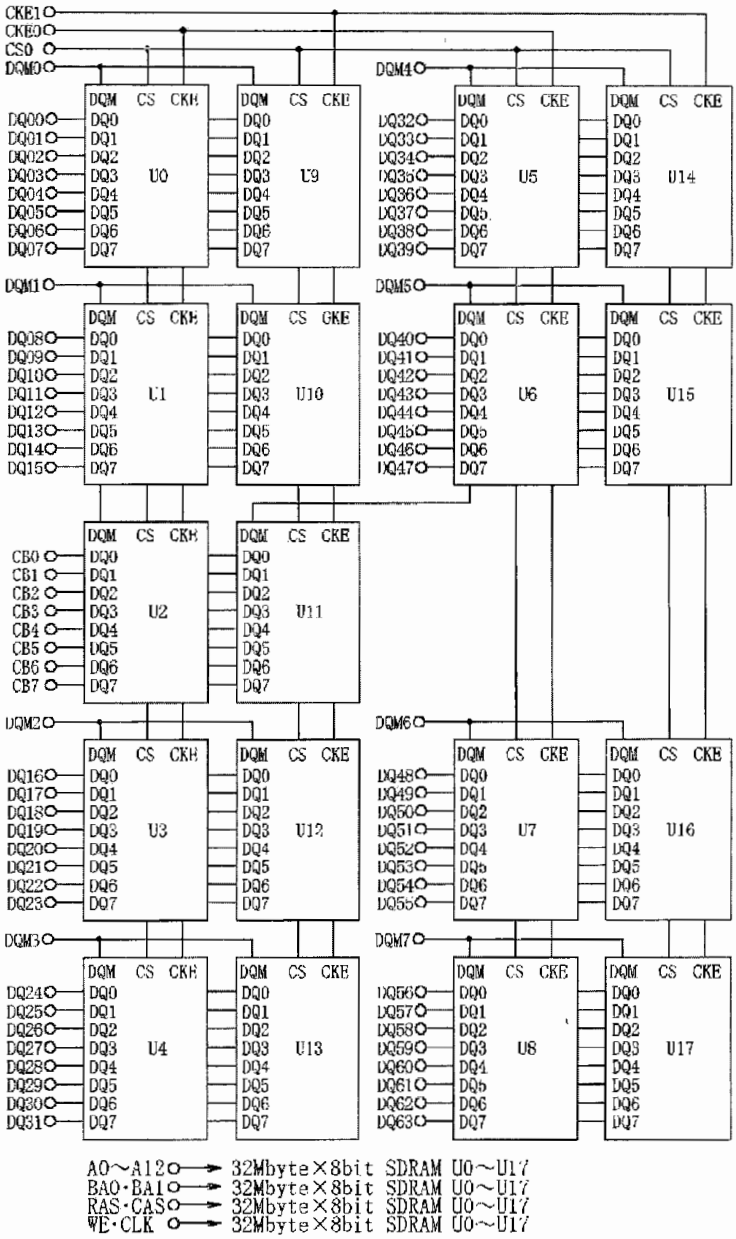
【図2】


CKE0
CS03
CS02
CS01
CS00
DQM0
DQM CS CKE
DQ0
DQ1
DQ2
DQ3
U0-0
U0-1
U0-2
U0-3
DQ00
DQ01
DQ02
DQ03
2
U0
AO～A11 → 16Mbyte×4bit SDRAM U0-0～U0-3
BAQ·BA1 → 16Mbyte×4bit SDRAM U0-0～U0-3
RAS·CAS → 16Mbyte×4bit SDRAM U0-0～U0-3
WE·CLK → 16Mbyte×4bit SDRAM U0-0～U0-3

【図3】


CLK
CKE
A12
CKE0
CS0
CS00
CS02
(a)
(b)
(c)
(d)
(e)
(f)

【図5】


CKE1
CKE0
CS0
DQM0
DQM4
DQM1
DQM5
DQM2
DQM6
DQM3
DQM7
U0
U9
U5
U14
U1
U10
U6
U15
U2
U11
U7
U16
U3
U12
U8
U17
U4
U13
AO～A12 → 32Mbyte×8bit SDRAM U0～U17
BAO·BA1 → 32Mbyte×8bit SDRAM U0～U17
RAS·CAS → 32Mbyte×8bit SDRAM U0～U17
WE·CLK → 32Mbyte×8bit SDRAM U0～U17

(19) JAPANESE PATENT OFFICE (JP)
(12) Official Gazette for Unexamined Patent Publications (A)

(11) Japanese Patent Application Publication No.: 2002-184176
(P2002-184176A)
(43) Disclosure Date: June 28, 2002

(51) Int. Cl.[7]:	ID Symbols:	FI:		Theme Code (ref.)
G11C 11/407		G11C	11/34	362S 5B024
11/401				362H

Request for Substantive Examination: Submitted
Number of Claims: 8 OL
(Total of 7 pages [in the Japanese text])

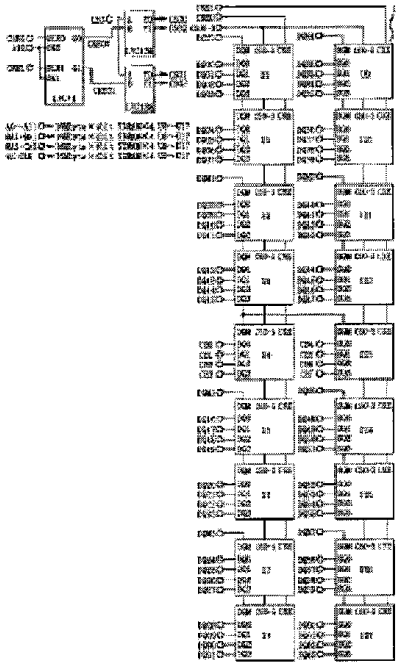
(21) Patent Application No.: 2000-378746
(P2000-378746)
(22) Filing Date: December 13, 2000
(71) Applicant: 000222060
NEC Tohoku Ltd., 1 Karakai, Ichinoseki-shi, Iwate-ken
(72) Inventor: Masashi KIKUCHI
c/o NEC Tohoku Ltd., 1 Karakai, Ichinoseki-shi, Iwate-ken
(74) Representative: 100082935 Naoki KYOMOTO, Patent Attorney (and two others)
F-Terms (reference) 5B024 AA15 BA21 CA16

(54) [Title of the Invention] SDRAM control circuit and SDRAM module

(57) [Abstract]

[Problem] Every year sees a demand for SDRAMs having greater capacity, and in order to respond to this demand it would be desirable to develop SDRAMs of greater capacity and to modularize these for use. However, the development of large-capacity SDRAMs entails problems such as very large development costs and a long development period.

[Means of Resolution] A control circuit for outputting a plurality of new chip select signals from a chip select signal, an address signal, and a clock enable signal is formed in an SDRAM module, and the number of controllable SDRAMs is increased. As a result, it is possible to use four inexpensive 16 Mbit × 4 bit SDRAMs instead of costly 32 Mbit × 8 bit SDRAMs.



[Patent Claims]

[Claim 1] An SDRAM control circuit for generating a control signal for an SDRAM module, characterized in that it comprises: a first circuit which receives as input an address signal and a clock enable signal and outputs a plurality of new clock enable signals; and a second circuit which receives as input a clock enable signal and a chip select signal output from said first circuit and outputs a plurality of new chip select signals.

[Claim 2] The SDRAM control circuit as claimed in claim 1, wherein the SDRAM module comprises a plurality of first SDRAM parts and each first SDRAM part is composed of a plurality of second SDRAM parts having a smaller memory capacity than said first SDRAM parts.

[Claim 3] The SDRAM control circuit as claimed in claim 2, wherein the chip select signals output by the second circuit are input to the second SDRAMs, respectively.

[Claim 4] The SDRAM control circuit as claimed in claim 1, 2 or 3, wherein the address signal is a bank address signal.

[Claim 5] An SDRAM module comprising a plurality of first SDRAM parts, characterized in that the first SDRAM parts are composed of a plurality of second SDRAM parts having a smaller memory capacity than said SDRAMs, and comprise an SDRAM control circuit which receives as input an address signal and outputs a signal for controlling the second SDRAMs.

[Claim 6] The SDRAM module as claimed in claim 5, in which the signal output by the SDRAM control circuit is a chip select signal for controlling the second SDRAMs.

[Claim 7] The SDRAM module as claimed in claim 5, in which the SDRAM control circuit comprises: a first

circuit which receives as input an address signal and a clock enable signal and outputs a plurality of new clock enable signals; and a second circuit which receives as input a clock enable signal and a chip select signal output from said first circuit and outputs a plurality of new chip select signals.

[Claim 8] The SDRAM module as claimed in claim 5, 6 or 7, which is a dual in-line memory module.

[Detailed Description of the Invention]

[0001]

[Technical Field of the Invention] The present invention relates to a Synchronous Dynamic Random Access Memory (referred to below as an "SDRAM") control circuit, and to a memory module employing SDRAMs.

[0002]

[Prior Art] SDRAMs are used in memories of computers such as CPUs and processor modules, or communication unit systems. SDRAMs generally require a memory capacity which is compatible with the system, but as access data and communication speed tend to increase year on year, there is a need for increased SDRAM capacity.

[0003] Figure 5 is a block diagram showing a 512 Mbyte-capacity dual in-line memory module (referred to below as a "DIMM") as one example of an SDRAM module comprising a plurality of combined SDRAMs, whereof the capacity is to be increased. This DIMM is supplied with and controlled by one each of the following signals: 13 address signals (A0 - A12), and a chip select signal (CS0) and clock enable signals (CKE0 - CKE1). A clock signal (CLK), bank address signals (BA0 - BA1), row-address strobe signal (RAS), column-address strobe signal (CAS), and a write enable signal are also input. Furthermore, data mask signals (DQM0 - DQM7), data

signals (DQ00 - DQ63) and check bit signals (CB0 - CB7) are supplied.

[0004]

[Problem to be Resolved by the Invention] In order to respond to the abovementioned need for increased SDRAM capacity, it would be desirable to develop SDRAMs of greater capacity and to modularize these for use. However, the development of large-capacity SDRAMs entails problems such as very large development costs and a long development period. In addition, finished systems are costly and risk becoming impractical. Furthermore, there is also a need for development such as LSI for controlling these systems, so there are problems in terms of development periods becoming longer and the costs of the systems increasing.

[0005] The aim of the present invention lies in providing an inexpensive, practical and large-capacity memory module by constructing a memory module (DIMM) using inexpensive SDRAMs which are currently utilized. Specifically, for each SDRAM which forms part of the DIMM, a plurality of SDRAMs having a smaller capacity than the capacity of said SDRAMs are employed, with the aim of forming a DIMM having the same capacity.

[0006]

[Means for Resolving the Problem] The SDRAM control circuit according to the present invention is an SDRAM control circuit for controlling an SDRAM module, which comprises: a first circuit which receives as input an address signal and a clock enable signal and outputs a plurality of new clock enable signals; and a second circuit which receives as input a clock enable signal and a chip select signal output from said first circuit and outputs a plurality of new chip select signals.

[0007] Furthermore, the SDRAM module according to the present invention comprises a plurality of first SDRAM

parts, where each of the first SDRAMs is composed of a plurality of SDRAMs having a smaller memory capacity than said SDRAMs, and comprises an SDRAM control circuit which receives as input an address signal and outputs a signal for controlling the SDRAMs having a small memory capacity.

[0008] According to the present invention described above, a large-capacity SDRAM can be constructed using inexpensive, small-capacity SDRAMs which are currently available.

[0009]

[Mode of Embodiment of the Invention] Figure 1 is a block diagram showing an example of the inventive SDRAM control circuit and SDRAM module (a DIMM, to be specific). In the DIMM according to the present invention shown in Figure 1, the circuit portion on the right-hand side comprises eighteen 32 Mbit × 8 bit SDRAMs (U0 - U17) 1, and the circuit portion on the left-hand side is an SDRAM control circuit for controlling the SDRAMs on the right-hand side. The SDRAM module also employs the same signals as described in the conventional example above as input signals for control. However, according to the present invention, each of the SDRAMs (U0 - U17) is formed by four 16 Mbit × 4 bit SDRAMs (U0-0 to U0-3) 2 shown in Figure 2. One 16 Mbit × 4 bit SDRAM 2 requires one chip select signal (CS) and 12 address signals (A0 - A11).

[0010] That is to say, the A12 address signal used in a conventional DIMM employing the 32 Mbit × bit DRAMs 1 is a control signal for designating the column address number 12, and is not required to control the 16 Mbit × 4 bit SDRAMs 2 according to the present invention. In the SDRAM module according to the present invention, there is therefore one extra address signal compared with a conventional DIMM. Using this extra A12 signal, chip select signals (CS00 - CS03) for controlling the

four 16 Mbit x 4 bit SDRAMs are produced by means of the SDRAM control circuit.

[0011] The A12 signal is first of all input to a D-flip flop circuit (LVC74) together with clock enable signals (CKE0, CKE1), processing (clear/preset) is carried out, and new clock enable signals (CKEO0, CKEO1) are output. The clock enable signals (CKEO0, CKEO1) are input to two decoder circuits (LVC138) together with a chip select signal (CS0), respectively, this circuit carries out processing (decoding), and four new chip select signals (CSO0 - CSO3) are output. The chip select signals (CSO0 - CSO3) obtained in this way are signals which operate synchronously with the chip select signal (CS0), address signal (A12) and clock enable signal (CKE) which are input to said module, and constitute chip select signals of the 16 Mbit x 4 bit SDRAMs (U0-0 to U0-3) 2 shown in Figure 2, controlling read/write, refresh, precharge and the like for these SDRAMs. The same also applies to the other SDRAMs (U1 - U17) shown in Figure 2. It should be noted that the D-flip flop circuit and decoder circuit which output the abovementioned chip select signals (CSO0 - CSO3) may also comprise gate arrays which provide high speed and little delay in output.

[0012] The seventy-two 16 Mbit x 4 bit SDRAMs 2 can therefore be controlled by means of these signals, and it is possible to construct a 512 Mbyte-DIMM having the same capacity as the module shown in Figure 5.

[0013] The operation of this SDRAM module will be described next with the aid of the timing charts shown in Figures 3(a) - (f). The signals are controlled synchronously with the CLK signal.

[0014] In Figure 3(a), when the clock enable signal (CKE) input to the CLR terminal is high level and the address signal (A12) input to the CLK terminal is low

level, the output terminal of the D-flip flop circuit (LVC74) outputs the low-level clock enable signal (CKEO). This low-level clock enable signal (CKEO) and the chip select signal (CS0) are input to the decoder circuit (LVC138), and when the chip select signal (CS0) is low level, the low-level chip select signal (CSO0) synchronous with the CS0 signal is output at the output terminal, the SDRAM (U0-0) becomes active, and a high-level chip select signal (CSO2) continues to be output, and the SDRAM (U0-2) is in a standby state.

[0015] In Figure 3(b), when the clock enable signal (CKE) input to the CLR terminal of the D-flip flop circuit (LVC74) is high level and the address signal (A12) input to the CLK terminal is high level, the new clock enable signal (CKEO) output from said circuit is subsequently output at high level until the CKE signal becomes low level. When the high-level clock enable signal (CKEO) and the chip select signal (CS0) are input to the decoder circuit (LVC138) and the chip select signal (CS0) becomes low level, the new chip select signal (CSO0) output from said decoder circuit continues to be output at high level, and the SDRAM (U0-0) is in a standby state, while the chip select signal (CSO2) is a low-level signal synchronous with the CS0 signal, and the SDRAM (U0-2) becomes active.

[0016] In Figure 3(c), the input signal and output signal of the D-flip flop circuit (LVC74) are the same as in Figure 3(a). In this case, when the chip select signal (CS0) input to the decoder circuit (LVC138) remains high level, the CSO0 signal and CSO2 signal which are output remain high level.

[0017] In Figure 3(d), the input signal and output signal of the D-flip flop circuit (LVC74) are the same as in Figure 3(b). In this case, when the chip select signal (CS0) input to the decoder circuit (LVC138)

remains high level, the CSO0 signal and CSO2 signal which are output remain high level.

[0018] In Figure 3(e) and Figure 3(f), if the clock enable signal (CKE) input to the D-flip flop circuit (LVC74) is low level, the address signal (A12) is low level, and even if it is high level, the new clock enable signal (CKE0) which is output is low level. Likewise, if the chip select signal (CS0) input to the decoder circuit (LVC138) remains at high level, the CSO0 signal and CSO2 signal which are output remain at high level. The system operates in the same way as in the time charts described above for CSO1, CSO3 which are output from the other decoder circuit (LVC138).

[0019] The SDRAMs are read/write controlled and controlled in other ways by using the chip select signals (CSO0 - CSO3) produced by means of the control circuit provided with the D-flip flop circuit (LVC74) and decoder circuit (LVC138) and combining these signals with control signals conventionally used by a DIMM.

[0020] It should be noted that a bank address signal (BA0) may be used instead of the address signal (A12) as a signal which is input to the abovementioned D-flip flop circuit (LVC74). Figure 4 shows the construction of a control circuit in this case. The SDRAMs can still be controlled by means of the new CSO0 - CSO3 signals if the BA0 signal is used, and SDRAM components which do not have a BA signal may be used, so the range of components which can be used for the SDRAM components in the DIMM can be broadened.

[0021] The abovementioned 16 Mbit × 4 bit SDRAM components have a CSP (chip size package) structure arranged in a four-level stack to form one SDRAM.

[0022]

[Advantage of the Invention] In this way, according to the present invention, four output control signals, namely the new chip select signals (CSO0 - CSO3) are produced using a chip select signal (CS), clock enable signal (CKE), and address signal (A12), thereby increasing the SDRAMs which can be controlled, and small-capacity SDRAMs are used, and it is possible to construct a large-capacity DIMM. Furthermore, four 16 Mbit × 4 bit SDRAM components are fairly inexpensive compared with 32 Mbit × 8 bit SDRAM components, and therefore a large-capacity DIMM can be constructed inexpensively.

[Brief Description of the Figures]

[Figure 1] is a structural diagram showing one example of a DIMM in which the present invention has been applied;

[Figure 2] is a structural diagram of 32 Mbit × 8 bit SDRAMs which form a DIMM;

[Figure 3] is a time chart of DIMM control signals;

[Figure 4] is a structural diagram of a DIMM control unit to which a bank address signal is input; and

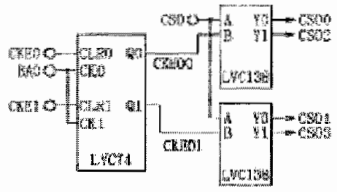
[Figure 5] is a structural diagram of a conventional DIMM.

[Key to Symbols]

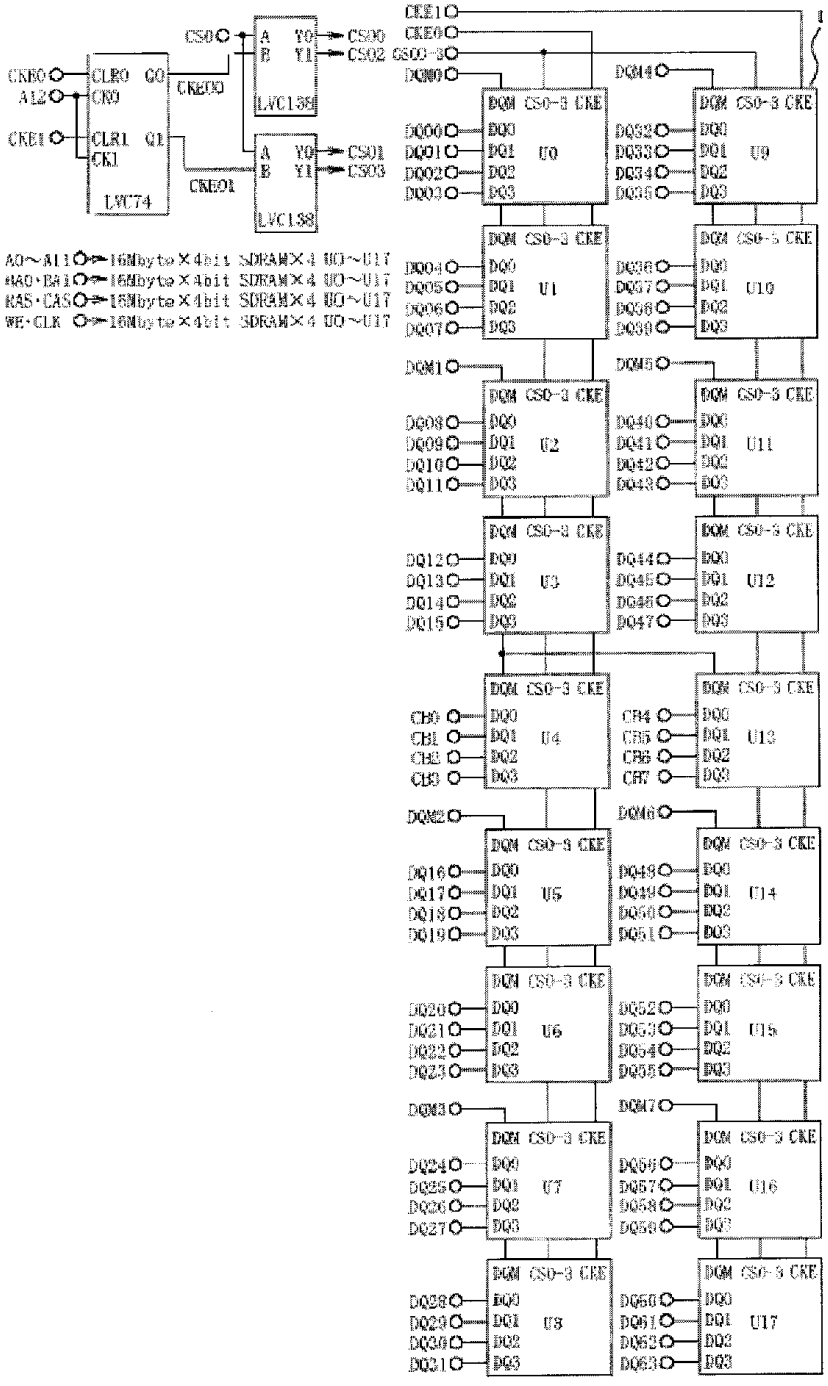
1 32 Mbit × 8 bit SDRAM

2 16 Mbit × 4 bit SDRAM

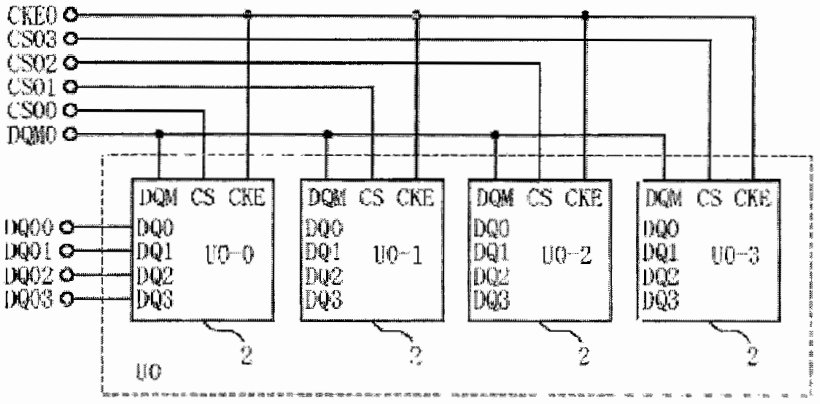
[Figure 4]



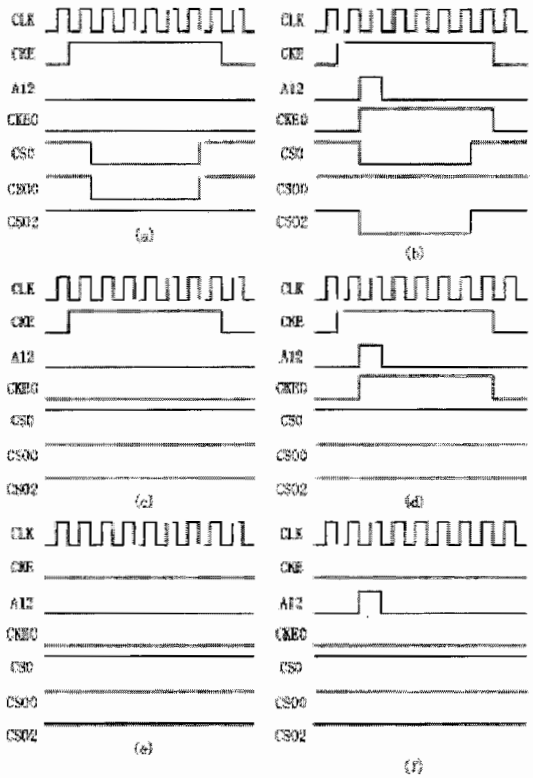
[Figure 1]


CS0
CKE0
A12
CKE1
CLR0
CK0
CLR1
CK1
LVC74
CKE00
CKE01
A
B
Y0
Y1
LVC138
CS00
CS02
CS01
CS03
A0~A11 16Mbyte×4bit SDRAM×4 U0~U17
BA0·BA1 16Mbyte×4bit SDRAM×4 U0~U17
RAS·CAS 16Mbyte×4bit SDRAM×4 U0~U17
WE·CLK 16Mbyte×4bit SDRAM×4 U0~U17
CKE1
CKE0
CS0-3
DQM0
DQM4
DQM1
DQM5
DQM2
DQM6
DQM3
DQM7
DQM CS0-3 CKE
DQ0
DQ1
DQ2
DQ3
U0
U1
U2
U3
U4
U5
U6
U7
U8
U9
U10
U11
U12
U13
U14
U15
U16
U17

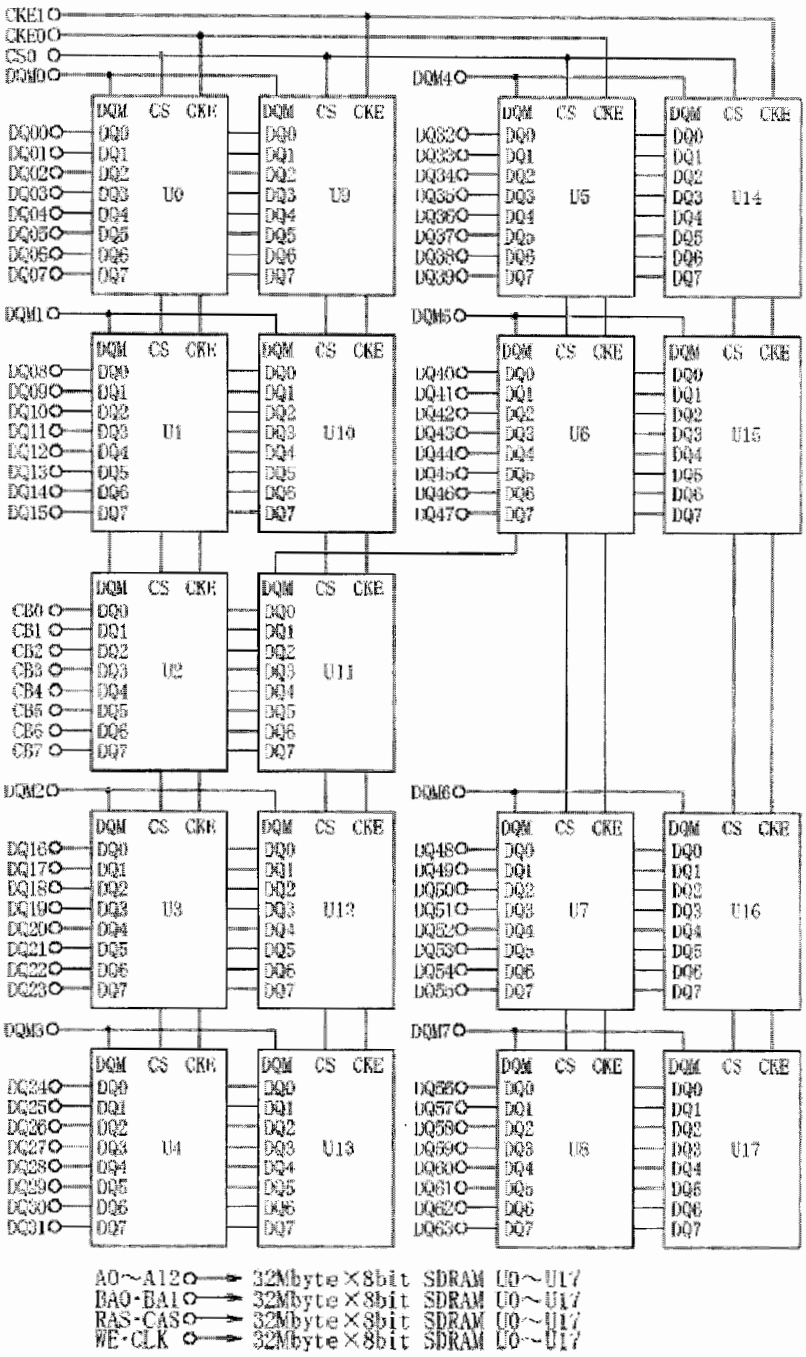
[Figure 2]



[Figure 3]



[Figure 5]


CKE1
CKE0
CS0
DQM0
DQM4
DQM1
DQM5
DQM2
DQM6
DQM3
DQM7
U0
U9
U5
U14
U1
U10
U6
U15
U2
U11
U7
U16
U3
U12
U8
U17
U4
U13
A0~A12 → 32Mbyte×8bit SDRAM U0~U17
BA0·BA1 → 32Mbyte×8bit SDRAM U0~U17
RAS·CAS → 32Mbyte×8bit SDRAM U0~U17
WE·CLK → 32Mbyte×8bit SDRAM U0~U17

Espacenet

Bibliographic data: JP2003007963 (A) — 2003-01-10

SEMICONDUCTOR MEMORY DEVICE AND ITS MANUFACTURING METHOD

Inventor(s): HONMA KAZUKI; ITO SHIN; WADA MASASHI ± (HONMA KAZUKI, ; ITO SHIN, ; WADA MASASHI)

Applicant(s): HITACHI LTD; AKITA DENSHI SYSTEMS KK ± (HITACHI LTD, ; AKITA DENSHI SYSTEMS:KK)

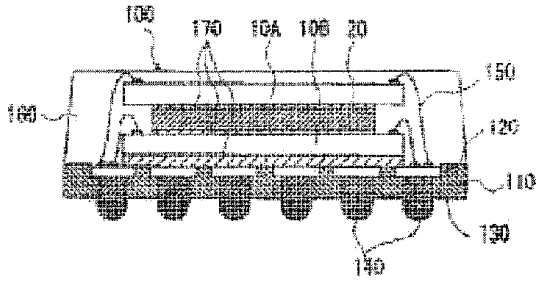
Classification: - international: ***G11C16/02; G11C29/00; G11C29/04; H01L25/065; H01L25/07; H01L25/18;*** (IPC1-7): G11C16/02; G11C29/00; H01L25/065; H01L25/07; H01L25/18

- European:

Application number: JP20010186288 20010620

Priority number (s): JP20010186288 20010620

Abstract of JP2003007963 (A)



Micron Technology Inc. et al.
Ex. 1010, p. 270

PROBLEM TO BE SOLVED: To overcome
such a problem of a conventional multichip
stack type semiconductor memory that since a plurality of chip enable terminals are provided, the number of terminals is increased as compared with a one chip semiconductor device when the semiconductor memory is regarded as one semiconductor device and thereby the package size is increased and compatibility with conventional products is lost. SOLUTION: The semiconductor memory device comprises a first terminal part (CAD) being connected with any one of first or second power supply voltage terminal, and chip selection control circuits (17, 20) for controlling internal chip selection signal based on the state of the first terminal part and the state of any one or more than one of a plurality of address input terminal parts (e.g. the input pad of a most significant bit A22). The chip selection control circuit controls the internal chip selection signal (CEB) to valid or invalid state based on any one or more than one bit of address signals inputted under a set state of the first terminal part and a state where an external chip selection signal is validated.

Last updated: 24.09.2012 Worldwide Database 5.7.42; 83p

Micron Technology Inc. et al.
Ex. 1010, p. 271

(19)日本国特許庁（ＪＰ）　(12)公開特許公報（Ａ）　(11)特許出願公開番号
特開2003－7963
（Ｐ2003－7963Ａ）
(43)公開日　平成15年１月10日(2003.1.10)

(51)Int.Cl.[7]		識別記号	ＦＩ		テーマコード（参考）
Ｈ０１Ｌ	25/065		Ｃ１１Ｃ　29/00	６０３Ｚ	５Ｂ０２５
Ｇ１１Ｃ	16/02		Ｈ０１Ｌ　25/08	Ｚ	５Ｌ１０６
	29/00	６０３	Ｃ１１Ｃ　17/00	６０１Ｚ	
Ｈ０１Ｌ	25/07				
	25/18				

審査請求　未請求　請求項の数10　ＯＬ　（全 14 頁）

(21)出願番号　特願2001－186288(Ｐ2001－186288)

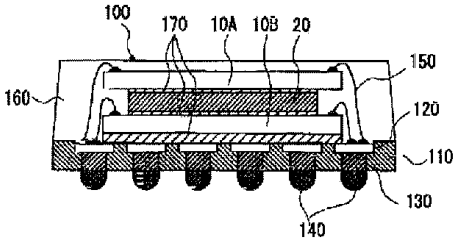
(22)出願日　平成13年６月20日(2001.6.20)

(71)出願人　000005108
株式会社日立製作所
東京都千代田区神田駿河台四丁目６番地
(71)出願人　000100997
株式会社アキタ電子システムズ
秋田県河辺郡雄和町相川字後野85番地
(72)発明者　本間　和樹
東京都小平市上水本町五丁目20番１号　株式会社日立製作所半導体グループ内
(74)代理人　100085811
弁理士　大日方　富雄

最終頁に続く

(54)【発明の名称】　半導体記憶装置および製造方法

(57)【要約】
【課題】　従来のマルチチップスタック型の半導体記憶装置は、チップイネーブル用端子が複数個あるため、これを１つの半導体装置として見たときにはワンチップの半導体装置に比べて端子数が増えることとなり、パッケージが大型化するとともに、従来の製品との互換性がなくなるという課題があった。
【解決手段】　第１の電源電圧端子または第２の電源電圧端子のいずれかに接続される第１端子部（ＣＡＤ）と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部（例えば最上位ビットＡ２２の入力パッド）の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路（１７，２０）とを備え、該チップ選択制御回路は、上記第１端子部の設定状態と、外部からのチップ選択信号が有効にされた状態で入力されたアドレス信号のいずれか１または２以上のビットとに基づいて内部のチップ選択信号（ＣＥＢ）を有効状態または無効状態に制御するように構成した。



Micron Technology Inc. et al.
Ex. 1010, p. 272

【特許請求の範囲】

【請求項１】 第１の電源電圧端子または第２の電源電圧端子のいずれかに接続される第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、該チップ選択制御回路は、上記第１端子部の設定状態と、外部からのチップ選択信号が有効にされた状態で入力されたアドレス信号のいずれか１または２以上のビットとに基づいて内部のチップ選択信号を有効状態または無効状態に制御するように構成されていることを特徴とする半導体記憶装置。

【請求項２】 複数の半導体メモリチップが１つのパッケージに収納され各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続されてなる半導体記憶装置であって、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が各々異なる状態に設定されている場合には、外部からのチップ選択信号が有効にされた状態で入力されたアドレス信号が同一であってもいずれか１のチップにおいて内部のチップ選択信号を有効にするように構成されていることを特徴とする半導体記憶装置。

【請求項３】 複数の半導体メモリチップが１つのパッケージに収納され、データ入出力端子部以外の各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続されてなる半導体記憶装置であって、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が同一の状態に設定されかつアドレス入力端子部のうち所定の１または２以上の端子部が同一の電位に固定されている場合には、外部からのチップ選択信号が有効にされたことに応じてそれぞれ内部のチップ選択信号を有効にするように構成されていることを特徴とする半導体記憶装置。

【請求項４】 上記アドレス信号のいずれか１または２以上のビットは、アドレスの最上位ビットまたは最上位側２ビット以上であることを特徴とする請求項２または３に記載の半導体記憶装置。

【請求項５】 複数の半導体メモリチップが１つのパッケージに収納され、該パッケージには各チップに対応してデータ入出力用の端子が設けられ各チップのデータ入出力端子部はパッケージの対応する上記端子に別々に接続されてなるとともに、パッケージには所定の端子が設けられている半導体記憶装置であって、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が同一の状態に設定されかつアドレス入力端子部のうち所定の１または２以上の端子部が同一の電位に固定されている場合には、外部からのチップ選択信号が有効にされたことに応じてそれぞれ内部のチップ選択信号を有効にするように構成され、上記複数の半導体メモリチップのいずれかの上記第１端子部は上記パッケージの上記所定の端子に接続され、他の半導体メモリチップの上記第１端子部は上記パッケージのいずれかの電源電圧端子に接続されていることを特徴とする半導体記憶装置。

【請求項６】 上記複数の半導体メモリチップは、それぞれ同一方向に端子部が現われるように積層され、各チップ間はこれらのチップの上記端子部よりも内側の部位に介在された接着剤層によって互いに接合されていることを特徴とする請求項５に記載の半導体記憶装置。

【請求項７】 複数の半導体メモリチップが１つのパッケージに収納され各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続され、上記複数のメモリチップ毎に各々異なる状態に設定可能な第１端子部を有してなる半導体記憶装置の製造方法であって、前工程終了後にウェハ状態で上記複数の半導体メモリチップのテストを行ない、該テストにより良品と判定された半導体メモリチップに対してウェハ状態でトリミング処理を行なった後、ウェハを切断して各チップに分割し、上記各半導体メモリチップの上記第１端子部に対する設定および各チップの端子部とパッケージの端子との接続を行なってからパッケージに封止することを特徴とする半導体記憶装置の製造方法。

【請求項８】 上記第１端子部に対する設定は、該第１端子部をパッケージに設けられているいずれか一の電源電圧端子にボンディングワイヤで接続する処理であることを特徴とする請求項７に記載の半導体記憶装置の製造方法。

【請求項９】 上記複数の半導体メモリチップはそれぞれ同一方向に端子部が現われるように積層され、上記ボンディングワイヤで接続する処理は各チップを接着剤で接合する処理の後に各チップ毎に行なわれることを特徴とする請求項８に記載の半導体記憶装置の製造方法。

【請求項１０】 上記ウェハ状態でのトリミング処理と共に上記テストにより検出された不良メモリセルを置換する救済処理を行なうことを特徴とする請求項７に記載の半導体記憶装置の製造方法。

Micron Technology Inc. et al.
Ex. 1010, p. 273

【発明の詳細な説明】

【0001】

【発明の属する技術分野】本発明は、半導体記憶装置さらには複数の半導体メモリチップが1つのパッケージに収納されたマルチチップスタック型の半導体記憶装置に適用して有効な技術に関し、例えばフラッシュメモリのような不揮発性半導体記憶装置の大容量化に利用して有効な技術に関する。

【0002】

【従来の技術】従来、半導体集積回路の分野では、1つのパッケージ内に複数の半導体チップを積層収納して実装密度の向上を図るようにした技術があり、半導体メモリにおいてもかかる技術により大容量化を図るようにしたものが実用化されつつある。このようなマルチチップスタック型の半導体記憶装置においては、アドレス信号や制御信号など大部分の入力信号に関しては、パッケージの端子を複数のチップ間で共有することで端子数を減らすことができる。また、データを入出力するI／O端子についてもチップ間で共有することによりパッケージの端子数を減らすことができる。

【0003】しかしながら、I／O端子を共有させた場合には、各チップを別々に選択して所望のチップに対してデータのリード／ライトを行なえるようにするために、一般にチップイネーブル信号／CEと呼ばれるチップ選択信号に関しては、チップごとにパッケージの別の端子を割り当てそれぞれのチップのCE信号用のパッドとパッケージのCE信号用端子とをボンディングワイヤ等で接続するように構成せざるを得なかった。

【0004】

【発明が解決しようとする課題】そのため、従来のマルチチップスタック型の半導体記憶装置は、チップイネーブル用端子が複数個あるため、これを1つの半導体装置として見たときにはワンチップの半導体装置に比べて端子数が増えることとなり、パッケージが大型化するとともに、従来の製品との互換性がなくなる。つまり、チップイネーブル用端子を複数個有しているため、メモリを使用したシステムを新たに開発したり、従来のメモリチップに代えてマルチチップスタック型の半導体記憶装置を用いたメモリモジュールを構成したりする場合に、すでにある基板の設計データの利用が図れず、新たに基板を設計し直す必要が生じる。

【0005】また、1つのパッケージに複数のCE信号用端子があるとアドレス信号をデコードして複数のチップイネーブル信号のいずれか1つを有効レベルにするためのデコーダ回路を外付け回路として設けたり、メモリをアクセスするためアドレスを出力する装置の側で複数のチップイネーブル信号を生成して出力できるように構成したりしなくてはならず、ユーザの設計負担が大きくなるといった不具合があった。

【0006】そこで、チップイネーブル端子を2つのチップで共有し、アドレス信号（最上位ビット）によってチップを選択できるようにしたマルチチップスタック型の半導体記憶装置の発明が提案されている（特願平11－207701号）。ただし、この先願発明においては、マルチチップスタック型の半導体記憶装置の概略構成およびデバイス構造は開示されているものの、アドレス最上位ビットによっていずれかのチップを選択できるようにする具体的な回路構成や具体的なボンディングの仕方など製造技術については開示されていない。

【0007】この発明の目的は、チップイネーブル端子を複数のチップで共有し、アドレス信号によってチップを選択できるようにしたマルチチップスタック型の半導体記憶装置を提供することにある。

【0008】この発明の他の目的は、チップイネーブル端子を複数のチップで共有し、アドレス信号によってチップを選択できるようにしたマルチチップスタック型の半導体記憶装置の有効な製造方法を提供することにある。

【0009】この発明の前記ならびにそのほかの目的と新規な特徴については、本明細書の記述および添付図面から明らかになるであろう。

【0010】

【課題を解決するための手段】本願において開示される発明のうち代表的なものの概要を説明すれば、下記のとおりである。すなわち、本発明に係る半導体記憶装置は、第1の電源電圧端子または第2の電源電圧端子のいずれかに接続される第1端子部と、該第1端子部の状態と複数のアドレス入力端子部のうちいずれか1または2以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、該チップ選択制御回路は、上記第1端子部の設定状態と、外部からのチップ選択信号が有効にされた状態で入力されたアドレス信号のいずれか1または2以上のビットとに基づいて内部のチップ選択信号を有効状態または無効状態に制御するように構成したものである。

【0011】上記した手段によれば、複数のメモリチップを組み合わせたときに、第1端子部の設定状態を変えることでアドレス信号に応じてチップが選択される記憶装置として動作したり、第1端子部の設定状態とアドレス端子部の設定状態の組合せにより入出力ビット数が異なる記憶装置としても動作したりすることができるようになる。

【0012】また、複数の半導体メモリチップが1つのパッケージに収納され各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続されてなる半導体記憶装置において、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第1端子部と、該第1端子部の状態と複数のアドレス入力端子部のうちいずれか1または2以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備

Micron Technology Inc. et al.
Ex. 1010, p. 274

え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が各々異なる状態に設定されている場合には、外部からのチップ選択信号が有効にされた状態で入力されたアドレス信号が同一であってもいずれか１のチップにおいて内部のチップ選択信号を有効にするように構成した。

【００１３】これにより、複数のメモリチップを組み合わせたときに、第１端子部の設定状態を変えることでアドレス信号に応じてチップが選択される記憶装置として動作することができ、チップが複数個あってもチップ選択信号は共通化することができ、装置にとっての外部端子数を減らすことができるとともに、通常の１個のメモリチップと同一に扱うことができるため、メモリを使用したシステムを開発する場合に既存の設計基板を利用することができるようになる。

【００１４】さらに、複数の半導体メモリチップが１つのパッケージに収納され、データ入出力端子部以外の各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続されてなる半導体記憶装置において、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が同一の状態に設定されかつアドレス入力端子部のうち所定の１または２以上の端子部が同一の電位に固定されている場合には、外部からのチップ選択信号が有効にされたことに応じてそれぞれ内部のチップ選択信号を有効にするように構成した。

【００１５】これにより、複数のメモリチップを組み合わせたときに、各チップの第１端子部の設定状態と所定のアドレス入力端子部の設定状態を一致させることにより、チップが複数個あっても入出力データのビット数が２倍である通常の１個のメモリチップと同一に扱うことができるため、メモリを使用したシステムを開発する場合に既存の設計基板を利用することができるようになる。

【００１６】また、望ましくは、上記アドレス信号のいずれか１または２以上のビットは、アドレスの最上位ビットまたは最上位側から２ビット以上とする。これにより、従来のメモリチップのマット構成やアドレスデコーダの構成を変えずにアドレス信号に応じてチップが選択されるマルチチップスタック型の半導体記憶装置を実現することができる。

【００１７】さらに、本願の他の発明は、複数の半導体メモリチップが１つのパッケージに収納され、該パッケージには各チップに対応してデータ入出力用の端子が設けられ各チップのデータ入出力端子部はパッケージの対応する上記端子に別々に接続されてなるとともに、パッケージには所定の端子が設けられている半導体記憶装置において、上記複数の半導体メモリチップは、各々異なる状態に設定可能な第１端子部と、該第１端子部の状態と複数のアドレス入力端子部のうちいずれか１または２以上の端子部の状態に基づいて内部のチップ選択信号を制御するチップ選択制御回路とを備え、上記複数の半導体メモリチップ内の各チップ選択制御回路は、上記第１端子部が同一の状態に設定されかつアドレス入力端子部のうち所定の１または２以上の端子部が同一の電位に固定されている場合には、外部からのチップ選択信号が有効にされたことに応じてそれぞれ内部のチップ選択信号を有効にするように構成され、上記複数の半導体メモリチップのいずれかの上記第１端子部は上記パッケージの上記所定の端子に接続され、他の半導体メモリチップの上記第１端子部は上記パッケージのいずれかの電源電圧端子に接続されるようにした。これにより、ユーザは上記所定の端子に印加される電圧を適宜設定することによって、当該半導体記憶装置の入出力データのビット数を選択することができる。

【００１８】また、望ましくは、上記複数の半導体メモリチップは、それぞれ同一方向に端子部が現われるように積層され、各チップ間はこれらのチップの上記端子部よりも内側の部位に介在された接着剤層によって互いに接合されるようにする。これにより、各チップの端子部とパッケージ側の対応する端子との電気的接続をボンディングワイヤで行なう場合に、それぞれのチップの端子部が同一方向に現われるように積層してもワイヤボンディング処理を行なうことができる。

【００１９】本発明に係る半導体記憶装置の製造方法は、複数の半導体メモリチップが１つのパッケージに収納され各チップの端子部は対応するもの同士がパッケージの対応する端子に共通に接続されてなる半導体記憶装置の製造方法において、前工程終了後にウェハ状態で上記複数の半導体メモリチップのテストを行ない、該テストにより良品と判定された半導体メモリチップに対してウェハ状態でトリミング処理を行なった後、ウェハを切断して各チップに分割し、上記各半導体メモリチップの上記第１端子部に対する設定および各チップの端子部とパッケージの端子との接続を行なってからパッケージに封止するようにしたものである。これにより、状態を設定する必要があるメモリチップを使用しながらその端子設定を、各チップの端子部とパッケージの端子とを接続する工程で行なうことができるため、従来の半導体メモリのプロセスを全く変更することなくマルチチップスタック型の半導体記憶装置を製造することができる。上記トリミング処理とともに不良ビットの救済処理を行うようにしても良い。

【００２０】また、上記第１端子部に対する設定は、該第１端子部をパッケージに設けられているいずれか一の電源電圧端子にボンディングワイヤで接続する処理とす

Micron Technology Inc. et al.
Ex. 1010, p. 275

る。ワイヤボンディング技術はきわめて信頼性の高い技術であるので、これを用いて第１端子部に対する設定を行なうことにより、設定ミスを防止できるとともに、従来の半導体メモリのプロセスを全く変更することなくマルチチップスタック型の半導体記憶装置を製造することができる。

【００２１】さらに、上記複数の半導体メモリチップはそれぞれ同一方向に端子部が現われるように積層され、上記ボンディングワイヤで接続する処理は各チップを接着剤で接合する処理の後に各チップ毎に行なわれるようにする。これにより、それぞれのチップの端子部が同一方向に現われるように積層しても既存の装置を用いてワイヤボンディング処理を行なうことができる。

【００２２】

【発明の実施の形態】以下、本発明の好適な実施例を図面に基づいて説明する。図１は、本発明を適用したマルチチップスタック型半導体記憶装置のパッケージ構造を、また図２は回路のブロック構成図を示す。

【００２３】図１に示されているように、この実施例のマルチチップスタック型半導体記憶装置は、２個のメモリチップ１０Ａ，１０Ｂが積層されてパッケージ１００内に封入されている。この実施例でパッケージ１００内に封入されるメモリチップ１０Ａ，１０Ｂは、例えば電気的にデータを書き込むことが可能でありかつデータを所定の単位で電気的に一括消去可能なフラッシュメモリのような不揮発性メモリである。

【００２４】特に制限されるものでないが、メモリチップ１０Ａと１０Ｂは、これらのチップよりも小さな絶縁体からなるスペーサ２０を挟んで互いに電気的に絶縁された状態で積層されている。スペーサ２０を設けているのは、２つのチップはサイズが同じであるため、直接重ねると下のチップのボンディングパッド部分が隠れてしまいワイヤボンディングができなくなるためである。メモリチップ１０Ａおよび１０Ｂとスペーサ２０とはペレット接着剤１７０により結合される。メモリチップ１０Ａ，１０Ｂは、１つのメモリセルに１ビットのデータを記憶可能な２値メモリであっても良いし、１つのメモリセルに２ビット以上のデータを記憶可能ないわゆる多値メモリであっても良い。

【００２５】１１０はセラミックなどからなる絶縁基板で、この絶縁基板１１０の一面（内面）にはチップの外延に沿って複数のボンディング用電極端子１２０が所定の間隔で配置され、また各ボンディング用電極端子１２０に対応して絶縁基板１１０には当該基板を貫通するように埋設された導電性プラグ１３０が設けられている。これらの導電性プラグ１３０は、一方の端部（図では上端）が上記いずれか一つのパッド１２０の下面に接触され、また他端すなわち基板１１０の他面（外面）に露出した部位（図では下端）には、半田ボール１４０がそれぞれ溶着されている。メモリチップ１０Ｂは絶縁基板１１０上にペレット接着剤１７０により結合される。

【００２６】また、上記基板側のボンディング用電極端子１２０にボンディングワイヤ１５０の一端が結合され、該ボンディングワイヤ１５０の他端はメモリチップ１０Ａ，１０Ｂの対応するボンディングパッドに結合されて、基板側のボンディング用電極端子１２０とメモリチップ１０Ａ，１０Ｂの対応するボンディングパッドとが電気的に接続されている。そして、絶縁基板１１０上にメモリチップ１０Ａと１０Ｂがボンディングワイヤ１５０で接続された状態で樹脂１６０によりモールドされてパッケージ１００が構成されている。

【００２７】図２に示されているように、各メモリチップ１０Ａ，１０Ｂは、複数のメモリセルがマトリックス状に配置されたメモリアレイ１１と、入力されたＸ系アドレスをデコードしてメモリアレイ１１内の一本のワード線を選択レベルにするＸデコーダ１２と、入力されたＹ系アドレスをデコードしてメモリアレイ１１内の対応するビット線を選択するＹデコーダ１３と、入力されたＸ系アドレスおよびＹ系アドレスを取り込んで上記Ｘデコーダ１２およびＹデコーダ１３に供給するアドレスバッファ１４と、選択されたビット線に読み出された信号を増幅したり書込みデータに応じてビット線の電位を与えるセンスアンプ＆書込み回路１５と、読出しデータを出力したり書込みデータをチップ外部から取り込む入出力回路１６と、チップ外部から入力されるチップイネーブル信号／ＣＥを取り込んで内部制御信号を生成するチップイネーブル制御回路１７と、ライトイネーブル信号／ＷＥやリセット信号ＲＥＳ、アウトイネーブル信号／ＯＥなどの外部制御信号を取り込む入力バッファ１８と、取り込まれた制御信号に応じてチップ内部の制御信号を生成して内部を制御するコントロール回路１９などにより構成されている。

【００２８】チップイネーブル信号／ＣＥは当該チップが選択状態にされることを示す信号、ライトイネーブル信号／ＷＥは当該チップが書込み状態にされることを示す信号、リセット信号ＲＥＳはチップ内部をリセット状態にする信号、アウトイネーブル信号／ＯＥは読出しデータ信号が出力される状態にあることを示す信号である。これらの制御信号およびアドレス信号が入力されるパッドのうち、以下に述べる２つのパッド（ＣＡＤ，Ａ２２）以外のパッドに関しては、メモリチップ１０Ａ，１０Ｂの対応するパッド同士が、それぞれボンディングワイヤでパッケージに設けられている対応する電極端子に共通に接続されている。また、各チップのデータ入力用パッドＩ／Ｏ０～Ｉ／Ｏ１５もパッケージに設けられている対応する電極端子に共通に接続されている。

【００２９】この実施例のメモリチップ１０Ａ，１０Ｂには、通常のメモリには設けられていないチップアドレスデータ設定用パッドＣＡＤがそれぞれ設けられている。また、各メモリチップ１０Ａ，１０Ｂには、チップ

Micron Technology Inc. et al.
Ex. 1010, p. 276

外部から供給される例えば２３ビットのようなアドレス信号Ａ０～Ａ２２のうち最上位ビットＡ２２と、上記チップアドレスデータ設定用パッドＣＡＤに印加される電位とに応じて、上記チップイネーブル制御回路１７を有効にするアドレスビット比較回路２０がそれぞれ設けられている。そして、メモリチップ１０Ａのチップアドレスデータ設定用パッドＣＡＤとメモリチップ１０Ｂのチップアドレスデータ設定用パッドＣＡＤは、それぞれ一方がパッケージの第１の電源電圧端子Ｖｃｃに、また他方がパッケージの第２の電源電圧端子Ｖｓｓに接続される。

【００３０】図３には、チップイネーブル制御回路１７とアドレスビット比較回路２０の具体的な構成例が示されている。チップイネーブル制御回路１７は、チップイネーブル信号／ＣＥとリセット信号ＲＥＳを入力とするＡＮＤゲートＧ１と、該ＡＮＤゲートＧ１の出力を反転してアドレスビット比較回路２０に供給するインバータＧ２と、ＡＮＤゲートＧ１の出力信号とアドレスビット比較回路２０からの信号とに基づいて内部チップイネーブル信号ＣＥＢを生成するＮＡＮＤゲートＧ３とからなる。

【００３１】一方、アドレスビット比較回路２０は、チップイネーブル制御回路１７の上記インバータＧ２の出力信号とアドレス最上位ビットＡ２２とを入力とするＮＯＲゲートＧ１１と、上記チップアドレスデータ設定用パッドＣＡＤに印加される電位に応じた信号を生成するインバータＧ１２と、該インバータＧ１２の出力信号と上記ＮＯＲゲートＧ１１の出力信号を入力とするイクスクルーシブＯＲゲートＧ１３と、該イクスクルーシブＯＲゲートＧ１３の出力を反転してチップイネーブル制御回路１７のＮＡＮＤゲートＧ３に供給するインバータＧ２４とからなる。

【００３２】以下、図３の回路の動作を説明する。リセット信号ＲＥＳはハイレベルが有効レベルの信号であり、リセット信号ＲＥＳがロウレベルのときにチップイネーブル制御回路１７はチップイネーブル信号／ＣＥを有効にする。すなわち、リセット信号ＲＥＳがロウレベルのときにＡＮＤゲートＧ１の出力がチップイネーブル信号／ＣＥに応じて変化する。チップイネーブル信号／ＣＥは、ロウレベルが有効レベルつまりチップ選択状態を示す信号であり、リセット信号ＲＥＳとチップイネーブル信号／ＣＥが共にロウレベルのときに、ＡＮＤゲートＧ１の出力がハイレベルとなる。

【００３３】そして、ＡＮＤゲートＧ１の出力はインバータＧ２により反転されてアドレスビット比較回路２０のＮＯＲゲートＧ１１に供給されるため、リセット信号ＲＥＳとチップイネーブル信号／ＣＥが共にロウレベルのときに、ＮＯＲゲートＧ１１の他方の入力であるアドレス最上位ビットＡ２２に応じてＮＯＲゲートＧ１１が変化する。リセット信号ＲＥＳまたはチップイネーブル信号／ＣＥのいずれか一方がハイレベルのときは、ＮＯＲゲートＧ１１の出力はロウレベルに固定される。また、リセット信号ＲＥＳまたはチップイネーブル信号／ＣＥのいずれか一方がハイレベルのときは、チップイネーブル制御回路１７の出力段のＮＡＮＤゲートＧ３の出力がハイレベルに固定され、内部チップイネーブル信号ＣＥＢはチップ非選択状態を示すこととなる。

【００３４】リセット信号ＲＥＳとチップイネーブル信号／ＣＥが共にロウレベルのとき、アドレスビット比較回路２０のＮＯＲゲートＧ１１はアドレス最上位ビットＡ２２を入力信号とするインバータとして動作し、アドレス最上位ビットＡ２２がロウレベルのときにはＮＯＲゲートＧ１１の出力がハイレベルに、またアドレス最上位ビットＡ２２がハイレベルのときにはＮＯＲゲートＧ１１の出力がロウレベルになる。そして、このＮＯＲゲートＧ１１の出力信号を入力とするイクスクルーシブＯＲゲートＧ１３は排他的論理和回路であるので、チップアドレスデータ設定用パッドＣＡＤに印加される電位に応じた信号を出力するバッファＧ１２の出力論理レベルとＮＯＲゲートＧ１１の出力論理レベルが同じである時はイクスクルーシブＯＲゲートＧ１３の出力はロウレベルに、またＧ１１とＧ１２の出力論理レベルが異なる時はイクスクルーシブＯＲゲートＧ１３の出力はハイレベルとなる。

【００３５】そして、イクスクルーシブＯＲゲートＧ１３の出力はインバータＧ１４により反転してチップイネーブル制御回路１７の出力段のＮＡＮＤゲートＧ３に供給されるため、Ｇ１１とＧ１２の出力論理レベルが異なる時はＮＡＮＤゲートＧ３の出力である内部チップイネーブル信号ＣＥＢはチップ非選択状態を示すハイレベルにされる。一方、Ｇ１１とＧ１２の出力論理レベルが同じである時は、ＮＡＮＤゲートＧ３がインバータとして動作し、その出力である内部チップイネーブル信号ＣＥＢは、ＮＡＮＤゲートＧ３の他方の入力信号（このときハイレベル）に応じてチップ選択状態を示すロウレベルにされる。

【００３６】以上のように、図３の回路はアドレス最上位ビットＡ２２とチップアドレスデータ設定用パッドＣＡＤに印加される電位に応じて内部チップイネーブル信号ＣＥＢをチップ選択状態にしたり、非選択状態にしたりする。次の表１に、図３の実施例回路を搭載したメモリチップを２個使用した半導体記憶装置におけるアドレス最上位ビットＡ２２とチップアドレスデータ設定用パッドＣＡＤに印加される電位との組合せと、チップの動作状態との関係を整理して示す。

【００３７】

【表１】

Micron Technology Inc. et al.
Ex. 1010, p. 277

	CAD	A22	内部状態	メモリ構成
(1)	Vss	H	非選択	×16
		L	選択	
(2)	Vcc	H	選択	
		L	非選択	
(3)	Vcc	Vcc	同時選択	×32

【００３８】表1に示されているように、チップアドレスデータ設定用パッドＣＡＤに印加される電位がＶｓｓであるチップ（１）は、アドレス最上位ビットＡ２２に応じて、Ａ２２がハイレベルのときは／ＣＥのいかんにかかわらず内部チップイネーブル信号ＣＥＢがハイレベルとなってチップは非選択状態となり、Ａ２２がロウレベルのときは／ＣＥがロウレベルとされると内部チップイネーブル信号ＣＥＢがロウレベルとなってチップは選択状態となる。

【００３９】一方、チップアドレスデータ設定用パッドＣＡＤに印加される電位がＶｃｃであるチップ（２）は、アドレス最上位ビットＡ２２に応じて、Ａ２２がハイレベルのときは／ＣＥがロウレベルとされると内部チップイネーブル信号ＣＥＢがロウレベルとなってチップは選択状態となり、Ａ２２がロウレベルのときは／ＣＥのいかんにかかわらず内部チップイネーブル信号ＣＥＢがハイレベルとなってチップは非選択状態となる。

【００４０】さらに、表1の（３）のように、アドレス最上位ビットＡ２２とチップアドレスデータ設定用パッドＣＡＤに印加される電位が共にＶｃｃ（または共にＶｓｓ）のときは、チップは外部からのチップイネーブル信号／ＣＥに応じて内部チップイネーブル信号ＣＥＢが同じように変化される。そのため、２チップともチップアドレスデータ設定用パッドＣＡＤに印加される電位と、アドレス最上位ビットＡ２２が入力されるべき端子をＶｃｃに固定することにより、２つのチップを同時に選択状態にさせることができる。

【００４１】従って、１つのパッケージに封入される２つのチップのデータ入出力パッドに対応してそれぞれ別のデータ入出力用電極端子をパッケージに設けるとともに、チップアドレスデータ設定用パッドＣＡＤとアドレス最上位ビットＡ２２が入力されるべきパッドをＶｃｃに固定することにより、２倍のデータ幅を有するメモリ（×２ｎ）として動作させることができる。例えば、２つのチップがそれぞれ６４Ｍビットの記憶容量を有し１６ビットのデータを並列入出力するものであれば、１２８Ｍビットの記憶容量を有し３２ビットのデータを並列入出力する半導体記憶装置として構成することができる。

【００４２】ただし、この場合、チップのアドレスＡ２２入力用パッドはパッケージ内でＶｃｃ端子に接続されるので、パッケージにはアドレス最上位ビットＡ２２に対応した電極端子は不要である。図２に示すような構成を有するチップは、マルチチップとしてではなく１つのチップを１つのパッケージに収納した半導体記憶装置として構成した場合には、アドレス最上位ビットＡ２２とチップアドレスデータ設定用パッドＣＡＤに印加される電位を共にパッケージ内部でＶｃｃに固定することにより、６４Ｍビットの記憶容量を有し１６ビットのデータを並列入出力する通常の半導体記憶装置として構成することもできる。

【００４３】一方、表１（１），（２）の場合は、２つのチップがそれぞれ６４Ｍビットの記憶容量を有し１６ビットのデータを並列入出力するものであれば、１２８Ｍビットの記憶容量を有し１６ビットのデータを並列入出力する半導体記憶装置として構成され、各チップはアドレス最上位ビットＡ２２によって選択的にアクセスされる。

【００４４】なお、図３の実施例においては、チップイネーブル制御回路１７からアドレスビット比較回路２０のＮＯＲゲートＧ１１にチップイネーブル信号／ＣＥに応じた制御信号を、インバータＧ２により送ってアドレスＡ２２の取り込みを制御しているが、この信号は必ずしも必要なものではなく、省略することも可能である。この制御信号によりＮＯＲゲートＧ１１を制御することで、チップ非選択状態においてＮＯＲゲートＧ１１の入力であるアドレスＡ２２が変化することで内部回路に貫通電流が流れて消費電量が増加するのを防止する効果がある。

【００４５】従って、アドレスＡ２２以外のビットＡ０～Ａ２１をチップ内部に取り込むバッファに関しても、同様にインバータＧ２によりチップ非選択状態では動作しないように制御するのが望ましい。また、図３の実施例においては、チップイネーブル制御回路１７にリセット信号ＲＥＳに応じてチップイネーブル信号／ＣＥの入力を制御するＡＮＤゲートＧ１が設けられているが、このＡＮＤゲートＧ１も省略することが可能である。

【００４６】図４には、チップイネーブル制御回路１７とアドレスビット比較回路２０の他の構成例が示されている。図４は、４つのチップを１組として半導体記憶装置を構成する場合に適した回路構成例である。この実施例においては、２つのチップアドレスデータ設定用パッドＣＡＤ１，ＣＡＤ２と、アドレス信号Ａ２３の入力端子が設けられている。そして、これに対応して、アドレスビット比較回路２０には、論理ゲートＧ１１～Ｇ１３と同様な論理ゲートＧ２１～Ｇ２３が設けられているとともに、インバータＧ１４の代わりにＡＮＤゲートＧ２４が設けられている。そして、ＮＯＲゲートＧ２１にはアドレス信号Ａ２３とチップイネーブル制御回路１７のインバータＧ２の出力が入力され、イクスクルーシブＯＲゲートＧ２３にはＮＯＲゲートＧ２１の出力とチップアドレスデータ設定用パッドＣＡＤ２に印加される電位に応じた信号を生成するインバータＧ２２の出力が入力されている。

Micron Technology Inc. et al.
Ex. 1010, p. 278

【００４７】さらに、イクスクルーシブＯＲゲートＧ１３とＧ２３の出力がＡＮＤゲートＧ２４に入力されている。チップイネーブル制御回路１７は、図３の実施例と同じである。この実施例においては、４つのチップにおいて、それぞれチップアドレスデータ設定用パッドＣＡＤ１，ＣＡＤ２に印加される電圧Ｖｃｃ，Ｖｓｓの組合せが異なるように設定（ボンディングワイヤの接続）がなされることにより、アドレス信号Ａ２２，Ａ２３に応じていずれか１つのチップの内部イネーブル信号ＣＥＢがロウレベルにされて、４つのうち１つのチップのみが選択状態とされる。ただし、図３の実施例と同様に、すべてのチップのチップアドレスデータ設定用パッドＣＡＤ１，ＣＡＤ２とアドレス信号Ａ２２，Ａ２３の入力パッドに印加される電圧をＶｃｃとすることにより、４倍のデータ幅を有するメモリとして動作させることができる。

【００４８】表２に、図４の実施例回路を搭載したメモリップを４個使用した半導体記憶装置におけるアドレス最上位ビットＡ２２，Ａ２３とチップアドレスデータ設定用パッドＣＡＤ１，ＣＡＤ２に印加される電位との組合せと、チップの動作状態との関係を整理して示す。

【００４９】

【表２】

CAD1	CAD2	A22	A23	内部状態	メモリ構成
Vss	Vss	H	H	非選択	×16
		L	H	非選択	
		H	L	非選択	
		L	L	選択	
Vcc	Vss	H	H	非選択	
		L	H	非選択	
		H	L	選択	
		L	L	非選択	
Vss	Vcc	H	H	非選択	
		L	H	選択	
		H	L	非選択	
		L	L	非選択	
Vcc	Vcc	H	H	選択	
		L	H	非選択	
		H	L	非選択	
		L	L	非選択	
Vcc	Vcc	Vcc	Vcc	同時選択	×64

【００５０】図５には、図３に示されているアドレスビット比較回路２０の他の構成例が示されている。図３においては、アドレス信号の最上位ビットＡ２２とチップアドレスデータ設定用パッドＣＡＤの状態とを比較するイクスクルーシブＯＲゲートＧ１３を用いているが、図５においては、イクスクルーシブＯＲゲートＧ１３の代わりにセレクタＳＥＬを設け、このセレクタＳＥＬをチップアドレスデータ設定用パッドＣＡＤに印加される電位に応じた信号を生成するバッファＧ１２の出力で制御するように構成されている。

【００５１】また、図５の実施例では、図３のＮＯＲゲートＧ１１の代わりにＯＲゲートＧ１１’を用いている。ＯＲゲートＧ１１’の入力はＮＯＲゲートＧ１１の入力と同じく、アドレスＡ２２とチップイネーブル制御回路１７からの信号であり、ＯＲゲートＧ１１’はチップイネーブル制御回路１７からの信号に応じてアドレスＡ２２を通過させたり遮断したりする伝送ゲートとして作用する。

【００５２】上記セレクタＳＥＬのデータ入力端子には、ＯＲゲートＧ１１’の出力と、該ＯＲゲートＧ１１’の出力をインバータＧ１６で反転した信号が入力されている。このようにして、パッドＣＡＤの電位に応じてＯＲゲートＧ１１’を通過したアドレス最上位ビットＡ２２またはその反転信号をセレクタＳＥＬで選択してチップイネーブル制御回路１７に供給することにより、同一のアドレスＡ２２が入力された場合でも、パッドＣＡＤの電位に応じて一方は内部チップイネーブル信号ＣＥＢが有効レベルとされ、他方は内部チップイネーブル信号ＣＥＢが無効レベルとされる。その結果、アドレスＡ２２によってチップを選択することができるようになる。

【００５３】なお、図５の実施例においては、パッドＣＡＤに印加される電位に応じた信号を生成するバッファＧ１２の出力でセレクタＳＥＬを制御してアドレスＡ２２またはその反転信号のいずれか一方を選択するようにしているが、逆にアドレスＡ２２すなわちＯＲゲートＧ１１’の出力でセレクタＳＥＬを制御してバッファＧ１２の出力またはその反転信号のいずれか一方を選択して出力するように構成しても良い。

【００５４】また、本実施例のアドレスビット比較回路２０は、４つのチップを１組として半導体記憶装置を構成する場合に適した図４に示されている実施例に対しても適用することができる。具体的には、図４のイクスクルーシブＯＲゲートＧ１３，Ｇ２３の代わりにセレクタを用いてアドレスビット比較回路２０を構成すれば良い。

Micron Technology Inc. et al.
Ex. 1010, p. 279

【００５５】図６は、本発明の他の実施例を示す。この実施例は、図１および図２に示されている２つのフラッシュメモリチップを積層した半導体記憶装置の変形例であって、フラッシュメモリチップ１０Ａ，１０Ｂの構成は、図２と同一である。図２と異なるのは、パッケージにビット構成切換え端子ＢＣが設けられメモリチップ１０Ａと１０Ｂのうち一方のチップ（例えば１０Ａ）のチップアドレス設定用端子ＣＡＤがこのビット構成切換え端子ＢＣに接続されている点と、パッケージのデータ入出端子が１６本（Ｉ／Ｏ０～Ｉ／Ｏ１５）でなく３２本（Ｉ／Ｏ０～Ｉ／Ｏ３１）とされ一方のチップ（例えば１０Ａ）のデータ入出力パッドはパッケージのデータ入出端子Ｉ／Ｏ０～Ｉ／Ｏ１５に接続され、他方のチップ（例えば１０Ｂ）のデータ入出力パッドはパッケージのデータ入出端子Ｉ／Ｏ１６～Ｉ／Ｏ３１に接続されている点である。

【００５６】この実施例の半導体記憶装置は、ユーザが１６ビット並列入出力のメモリとしても３２ビット並列入出力のメモリとしても利用できる点に特徴がある。具体的には、メモリチップ１０Ａのチップアドレス設定用端子ＣＡＤが接続されているビット構成切換え端子ＢＣを電源電圧Ｖｓｓに接続すると、図２の実施例の半導体記憶装置と同様にアドレス最上位ビットＡ２２によってチップが選択され１６ビットのデータを並列に入出力するメモリとして動作する。従って、この場合には、パッケージのデータ入出端子Ｉ／Ｏ０～Ｉ／Ｏ１５とＩ／Ｏ１６～Ｉ／Ｏ３１は、同一の１６ビットのバスに共通に接続される。

【００５７】一方、メモリチップ１０Ａのチップアドレス設定用端子ＣＡＤが接続されているビット構成切換え端子ＢＣを電源電圧Ｖｃｃに接続すると、２つのチップはチップイネーブル信号ＣＥによって同時に選択状態とされ、３２ビットのデータを並列に入出力するメモリとして動作する。従って、この場合には、パッケージのデータ入出端子Ｉ／Ｏ０～Ｉ／Ｏ１５とＩ／Ｏ１６～Ｉ／Ｏ３１は、３２ビットのバスの別の信号線に接続される。

【００５８】図７および図８は、本発明のさらに他の実施例を示す。この実施例は、図１および図２に示されている２つのフラッシュメモリチップを積層した半導体記憶装置の発展形態であって、積層された２つのフラッシュメモリチップ１０Ａ，１０Ｂの上にさらにスタティックＲＡＭチップ１０Ｃを積層したものである。フラッシュメモリチップ１０Ａ，１０Ｂの構成は、図２と同一である。この実施例では、スタティックＲＡＭチップ１０Ｃのサイズがフラッシュメモリチップ１０Ａ，１０Ｂのサイズよりも小さくフラッシュメモリチップ１０Ａの上にスタティックＲＡＭチップ１０Ｃを重ねてもフラッシュメモリチップ１０Ａのボンディングパッドが隠れないため、フラッシュメモリチップ１０Ａの保護膜の上にペレット接着剤１７０によりスタティックＲＡＭチップ１０Ｃが接合されている。

【００５９】これらのチップ１０Ａ，１０Ｂ，１０Ｃに設けられているパッドのうちライトイネーブル信号／ＷＥが入力されるパッドおよびアドレス信号Ａ０～Ａ１８の入力パッドは、図７に示すように、ボンディングワイヤ１５０を介してパッケージ１００に設けられている対応する電極端子１２０に共通に接続されている。

【００６０】また、図８に示すように、各チップのデータ入力用パッドＩ／Ｏ０～Ｉ／Ｏ１５もパッケージに設けられている対応する端子に共通に接続されている。チップイネーブル信号／ＣＥに関しては、フラッシュメモリチップ１０Ａ，１０Ｂのパッドは前記実施例と同様に、パッケージに設けられている対応する端子ＰＣＥ-Ｆに共通に接続されているが、ＳＲＡＭチップ１０Ｃのチップイネーブル信号／ＣＥが入力されるパッドは、パッケージに設けられている専用の端子ＰＣＥ-Ｓに接続されている。

【００６１】スタティックＲＡＭチップ１０Ｃは、メモリアレイ１１、Ｘデコーダ１２、Ｙデコーダ１３、アドレスバッファ１４、センスアンプ＆書込み回路１５、データ入出力回路１６、ライトイネーブル信号／ＷＥなど制御信号の入力バッファ１８、コントロール回路１９などを備えた汎用のＳＲＡＭと同じ構成にされており、前記実施例のフラッシュメモリチップ１０Ａ，１０Ｂに設けられているようなチップアドレスデータ設定用パッドＣＡＤやチップイネーブル制御回路１７、アドレスビット比較回路２０は設けられていない。

【００６２】本実施例の半導体記憶装置は１つのパッケージ内に２つのフラッシュメモリと１つのＳＲＡＭが積層収納されているため、別々のパッケージに収納されているものを使用する場合に比べてシステムの実装密度をさらに高めることができ、装置の小型化を図ることができる。

【００６３】次に、上記実施例のマルチチップスタック型の半導体記憶装置の開発から製造までの手順の一例を、図９を用いて説明する。まず、前記実施例で説明したチップアドレスデータ設定用パッドＣＡＤやチップイネーブル制御回路１７、アドレスビット比較回路２０を設けたメモリチップの論理設計とシミュレーションによる論理動作の確認を行なう（ステップＳ２００）。次に、上記論理設計データに基づいてレイアウトツールと呼ばれる設計支援プログラムを用いてレイアウト設計を行なう（ステップＳ２０１）。それから、このレイアウト設計データに基づいてプロセスに使用するマスクを作成する（ステップＳ２０２）。

【００６４】次に、上記マスクを用いて半導体ウェハ上に複数のメモリチップを形成する前工程を行なう（ステップＳ２０３）。そして、テスタによりウェハ状態で各チップのパッドにプローブを当てて試験するウェハテス

Micron Technology Inc. et al.
Ex. 1010, p. 280

トを行なう（ステップS204）。このウェハテストで不良と判定されたチップはマーキングが付されて各チップに切断されてから不良品として廃棄される。一方、ステップS204で良品と判定されたチップは、冗長回路による欠陥ビットの救済情報や内部電源回路の電圧値のトリミング情報をチップ上に設けられているヒューズや不揮発性記憶素子に設定する処理が行なわれる（ステップS205）。

【0065】ウェハ上のすべてのチップについてプローブ検査および救済等の処理が終了すると、ウェハは各チップに切断される（ステップS206）。それから、チップを絶縁基板に搭載して、チップのパッドと対応する基板側の電極端子とを接続するワイヤボンディングと樹脂でモールドするパッケージング処理を行なう（ステップS207）。このボンディング&パッケージング処理では、先ず下側のメモリチップ10Bを絶縁基板110上に接着剤ペレット170で接合してワイヤボンディングを行なった後、上側のメモリチップ10Aを接着剤ペレット170で下側のメモリチップ10B上に接合してこのチップ10Aに対してワイヤボンディングを行なうようにするのが良い。

【0066】そして、この実施例においては、上記ワイヤボンディングの際に、各チップのチップアドレスデータ設定用パッドCADをパッケージに設けられているいずれか一方の電源電圧端子に接続する処理を行なう。これにより、チップアドレスデータ設定用パッドCADに対する設定処理を何ら新たな工程を追加することなく行なうことができる。また、上側のメモリチップ10Aを接着剤ペレット170で下側のメモリチップ10B上に接合してから2つのチップ10A，10Bに対してまとめてワイヤボンディングを行なおうとすると、上のチップが邪魔になって下のチップに対するワイヤボンディングが困難になるが、下のチップ10Bに対するワイヤボンディングを行なった後、上側のメモリチップ10Aを接合してワイヤボンディングを行なうことにより、既存装置でワイヤボンディング処理を行なうことができる。

【0067】なお、チップアドレスデータ設定用パッドCADをいずれかの電源電圧端子に接続する方法としては、ワイヤボンディングの他に、予めチップアドレスデータ設定用パッドCADを電源電圧VccまたはVssに接続するヒューズを設けておいて、パッケージング処理の前にヒューズを切断する方法が考えられるが、ヒューズの場合にはチップアドレスデータ設定用パッドCADを電源電圧Vccに接続したチップとVssに接続したチップの2つが生じるので、各チップを個々に管理する必要があるが、前記実施例のようにパッケージング処理のワイヤボンディングで設定するようにすれば、チップを別々に管理する必要がないので、その分コストの低減が可能となる。

【0068】上記パッケージング処理の後は、各デバイスをテスタでテストする出荷前テストが実行される（ステップS208）。このテストで不良と判定されたチップはマーキングが付されて後の選別工程で不良品として廃棄される一方、ステップS208で良品と判定されたデバイスは製品として出荷される。

【0069】図10には、第1の実施例（図1および図2参照）の半導体記憶装置の応用例としてのメモリカードの構成を示す。図10において、CDは絶縁材料で形成されたカード本体、FLM0～FLM3は該カード本体に内蔵された不揮発性記憶装置、CNTはデータバスDB，アドレスバスABおよび制御バスCBを介してこれらの不揮発性記憶装置FLM0～FLM3に接続されたコントローラである。不揮発性記憶装置FLM0～FLM3は各々図2に示すような構成を有する2つのチップを内蔵し、または1のチップのみを内蔵していても3以上のチップを内蔵していても良い。コントローラCNTはシリアル入出力端子SIOを介して外部のマイクロプロセッサなどの制御装置と接続されて、制御装置からの指令に応じて、カード内部の不揮発性記憶装置FLM0～FLM3に対するデータの書込みや読出しを行なう。

【0070】コントローラCNTは各不揮発性記憶装置FLM0～FLM3に対してそれぞれ専用の選択信号／CE0，／CE1，／CE2，／CE3を生成していずれか1つを選択状態にする。各不揮発性記憶装置FLM0～FLM3は、図2の実施例で説明したように、一方メモリチップ（10A）のチップアドレスデータ設定用パッドCADはパッケージの第1の電源電圧Vssに接続され、他方のメモリチップ（10B）のチップアドレスデータ設定用パッドCADはパッケージの第2の電源電圧端子Vccに接続されており、アドレス最上位ビット（例えばA22）に応じていずれか一方のチップが選択状態にされる。

【0071】この応用例からも分かるように、前記実施例の半導体記憶装置を使用すると、従来の半導体メモリと全く同等に扱えしかも記憶容量が2倍のメモリカードを構成することができる。なお、この応用例のメモリカードにおいては、図1に示されている絶縁基板110を、各記憶装置FLM0～FLM3とコントローラCNTに共通の基板として構成して、全体を樹脂等でモールドして1つのパッケージに封入しても良い。つまり、図1に示されているような構造を有する半導体記憶装置として製造されたものを使用してメモリカードを構成するとパッケージが2重構造となるが、各記憶装置FLM0～FLM3とコントローラCNTを共通の基板上に搭載すれば、1つのパッケージに封入させることができる。

【0072】以上本発明者によってなされた発明を実施例に基づき具体的に説明したが、本発明は上記実施例に限定されるものではなく、その要旨を逸脱しない範囲で種々変更可能であることはいうまでもない。例えば図1

Micron Technology Inc. et al.
Ex. 1010, p. 281

の実施例においては、複数のメモリチップが積層状態でパッケージに封入されているものについて説明したが、複数のメモリチップが１枚の絶縁基板上に横に並べた状態でパッケージに封入されている場合にも適用することができる。また、チップの選択のためのアドレスビットは、アドレス最上位ビットに限定されるものでなく、他のビット例えば最下位ビット等であってもよい。

【００７３】以上の説明では主として本発明者によってなされた発明をその背景となった利用分野であるフラッシュメモリを主体とするマルチチップスタック型の半導体記憶装置について説明したが、フラッシュメモリ以外の例えばＥＥＰＲＯＭを複数個内蔵した半導体記憶装置などにも利用することができる。

【００７４】

【発明の効果】本願において開示される発明のうち代表的なものによって得られる効果を簡単に説明すれば下記のとおりである。すなわち、チップイネーブル端子を複数のチップで共有し、アドレス信号によってチップを選択できるようになり、これによって半導体記憶装置としての外部端子数を減らすことができるとともに、通常の１個の半導体メモリと同等に扱うことができるため、メモリを使用したシステムを開発する場合に既存の設計基板を利用することができるようなマルチチップスタック型の半導体記憶装置を実現できる。さらに、そのような半導体記憶装置を、従来の半導体メモリの製造プロセスを変更したり製造装置を改造したりすることなく製造することができる。

【図面の簡単な説明】

【図１】本発明を適用したマルチチップスタック型半導体記憶装置の一実施例のパッケージ構造を示す断面図である。

【図２】本発明を適用したマルチチップスタック型半導体記憶装置の一実施例を示すブロック構成図である。

【図３】チップイネーブル信号制御回路およびアドレスビット比較回路の具体例を示す論理構成図である。

【図４】チップイネーブル信号制御回路およびアドレスビット比較回路の他の具体例を示す論理構成図である。

【図５】チップイネーブル信号制御回路およびアドレスビット比較回路のさらに他の具体例を示す論理構成図である。

【図６】本発明を適用したマルチチップスタック型半導体記憶装置の他の実施例を示すブロック構成図である。

【図７】本発明を適用したマルチチップスタック型半導体記憶装置のさらに他の実施例のパッケージ構造を示す断面図である。

【図８】図７の実施例の半導体記憶装置のうちＳＲＡＭの構成例を示すブロック構成図である。

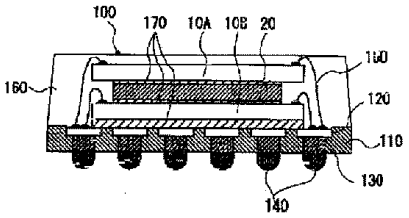
【図９】本発明に係るマルチチップスタック型半導体記憶装置の製造方法の一例を工程順に示すフローチャートである。

【図１０】本発明に係る半導体記憶装置の応用例としてのメモリカードの構成例を示すブロック図である。

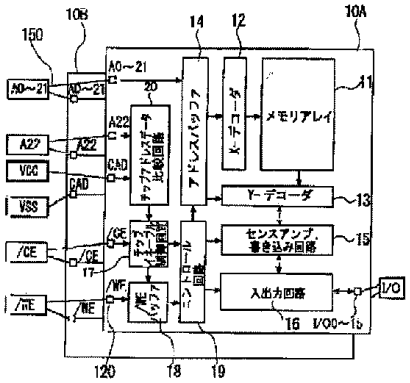
【符号の説明】

１０Ａ，１０Ｂ　フラッシュメモリチップ
１０Ｃ　ＳＲＡＭチップ
１７　チップイネーブル信号制御回路
２０　アドレスビット比較回路
１００　パッケージ
１１０　絶縁基板
１２０　電極端子
１３０　導電性プラグ
１４０　半田ボール
１５０　ボンディングワイヤ
１６０　モールド樹脂

【図１】

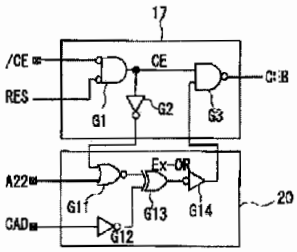


【図２】

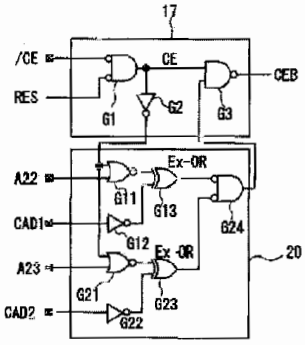


Micron Technology Inc. et al.
Ex. 1010, p. 282

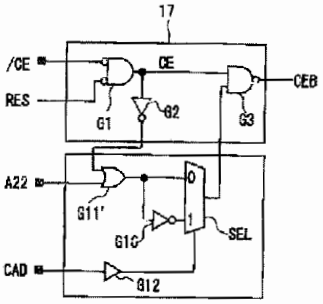
【図3】



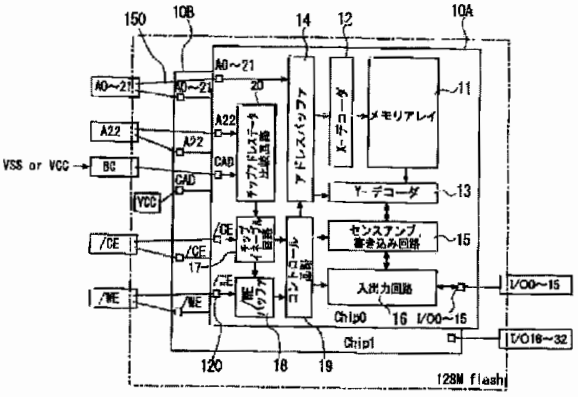
【図4】



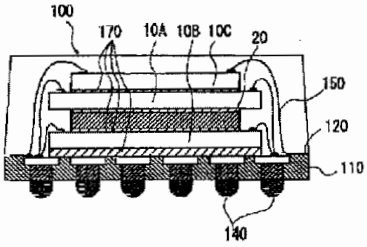
【図5】



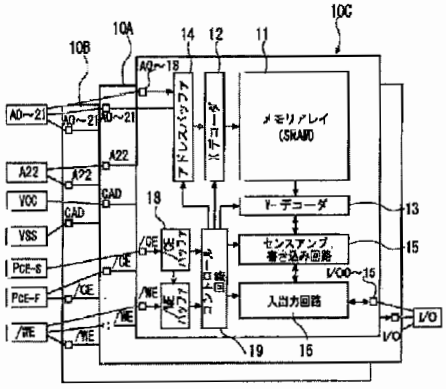
【図6】



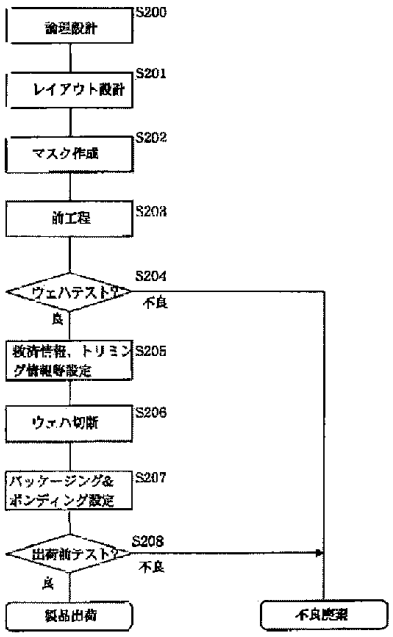
【図7】



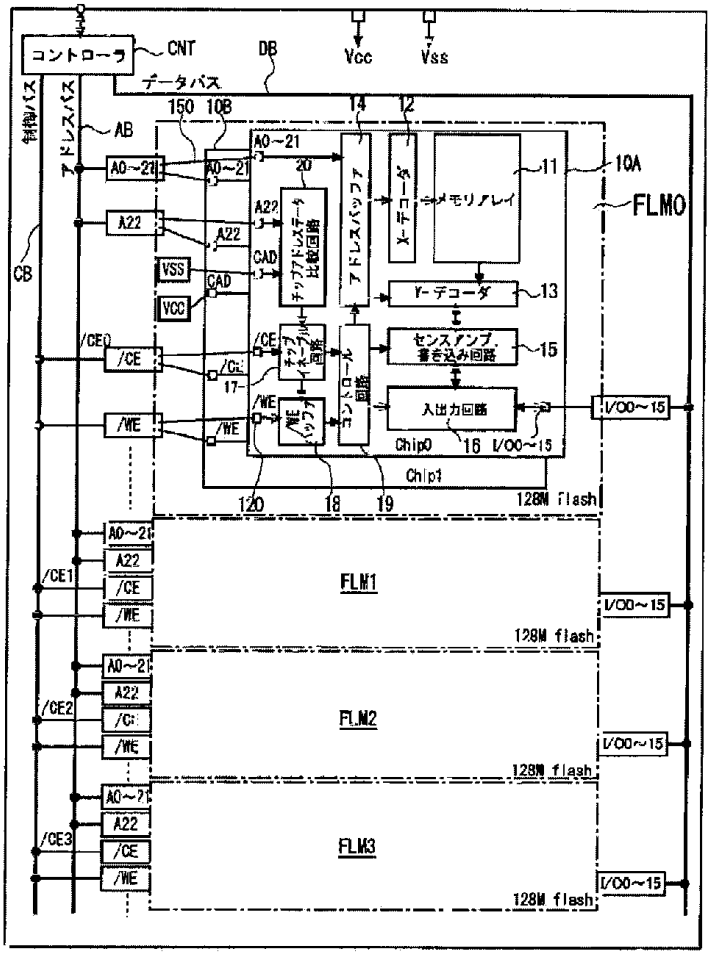
【図8】



【図9】


論理設計
S200
レイアウト設計
S201
マスク作成
S202
前工程
S203
ウェハテスト
S204
不良
良
救済情報、トリミング情報等設定
S205
ウェハ切断
S206
パッケージング&ボンディング設定
S207
出荷前テスト
S208
不良
良
製品出荷
不良廃棄

【図10】



フロントページの続き

(72)発明者　伊藤　紳
秋田県南秋田郡天王町字長沼64　アキタ電子株式会社内

(72)発明者　和田　正志
東京都小平市上水本町五丁目20番１号　株式会社日立製作所半導体グループ内

Fターム(参考)　5B025 AD13 AD16 AE02 AE09
5L106 AA10 CC04 CC05 CC08 CC09
CC13 CC21 CC31 GG00 GG06

Micron Technology Inc. et al.
Ex. 1010, p. 285

* NOTICES *

JPO and INPIT are not responsible for any
damages caused by the use of this translation.

1.This document has been translated by computer. So the translation may not reflect the original precisely.
2.**** shows the word which can not be translated.
3.In the drawings, any words are not translated.

CLAIMS

[Claim(s)]
[Claim 1]The 1st terminal area connected to either a first source voltage terminal or a second source voltage terminal, It has a chip-select control circuit which controls Uchibe's chip selection signal based on a state of two or more terminal areas among [1 / any] a state of this 1st terminal area, and two or more address input terminal areas, This chip-select control circuit An established state of the 1st terminal area of the above, A semiconductor memory device constituting so that a chip selection signal from the outside may control Uchibe's chip selection signal to an effective state or an invalid state based on any 1 or two or more bits of an address signal input in the state where it validated.
[Claim 2]Have the following and each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above is set as the state of differing respectively, A semiconductor memory device constituting so that Uchibe's chip selection signal may be validated in any 1 chip, even if an address signal input where a chip selection signal from the outside is validated is the same.
The 1st terminal area that two or more semiconductor memory chips are stored by one package, and a terminal area of each chip is a semiconductor memory device to which it comes to connect corresponding things common to a terminal in which a package corresponds, and can be set as the state where two or more above-mentioned semiconductor memory chips differ respectively.
A chip-select control circuit which controls Uchibe's chip selection signal based on a state of two or more terminal areas among [1 / any] a state of this 1st terminal area, and two or more address input terminal areas.

[Claim 3]Have the following and each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above is set as the same state and 1 predetermined or two or more terminal areas are fixed to the same electric potential among address input terminal areas, It embraces that a chip selection signal from the outside was validated, and is each the semiconductor memory device constituting so that an internal chip selection signal may be validated.
What two or more semiconductor memory chips are stored by one package, and correspond in a terminal area of each chip other than a data input/output terminal part is the semiconductor memory devices which it comes to connect common to a terminal whose package corresponds,

Micron Technology Inc. et al.
Ex. 1010, p. 286

The 1st terminal area that can be set as the state where two or more above-mentioned semiconductor memory chips differ respectively.
A chip-select control circuit which controls an internal chip selection signal based on a state of two or more terminal areas among [1 / any] a state of this 1st terminal area, and two or more address input terminal areas.

[Claim 4]The semiconductor memory device according to claim 2 or 3, wherein any 1 or two or more bits of a most significant bit [of an address] or top side are 2 bits or more. [of the above-mentioned address signal]
[Claim 5]Have the following and each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above is set as the same state and 1 predetermined or two or more terminal areas are fixed to the same electric potential among address input terminal areas, It is constituted so that an internal chip selection signal may be validated according to a chip selection signal from the outside having been validated, respectively, A semiconductor memory device, wherein one 1st terminal area of the above of two or more above-mentioned semiconductor memory chips is connected to the above-mentioned predetermined terminal of the above-mentioned package and the 1st terminal area of the above of other semiconductor memory chips is connected to one source voltage terminal of the above-mentioned packages.
While two or more semiconductor memory chips are stored by one package, and a terminal for data input/outputs is provided by this package corresponding to each chip and coming independently to connect a data input/output terminal part of each chip to the above-mentioned terminal in which a package corresponds, The 1st terminal area that is the semiconductor memory device with which a predetermined terminal is provided by package, and can be set as the state where two or more above-mentioned semiconductor memory chips differ respectively.
A chip-select control circuit which controls an internal chip selection signal based on a state of two or more terminal areas among [1 / any] a state of this 1st terminal area, and two or more address input terminal areas.

[Claim 6]The semiconductor memory device according to claim 5 laminating two or more above-mentioned semiconductor memory chips so that a terminal area may appear in a uniform direction, respectively, and being mutually joined between each chip by adhesion agent layer by which it was interposed between parts inside the above-mentioned terminal area of these chips.
[Claim 7]Two or more semiconductor memory chips are stored by one package, and, as for a terminal area of each chip, corresponding things are connected common to a terminal in which a package corresponds, It is a manufacturing method of a semiconductor memory device which has the 1st terminal area that can be set as the state of differing respectively for two or more above-mentioned memory chips of every, Two or more above-mentioned semiconductor memory chips are tested by a wafer state after an end of a previous process, After performing trimming treatment by a wafer state to a semiconductor memory chip judged by this test to be a good article, A manufacturing method of a semiconductor memory device sealing in a package after cutting a wafer, dividing into each chip and making setting out to the 1st terminal area of the above of each above-mentioned semiconductor memory chip, and connection between a terminal area of each chip, and a terminal of a package.
[Claim 8]A manufacturing method of the semiconductor memory device according to claim 7, wherein setting out to the 1st terminal area of the above is processing which connects this 1st

Micron Technology Inc. et al.
Ex. 1010, p. 287

terminal area to any 1 source voltage terminal currently provided by package by a bonding wire.
[Claim 9]A manufacturing method of the semiconductor memory device according to claim 8, wherein processing which two or more above-mentioned semiconductor memory chips are laminated, respectively so that a terminal area may appear in a uniform direction, and is connected by the above-mentioned bonding wire is performed for every chip after processing which joins each chip with adhesives.
[Claim 10]A manufacturing method of the semiconductor memory device according to claim 7 performing relief processing which replaces a defective memory cell detected by the above-mentioned test with trimming treatment in the above-mentioned wafer state.

DETAILED DESCRIPTION

[Detailed Description of the Invention]
[0001]
[Field of the Invention]A semiconductor memory device and also two or more semiconductor memory chips apply the present invention to the multichip stack type semiconductor memory device stored by one package, and concerning effective technology, it is used for large scale-ization of a nonvolatile semiconductor memory like a flash memory, and relates to effective technology.
[0002]
[Description of the Prior Art]Conventionally, in the field of a semiconductor integrated circuit, there is technology which carries out lamination storage of two or more semiconductor chips into one package, and aimed at improvement in packaging density, and what attained large scale-ization by this technology also in semiconductor memory is being put in practical use. In such a multichip stack type semiconductor memory device, a terminal number can be reduced by sharing the terminal of a package among two or more chips about most input signals, such as an address signal and a control signal. The terminal number of a package can be reduced by sharing between chips also about the I/O terminal which outputs and inputs data.
[0003]However, when an I/O terminal is made to share. In order to choose each chip independently and to be able to perform read/write of data to a desired chip, about the chip selection signal generally called a chip enable signal/CE, It could not but constitute so that another terminal of a package might be assigned for every chip and the pad for CE signals of each chip and the terminal for CE signals of a package might be connected by a bonding wire etc.
[0004]
[Problem to be solved by the invention]Therefore, a conventional multichip stack type semiconductor memory device, While a terminal number will increase compared with the semiconductor device of a one chip and a package is enlarged when this is seen as one semiconductor device since there are two or more terminals for chip enable, compatibility with the conventional product is lost. That is, since it has two or more terminals for chip enable, newly develop the system which uses a memory, or. When it replaces with the conventional memory chip and constitutes the memory module using a multichip stack type semiconductor memory device, use of the design data of a certain substrate cannot already be aimed at, but it will be necessary to newly redesign a substrate.

Micron Technology Inc. et al.
Ex. 1010, p. 288

[0005]Provide the decoder circuit for decoding an address signal and making any one of two or more of the chip enable signals into a valid level, if one package has two or more terminals for CE signals as an external circuit, or, In order to access a memory, it had to constitute so that two or more chip enable signals could be generated and outputted by the equipment side which outputs an address, and there was fault that a user's design burden became largely.
[0006]Then, two chips share a chip enable terminal and invention of the multichip stack type semiconductor memory device which enabled it to choose a chip by an address signal (most significant bit) is proposed (Tokuganhei11-207701). However, in this prior invention, although multichip stack type the schematic structure and device structure of the semiconductor memory device are disclosed, It does not disclose about production technology, such as a specific circuit configuration which enables it to choose one of chips by the address most significant bit, and the method of specific bonding.
[0007]The purpose of this invention shares a chip enable terminal between two or more chips, and there is in providing the multichip stack type semiconductor memory device which enabled it to choose a chip by an address signal.
[0008]Other purposes of this invention share a chip enable terminal between two or more chips, and there are in providing the effective manufacturing method of the multichip stack type semiconductor memory device which enabled it to choose a chip by an address signal.
[0009]About the above of this invention and the other purposes, and the new characteristics, it will become clear from description and the accompanying drawing of this Description.
[0010]
[Means for solving problem]It will be as follows if the outline of a typical thing is described among invention disclosed in an application concerned. Namely, the 1st terminal area by which the semiconductor memory device concerning the present invention is connected to either a first source voltage terminal or a second source voltage terminal, It has a chip-select control circuit which controls an internal chip selection signal based on the state of two or more terminal areas among [1 / any] the state of this 1st terminal area, and two or more address input terminal areas, This chip-select control circuit is constituted so that an internal chip selection signal may be controlled to an effective state or an invalid state based on any 1 or two or more bits of the address signal input where the chip selection signal from the outside is validated, the established state of the 1st terminal area of the above, and.
[0011]When two or more memory chips are combined according to the above-mentioned means, It can operate as memory storage with which a chip is chosen according to an address signal by changing the established state of the 1st terminal area, or can operate now also as memory storage from which the input-and-output number of bits differs with the combination of the established state of the 1st terminal area, and the established state of an address terminal part.
[0012]In the semiconductor memory device to which it comes to connect those to which two or more semiconductor memory chips are stored by one package, and the terminal area of each chip corresponds common to the terminal in which a package corresponds, The 1st terminal area that can be set as the state where two or more above-mentioned semiconductor memory chips differ respectively, It has a chip-select control circuit which controls an internal chip selection signal based on the state of two or more terminal areas among [1 / any] the state of this 1st terminal area, and two or more address input terminal areas, Each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above was set as the state of differing respectively, even if the address signal input where the chip selection signal from the outside is validated was the same, it constituted so that an internal chip

Micron Technology Inc. et al.
Ex. 1010, p. 289

selection signal might be validated in any 1 chip.
[0013]By this, when two or more memory chips are combined, it can operate as memory storage with which a chip is chosen according to an address signal by changing the established state of the 1st terminal area, Even if there are two or more chips, while being able to commonalize a chip selection signal and being able to reduce the number of external terminals for equipment, Since it can treat identically to one usual memory chip, when developing the system which uses a memory, the existing design board can be used.
[0014]In the semiconductor memory device to which two or more semiconductor memory chips are stored by one package, and, as for the terminal area of each chip other than a data input/output terminal part, it comes to connect corresponding things common to the terminal in which a package corresponds, The 1st terminal area that can be set as the state where two or more above-mentioned semiconductor memory chips differ respectively, It has a chip-select control circuit which controls an internal chip selection signal based on the state of two or more terminal areas among [1 / any] the state of this 1st terminal area, and two or more address input terminal areas, Each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above was set as the same state and 1 predetermined or two or more terminal areas were fixed to the same electric potential among address input terminal areas, it constituted so that an internal chip selection signal might be validated according to the chip selection signal from the outside having been validated, respectively.
[0015]By this, when two or more memory chips are combined, by making the established state of the 1st terminal area of each chip, and the established state of a predetermined address input terminal area correspond, Since the number of bits of an I/O data can treat identically to double the one usual memory chip even if there are two or more chips, when developing the system which uses a memory, the existing design board can be used.
[0016]Any 1 or two or more bits of the above-mentioned address signal shall be 2 bits or more from the most significant bit [of an address], or top side desirably. The multichip stack type semiconductor memory device with which a chip is chosen by this according to an address signal without changing the mat composition of the conventional memory chip and the composition of an address decoder is realizable.
[0017]As for invention of everything but an application concerned, two or more semiconductor memory chips are stored by one package, While the terminal for data input/outputs is provided by this package corresponding to each chip and coming independently to connect the data input/output terminal part of each chip to the above-mentioned terminal in which a package corresponds, In a package, a predetermined terminal in the semiconductor memory device currently provided two or more above-mentioned semiconductor memory chips, It has a chip-select control circuit which controls an internal chip selection signal based on the state of two or more terminal areas among [1 / any] the 1st terminal area that can be set as the state of differing respectively, and the state of this 1st terminal area and two or more address input terminal areas, Each chip-select control circuit in two or more above-mentioned semiconductor memory chips, When the 1st terminal area of the above is set as the same state and 1 predetermined or two or more terminal areas are fixed to the same electric potential among address input terminal areas, embracing that the chip selection signal from the outside was validated -- each -- it being constituted so that an internal chip selection signal may be validated, and it being connected to the above-mentioned predetermined terminal of the above-mentioned package, and one 1st terminal area of the above of two or more above-mentioned semiconductor memory chips, The

Micron Technology Inc. et al.
Ex. 1010, p. 290

1st terminal area of the above of other semiconductor memory chips was connected to one source voltage terminal of the above-mentioned packages. Thereby, the user can choose the number of bits of the I/O data of the semiconductor memory device concerned by setting up suitably the voltage applied to the above-mentioned predetermined terminal.

[0018]Desirably, two or more above-mentioned semiconductor memory chips are laminated so that a terminal area may appear in a uniform direction, respectively, and it is mutually joined between each chip by the adhesion agent layer by which it was interposed between the parts inside the above-mentioned terminal area of these chips. When this performs the electrical link of the terminal area of each chip, and the corresponding terminal by the side of a package by a bonding wire, wirebonding processing can be performed, even if it laminates so that the terminal area of each chip may appear in a uniform direction.

[0019]In the manufacturing method of the semiconductor memory device to which it comes to connect those to which, as for the manufacturing method of the semiconductor memory device concerning the present invention, two or more semiconductor memory chips are stored by one package, and the terminal area of each chip corresponds common to the terminal in which a package corresponds, Two or more above-mentioned semiconductor memory chips are tested by a wafer state after the end of a previous process, After performing trimming treatment by a wafer state to the semiconductor memory chip judged by this test to be a good article, A wafer is cut and it divides into each chip, and after making setting out to the 1st terminal area of the above of each above-mentioned semiconductor memory chip, and connection between the terminal area of each chip, and the terminal of a package, it is made to seal in a package. Since the terminal setting out can be performed at the process of connecting the terminal area of each chip, and the terminal of a package while this uses the memory chip which needs to set up a state, A multichip stack type semiconductor memory device can be manufactured without completely changing the process of the conventional semiconductor memory. It may be made to perform relief processing of a defective bit with the above-mentioned trimming treatment.

[0020]Setting out to the 1st terminal area of the above is considered as the processing which connects this 1st terminal area to any 1 source voltage terminal currently provided by the package by a bonding wire. Since wirebonding technology is very reliable technology, while being able to prevent a setting error by performing setting out to the 1st terminal area using this, A multichip stack type semiconductor memory device can be manufactured without completely changing the process of the conventional semiconductor memory.

[0021]Two or more above-mentioned semiconductor memory chips are laminated, respectively so that a terminal area may appear in a uniform direction, and the processing connected by the above-mentioned bonding wire is made to be performed after the processing which joins each chip with adhesives for every chip. Even if this laminates so that the terminal area of each chip may appear in a uniform direction, wirebonding processing can be performed using existing equipment.

[0022]

[Mode for carrying out the invention]Hereafter, the preferable working example of the present invention is described based on Drawings. Fig.2 shows the Brock configuration diagram of a circuit for the package structure of the multichip stack type semiconductor memory device with which Fig.1 applied the present invention again.

[0023]The two memory chips 10A and 10B are laminated, and the multichip stack type semiconductor memory device of this working example is enclosed in the package 100 as shown in Fig.1. The memory chips 10A and 10B enclosed in the package 100 in this working example

Micron Technology Inc. et al.
Ex. 1010, p. 291

can write in data electrically, for example, and are nonvolatile memory like the flash memory in which collective erasure is possible electrically in a predetermined unit about data.
[0024]Although not restricted in particular, the memory chips 10A and 10B are laminated in the state where it was mutually insulated from these chips electrically on both sides of the spacer 20 which consists of a small insulator. The spacer 20 is provided because the bonding pad part of Shimo's chip hides and wirebonding becomes impossible, when it piles up directly, since two chips have the same size. The memory chips 10A and 10B and the spacer 20 are combined by the pellet adhesives 170. The memory chips 10A and 10B may be the binary memories which can memorize 1-bit data to one memory cell, and may be what is called multi level memories that can memorize data of 2 bits or more to one memory cell.
[0025]110 is an insulating substrate which consists of ceramics etc., and two or more electrode terminals 120 for bonding are arranged at the predetermined intervals along with the extent of a chip at the whole surface (inner surface) of this insulating substrate 110, Corresponding to each electrode terminal 120 for bonding, the conductive plug 130 embedded so that the substrate concerned might be penetrated is provided by the insulating substrate 110. The solder ball 140 is welded on the part (a figure lower end) to which one end (a figure upper end) was contacted on the lower surface of any one above-mentioned pad 120, and the other end 110, i.e., a substrate, on the other hand (outer surface), exposed these conductive plugs 130, respectively. The memory chip 10B is combined by the pellet adhesives 170 on the insulating substrate 110.
[0026]The end of the bonding wire 150 is combined with the electrode terminal 120 for bonding by the side of the above-mentioned substrate, The other end of this bonding wire 150 is combined with the bonding pad in which the memory chips 10A and 10B correspond, and the electrode terminal 120 for bonding by the side of a substrate and the bonding pad in which the memory chips 10A and 10B correspond are electrically connected. And where the memory chips 10A and 10B are connected by the bonding wire 150 on the insulating substrate 110, a mold is carried out with the resin 160, and the package 100 is constituted.
[0027]As shown in Fig.2 each memory chip 10A and 10B, The memory array 11 by which two or more memory cells are arranged at matrix form, Input X decoder 12 which decodes an address X system and makes a selection level one word line in the memory array 11, The Y decoder 13 which chooses the input bit line which decodes an address Y system, and with which it corresponds in the memory array 11, input X system -- an address and Y system -- with the address buffer 14 which incorporates an address and is supplied to above-mentioned X decoder 12 and the Y decoder 13, The sense amplifier & writing circuit 15 which gives the electric potential of a bit line corresponding to write data in amplifying the signal read to the selected bit line ****, The input output circuit 16 which outputs read-out data or incorporates write data from chip outside, The chip enable control circuit 17 which incorporates the chip enable signal/CE input from the chip outside, and generates an internal control signal, and the input buffer 18 which incorporates external control signals, such as a write enable signal/WE, reset-signal RES, an out enable signal / OE, It is constituted by the control circuit 19 etc. which generate the control signal inside a chip according to the incorporated control signal, and control an inside.
[0028]The signal with which a chip enable signal/CE shows that the chip concerned is made into a selective state, They are a signal which shows that the chip concerned writes in a write enable signal/WE, and it changes it into a state, a signal with which reset-signal RES makes the inside of a chip a reset state, and a signal which shows that an out enable signal / OE is in the state where a read-out data signal is outputted. About pads other than two pads (CAD, A22) described

Micron Technology Inc. et al.
Ex. 1010, p. 292

below among the pads into which these control signals and address signals are input, The pads to which the memory chips 10A and 10B correspond are connected common to the corresponding electrode terminal currently provided by the package by the bonding wire, respectively. Pad I/O0 for data input of each chip - I/O15 are connected common to the corresponding electrode terminal currently provided by the package.

[0029]Pad CAD for chip address information setting out which is not provided is provided by the usual memory at the memory chips 10A and 10B of this working example, respectively. For example, each memory chip 10A and 10B is supplied from the chip outside, among the address signals A0-A22 like 23 bits The most significant bit A22, According to the electric potential applied to the above-mentioned pad CAD for chip address information setting out, the address bit comparison circuit 20 which validates the above-mentioned chip enable control circuit 17 is provided, respectively. and pad CAD for chip address information setting out of the memory chip 10A and pad CAD for chip address information setting out of the memory chip 10B -- respectively -- on the other hand, the first source voltage terminal Vcc of a package -- another side is connected to the second source voltage terminal Vss of a package.

[0030]The specific configuration example of the chip enable control circuit 17 and the address bit comparison circuit 20 is shown in Fig.3. AND gate G1 to which the chip enable control circuit 17 considers a chip enable signal/CE, and reset-signal RES as an input, It consists of NAND gate G3 which generates the internal chip enable signal CEB based on the inverter G2 which reverses the output of this AND gate G1, and is supplied to the address bit comparison circuit 20, and the output signal of AND gate G1 and the signal from the address bit comparison circuit 20.

[0031]NOR gate G11 to which the address bit comparison circuit 20, on the other hand, considers the output signal and the address most significant bit A22 of the above-mentioned inverter G2 of the chip enable control circuit 17 as an input, The inverter G12 which generates the signal according to the electric potential applied to the above-mentioned pad CAD for chip address information setting out, It consists of IKUSUKURUSHIBU OR gate G13 which considers the output signal of this inverter G12, and the output signal of above-mentioned NOR gate G11 as an input, and the inverter G24 which reverses the output of this IKUSUKURUSHIBU OR gate G13, and is supplied to NAND gate G3 of the chip enable control circuit 17.

[0032]Hereafter, operation of the circuit of Fig.3 is described. The high level of reset-signal RES is a signal of a valid level, and when reset-signal RES is a low level, the chip enable control circuit 17 validates a chip enable signal/CE. That is, when reset-signal RES is a low level, the output of AND gate G1 changes according to a chip enable signal/CE. A low level is a signal which shows a valid level, i.e., a chip-select state, and when both reset-signal RES, and a chip enable signal/CE are low levels, it becomes high-level outputting a chip enable signal/CE of AND gate G1.

[0033]And since it is reversed with the inverter G2 and the output of AND gate G1 is supplied to NOR gate G11 of the address bit comparison circuit 20, When both reset-signal RES, and a chip enable signal/CE are low levels, NOR gate G11 changes according to the address most significant bit A22 which is an input of another side of NOR gate G11. When either one of reset-signal RES, or a chip enable signal/CE is high-level, the output of NOR gate G11 is fixed to a low level. When either one of reset-signal RES, or a chip enable signal/CE is high-level, the output of NAND gate G3 of the output stage of the chip enable control circuit 17 will be fixed high-level, and, as for the internal chip enable signal CEB, a chip non selection state will be

Micron Technology Inc. et al.
Ex. 1010, p. 293

shown.
[0034]When both reset-signal RES, and a chip enable signal/CE are low levels, NOR gate G11 of the address bit comparison circuit 20 operates as an inverter which makes an input signal the address most significant bit A22, When the address most significant bit A22 is a low level, in the output of NOR gate G11, when the address most significant bit A22 is high-level, the output of NOR gate G11 becomes a low level high-level. And since IKUSUKURUSHIBU OR gate G13 which considers the output signal of this NOR gate G11 as an input is an exclusive "or" circuit, When the output logical level of the buffer G12 and the output logical level of NOR gate G11 which output the signal according to the electric potential applied to pad CAD for chip address information setting out are the same, the output of IKUSUKURUSHIBU OR gate G13 to a low level, When the output logical levels of G11 and G12 differ, the output of IKUSUKURUSHIBU OR gate G13 becomes high-level.
[0035]And since it is reversed with the inverter G14 and the output of IKUSUKURUSHIBU OR gate G13 is supplied to NAND gate G3 of the output stage of the chip enable control circuit 17, When the output logical levels of G11 and G12 differ, the internal chip enable signal CEB which is an output of NAND gate G3 is made into the high level which shows a chip non selection state. On the other hand, when the output logical level of G11 and G12 is the same, NAND gate G3 operates as an inverter and the internal chip enable signal CEB which is that output is made into the low level which shows a chip-select state according to the input signal (at this time high level) of another side of NAND gate G3.
[0036]As mentioned above, the circuit of Fig.3 changes the internal chip enable signal CEB into a chip-select state according to the electric potential applied to the address most significant bit A22 and pad CAD for chip address information setting out, or changes it into a non selection state. The relation between the combination of the address most significant bit A22 in the semiconductor memory device which uses two MEMORIPPU which carries the working-example circuit of Fig.3 in the next table 1, and the electric potential applied to pad CAD for chip address information setting out, and the operating state of a chip is arranged and shown.
[0037]
[Table 1]

[0038]The electric potential applied to pad CAD for chip address information setting out the chip (1) which is Vss as shown in Table 1, According to the address most significant bit A22, when A22 is high-level, the internal chip enable signal CEB becomes high-level regardless of /CE, and a chip will be in a non selection state, If /CE is made into a low level when A22 is a low level, the internal chip enable signal CEB will serve as a low level, and a chip will be in a selective state.
[0039]On the other hand, the electric potential applied to pad CAD for chip address information setting out the chip (2) which is Vcc, According to the address most significant bit A22, if /CE is made into a low level when A22 is high-level, the internal chip enable signal CEB will serve as a low level, and a chip will be in a selective state, When A22 is a low level, regardless of /CE, the internal chip enable signal CEB becomes high-level, and a chip will be in a non selection state.
[0040]As shown in (3) of Table 1, when both the electric potential applied to the address most

Micron Technology Inc. et al.
Ex. 1010, p. 294

significant bit A22 and pad CAD for chip address information setting out is Vcc(s) (or both Vss), In a chip, according to the chip enable signal/CE from the outside, the internal chip enable signal CEB changes in a similar manner. Therefore, two chips can be simultaneously made into a selective state by fixing to Vcc the electric potential by which both chips are applied to pad CAD for chip address information setting out, and the terminal into which the address most significant bit A22 should be input.

[0041]Therefore, while providing the respectively different electrode terminal for data input/outputs in a package corresponding to the data input/output pad of two chips enclosed with one package, It can be made to operate as a memory (x2n) which has twice as many data width as this by fixing to Vcc the pad into which pad CAD for chip address information setting out and the address most significant bit A22 should be input. For example, if two chips have a storage capacity which is 64M bit, respectively and carry out the parallel input/output of the 16-bit data, it can constitute as a semiconductor memory device which has a 128M bit storage capacity and carries out the parallel input/output of the 32-bit data.

[0042]However, since the pad for address A22 input of a chip is connected to a Vcc terminal within a package in this case, in a package, the electrode terminal corresponding to the address most significant bit A22 is unnecessary. When constituted as a semiconductor memory device which stored one chip in one package as multichip, the chip which has composition as shown in Fig.2, By fixing to Vcc both the electric potential applied to the address most significant bit A22 and pad CAD for chip address information setting out inside a package, it can also constitute as a usual semiconductor memory device which has a 64M bit storage capacity and carries out the parallel input/output of the 16-bit data.

[0043]On the other hand, in the case of Table 1 (1) and (2), if two chips have a storage capacity which is 64M bit, respectively and carry out the parallel input/output of the 16-bit data, It is constituted as a semiconductor memory device which has a 128M bit storage capacity and carries out the parallel input/output of the 16-bit data, and each chip is selectively accessed by the address most significant bit A22.

[0044]Although the control signal according to a chip enable signal/CE is sent to NOR gate G11 of the address bit comparison circuit 20 with the inverter G2 from the chip enable control circuit 17 and incorporation of the address A22 is controlled in the working example of Fig.3, This signal is not necessarily required and omitting is also possible. It is effective in preventing consumption electrical quantity from penetration current flowing into an internal circuit because the address A22 which is an input of NOR gate G11 changes in a chip non selection state, and increasing by controlling NOR gate G11 by this control signal.

[0045]Therefore, it is desirable to control also about the buffer which incorporates the bits A0-A21 other than address A22 into the inside of a chip not to operate in the state of chip non selection with the inverter G2 similarly. In the working example of Fig.3, although AND gate G1 which controls the input of a chip enable signal/CE according to reset-signal RES to the chip enable control circuit 17 is provided, it is possible to also omit this AND gate G1.

[0046]Other configuration examples of the chip enable control circuit 17 and the address bit comparison circuit 20 are shown in Fig.4. Fig.4 is the example of a circuit configuration for which it was suitable when four chips were made into 1 set and a semiconductor memory device was constituted. In this working example, two pad CAD1 for chip address information setting out, CAD2, and the input terminal of the address signal A23 are provided. And while the logic gates G11-G13 and the same logic gates G21-G23 are provided by the address bit comparison circuit 20 corresponding to this, AND gate G24 is provided instead of the inverter G14. And the

Micron Technology Inc. et al.
Ex. 1010, p. 295

output of the inverter G2 of the address signal A23 and the chip enable control circuit 17 is input into NOR gate G21, The output of NOR gate G21 and the output of the inverter G22 which generates the signal according to the electric potential applied to pad CAD2 for chip address information setting out are input into IKUSUKURUSHIBU OR gate G23.
[0047]The output of IKUSUKURUSHIBU OR gates G13 and G23 is input into AND gate G24. The chip enable control circuit 17 is the same as the working example of Fig.3. By making setting out (connection of a bonding wire) so that the combination of the voltage VccVss applied to pad CAD1 for chip address information setting out and CAD2, respectively may differ in four chips in this working example, It has responded to the address signal A22 and A23, the internal enable signal CEB of a gap or one chip is made into a low level, and only one chip is made into a selective state among four. However, it can be made to operate as a memory which has one 4 times the data width of this by setting to Vcc pad CAD1 for chip address information setting out of all the chips, CAD2, the address signal A22, and voltage applied to the input pad of A23 like the working example of Fig.3.
[0048]The relation between the combination of the address most significant bit A22 in the semiconductor memory device which uses four MEMORIPPU which carries the working-example circuit of Fig.4 in Table 2, A23, and the electric potential applied to pad CAD1 for chip address information setting out and CAD2, and the operating state of a chip is arranged and shown.
[0049]
[Table 2]

[0050]Other configuration examples of the address bit comparison circuit 20 shown in Fig.3 are shown in Fig.5.Although IKUSUKURUSHIBU OR gate G13 which compares the most significant bit A22 of an address signal with the state of pad CAD for chip address information setting out is used in Fig.3, In Fig.5, the selector SEL is provided instead of IKUSUKURUSHIBU OR gate G13, and it is constituted so that it may control by the output of the buffer G12 which generates the signal according to the electric potential applied to pad CAD for chip address information setting out in this selector SEL.
[0051]In the working example of Fig.5, OR gate G11' is used instead of NOR gate G11 of Fig.3.The input of OR gate G11' is a signal from the address A22 and the chip enable control circuit 17 as well as the input of NOR gate G11, OR gate G11' acts as a transmission gate passed

Micron Technology Inc. et al.
Ex. 1010, p. 296

the address A22 or intercepted according to the signal from the chip enable control circuit 17.
[0052]The signal which reversed the output of OR gate G11' and the output of this OR gate G11' with the inverter G16 is input into the data input terminal of the above-mentioned selector SEL. Thus, by choosing the address most significant bit A22 which passed OR gate G11' according to the electric potential of pad CAD, or its inversion signal by the selector SEL, and supplying the chip enable control circuit 17, Even when the same address A22 is input, as for one side, the internal chip enable signal CEB is made into a valid level according to the electric potential of pad CAD, and, as for another side, let the internal chip enable signal CEB be an invalid level. As a result, a chip can be chosen now with the address A22.
[0053]Although the selector SEL is controlled by the output of the buffer G12 which generates the signal according to the electric potential applied to pad CAD in the working example of Fig.5 and he is trying to choose either one of the address A22 or its inversion signal, Conversely, the selector SEL may be controlled by the output of the address A22, i.e., OR gate G11', and it may constitute so that the output of the buffer G12 or either of the inversion signal may be chosen and outputted.
[0054]The address bit comparison circuit 20 of this example is applicable also to the working example shown in the Fig.4 which was suitable when four chips were made into 1 set and a semiconductor memory device was constituted. What is necessary is specifically, to use a selector instead of IKUSUKURUSHIBU OR gate [of Fig.4] G13, and G23, and just to constitute the address bit comparison circuit 20.
[0055]Fig.6 shows other working examples of the present invention. This working example is a modification of the semiconductor memory device which laminated two flash memory chips shown in Fig.1 and Fig.2, and the composition of the flash memory chips 10A and 10B is the same as that of Fig.2.The point that, as for differing from Fig.2, bit configuration change terminal BC is provided by the package, and terminal CAD for chip address selections of one chip (for example, 10A) is connected to this bit configuration change terminal BC among the memory chips 10A and 10B, While is carried out to 32 (I/O0-I/O31) instead of 16 (I/O0-I/O15), and, as for the data input/output pad of a chip (for example, 10A), the data ON appearance terminal of a package is connected to data ON appearance terminal I/O0 of a package - I/O15, The data input/output pad of the chip (for example, 10B) of another side is a point connected to data ON appearance terminal I/O16 of a package - I/O31.
[0056]The semiconductor memory device of this working example has the characteristics in the point which a user can use also as a memory of a 32-bit parallel input/output also as a memory of a 16-bit parallel input/output. If bit configuration change terminal BC to which terminal CAD for chip address selections of the memory chip 10A is connected is specifically connected to the power supply voltage Vss, It operates as a memory which a chip is chosen by the address most significant bit A22 like the semiconductor memory device of the working example of Fig.2, and outputs and inputs in parallel the data which is 16 bits. Therefore, data ON appearance terminal I/O0 of a package - I/O15, I/O16 - I/O31 are connected common to the 16-bit same Buss in this case.
[0057]On the other hand, if bit configuration change terminal BC to which terminal CAD for chip address selections of the memory chip 10A is connected is connected to the power supply voltage Vcc, Two chips are simultaneously made a selective state by chip enable signal CE, and it operates as a memory which outputs and inputs 32-bit data in parallel. Therefore, data ON appearance terminal I/O0 of a package - I/O15, I/O16 - I/O31 are connected to Buss's 32-bit another signal wire in this case.

Micron Technology Inc. et al.
Ex. 1010, p. 297

[0058]Fig.7 and Fig.8 show the working example of further others of the present invention.This working example is a development form of the semiconductor memory device which laminated two flash memory chips shown in Fig.1 and Fig.2, and laminates the static RAM chip 10C further on the two laminated flash memory chips 10A and 10B.The composition of the flash memory chips 10A and 10B is the same as that of Fig.2. In this working example, Even if the size of the static RAM chip 10C piles up the static RAM chip 10C on the flash memory chip 10A smaller than the size of the flash memory chips 10A and 10B, in order that the bonding pad of the flash memory chip 10A may not hide, The static RAM chip 10C is joined by the pellet adhesives 170 on the protective film of the flash memory chip 10A.
[0059]The pad and the input pad of the address signals A0-A18 into which a write enable signal/WE is input among the pads currently provided by these chips 10A, 10B, and 10C, As shown in Fig.7, it is connected to the corresponding electrode terminal 120 currently provided by the package 100 via the bonding wire 150 in common.
[0060]As shown in Fig.8, pad I/O0 for data input of each chip - I/O15 are connected common to the corresponding terminal currently provided by the package. Although the pad of the flash memory chips 10A and 10B is connected common to corresponding terminal PCE-F currently provided by the package like the aforementioned working example about the chip enable signal/CE, The pad into which the chip enable signal/CE of the SRAM chip 10C are input is connected to terminal PCE-S for exclusive use currently provided by the package.
[0061]The static RAM chip 10C, The input buffer 18 of control signals, such as the memory array 11, X decoder 12, the Y decoder 13, the address buffer 14, the sense amplifier & writing circuit 15, the data output circuit 16, a write enable signal/WE, the control circuit 19, etc. It has the same composition as general-purpose SRAM which it had, and pad CAD for chip address information setting out, the chip enable control circuit 17, and the address bit comparison circuit 20 which are provided by the flash memory chips 10A and 10B of the aforementioned working example are not provided.
[0062]Since lamination storage of two flash memories and the one SRAM is carried out into one package, the semiconductor memory device of this example can raise the packaging density of a system further compared with the case where what is stored by the separate package is used, and can attain the miniaturization of equipment.
[0063]Next, an example of the procedure from development of the multichip stack type semiconductor memory device of the above-mentioned working example to manufacture is described using Fig.9. First, the logic operation by the logic design and simulation of a memory chip which provided pad CAD for chip address information setting out, the chip enable control circuit 17, and the address bit comparison circuit 20 which were described in the aforementioned working example is checked (Step S200). Next, a layout design is performed using the design support program called a layout tool based on the above-mentioned logic design data (Step S201). And the mask used for a process based on this layout design data is created (Step S202).
[0064]Next, the previous process which forms two or more memory chips on a semiconductor wafer using the above-mentioned mask is performed (Step S203). And the wafer tests which put a probe in a wafer state to the pad of each chip with a circuit tester, and are examined are done (Step S204). The chip judged that is poor at these wafer tests is discarded as a defective article, after marking is attached and being cut by each chip. On the other hand, processing which sets the relieving information of a defective bit and the trimming information of the pressure value of an internal electrical power source circuit according [the chip judged at Step S204 to be a good article] to a redundant circuit as the fuse currently provided on the chip or a nonvolatile memory

Micron Technology Inc. et al.
Ex. 1010, p. 298

element is performed (Step S205).
[0065]A wafer is cut by each chip after processing of probe inspection, relief, etc. is completed about all the chips on a wafer (Step S206). And a chip is carried in an insulating substrate and packaging processing which carries out a mold to wirebonding which connects the pad of a chip and the electrode terminal by the side of a corresponding substrate by resin is performed (Step S207). In this bonding & packaging processing, After joining the lower memory chip 10B with the adhesives pellet 170 on the insulating substrate 110 first and performing wirebonding, It is good to join the upper memory chip 10A on the lower memory chip 10B with the adhesives pellet 170, and to be made to perform wirebonding to this chip 10A.
[0066]And in this working example, processing which connects pad CAD for chip address information setting out of each chip to one which is provided by the package of source voltage terminals in the case of the above-mentioned wirebonding is performed. Thereby, setting processing to pad CAD for chip address information setting out can be performed, without adding a new process in any way. If it collects to the two chips 10A and 10B and tries to perform wirebonding after joining the upper memory chip 10A on the lower memory chip 10B with the adhesives pellet 170, Although the upper chip becomes obstructive and wirebonding to a lower chip becomes difficult, After performing wirebonding to the lower chip 10B, an existing apparatus can perform wirebonding processing by joining the upper memory chip 10A and performing wirebonding.
[0067]As a method of connecting pad CAD for chip address information setting out to one of source voltage terminals, Although the fuse which connects pad CAD for chip address information setting out other than wirebonding to the power supply voltage Vcc or Vss previously is provided and how to cut a fuse before packaging processing can be considered, Since two, the chip connected to the power supply voltage Vcc and the chip connected to Vss, produce pad CAD for chip address information setting out in the case of a fuse, it is necessary to manage each chip separately but, and. Since it is not necessary to manage a chip independently if it is made to set up by wirebonding of packaging processing like the aforementioned working example, reduction of the part cost is attained.
[0068]After the above-mentioned packaging processing, the test before shipment which tests each device with a circuit tester is performed (Step S208). While marking is attached and the chip judged that is poor by this test is discarded as a defective article by a next sorting process, the device judged at Step S208 to be a good article is shipped as a product.
[0069]The composition of the memory card as an application of the semiconductor memory device of a first working example (refer to Fig.1 and Fig.2) is shown in Fig.10.A card body in which CD was formed by the insulating material in Fig.10, a nonvolatile storage with which FLM0-FLM3 were built in this card body, CNT is the controller connected to these nonvolatile storages FLM0-FLM3 via the data bus DB, the address bus AB, and the control bus CB. The nonvolatile storages FLM0-FLM3 may contain two chips which have composition as respectively shown in Fig.2, may contain only the chip of 1, or may contain three or more chips. The controller CNT is connected with control devices, such as an external microprocessor, via serial input/output terminal SIO, and the writing and read-out of data to the nonvolatile storages FLM0-FLM3 inside a card are performed according to the instructions from a control device.
[0070]The controller CNT is generating a selection signal /CE0 for exclusive use, /CE1, /CE2, and /CE3 to each nonvolatile storage FLM0-FLM3, respectively, and makes a gap or one a selective state. As described in the working example of Fig.2, each nonvolatile storage FLM0-FLM3. On the other hand, pad CAD for chip address information setting out of a memory chip

Micron Technology Inc. et al.
Ex. 1010, p. 299

(10A) is connected to the first power supply voltage Vss of a package, Pad CAD for chip address information setting out of the memory chip (10B) of the other is connected to the second source voltage terminal Vcc of a package, it has responded to the address most significant bit (for example, A22), and a gap or one chip is made into a selective state.

[0071]If the semiconductor memory device of the aforementioned working example is used so that this application may also show, it can treat on a par with the conventional semiconductor memory completely, and the memory card it is moreover twice whose storage capacity of this can be constituted. In the memory card of this application, the insulating substrate 110 shown in Fig.1 may be constituted as a substrate common to each memory storage FLM0-FLM3 and the controller CNT, the mold of the whole may be carried out by resin etc., and it may enclose with one package. That is, if a memory card is constituted using what was manufactured as a semiconductor memory device which has structure as shown in Fig.1, a package will serve as double structure, but. If each memory storage FLM0-FLM3 and the controller CNT are carried on a common substrate, it can be made to enclose with one package.

[0072]Although invention made by the inventor above was specifically described based on the working example, it cannot be overemphasized that it can change variously in the range which the present invention is not limited to the above-mentioned working example, and does not deviate from the summary. For example, in the working example of Fig.1, although two or more memory chips described about what is enclosed with the package by the laminating condition, also when two or more memory chips are enclosed with the package in the state where it arranged horizontally on the insulating substrate of one sheet, it can apply. The address bit for selection of a chip may not be limited to the address most significant bit, and may be, other bits, for example, least significant bit etc., etc.

[0073]Although described in the above description about the multichip stack type semiconductor memory device which makes a subject the flash memory which is the field of the invention which became the background about invention mainly made by the inventor, EEPROMs other than a flash memory can be used for the semiconductor memory device etc. which were built in.

[0074]

[Effect of the Invention]It will be as follows if the effect acquired by the typical thing among invention disclosed in an application concerned is described simply. Namely, while two or more chips share a chip enable terminal, being able to choose a chip now by an address signal and being able to reduce the number of external terminals as a semiconductor memory device by this, Since it can treat on a par with one usual semiconductor memory, when developing the system which uses a memory, the multichip stack type semiconductor memory device which can use the existing design board can be realized. It can manufacture without changing the manufacturing process of the conventional semiconductor memory for such a semiconductor memory device, or converting a manufacturing installation.

[Translation done.]

Micron Technology Inc. et al.
Ex. 1010, p. 300

Espacenet

Bibliographic data: JP2008046989 (A) — 2008-02-28

MEMORY CONTROL DEVICE

Inventor(s): TAKAKU KAZUYA; HONDA IKUSHI; SUZUKI KENJI ± (TAKAKU KAZUYA, ; HONDA IKUSHI, ; SUZUKI KENJI)

Applicant(s): FUJITSU LTD ± (FUJITSU LTD)

Classification: - international: ***G06F12/00; G06F12/06; G11C11/403; G11C5/00***
- European: G06F13/16D

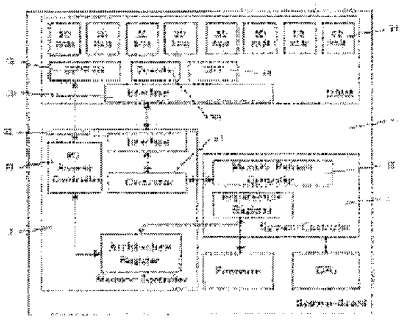
Application number: JP20060223567 20060818

Priority number (s): JP20060223567 20060818

Also published as: US2008046631 (A1)

Abstract of JP2008046989 (A)

PROBLEM TO BE SOLVED: To provide a technique for increasing the number of ranks of a memory module by a small change in configuration. ; SOLUTION: A memory control device which accesses the memory module having a plurality of ranks includes: an interface part having a select signal line for a select signal for rank selection and an address signal line for an address signal indicating an address on a selected rank; and a control part which transmits the signal for rank selection through a part of an address signal line and the select signal line. ; COPYRIGHT: (C) 2008,JPO&INPIT



Micron Technology Inc. et al.
Ex. 1010, p. 301

Last updated: 24.09.2012 Worldwide
Database 5.7.42.93p

Micron Technology Inc. et al.
Ex. 1010, p. 302

(19)日本国特許庁(JP)　(12)公開特許公報(A)　(11)特許出願公開番号
特開2008-46989
(P2008-46989A)
(43)公開日　平成20年2月28日(2008.2.28)

(51)Int.Cl.		F I			テーマコード（参考）
G06F 12/06	(2006.01)	G06F	12/06	515H	5B060
G06F 12/00	(2006.01)	G06F	12/00	564A	5M024
G11C 11/403	(2006.01)	G06F	12/00	597D	
G11C 5/00	(2006.01)	G11C	11/34	363M	
		G11C	5/00	303	

審査請求 未請求　請求項の数 10　OL　（全 13 頁）

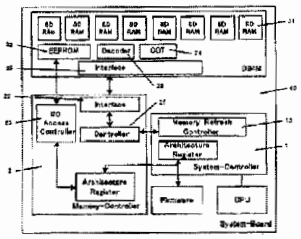
(21)出願番号　特願2006-223567 (P2006-223567)
(22)出願日　平成18年8月18日 (2006.8.18)

(71)出願人　000005223
富士通株式会社
神奈川県川崎市中原区上小田中４丁目１番１号
(74)代理人　100090516
弁理士　松倉　秀実
(74)代理人　100113608
弁理士　平川　明
(74)代理人　100105407
弁理士　高田　大輔
(74)代理人　100089244
弁理士　遠山　勉
(72)発明者　高久　和也
神奈川県川崎市中原区上小田中４丁目１番１号　富士通株式会社内

最終頁に続く

(54)【発明の名称】メモリ制御装置

(57)【要約】
【課題】少ない構成変更でメモリモジュールのランク数を増加させる技術を提供する。
【解決手段】複数のランクを有するメモリモジュールにアクセスするメモリ制御装置であって、ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するインターフェイス部と、前記アドレス信号線の一部及びセレクト信号線を介して前記ランクを選択する信号を送信する制御部とを備えた。
【選択図】図３



システム全体図

Micron Technology Inc. et al.
Ex. 1010, p. 303

【特許請求の範囲】
【請求項1】
複数のランクを有するメモリモジュールにアクセスするメモリ制御装置であって、
ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するインターフェイス部と、
前記アドレス信号線の一部及びセレクト信号線を介して前記ランクを選択する信号を送信する制御部と、
を備えたメモリ制御装置。
【請求項2】
前記制御部が、前記セレクト信号線数よりも多い数のランクを選択する信号を送信する請求項１に記載のメモリ制御装置。
【請求項3】
前記インターフェイス部が、ＤＤＲ２インターフェイスである請求項１又は２に記載のメモリ制御装置。
【請求項4】
前記制御部からの命令に基づいて、前記ランク数に応じたメモリリフレッシュ処理を行うメモリリフレッシュコントロール回路を備えた請求項１から３の何れかに記載のメモリ制御装置。
【請求項5】
前記制御部は、異ランク間に対するアクセスが連続した場合に、該アクセス間に所定の待機時間を挿入する請求項１から４の何れかに記載のメモリ制御装置。
【請求項6】
前記制御部が、前記メモリモジュールのランク数及び／又はアクセス命令の種別に応じて、各ランクが備えるＯＤＴをそれぞれ制御する請求項１から５の何れかに記載のメモリ制御装置。
【請求項7】
複数ランクに区分され、各ランクが選択的にアクセス可能に接続されたメモリ郡と、
ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するインターフェイス部と、
前記インターフェイス部を介してアクセス命令を受信した際のアドレス信号線を介して得た信号及びセレクト信号に基づいて前記ランクを選択してアクセスさせるデコード部と、
を備えたメモリモジュール。
【請求項8】
前記セレクト信号線数よりも多い数のランクを備えた請求項７に記載のメモリモジュール。
【請求項9】
前記デコード部をレジスタ間に配置した請求項７又は８に記載のメモリモジュール。
【請求項10】
複数ランクのメモリモジュールと、
前記メモリモジュールに格納される情報を用いて処理を実行するプロセッサと、
前記メモリモジュールへのアクセスを制御するメモリ制御装置と、を有する情報処理装置において、
前記メモリ制御装置が、
メモリモジュールのランクを選択するためのセレクト信号を介するセレクト信号線と、選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線とが接続されるインターフェイス部と、
前記アドレス信号及び前記セレクト信号から、ランクを選択する信号を送信する制御部と、を備えた情報処理装置。

【発明の詳細な説明】

Micron Technology Inc. et al.
Ex. 1010, p. 304

【技術分野】
【0001】
本発明は、大容量のメモリモジュールを実現する技術に関する。
【背景技術】
【0002】
バイトスライス型メモリコントローラを搭載したサーバ等のコンピュータの構成を例として図１に示す。
装置を制御するシステムコントローラ１００に、ＣＰＵ、メモリコントローラ１０１が接続されている。各メモリコントローラ１０１に少なくとも１つのＤＩＭＭ(Dual Inline Memory Module)が接続され、４つのメモリコントローラが同期してメモリアクセスを行う。（図１）
【0003】
図２は、該ＤＩＭＭの概略構成図であり、ＤＤＲ２インターフェイスを用いた２ランク(RANK)のＤＩＭＭの例を示している。
図２のＤＩＭＭは、１ランクに１８個のＳＤＲＡＭ(Synchronous Dynamic Random Access Memory)を備え、２ランク有しているので、ＤＩＭＭ全体でのＳＤＲＡＭ総数は３６個となっている。
複数のランクで構成されたＤＩＭＭは、ランク数分のチップセレクト信号を用い、各ランクの読み書きを選択的に行っている。現在、デファクトスタンダードの一つとなっているＤＤＲ２インターフェイスでは、チップセレクト信号線２本、アドレス信号線がＡ０～Ａ１５の１６本となっているため、このインターフェイスを採用した場合、ランク数は最大２であり、２ランクで１６本のアドレス信号線を用いて読み書きできる最大の容量は１６ＧＢである。
【0004】
また、本願発明に関連する先行技術として、例えば、下記の特許文献１に開示される技術がある。
【特許文献1】特開２００３－７９６３号公報
【発明の開示】
【発明が解決しようとする課題】
【0005】
（１）問題点１
従来のＤＤＲ２インターフェイスでは、制御できるランク数が２ランクまでとなっており、装置に搭載されるメモリ容量を増やすには、ＤＩＭＭアクセスチャンネルを増やすか、ＤＩＭＭに搭載される各ＳＤＲＡＭの容量を増加させるしか方法がなかった。
しかし、ＳＤＲＡＭの容量の上限は、製造技術に依存するため、任意に上限を超えて容量を増加できるものではない。
【0006】
また、ランク数を増やす場合、１ランク毎にチップセレクト信号線１本を追加する必要があり、メモリコントローラのピン数を増やすなどボードの再開発も行わなくてはならず、現実的ではない。
【0007】
（２）問題点２
公知例「特開２００２－１８４１７６」では、２ランクから４ランクに増量する例が提案されているが、ＤＤＲ２のような高速インターフェイスには適用できない。また、ランク数を増やした場合、メモリコントローラ側を変更しなければ、公知例に挙げられたＤＩＭＭの制御ができないはずであるが、この制御については開示されておらず、実現性に乏しいという問題があった。
【0008】
例えば、ランクを増やすと言うことは、物理的なＳＤＲＡＭが増えるということなので、増えたＳＤＲＡＭを初期化するための回路が必要であるし、メモリリフレッシュの回路

Micron Technology Inc. et al.
Ex. 1010, p. 305

も必要である。
【0009】
そこで本発明は、少ない構成変更でメモリモジュールのランク数を増加させる技術を提供する。
【課題を解決するための手段】
【0010】
上記課題を解決するため、本発明は、以下の構成を採用した。
即ち、本発明のメモリ制御装置は、
複数のランクを有するメモリモジュールにアクセスする装置であって、
ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するインターフェイス部と、
前記アドレス信号線の一部及びセレクト信号線を介して前記ランクを選択する信号を送信する制御部と、を備えた。
【0011】
また、本発明の情報処理装置は、
複数ランクのメモリモジュールと、
前記メモリモジュールに格納される情報を用いて処理を実行するプロセッサと、
前記メモリモジュールへのアクセスを制御するメモリ制御装置と、を有する情報処理装置において、
前記メモリ制御装置が、
メモリモジュールのランクを選択するためのセレクト信号を介するセレクト信号線と、選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線とが接続されるインターフェイス部と、
前記アドレス信号及び前記セレクト信号から、ランクを選択する信号を送信する制御部と、を備えた。
【0012】
前記制御部は、前記セレクト信号線数よりも多い数のランクを選択する信号を送信しても良い。
【0013】
前記インターフェイス部は、ＤＤＲ２インターフェイスであっても良い。
【0014】
前記メモリ制御部は、前記制御部からの命令に基づいて、前記ランク数に応じたメモリリフレッシュ処理を行うメモリリフレッシュコントロール回路を備えても良い。
【0015】
前記制御部は、異ランク間に対するアクセスが連続した場合に、該アクセス間に所定の待機時間を挿入しても良い。
【0016】
前記制御部は、前記メモリモジュールのランク数及び／又はアクセス命令の種別に応じて、各ランクが備えるＯＤＴをそれぞれ制御しても良い。
【0017】
また、本発明のメモリモジュールは、
複数ランクに区分され、各ランクが選択的にアクセス可能に接続されたメモリ郡と、
ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するインターフェイス部と、
前記インターフェイス部を介してアクセス命令を受信した際のアドレス信号線を介して得た信号及びセレクト信号に基づいて前記ランクを選択してアクセスさせるデコード部と、を備えた。
【0018】
前記メモリモジュールは、前記セレクト信号線数よりも多い数のランクを備えても良い。

Micron Technology Inc. et al.
Ex. 1010, p. 306

【0019】
前記メモリモジュールは、前記デコード部をレジスタ間に配置しても良い。
【発明の効果】
【0020】
本発明によれば、少ない構成変更でメモリモジュールのランク数を増加させる技術を提供できる。
【発明を実施するための最良の形態】
【0021】
以下、図面を参照して本発明を実施するための最良の形態について説明する。以下の実施の形態の構成は例示であり、本発明は実施の形態の構成に限定されない。
【0022】
§１．全体構成
図３は、本発明の一実施例に係る情報処理システムの概略図である。
本例の情報処理システム１０は、システムボード上に、ＣＰＵや、システムコントローラ１、メモリコントローラ２、メモリモジュール３を備えた汎用のコンピュータである。
情報処理システム１０は、メモリコントローラ（メモリ制御装置）２やメモリモジュール３を複数備えても良い。例えば図１と同様に各メモリコントローラ２に少なくとも１つのＤＩＭＭが接続され、４つのメモリコントローラ２が同期してメモリアクセスを行う。
【0023】
§２．各構成部分の説明
本実施形態のメモリモジュール３は、図４に示すように、ＤＤＲ２インターフェイスを利用したＤＩＭＭであり、以下ＤＩＭＭ３とも称する。該ＤＩＭＭ３は、１ランクに１８個のＳＤＲＡＭ３１を持ち、これを４ランク、即ち７２個のＳＤＲＡＭ３１（メモリ群）を備えている。
該ＤＩＭＭ３は、１ＧＢのＳＤＲＡＭ２つ分で従来の１ランク相当を構成しており、計８ＧＢの容量を有している。
【0024】
ＤＩＭＭ３上には記憶素子としてのＳＤＲＡＭ３１の他に、ＥＥＰＲＯＭ３２、デコード回路（デコード部）３３、ＯＤＴ（On Die Termination）３４、インターフェイス３５を備えている。
【0025】
ＥＥＰＲＯＭ３２は、高速伝送実現の為にアドレス信号線等をラッチするレジスタや、ＰＬＬ、ＤＩＭＭのスペックを記録している。
【0026】
ＯＤＴ３４は、ランク毎のＳＤＲＡＭに搭載されたバス終端抵抗であり、メモリコントローラ２からON/OFF制御される。
【0027】
インターフェイス部３５は、ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有するＤＤＲ２インターフェイスである。
【0028】
デコード回路３３は、前記インターフェイス部３５を介してアクセス命令を受信した際のアドレス信号線を介して得た信号及びランクセレクト信号に基づいて、増加したランク２，３を含めたランクセレクト信号を生成することにより、前記ランクを選択及びアクセスを可能にしている。
【0029】
システムコントローラ１は、ＣＰＵ及び各メモリコントローラ２を制御する。システムコントローラ１は各メモリコントローラ２に対し、同期制御とメモリアクセスを命令する。各メモリコントローラ２はシステムコントローラ１からの命令に従い、メモリアクセスを行う。

Micron Technology Inc. et al.
Ex. 1010, p. 307

【0030】
システムコントローラ１は、メモリリフレッシュコントロール回路１３を備えている。
【0031】
メモリコントローラ２は、制御部２１や、インターフェイス部２２、ＥＥＰＲＯＭアクセスコントロール回路２３を備えている。
【0032】
制御部２１は、前記アドレス信号線の一部及びセレクト信号線を介して前記ランクを選択する信号を送信するなど、ＤＩＭＭ３に対するアクセス制御を行う。
【0033】
インターフェイス部２２は、ランクを選択するためのセレクト信号を介するセレクト信号線及び選択されたランク上のアドレスを示すアドレス信号を介するアドレス信号線を有したＤＤＲ２インターフェイスである。
【0034】
各メモリコントローラ２は、図１と同様にＤＤＲ２バスインターフェイスを利用したＤＩＭＭバスチャンネルを４つ搭載し、各ＤＩＭＭバスチャンネルには２つのＤＩＭＭスロットがバス接続されている。
【0035】
図５にＤＩＭＭに搭載されるＳＤＲＡＭ３１のダイヤグラムを示す。
メモリコントローラ２から接続されるコマンド信号（／ＲＡＳ，／ＣＡＳ，／ＷＥ）と／ＣＳによりＳＤＲＡＭ３１のセレクト及び動作が決定される。
【0036】
アドレス選択用として接続される信号はＡ０～Ａ１４、ＢＡ０～ＢＡ２があり、ＢＡ０～ＢＡ２は合計８つのメモリバンクの中の１つを選択し、Ａ０～Ａ１４はバンク内の記憶セルの選択に使用される。
【0037】
記憶セルの選択には、Ａ０～Ａ１４を時分割した、ロウアドレスとカラムアドレスの２つのアドレスが使用さる。記憶セル内のデータは、４本のデータ信号（ＤＱ）で入出力される。ＤＱ信号は、データストローブ信号対（ＤＱＳ、／ＤＱＳ）によりラッチされる。
【0038】
ＤＩＭＭ３上に搭載されるＥＥＰＲＯＭ３２には、コマンドのタイミングやランクなどの構成情報が記録されている。メモリコントローラ２のＩ２Ｃアクセスコントロール回路２３は、Ｉ２Ｃと呼ばれるシリアル規格に従ってこの情報を読み出し、ＤＩＭＭ３をコントロールする為の初期化をファームウェアに従って行う。
【0039】
ＳＤＲＡＭ３１は、記憶保持の為のメモリリフレッシュ機能を備えている。メモリリフレッシュはメモリリードやメモリライトと同様に、メモリコントローラ２からの命令により行われる。
【0040】
メモリコントローラ２とＳＤＲＡＭ３１間は、ＤＤＲ２インターフェイスを利用したバスで構成されている。バス上にはスロット及びランクによりドット接続される。
【0041】
システムコントローラ１はＣＰＵからの物理アドレスをメモリアドレスに変更し、リード／ライトコマンドと共に各メモリコントローラ２に送出する。メモリライトの際には、ライトデータがメモリコントローラ２に送出される。
【0042】
上記のメモリアドレスとは、メモリインターフェイス上のＳＤＲＡＭ郡のランクを選択するチップセレクト信号（ＣＳ［３：０］）、ＳＤＲＡＭ内のＢＡＮＫを選択するバンクアドレス（ＢＡ［２：０］）、ＢＡＮＫ内の記憶セル郡を選択するロウアドレスと、ロウアドレスで選択された記憶セル郡からデータビットを選択するカラムアドレスで構成される。

Micron Technology Inc. et al.
Ex. 1010, p. 308

【0043】
メモリコントローラ２は初期化動作時にＩ２Ｃアクセスコントロール回路２３がＥＥＰＲＯＭ３２を参照し、搭載されたＤＩＭＭの構成やコマンドタイミングといったスペックを知ることができる。
【0044】
本実施形態の初期化動作の例を図３を用いて説明する。情報処理システム１０への電源投入後にファームウェアが起動され、ＤＩＭＭ３のＥＥＰＲＯＭ３２から情報の読み出しを行う。読み出された情報を元にメモリコントローラ２内のアーキテクチャレジスタをセットする。例えば、ＤＩＭＭ搭載情報によりメモリマップの決定、動作タイミングの決定、ランク数の決定を行う。また、構成情報を元にＤＩＭＭ３そのものの初期化や、ＤＩＭＭ３用の設定レジスタのセット、ＯＤＴ（On Die Termination）３４の抵抗値の設定を行う。
【0045】
本実施形態では、制御部２１が増加させたランク２，３用のレジスタセットを備えており、参照したランク数が４であった場合には、通常どおりランク０，１の初期化を行うのと共に、このレジスタセットを用いてランク２，３の初期化も行う。
【0046】
また、情報処理システム１０におけるメモリリフレッシュの例を図８に示す。前述した初期化動作により、システムコントローラ内のアーキテクチャレジスタにリフレッシュ間隔が設定される。この値はメモリコントローラ２に搭載されるＤＩＭＭ３のランク総数によって決定される。システムコントローラ１内のリフレッシュカウンターが、設定されたリフレッシュ間隔になるとメモリコントローラ２を介してリフレッシュ命令がＤＩＭＭ３に対して発行される。リフレッシュはランク毎に発行され、対象アドレスはＤＩＭＭ内で自動的に決定される。
【0047】
そして、メモリコントローラ２がシステムコントローラ１からリード／ライトコマンドを受信した場合、該メモリコントローラ２のシーケンサによってＤＩＭＭ３に対するリード／ライトを行い、データの送受信を行う。
【0048】
§３．メモリコントローラ２及びＤＩＭＭ３の特徴部分の説明
本実施形態のＤＩＭＭ３では、A0～A14の15本持っているアドレス信号線のうち、最上位の1本（A14)をチップセレクト信号の生成に用いる為、図６に示すようにデコード回路３３を備えている。
【0049】
デコーダ回路３３は、アドレス信号線A14からの信号と既存のチップセレクト信号線2本からのセレクト信号に基づき、図７の真理値表に示す４つの状態のチップセレクト信号を生成する。
【0050】
即ち、アドレス信号線Ａ１４の信号がアサートされておらず、チップセレクト信号CS0がアサートされた場合にはランク０が選択され、チップセレクト信号ＣＳ１がアサートされた場合にはランク１が選択される。また、アドレス信号線Ａ１４の信号がアサートされており、チップセレクト信号ＣＳ０がアサートされた場合にはランク２が選択され、チップセレクト信号ＣＳ１がアサートされた場合にはランク３が選択される。
【0051】
ここで使用されるアドレス信号線Ａ１４は、ＤＩＭＭ３に入力されるアドレスの最上位ビットを示すアドレス信号を伝送する為の既存の線であるので、メモリコントローラ２とＤＩＭＭ３の間の伝送路上に信号の増減が発生せず、ボードの再開発が必要ない。
【0052】
本実施形態のＤＩＭＭ３には１ＧＢのＳＤＲＡＭが搭載されており、アドレス信号線はＡ０～Ａ１３しか使用しないため、最上位ビットであるＡ１４をチップセレクトに使用し

Micron Technology Inc. et al.
Ex. 1010, p. 309

てもシステムとして問題ない。

【0053】

図９にＤＩＭＭ３とＤＤＲ２インターフェイスで規定されているメモリ空間イメージの比較を示す。

空間イメージ４１は、ＤＤＲ２インターフェイスで規定されているもので、２ＧＢのＳＤＲＡＭを搭載した２ランク構成とすることで、最大８ＧＢの容量が得られる。

【0054】

これに対し、本実施形態のＤＩＭＭ３は、１ＧＢのＳＤＲＡＭを搭載した４ランク構成で８ＧＢの容量を得ている。このため、ロウアドレスの表現にＡ０～Ａ１３、カラムアドレスの表現にＡ０～Ａ９、Ａ１１を使用しており、Ａ１４は未使用である。

【0055】

従って、メモリコントローラ２の制御部２１は、システムコントローラ１からＤＩＭＭ３に対するリード／ライトコマンドを受信した場合、選択するランクに応じて前記真理値表のとおりＡ１４，ＣＳ０，ＣＳ１に信号をアサートする。

【0056】

なお制御部２１は、初期化時にＥＥＰＲＯＭ３２から読み出した情報に従い、アクセスするメモリモジュールが、本実施形態の４ランク構成であればＡ１４，ＣＳ０，ＣＳ１を用いてランク選択の為の信号をアサートし、２ランク構成であれば規定どおりＣＳ０，ＣＳ１を用いてランク選択の為の信号をアサートしてＡ１４をアドレス信号線として使用する。

【0057】

これにより、本実施形態のシステムボードは、ＤＤＲ２インターフェイスに従う２ＧＢのＳＤＲＡＭを搭載した２ランクのＤＩＭＭとの互換性を維持しつつ、１ＧＢのＳＤＲＡＭを搭載した４ランクのＤＩＭＭ３も利用可能にしている。

【0058】

なお、本実施形態では、既存の２ランク構成よりもランク数を増加させているため、図１０（Ａ）のようにランク０，１に対するメモリリフレッシュのフローを図１０（Ｂ）のように２ＦＩｏｗにして行う。

【0059】

メモリコントローラ２内の制御部２１は、増加させたランク数に応じたシーケンサを備え、ランクを増加させたＤＩＭＭ３に対するリフレッシュコマンドを受けた場合、メモリリフレッシュの処理をランク数分増加させて各ランクのメモリリフレッシュを行わせる。

【0060】

例えば、図１１に示すパターン１のように、ランク０，１に対して、ＰＡＬＬからｔＲＡＳを待ってＲＥＦまでの処理のフロー１を行った後に同様のフロー２をランク２，３に行う。しかし、この場合、メモリリフレッシュにかかる時間がランク数に応じて単純に増加してしまい、リード/ライトのアクセスを圧迫してしまう可能性がある。

【0061】

このため、パターン２のように、ランク０，１に対するフロー１のＰＡＬＬの後、ランク２，３に対するフロー２のＰＡＬＬを行い、メモリリフレッシュ処理を並行して行うのが望ましい。

【0062】

また、本実施形態のように、ランク数を増加させた場合、スタブの増加も考えられるが、制御部２１が、ランク数に応じて、所定ランクのＯＤＴをオン又はオフに制御することで抑えることができる。また、このＯＤＴのオン／オフ制御は、リード／ライト等の命令の種別に応じて行っても良い。このとき、何れのランクのＯＤＴをオンにし、何れのランクのＯＤＴをオフにするかは、ＥＥＰＲＯＭ３２の内容で判断するかファームウェアに予め設定しておけば良い。

【0063】

更に、増設ランクのデコード回路３３の為に回路段数が増え、伝送特性が悪化して伝送

Micron Technology Inc. et al.
Ex. 1010, p. 310

速度の高速化の妨げになることが考えられるが、本実施形態では、図１２に示すようにデコード回路３３をレジスタで挟みこみ、この伝送特性を改善している。

【0064】

このように、レジスタ４３の直後にデコード回路３３を配置し、デコード回路３３の直後にレジスタ４４を配置したことにより、デコード回路３３の直前で信号のタイミングが揃い、デコード回路３３の直後でも信号のタイミングが揃うので、伝送特性の悪化が防止できる。

【0065】

また、ＤＩＭＭ３は、ランクを増加したことによってバスファイト（Bus Fight）が増える可能性がある。

【0066】

メモリコントローラ２の制御部２１は、同一ランク内の前半部分と後半部分のアクセスが連続した場合、図１３に示すパターン１のように、連続してアクセスを行う。しかし、異なるランクに対するアクセスが連続した場合、図１３に示すパターン２のように、該アクセス間にウェイトを挿入する。

【0067】

即ち、制御部２１は、増設したランク２，３も含め各ランクへのアクセスを排他的に管理し、異なるランクに対するアクセスの間には、ウェイトを挿入する。

【0068】

以上のように本実施形態では、既存信号線を用いてランク選択の為の信号を送信するので、ランク数を増やしても従来のインターフェイスを用いることが可能となった。

【0069】

従って、ＳＤＲＡＭの容量増加が望めない場合であっても、ランクを増やすことでＤＩＭＭの容量を増やすことが出来るようになった。例えば、現在入手が困難な２ＧＢのＳＤＲＡＭを用いなくても１ＧＢのＳＤＲＡＭで高容量のＤＩＭＭを達成できる。このため、普及している安価なＳＤＲＡＭを選択可能で、高容量のＤＩＭＭを安価に製造しやすい。

【0070】

また、増やしたランクに対する初期化やリフレッシュの制御が可能となった。

ＤＩＭＭに渡されるアドレス自体はそのまま使用できるので、ＣＰＵからの物理アドレスの制御に変更を加えなくともＤＩＭＭの容量を増やすことが可能となった。

【0071】

更に、情報処理システムに搭載するＤＩＭＭの数を増やさずともメモリ総容量を増やすことが出来るようになる為、システムの開発コストを軽減できる。

【図面の簡単な説明】

【0072】

【図1】関連技術の概略図
【図2】関連技術のメモリ構成を示す図
【図3】情報処理装置の全体図
【図4】本実施形態のメモリ構成を示す図
【図5】ＳＤＲＡＭの説明図
【図6】デコード部の説明図
【図7】ランク選択のための真理値表
【図8】メモリリフレッシュ処理の説明図
【図9】メモリ空間イメージを示す図
【図10】メモリリフレッシュ処理の説明図
【図11】メモリリフレッシュ処理の説明図
【図12】デコード回路の説明図
【図13】バスファイト防止の説明図

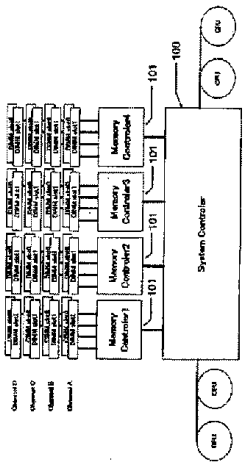
【符号の説明】

【0073】

Micron Technology Inc. et al.
Ex. 1010, p. 311

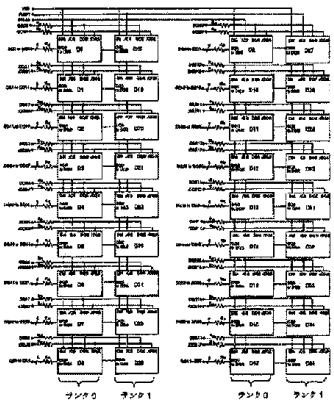
１０　情報処理システム
１　システムコントローラ
１３　メモリリフレッシュコントロール回路
２　メモリコントローラ
２１　制御部
２２　インターフェイス部
２３　Ｉ２Ｃアクセスコントロール回路
３　メモリモジュール（ＤＩＭＭ）
３１　ＳＤＲＡＭ
３２　ＥＥＰＲＯＭ
３３　デコード回路（デコード部）
３４　ＯＤＴ（On Die Termination）
３５　インターフェイス

【図1】



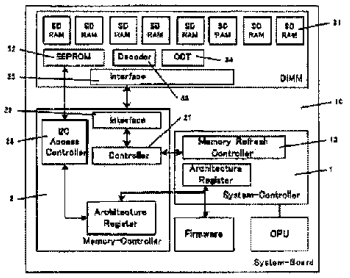
関連技術

【図2】



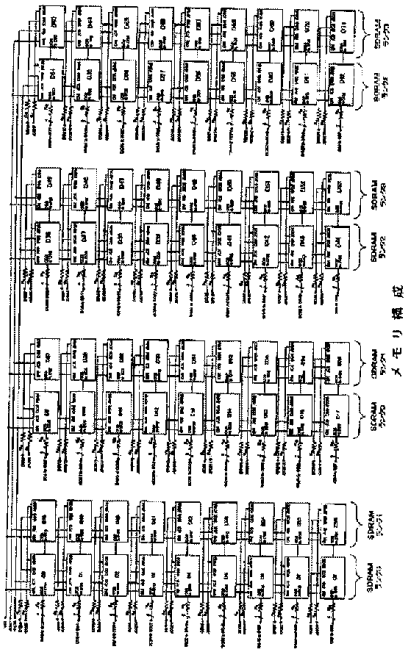
メモリ構成

【図3】



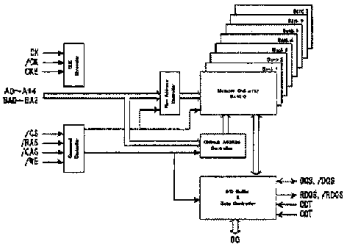
システム全体図

【図4】



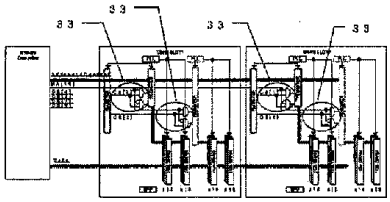
メモリ構成

【図5】



SDRAMの説明図

【図6】



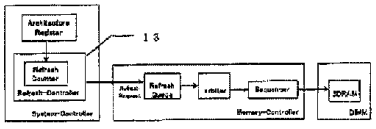
デコード部の説明図

【図7】

	A14	CS1	CS0
RANK0	0	0	1
RANK1	0	1	0
RANK2	1	0	1
RANK3	1	1	0

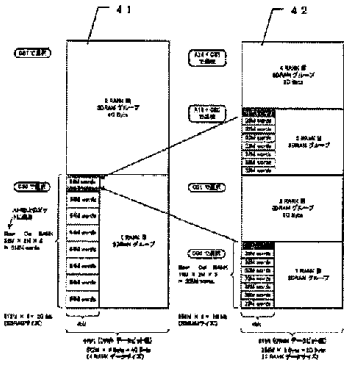
真理値表

【図8】



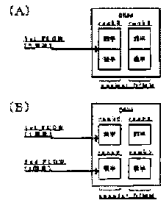
メモリリフレッシュ処理の説明図

【図9】



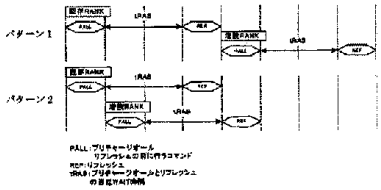
メモリ空間イメージを示す図

【図10】



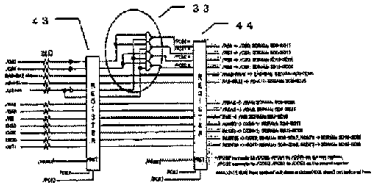
メモリリフレッシュ処理の説明図

【図11】



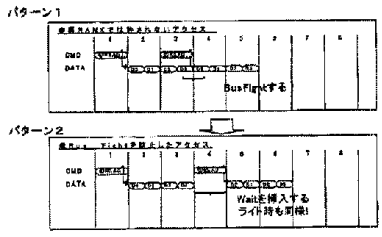
メモリリフレッシュ処理の説明図

【図12】



デコード回路の説明図

【図13】



バスファイト防止の説明図

(72)発明者　本田　育史
神奈川県川崎市中原区上小田中４丁目１番１号　富士通株式会社内
(72)発明者　鈴木　賢司
神奈川県川崎市中原区上小田中４丁目１番１号　富士通株式会社内
Ｆターム(参考) 5B060 AC01　AC11　CA11　CC06　MB09　MH16　MM17
5M024 AA54　EE02　JJ03　LL20　PP01　PP07　PP10

Micron Technology Inc. et al.
Ex. 1010, p. 315

PATENT ABSTRACTS OF JAPAN

(11)Publication number : **10-092169**
(43)Date of publication of application : **10.04.1998**

(51)Int.Cl. **G11C 5/00**

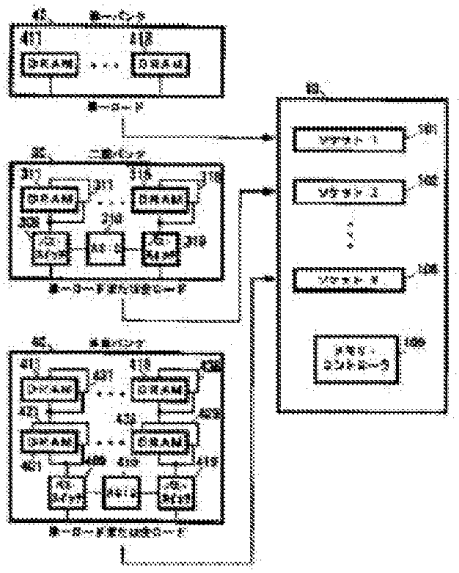
(21)Application number : **09-181302**
(22)Date of filing : **07.07.1997**

(71)Applicant : **INTERNATL BUSINESS MACH CORP <IBM>**
(72)Inventor : **BRIAN J CONNOLLY**
MARK W KELLOGG
BRUCE G HEEZERUZETSUTO

(30)Priority
Priority number : **96 676609** Priority date : **08.07.1996** Priority country : **US**

(54) MEMORY MODULE

(57)Abstract:
PROBLEM TO BE SOLVED: To minimize a data line load by providing a bus switch between a system and a RAM and then comprising a logic circuit to produce a signal for enabling this switch into an ASIC.
SOLUTION: DRAM chips 301 to 304 and DRAM chips 311 to 314 are connected to a bus switch 309, while DRAM chips 305 to 308 and DRAM chips 315 to 318 are connected to a bus switch 319. A couple of bus switches are connected to ASIC 310 which is connected to a RAS pin and a CAS pin. A data bus load can be minimized by forming a bank DIMM 30 or DIMM 40 by utilizing bus switches 309, 319 or bus switches 409, 419.



(19)日本国特許庁（ＪＰ）　(12) **公開特許公報（Ａ）**　(11)特許出願公開番号

特開平10－92169

(43)公開日　平成10年(1998)４月10日

(51)Int.Cl.[6]	識別記号	Ｆ Ｉ	
Ｇ１１Ｃ　5/00	３０３	Ｇ１１Ｃ　5/00	３０３Ａ

審査請求　未請求　請求項の数４　ＯＬ　（全 ７ 頁）

(21)出願番号　特願平9－181302

(22)出願日　平成９年(1997)７月７日

(31)優先権主張番号　０８／６７６６０９

(32)優先日　1996年７月８日

(33)優先権主張国　米国（ＵＳ）

(71)出願人　390009531
インターナショナル・ビジネス・マシーンズ・コーポレイション
ＩＮＴＥＲＮＡＴＩＯＮＡＬ　ＢＵＳＩＮＥＳＳ　ＭＡＳＣＨＩＮＥＳ　ＣＯＲＰＯＲＡＴＩＯＮ
アメリカ合衆国10504、ニューヨーク州　アーモンク　（番地なし）

(72)発明者　ブライアン・ジェイ・コノリー
アメリカ合衆国05495　バーモント州ウィリストンヒルクレスト・レーン　３

(74)代理人　弁理士　坂口　博　（外１名）

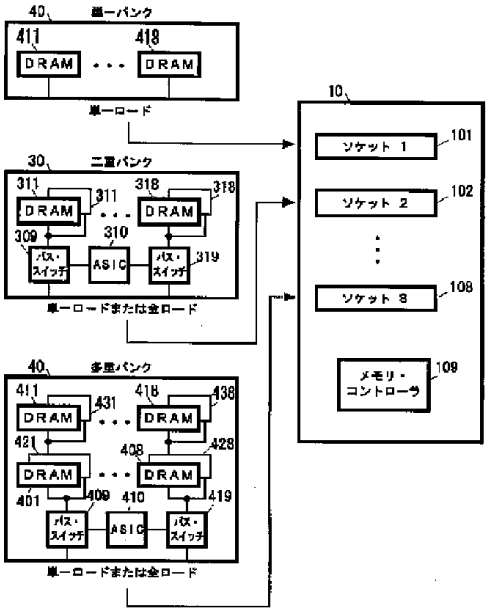
最終頁に続く

(54)【発明の名称】　メモリ・モジュール

(57)【要約】

【課題】　複数のＤＲＡＭを有する高密度モジュール上のデータ線負荷を最小にし、データ線容量性負荷に起因する性能低下なしに、システムの最大のメモリ密度を可能にすること。

【解決手段】　システム・ボード１０上のコネクタ１０１～１０８と相互接続されるタブを有するプリント回路カード２０～４０と、プリント回路カードに取り付けられた複数のバンクのランダム・アクセス・メモリ・デバイス３０１～３１８、４０１～４３８と、選択されたランダム・アクセス・メモリ・デバイスと接続された、イネーブル信号に応答する１つまたは複数のバス・スイッチ３０９、３１９、４０９、４１９と、システムの行アドレス・ストローブ信号および列アドレス・ストローブ信号に応答して、１つまたは複数のバス・スイッチのイネーブル信号を生成する、論理手段とを含む。



Micron Technology Inc. et al.
Ex. 1010, p. 317

【特許請求の範囲】
【請求項１】システム・ボード上のコネクタと相互接続される接点タブを有し、カード・データ・バスを有するプリント回路カードと、
前記プリント回路カードに取り付けられ、前記カード・データ・バスに接続された、複数バンクのランダム・アクセス・メモリ・デバイスと、
前記複数バンクのランダム・アクセス・メモリ・デバイスのうちの選択されたものが前記カード・データ・バスを介して接続された、イネーブル信号に応答する、１つまたは複数のバス・スイッチと、
システムの行アドレス・ストローブ信号および列アドレス・ストローブ信号に応答して、前記１つまたは複数のスイッチの前記イネーブル信号を生成する、論理手段とを含む、メモリ・モジュール。
【請求項２】前記１つまたは複数のバス・スイッチおよび前記論理手段が、前記プリント回路カードに取り付けられることを特徴とする、請求項１に記載の高密度メモリ・モジュール。
【請求項３】第２のプリント回路カードを含み、前記第１および第２のプリント回路カードが異なる数のメモリ・バンクを有することを特徴とする、請求項１に記載の高密度メモリ・モジュール。
【請求項４】前記１つまたは複数のバス・スイッチのそれぞれが、前記カード・データ・バスの信号線１本につき１つの複数の電界効果トランジスタ（ＦＥＴ）と、前記イネーブル信号に応答して前記ＦＥＴをバイアスするための手段とを含むことを特徴とする、請求項１に記載の高密度メモリ・モジュール。
【発明の詳細な説明】
【０００１】
【発明の属する技術分野】本発明は、全般的にはコンピュータ用の高密度メモリ・モジュールに関し、具体的には、最小限のメモリ・バス負荷を有する多重バンク・メモリ・モジュールに関する。
【０００２】
【従来の技術】ＳＩＭＭ（シングル・インライン・メモリ・モジュール）やＤＩＭＭ（デュアル・インライン・メモリ・モジュール）などの高密度メモリ・モジュールは、多くの応用分野で追加のシステム性能を達成できるので、ハイ・エンド・パーソナル・コンピュータ（ＰＣ）、ネットワーク・サーバおよびワークステーションで需要がある。しかし、最大システム密度は、下記の１つまたは複数によって不自然に制限されることがしばしばである。
・システムが有するメモリ・モジュール「スロット」の数が限られている。
・システムが有するメモリ「バンク」の数が（メモリ・コントローラからの選択信号線の不足が原因で）限られている。
・高密度メモリ・チップ（たとえば６４メガビット（Ｍｂ）、２５６Ｍｂおよびそれ以上）は、非常に高価であるか、簡単に入手できないか、その両方である。
・高密度メモリ・チップは、システムのメモリ・インターフェース電圧より低い動作電圧を有する場合がある。
・標準ランダム・アクセス・メモリ・デバイス（ＲＡＭ）（ダイナミックＲＡＭ（ＤＲＡＭ）、シンクロナスＤＲＡＭ（ＳＤＲＡＭ）およびスタティックＲＡＭ（ＳＲＡＭ）を含む）の積層化／立体化または大型モジュールの使用は、システム限界をはるかに超えるデータ線容量をもたらす。以前の解決は、最後の問題以外の上記の問題のすべてについて提供されてきたが、これらの解決は、データ線負荷の問題が原因で応用性が制限されている。メモリ拡張を制限されているシステムで安価なＲＡＭチップをよりよく利用するために必要なものは、データ線の容量性負荷を最小にし、その結果、ＲＡＭチップのバンクを有する大型メモリ・モジュールをシステムに追加できるようにする方法である。
【０００３】
【発明が解決しようとする課題】したがって、本発明の目的は、データ線ごとに複数のランダム・アクセス・メモリ・デバイスを有する高密度モジュールでのデータ線負荷を最小にするための解決を提供することである。
【０００４】本発明のもう１つの目的は、データ線の容量性負荷に起因する性能低下なしに、密度を制限されているシステムのメモリ密度を最大にすることである。
【０００５】
【課題を解決するための手段】本発明によれば、データ線容量を許容可能なシステム限界まで減らす解決法には、２つの部分がある。第１の部分には、インライン・バス・スイッチを有するメモリ・モジュールの設計が含まれる。バス・スイッチは、モジュール・タブ（システム）とランダム・アクセス・メモリ・デバイス（ＲＡＭ）の間に置かれ、高インピーダンス（オフ）またはアクティブ（活動）状態のいずれかになる。高インピーダンス状態の時には、メモリ・モジュールの実効負荷は、ビット・スイッチ・デバイスの負荷になる。アクティブ（たとえば、リード／ライト・サイクル）の時には、ＲＡＭ負荷の他に、最小限のキャパシタンス／抵抗がメモリ・バスに追加される。一時に１つのメモリ・モジュールだけがアクティブになる。
【０００６】本発明による解決の第２の部分は、バス・スイッチをイネーブルする信号を生成する論理回路を特定用途向け集積回路（ＡＳＩＣ）に組み込むことである。バス・スイッチは、メモリ・モジュールへの行アドレス・ストローブ（ＲＡＳ）選択線の立ち下がりエッジでアクティブになり、ＲＡＳ選択線または列アドレス・ストローブ（ＣＡＳ）選択線のうちの最後の線がインアクティブ（非活動）状態になるまで活動状態にとどまり、これによって、高速ページ・モード（ＦＰＭ）動作

Micron Technology Inc. et al.
Ex. 1010, p. 318

とＥＤＯ（拡張データ出力）動作の両方がサポートされる。ＡＳＩＣの回路は、システムのＲＡＳ選択線とＣＡＳ選択線をデコードし、信号を駆動してバス・スイッチをイネーブルすることによって、この作業を実行する。論理回路は、リード／ライト・サイクル、ＲＯＲ（RAS only refresh）サイクル、ＣＢＲ（CAS before RAS refresh）サイクルおよびヒドン・リフレッシュ（hidden refresh）サイクルを区別し、それ相応にバス・スイッチをアクティブにしなければならない。ＣＢＲサイクルが発生する時には、バス・スイッチはインアクティブのままになる。ヒドン・リフレッシュ・サイクルでは、バス・スイッチはアクティブにされ、リフレッシュ・サイクルが完了するまでアクティブのままになる。これは、リード動作またはライト動作が存在する可能性があるので必要である。

【０００７】

【発明の実施の形態】本発明を、ＤＩＭＭを使用する特定の実施例に関して説明するが、当業者であれば、ＳＩＭＭなどの他のメモリ・モジュールを本発明の実施に使用できることを諒解するであろう。このようなモジュールは、通常は、コンピュータ・システム基板上のコネクタに挿入するための接点タブを有するプリント回路カードとして実施される。

【０００８】ここで図面、具体的には図１を参照すると、８つのメモリ・モジュール・ソケット１０１～１０８とメモリ・コントローラ１０９を有する通常のシステム・ボード１０が示されている。たとえば８つの４MbのＤＲＡＭチップ２０１～２０８を有する通常の単一バンクＤＩＭＭ　２０は、通常はメモリ・コントローラ１０９に最も近いメモリ・モジュール・ソケット１０１から始めて、メモリ・モジュール・ソケット１０１～１０８のすべてに単一バンクＤＩＭＭ　２０が挿入されるまで順番に、メモリ・モジュール・ソケット１０１～１０８のうちの１つに挿入することができる。したがって、この例では、システム・ボードは、４メガバイト（ＭＢ）の最小構成と３２ＭＢの最大構成を有する。

【０００９】本発明によれば、二重バンクＤＩＭＭ　３０または多重バンクＤＩＭＭ　４０によって、単一バンクＤＩＭＭ　２０を置換することができる。二重バンクＤＩＭＭ　３０の場合、バス・スイッチ３０９および３１９に接続された、８つのＤＲＡＭチップ３０１～３０８の第１のグループと、８つのＤＲＡＭチップ３１１～３１８の第２のグループがある。ＤＲＡＭチップ３０１～３０４およびＤＲＡＭチップ３１１～３１４はバス・スイッチ３０９に接続され、ＤＲＡＭチップ３０５～３０８およびＤＲＡＭチップ３１５～３１８はバス・スイッチ３１９に接続される。２つのバス・スイッチは、ＡＳＩＣ　３１０に接続され、ＡＳＩＣ３１０は、このモジュールのＲＡＳピンおよびＣＡＳピンに接続される。この概念は、多重バンクＤＩＭＭ　４０の場合に拡張され、この例では、ＤＲＡＭチップの４つのグループすなわち、ＤＲＡＭチップ４０１～４０８、ＤＲＡＭチップ４１１～４１８、ＤＲＡＭチップ４２１～４２８およびＤＲＡＭチップ４３１～４３８がある。ＤＲＡＭチップ４０１～４０４、ＤＲＡＭチップ４１１～４１４、ＤＲＡＭチップ４２１～４２４およびＤＲＡＭチップ４３１～４３４は、バス・スイッチ４０９に接続され、ＤＲＡＭチップ４０５～４０８、ＤＲＡＭチップ４１５～４１８、ＤＲＡＭチップ４２５～４２８およびＤＲＡＭチップ４３５～４３８は、バス・スイッチ４１９に接続される。二重バンクＤＩＭＭ　３０の場合と同様に、バス・スイッチ４０９および４１９は、ＡＳＩＣ　４１０に接続され、ＡＳＩＣ　４１０は、このモジュールのＲＡＳピンおよびＣＡＳピンに接続される。

【００１０】二重バンクＤＩＭＭ　３０を使用する場合、システムの最小構成は８ＭＢに増加し、最大構成は６４ＭＢに増加する。多重バンクＤＩＭＭ　４０を使用すると、図示の４バンクのＤＲＡＭチップの例では、最小構成は１６ＭＢ、最大構成は１２８ＭＢに増加する。より高密度のデバイス（たとえば１６Ｍｂ、６４Ｍｂなど）を使用すると、最大メモリ密度のかなりの増加がもたらされるはずである。これは、下の表からわかるように、最小限のデータ・バス負荷で達成される。

【００１１】

【表１】

データ・バス負荷の比較＊（最大値）

	１バンク	２バンク	４バンク
従来技術	８	１６	３２
本発明	８	９	１１

＊ＤＩＭＭ当たりのＤＲＡＭ等価負荷の数として表す（通常は７ｐＦ／ＤＲＡＭ）

【００１２】これは、二重バンクＤＩＭＭ　３０の場合にはバス・スイッチ３０９および３１９、多重バンクＤＩＭＭ　４０の場合にはバス・スイッチ４０９および４１９を使用することによって達成される。ＡＳＩＣ　３

Micron Technology Inc. et al.
Ex. 1010, p. 319

１０および４１０は、バス・スイッチとシステム・ボードの間のインターフェースを提供する。

【００１３】ここで図２に移ると、バス・スイッチの一般化された論理図が示されている。バス・スイッチには、信号線ごとに１つずつ、複数のＦＥＴが含まれる。１つまたは複数のインバータ５０が、ＦＥＴをオンにバイアスするのに使用される。インバータ５０への入力ＢＥＮ（ＮはＢＥの反転を表す）は、バス・スイッチのイネーブル入力である。イネーブル入力がロウになると、インバータの出力はハイになり、その結果、正バイアスがＦＥＴのゲートに印加され、ＦＥＴが導通する。このバス・スイッチを介する遅延は、１ナノ秒（ｎｓ）未満であり、これによって、正しいメモリ動作が保証される。

【００１４】図３のブロック図からわかるように、ＡＳＩＣ　６０（図１のＡＳＩＣ　３１０および４１０に対応する）は、システムのＲＡＳ信号（ＳＹＳ＿ＲＡＳ）とＣＡＳ信号（ＳＹＳ＿ＣＡＳ）を受け取り、バス・スイッチ６１および６２のイネーブル入力へのＲＣ　ＳＥＬＥＣＴ信号を生成する。バス・スイッチ６１および６２は、カード・データ・バス６３とプリント回路カードの接点タブ６４の間のインターフェースを提供する。接点タブ６４はシステム・ボードのソケット１～８に挿入されるモジュール２０、３０、４０のプリント回路カードのエッジに設けられている。

【００１５】ＲＣ＿ＳＥＬＥＣＴ信号は、アクティブ・ロウの信号であり、独自の機能を有する。通常のメモリ・リード・サイクルおよびメモリ・ライト・サイクルの間、ＲＣ＿ＳＥＬＥＣＴ信号は、図４に示されるように、システムのＲＡＳ選択線のアクティブ（立ち下がり）エッジでアクティブにされ、システムのＲＡＳ選択線とＣＡＳ選択線の両方がインアクティブになるまでアクティブのままになる。

【００１６】ヒドン・リフレッシュ・サイクル中は、図５からわかるように、ヒドン・リフレッシュ・サイクルの性質が原因で、ＲＣ＿ＳＥＬＥＣＴ信号はアクティブになる。ヒドン・リフレッシュ・サイクルの第１の部分は、ＲＡＳが立ち上がってから立ち下がる（ＣＢＲリフレッシュ・サイクルの開始）までシステムのＣＡＳ選択線がアクティブのままになるメモリ・サイクルである。ＲＣ＿ＳＥＬＥＣＴ論理回路は、ＣＡＳ外乱をフィルタリングし、インアクティブのままになる。

【００１７】ＣＢＲリフレッシュ中は、図６からわかるように、システムのＣＡＳ選択線が、システムのＲＡＳ選択線がアクティブになるより前にアクティブになる。ＲＣ＿ＳＥＬＥＣＴ信号は、ＣＢＲサイクル中はインアクティブのままになる。

【００１８】ＲＯＲサイクル中は、図７からわかるように、システムのＲＡＳ選択線がアクティブになり、ＲＣ＿ＳＥＬＥＣＴ信号上に駆動される。バス・スイッチがアクティブになっても、データがその時点で駆動されていないので、システム動作には影響しない。

【００１９】所与の応用分野で必要であれば、追加回路を追加して、ＲＯＲサイクル中にバス・スイッチがアクティブにならないようにすることができる。

【００２０】図８のＡＳＩＣの内部ロジックには、システムのＣＡＳ選択線ＣＡＳ＿Ｎ（Ｎは反転を表わす）のアクティブ（すなわち立ち下がり）エッジでシステムＲＡＳ選択線（ＡＮＹ＿ＲＡＳ）の状態をラッチするＤ型のフリップフロップ７０が含まれる。ＡＮＹ＿ＲＡＳ信号は、ＡＮＤゲート７１によって生成され、バッファ７２を介してフリップフロップ７０のＤ入力に供給されるが、ＡＮＹ＿ＣＡＳ信号は、ＡＮＤゲート７３によって生成され、インバータ７４を介してフリップフロップ７０のクロック入力に供給される。フリップフロップ７０の出力は、インバータ７６を介して２入力ＮＡＮＤゲート７５に接続され、マルチプレクサ７７の入力「０」にも接続される。インバータ７６の出力は、信号ＣＢＲ＿ＢＩＴ　Ｎである。信号ＡＮＹ　ＲＡＳは、インバータ７８によって反転され、インバータ７８の出力は、２入力ＮＡＮＤゲート７５のもう一方の入力に接続される。２入力ＮＡＮＤゲート７５の出力は、信号ＲＣ＿ＡＣＴＩＶＥであり、マルチプレクサ７７の入力「１」に接続される。フリップフロップ７０は、システムのＲＡＳ選択線とＣＡＳ選択線がインアクティブであり、フリップフロップ７０の「Ｑ」出力が論理「１」である時に、ＮＡＮＤゲート７９の出力によってリセットされる。マルチプレクサ７７の選択線は、４つのゲート（インバータ８０と３つのバッファ８１～８３）によって遅延されたシステムのＣＡＳ選択線ＤＥＬＹ＿ＣＡＳによって選択される。これは、フリップフロップ７０が信号ＡＮＹ＿ＲＡＳの状態をラッチする時間を与えるために行われる。

【００２１】各サイクルの初期の開始時に、マルチプレクサ７７は、「１」入力位置にセットされる。これによって、システムのＲＡＳ選択線をすばやく流すことができるようになる。システムのＣＡＳ選択線がアクティブになった時には、マルチプレクサ７７は、「０」入力位置に切り替わり、これによって、フリップフロップ７０の出力が選択される。

【００２２】図１に戻って、本発明の実装は、ＤＲＡＭチップのバンク群を１つのプリント回路カードに取り付け、ＡＳＩＣとバス・スイッチを別のキャリアに取り付けることによって、変更することができる。たとえば、ＡＳＩＣとバス・スイッチは、システム・ボード１０に直接に取り付けることができるが、これにはシステム・ボードの変更が必要である。代替案は、ＤＲＡＭカードを受けるコネクタを備えた別のプリント回路カードにＡＳＩＣとバス・スイッチを取り付けることである。どちらの代替案でも、ＤＲＡＭチップだけでＤＲＡＭカード

Micron Technology Inc. et al.
Ex. 1010, p. 320

を製造することができる。

【図面の簡単な説明】

【図１】システム・メモリ・ボードと、従来の単一バンクＤＩＭＭと本発明による多重バンクＤＩＭＭの両方を示すブロック図である。

【図２】図１に示された本発明の実施例に使用されるバス・スイッチを一般化した論理ブロック図である。

【図３】図１の実施例でのＡＳＣＩとバス・スイッチの関係を示すブロック図である。

【図４】通常のリード／ライト・サイクルの場合の、ＲＣ＿ＳＥＬＥＣＴ出力信号を生成するＡＳＩＣ論理回路の動作を示すタイミング図である。

【図５】ヒドン・リフレッシュの場合の、ＲＣ＿ＳＥＬＥＣＴ出力信号を生成するＡＳＩＣ論理回路の動作を示すタイミング図である。

【図６】ＣＢＲの場合の、ＲＣ＿ＳＥＬＥＣＴ出力信号を生成するＡＳＩＣ論理回路の動作を示すタイミング図である。

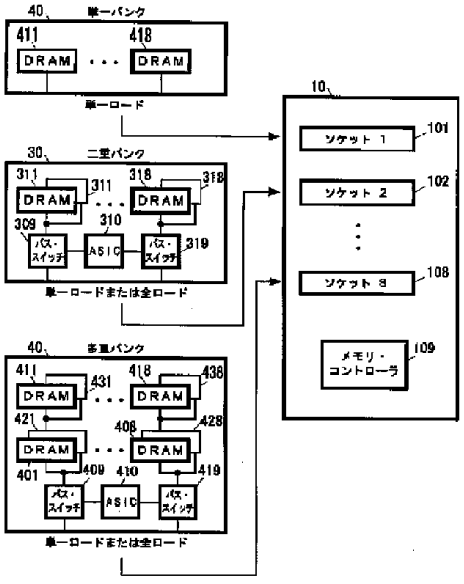
【図７】ＲＯＲの場合の、ＲＣ　ＳＥＬＥＣＴ出力信号を生成するＡＳＩＣ論理回路の動作を示すタイミング図である。

【図８】本発明の好ましい実施例によるＡＳＣＩの論理回路を示す論理ブロック図である。

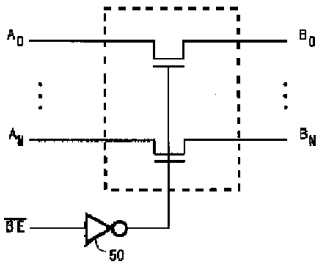
【符号の説明】

６０　ＡＳＩＣ
６１　バス・スイッチ
６２　バス・スイッチ
６３　カード・データ・バス
６４　タブ
７２　バッファ
７３　ＡＮＤゲート
７４　インバータ
７５　２入力ＮＡＮＤゲート
７６　インバータ
７７　マルチプレクサ
７８　インバータ
７９　ＮＡＮＤゲート
８０　インバータ
８１　バッファ
８２　バッファ
８３　バッファ

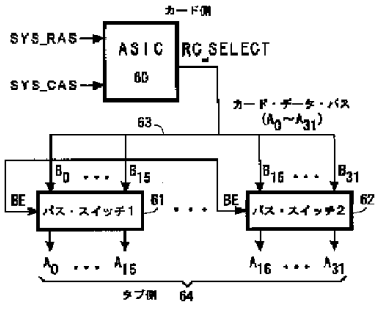
【図１】



【図２】

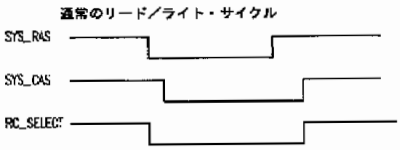


【図３】

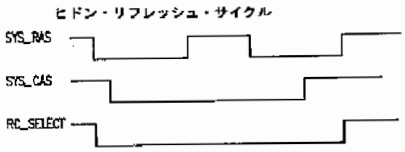


Micron Technology Inc. et al.
Ex. 1010, p. 321

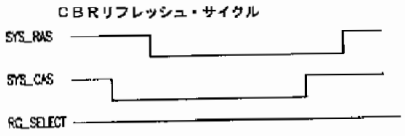
【図４】



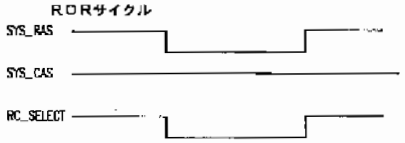
【図５】



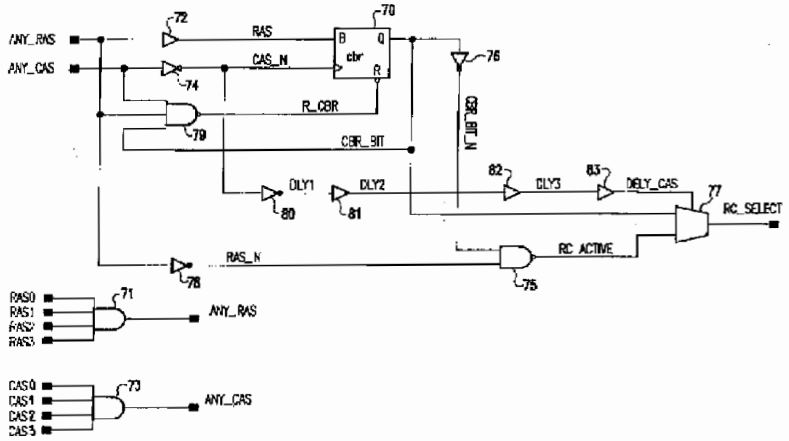
【図６】



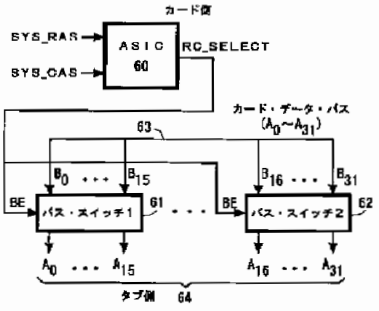
【図７】



【図８】



【手続補正書】
【提出日】平成９年７月２３日
【手続補正１】
【補正対象書類名】図面
【補正対象項目名】図３
【補正方法】変更
【補正内容】
【図３】



フロントページの続き

(72)発明者　マーク・ダブリュー・ケロッグ
アメリカ合衆国05452　バーモント州エセックス・ジャンクション　コーデュロイロード　29

(72)発明者　ブルース・ジー・ヘーゼルゼット
アメリカ合衆国05452　バーモント州エセックス・ジャンクション　グリーンフィールド・コート　8

Micron Technology Inc. et al.
Ex. 1010, p. 323

(19) JAPANESE PATENT OFFICE (JP)
(12) Official Gazette for Unexamined Patent Publications (A)

(11) Japanese Patent Application Publication No.: H10-320270
(43) Disclosure Date: December 4, 1998

(51) Int. Cl.[6]:	ID Symbols:	FI:	
G06F 12/06	515	G06F 12/06	515H

Request for Substantive Examination: Not yet submitted
Number of Claims: 3 OL
(Total of 5 pages [in the Japanese text])

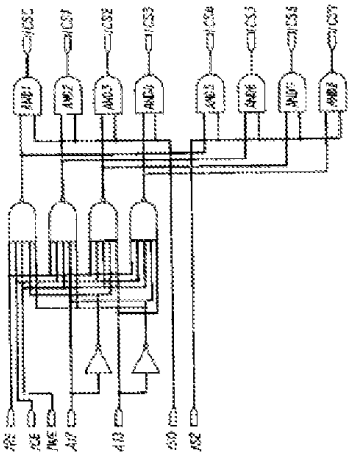
(21) Patent Application No.: H9-125328
(22) Filing Date: May 15, 1997
(71) Applicant: 000005821
Matsushita Electric Industrial Co. Ltd., 1006 Oaza-Kadoma, Kadoma-shi, Osaka-fu
(72) Inventor: Kazuo TAKEDA
c/o Matsushita Electric Industrial Co. Ltd., 1006 Oaza-Kadoma, Kadoma-shi, Osaka-fu
(74) Representative: Tomoyuki TAKIMOTO, Patent Attorney (and one other)

(54) [Title of the Invention] Memory module

(57) [Abstract]

[Problem] The aim is to provide a memory module which achieves lower costs, reduced power consumption and low temperature rises using current-generation SDRAMs, by resolving a conventional problem, namely that if the capacity of an SDRAM memory module is increased, there is excessively large power consumption or an excessively large rise in temperature with current-generation ×1 configuration SDRAMs unless next-generation SDRAMs are present.

[Means of Resolution] It is possible to reduce power consumption and produce low temperature rises which cannot be achieved using only next-generation SDRAMs, by adopting a configuration composed of a plurality of banks comprising a plurality of current-generation SDRAMS, and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks.



[Patent Claims]

[Claim 1] A memory module composed of: a plurality of banks comprising a plurality of current-generation SDRAMS for achieving increased memory capacity with current-generation SDRAMs, which cannot be achieved using only next-generation SDRAMs; and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks.

[Claim 2] The memory module as claimed in claim 1, composed of: a plurality of banks which are installed using an installation method enabling a plurality of current-generation SDRAMs to be installed in multiple stages, such as TCP (tape carrier packaging); and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks.

[Claim 3] The memory module as claimed in claim 1 or 2, in which the bank control unit is a one-chip unit.

[Detailed Description of the Invention]

[0001]

[Technical Field of the Invention] The present invention relates to a memory module which is used in a computer or the like.

[0002]

[Prior Art] Known configurations for memory modules which are used in computers and the like have been disclosed in JEDEC documents in the past. Figure 4 shows a conventional memory module structure, which is a 128 MB SDRAM module comprising sixteen 64 Mbit, ×4 configuration SDRAMs.

[0003]

[Problem to be Resolved by the Invention] However, there are problems with memory modules having the abovementioned conventional configuration in that it is necessary to use 64 Mbit SDRAMs (referred to here as "next-generation" components), and compared with configurations employing 16 Mbit SDRAMs (referred to here as "current-generation" components), configurations employing next-generation components are more costly until the value of SDRAMs is bit crossed. There are further problems in that when use is made of ×1 bit configuration current-generation components, the number of SDRAMs driven at the same time increases so a greater amount of power is consumed, and the heat generated by DRAMs causes the temperature of the module to rise.

[0004] The aim of the present invention is to resolve these kinds of conventional problems by providing a large-capacity memory module which employs a multibit configuration of current-generation components, and therefore makes it possible to reduce costs compared with using next-generation components, and also makes it possible to reduce the number of DRAMs driven at the same time compared with a system employing a current-generation ×1 bit configuration, so power consumption is reduced and temperature rises are lowered.

[0005]

[Means for Resolving the Problem] In order to resolve the problem, the memory module according to the present invention is composed of: a plurality of banks comprising a plurality of current-generation SDRAMS for achieving increased memory capacity with current-generation SDRAMs, which cannot be achieved using only next-generation SDRAMs; and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks.

[0006]
By virtue of the configuration according to the present invention, it is possible to obtain a large capacity memory module at lower cost than a system employing temperature next-generation components, and with lower temperature rises than a system employing a ×1 configuration with current-generation components, so it is a simple matter to install a large capacity memory in a set such as a computer, and an inexpensive memory module can be obtained.

[0007] Furthermore, it is also possible to reduce problematic temperature rises if the DRAMs are installed using an installation method enabling multiple stage installation, such as TCP (tape carrier packaging), and therefore the installation density can be increased, the size of the module can be reduced, and it is possible to obtain a compact, large-capacity memory module.

[0008]
[Mode of Embodiment of the Invention] The invention disclosed in claim 1 of the present invention is a memory module composed of: a plurality of banks comprising a plurality of current-generation SDRAMS for achieving increased memory capacity with current-generation SDRAMs, which cannot be achieved using only next-generation SDRAMs; and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks, so an action is afforded such that it is possible to obtain a large capacity memory module at lower cost than a system employing temperature next-generation components, and with lower temperature rises than a system employing a ×1 configuration with current-generation components, so it is a simple matter to install a large capacity memory in a set such as a

computer, and an inexpensive memory module can be obtained.

[0009] The invention disclosed in claim 2 is the memory module as claimed in claim 1, composed of: a plurality of banks which are installed using an installation method enabling a plurality of current-generation DRAMs to be installed in multiple stages, such as TCP (tape carrier packaging); and a bank control unit for converting drive signals from outside the module, which are sent in order to control the plurality of banks, to signals for controlling the plurality of banks, so an action is afforded such that the memory module can be made compact.

[0010] The invention disclosed in claim 3 is the memory module as claimed in claim 1 or 2, in which the bank control unit is a one-chip unit, so an action is afforded such that it is possible to obtain a compact memory module.

[0011] A mode of embodiment of the present invention will be described below with the aid of Figures 1 to 3. Figure 1 is a block diagram showing a plurality of banks composed of a plurality of multibit DRAMs according to a first mode of embodiment of the present invention; Figure 2 is a block diagram of the bank control unit for converting drive signals from outside the module into signals for controlling the plurality of banks; and Figure 3 is a timing chart showing the operation of the bank control unit in Figure 2. In Figure 1, D0-D63 are SDRAMs in four banks, where D0 - D15 form bank 0, D16 - D31 form bank 1, D32 - D47 form bank 2, and D48 - D63 form bank 3.

[0012] The operation of this mode of embodiment will be described below with the aid of Figures 1, 2 and 3. Drive signals /RE, /CE, /WE, A12, A13 etc. from outside the module are input to the memory module. The bank

control unit shown in Figure 2 identifies SDRAM bank active commands (when /RE; "L", /CE; "H", /WE; "H") using the signals /RE, /CE, /WE, decodes A12, A13, performs bank selection, and sets /CS (chip select signal), which corresponds to a selected bank, to active ("L").

[0013] Period (1) in Figure 3 involves active commands, as described above, and A12 is "L" and A13 is "L", so two from among /CS0 to 7, namely /CS0 and /CS4 are set to "L" in order to select bank 0 which is composed of SDRAMs D0 - D15 shown in Figure 1, but the other banks are in an inactive state.

[0014] Period (2) in Figure 3 involves write commands to the SDRAMs, but writing is performed only to the selected bank 0 and the other banks 1 to 3 do not operate. Period (3) in Figure 3 involves precharge commands to the SDRAMs, and precharging is performed only for the bank 0 and the other banks do not operate, in the same way as in period (2).

[0015] In the same way subsequently, in period (4) in Figure 3, only bank 1 is active, in period (7) in Figure 3, bank 2 is active, and in period (10) in Figure 3, bank 3 is active. As a result, it is possible to achieve equivalent operation to that of a system in which a 64 bit memory is made up of the conventional 64 Mbit SDRAMs shown in Figure 4. Furthermore, the banks are active one at a time, so in other words the active SDRAMs constitute 1/4 of the total number of SDRAMs in the memory module. Power consumption for the inactive SDRAMs is close to the power consumed during standby and is very low compared with the power consumed during operation. It is therefore possible for the memory module to operate at approximately 1/4 power consumption.

[0016]

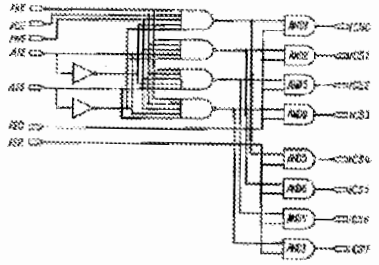
[Advantage of the Invention] The memory module obtained in the manner described above produces an advantageous effect in that low power consumption (low temperature rises) which can be achieved by using 64 Mbit SDRAMs can be achieved using 16 Mbit SDRAMs, so costs can be reduced.

[0017] It should be noted that in this mode of embodiment, the current-generation components were 16 Mbit SDRAMs while the next-generation components were 64 Mbit SDRAMs, but the same effect can still be achieved if the generations are advanced so that the current-generation components are 64 Mbit SDRAMs while the next-generation components are 256 Mbit SDRAMs.

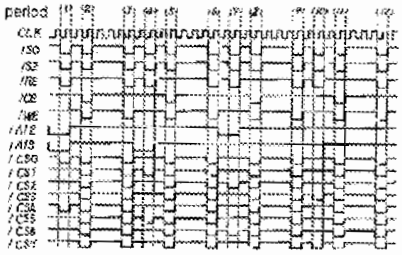
[Detailed Description of the Invention]
[Figure 1] is a block diagram showing the configuration of the plurality of banks composed of a plurality of multibit DRAMs, in accordance with a mode of embodiment of the present invention;
[Figure 2] is a block diagram of the bank control unit for converting drive signals from outside the module into signals for controlling the plurality of banks, in accordance with a mode of embodiment of the present invention;
[Figure 3] is a timing chart showing the operation of the bank control unit in Figure 2, in accordance with a mode of embodiment of the present invention; and
[Figure 4] is a block diagram of a memory module, which is a conventional 128 MB SDRAM module comprising sixteen 64 Mbit, ×4 configuration SDRAMs.

[Key to Symbols]
D0 - D63 SDRAM (synchronous DRAM)

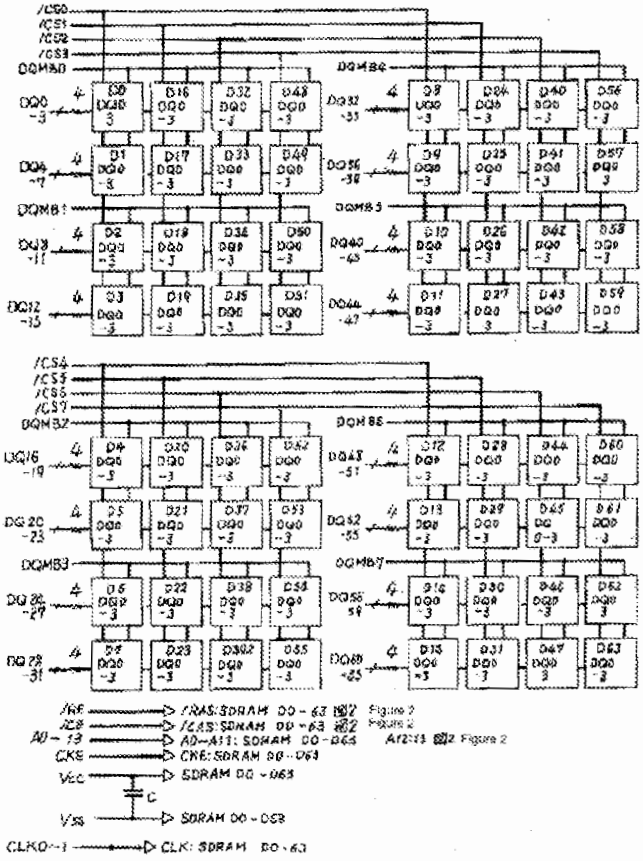
[Figure 2]



[Figure 3]

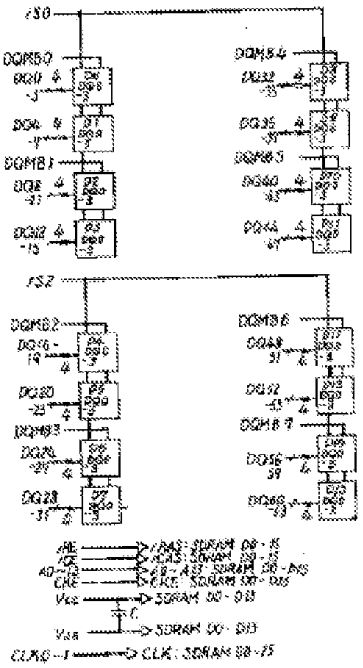


[Figure 1]



5

[Figure 4]



PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 :		(11) International Publication Number:	WO 92/02879
G06F 12/10, G09G 1/16	A1	(43) International Publication Date:	20 February 1992 (20.02.92)

(21) International Application Number: PCT/GB90/01213

(22) International Filing Date: 3 August 1990 (03.08.90)

(71) Applicant *(for all designated States except US)*: DU PONT PIXEL SYSTEMS LIMITED [GB/GB]; Wedgewood Way, Stevenage, Hertfordshire SG1 4QN (GB).

(72) Inventors; and

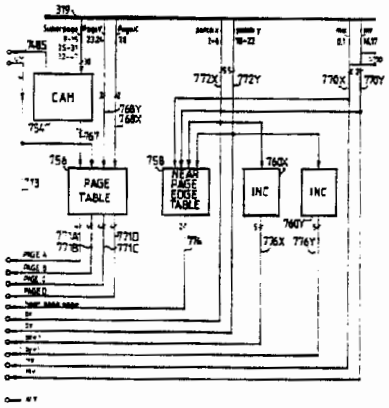
(75) Inventors/Applicants *(for US only)* : MAUND, Andrew, Peter [GB/GB]; 38 Laughton Road, Northolt, Middlesex UB5 5LL (GB). NEAVE, John, Walter [GB/GB]; 49 Bennetts Wood, Capel, Dorking, Surrey RH5 5EX (GB). TREVETT, Neil, Francis [GB/GB]; 16 Manorgate Road, Kingston-Upon-Thames, Surrey KT13 0RJ (GB). MOORE, Simon, James [GB/GB]; 61 Swallowtail Road, Horsham, West Sussex RH12 4YG (GB). WILSON, Malcolm, Eric [GB/GB]; 37 Manor Crescent, Byfleet, Weybridge, Surrey KT14 7EN (GB).

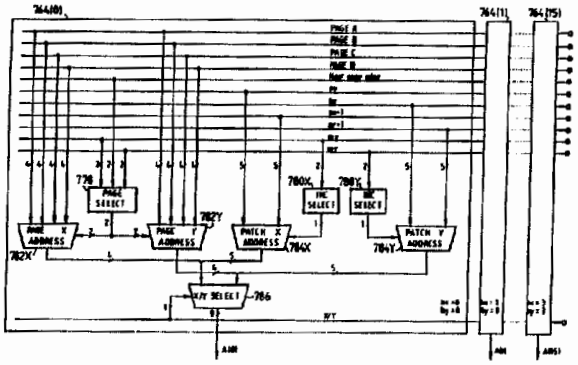
(74) Agents: CRAMPTON, Keith, John, Allen et al.; D. Young & Co., 10 Staple Inn, London WC1V 7RD (GB).

(81) Designated States: AT (European patent), BE (European patent), CH (European patent), DE (European patent)*, DK (European patent), ES (European patent), FR (European patent), GB (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent), US.

Published
With international search report.

(54) Title: VIRTUAL MEMORY SYSTEM





(57) Abstract

In a demand-paged virtual memory system, the pages are arranged in the virtual memory space in groups. In order to translate an address from the virtual address space to a physical memory address space, the virtual group address component is input to a contents addressable memory (CAM)(745) which outputs a group code (767), and the group code and virtual page address component (768X,Y) are input to a RAM page table (750) which outputs the page address. When the physical memory capacity is substantially smaller than the virtual address space, the CAM provides a large saving in page table size. In the case where the data-elements provide a plural-dimensional representation, for example as in pixel data, the pages include data elements which are contiguous in each of the plural dimensions in order to reduce the amount of page-swapping between the physical memory and a paging memory. The data-elements in the physical memory are accessible in parallel as contiguous patches. In order to allow parallel access of patches which are not aligned with respect to a page and therefore may fall partly on the edge of one page and partly on the edge of a page which is contiguous therewith in the virtual memory but not in the physical memory, the addressing means includes means (758) to determine whether a patch to be accessed has data-elements in different pages, and means (756, 778, 782X,Y) to modify the addresses to the physical memory.

* See back of page

Micron Technology Inc. et al.
Ex. 1010, p. 334

*** DESIGNATIONS OF "DE"**

Until further notice, any designation of "DE" in any international application whose international filing date is prior to October 3, 1990, shall have effect in the territory of the Federal Republic of Germany with the exception of the territory of the former German Democratic Republic.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	**ES**	Spain	**MG**	Madagascar
AU	Australia	**FI**	Finland	**ML**	Mali
BB	Barbados	**FR**	France	**MN**	Mongolia
BE	Belgium	**GA**	Gabon	**MR**	Mauritania
BF	Burkina Faso	**GB**	United Kingdom	**MW**	Malawi
BG	Bulgaria	**GN**	Guinea	**NL**	Netherlands
BJ	Benin	**GR**	Greece	**NO**	Norway
BR	Brazil	**HU**	Hungary	**PL**	Poland
CA	Canada	**IT**	Italy	**RO**	Romania
CF	Central African Republic	**JP**	Japan	**SD**	Sudan
CG	Congo	**KP**	Democratic People's Republic of Korea	**SE**	Sweden
CH	Switzerland			**SN**	Senegal
CI	Côte d'Ivoire	**KR**	Republic of Korea	**SU+**	Soviet Union
CM	Cameroon	**LI**	Liechtenstein	**TD**	Chad
CS	Czechoslovakia	**LK**	Sri Lanka	**TG**	Togo
DE*	Germany	**LU**	Luxembourg	**US**	United States of America
DK	Denmark	**MC**	Monaco		

+ It is not yet known for which States of the former Soviet Union any designation of the Soviet Union has effect.

VIRTUAL MEMORY SYSTEM

The present invention relates to a demand-paged virtual memory system.

More particularly, the present invention is concerned with a memory system for storing an array of data elements, such as pixel data, which form in a virtual memory space contiguous aligned pages of contiguous data elements and comprising a first memory, or paging memory, for storing the data-array, a second memory, such as a video RAM, for storing part of the data-array, and means for transferring data-elements page-by-page between the first and second memories.

In a simple non-demand-paged system, the whole of the second memory could be filled with contiguous pages of the data elements. In this case, in order to translate between a page address in the virtual memory address space and a page address in the second memory address space, all that would be required would be to know the virtual page address of page zero in the second memory, and to use this as an offset between a second memory page address and the corresponding virtual memory page address. However, such a simple system is inflexible, and requires the whole content of the second memory to be swapped each time a new page is to be transferred from the first memory to the second memory.

In a more flexible system, that is a demand-paged system, pages of data-elements are transferred from the first memory to the second memory as required and placed in any available page location in the second memory. As a result, pages which are contiguous in the virtual memory address space are not necessarily stored contiguously in the second memory address space, and it is therefore necessary to provide a means to keep track of the pages in the second memory and to translate between the virtual page address and the corresponding second memory page address. This means could be provided by a RAM page table which has an address for each page in the virtual memory space and which stores

Micron Technology Inc. et al.
Ex. 1010, p. 336

at that address in the page table a flag indicating whether that page is also stored in the second memory and if so also the address in the second memory of that page.

In the system with which the preferred embodiment of this invention is concerned, the virtual memory space contains 2^{34} pages (or 16 Gigapages) and the second memory space contains 2^{8} (or 256) pages. Therefore, if a page table as described above were used in a system of such size, the page table would require a RAM of 2^{42} bits (or 4 Terabits) for the addresses and a RAM of 2^{34} (or 16 Gigabits) for the flags.

A first aspect of the present invention is concerned with providing an address translation means which does not require tables of such massive size.

In the system in accordance with the first aspect of the present invention, the virtual memory address space is considered to be formatted as aligned groups of the pages, with the pages in each group being aligned and contiguous, and the translating means comprises a contents addressable memory (CAM) which receives the group component of the virtual address and outputs a group code corresponding thereto, and a page table which receives the group code and the page component of the virtual address and outputs the corresponding page address in the second memory address space.

Comparing the system of the first aspect of the invention with the example given above, if for example each group is chosen to contain 2^{4} pages, then the virtual memory space will contain 2^{30} groups (or 1 Gigagroup). A CAM may then be chosen which has, say 2^{7} (or 128) addresses, each with a 30 bit capacity. Thus the CAM will have a 30-bit input and a 7-bit group code output. The page table will therefore be addressed by 11 bits (comprising the 7-bit group code and then 4-bit page component) and will output the 8-bit page address in the second memory. Accordingly, the CAM stores 30 x 2^{7} (or 3840 bits), and the page table stores 2^{11} x 2^{8} bits (or $\frac{1}{2}$ Megabit), which in combination compare favourably with the 4 Terabits required by the earlier described system.

Micron Technology Inc. et al.
Ex. 1010, p. 337

It will be appreciated that, for good advantage, the number of group codes which can be stored in the CAM is preferably substantially less than the number of the groups which can be stored in the first memory, or alternatively stated that each group code has a number of bits (e.g. 7) which is substantially less than the number of bits (e.g. 30) of the virtual group component.

In the preferred embodiment, the CAM is operable to compare the received virtual group component with the contents at the addresses of the CAM and to output as the group code the address of the CAM whose content matches the virtual group component. In this case, in order to flag locations in the CAM at which a virtual group component is not stored the CAM contents preferably have a bit (e.g. msb) which is not used in the virtual group component, and which is set at any location in the CAM at which a virtual group component is not stored.

In the preferred embodiment, the CAM is preferably operable to supply a group fault signal if a match is not made with the received virtual group component, and the system preferably further comprises processor means operable in response to the group fault signal to store the received virtual group component at an available location in the CAM. In this case the processor means is preferably operable in response to the group fault signal to determine whether there is any available location in the CAM, and if not to select a virtual group component to delete from the CAM, and to determine whether any pages of data-elements of the selected group are stored in the second memory, and if so to cause the transfer means to transfer those pages of data-elements from the second memory to the first memory. In order to do this, the system preferably further comprises means for determining the least recently used group having a page stored in the second memory, and the processor means is operable to select the virtual group component of the determined least recently used group as the virtual group component for deletion from the CAM.

In a conventional demand-paged virtual memory system, the virtual memory address space is one-dimensional, and the locations in each page

Micron Technology Inc. et al.
Ex. 1010, p. 338

are adjacent in one dimension. The present invention is concerned more particularly, but not exclusively, concerned with processing of a data-array in which the data-elements form a plural (such as 2- or 3-) dimensional representation and in which the relative positions of the data-elements in the array are significant, in addition to their values, for example as in pixel data or vector data. In this case, if processing concerns a plurality of adjacent data-element arranged in the one dimension of the address space, there is a high probability that these data-elements will belong to the same page. However if processing concerns data-elements which are adjacent in some other direction, they will belong to different pages. Accordingly, in the case where the second memory is of substantially smaller capacity than the first memory, a vast amount of page swapping may be required between the first and second memories.

A second aspect of the present invention is concerned with a demand-paged virtual memory system in which the amount of page-swapping, when dealing with plural-dimensional data-arrays, is reduced. In accordance with this second aspect of the present invention, the pages which are transferred between the memories include data-elements which are contiguous in each of the plural dimensions of the representation provided by the data elements. For example, in the case of pixel data representing a two-dimensional image, each transferred page of pixel data is a two-dimensional portion of the image, for example of 128 pixels x 128 pixels. Similarly, in the case of three-dimensional pixel data, each transferred page would be a three dimensional portion of the image, for example of 32 pixels x 32 pixels x 16 pixels.

Preferably, the virtual address space and the second memory address space are plural dimensional and the addresses include components in each of the plural dimensions of the data-array.

In order to increase memory access speed and processing speed, it is desirable that a set or "patch" of contiguous data-element locations in the second memory can be accessed in parallel, for example so that the parallel-accessed data-elements can be processed in parallel by respective

Micron Technology Inc. et al.
Ex. 1010, p. 339

ones of a set of processors. Conveniently, t ..se patches are aligned with respective to the page so that the patches form a sub-array in each page. Despite this, the memory is preferably accessible so that a non-aligned patch can be accessed having some data-elements in one aligned patch and other data-elements in one or more adjacent aligned patches. Throughout most of the page, the data-elements of such misaligned patches will all be in the same page, but for patches on the edge of the page, some of the data-elements of the misaligned patch may belong to one or more different page. As described above with respect to the first aspect of the invention, contiguous pages in the virtual address space are not necessarily stored contiguously in the second memory address space, and therefore these misaligned patches on the page boundaries present a problem, because the page components of the addresses in the second memory of some of the data-elements in the patch will be entirely different to those of others in the patch.

In order to deal with this problem, in accordance with a third aspect of the present invention, the means for addressing the second memory to access in parallel a patch of data-elements which is not necessarily an aligned patch includes page edge means for determining whether the patch to be accessed includes data-elements in different pages in the second memory, and means to modify the addresses to the second memory in response to such a determination.

Preferably, the address of a patch to be accessed in the second memory has an aligned page address component, an aligned patch address component and a misalignment component indicating a misalignment of the patch to be accessed relative to an aligned patch having an address defined by the aligned page and aligned patch components, and in this case the addressing means preferably includes page determining means operable to determine, upon addressing, an aligned page add: .ss component for a basic aligned page for the patch to be accessed and an aligned page address component of at least one other page which is contiguous with the basic aligned page in the virtual address space.

In the case where the data-array is a plural-dimensional representation having D-dimensions, the page determining means is

preferably operable to determine the aligned page addresses in the second memory for all aligned pages which differ by one from the basic aligned page address in one direction in each of the D-dimensions in the virtual address space and in any combinations of those one directions.

In order to deal with the problem that not all required pages may be stored in the second memory, the system preferably further comprises page fault determining means for determining, upon addressing, whether the or each aligned page which is to be accessed is not stored in the second memory, the transferring means being operable to transfer the missing page from the first memory to the second memory in response to such a determination. Preferably, the page fault determining means is operable to indicate a page fault if the basic page is not stored in the second memory, or if any said other aligned page is not stored in the second memory and that page is required to be accessed. Preferably, the page edge determining means is operable to determine when the aligned patch address component has a maximum value for a patch in a page and the misalignment component together with said maximum value indicate that the patch to be accessed extends outside the page defined by the aligned page address component.

Preferred embodiments of the invention will now be described, by way of non-limiting example, with reference to the drawings, in which Figures 4 to 16, 20 and 21 are of major relevance to the present invention, and in which:

Figure 1 is a high-level schematic illustration of a computer system embodying the invention;

Figures 2 and 3 illustrations of modified forms of the system of Figure 1;

Figure 4 is an illustration in greater detail of a renderer employed in the systems of Figures 1 to 3;

Figure 5 is an illustration in greater detail of a front-end processor

Micron Technology Inc. et al.
Ex. 1010, p. 341

board employed in the systems of Figures 1 to 3;

Figure 6A and 6B show how patches of pixel data are made up;

Figures 7A and 7B show how pages of patch data, and groups or 'superpages' of page data are made up;

Figure 8 is a schematic illustration of a physical image memory and the address lines therefor, used in the renderer of Figure 4;

Figure 9A is a 3-D representation of an aligned patch of data within a single page in the image memory;

Figure 9B is a 2-D representation of a page, showing the patch of Figure 9A;

Figure 10A is a 3-D representation of a non-aligned patch of data within a single page in the image memory;

Figure 10B is a 2-D representation of a page, showing the patch of Figure 10A;

Figure 11A is a 2-D representation of four pages in a virtual memory, showing a non-aligned patch which crosses the page boundaries and an enlargement of the circled part of the page boundary intersection;

Figure 11B is a 2-D representation of the physical memory illustrating locations of the four pages shown in Figure 11A;

Figure 11C is a 3-D representation of the non-aligned patch of Figure 11A;

Figure 12 is a truth table showing how page selection is made for patches which cross page boundaries;

Figure 13 shows two truth tables for selecting, respectively, X and Y patch address incrementation;

Figure 14 is a schematic illustration in greater detail of part of the renderer of Figure 4;

Figure 15 is a schematic illustration in greater detail of an address translator of Figure 14;

Figure 16 is an illustration of the operation of a contents addressable memory used in the address translator of Figure 15;

Figure 17 is a schematic illustration in greater detail of a read surface shifter used Figure 14;

Figure 18 shows in greater detail an array of multiplexers forming part of the surface shifter of Figure 17;

Figure 19 illustrates the translation made by the surface shifter of Figure 17;

Figure 20 is an illustration of the operation a least-recently-used superpage table which may be used with the address translator of Figure 15;

Figure 21 is a schematic diagram showing a page fault table which may be used with the address translator of Figure 15;

Figure 22 is a schematic diagram of an exchange and grid processor of the renderer of Figure 4;

Figure 23 is a flow diagram illustrating the operation of the processors and a priority encoder of the grid processor of Figure 22;

Figure 24 is a table giving an example of the operation of the priority encoder of Figure 22;

Micron Technology Inc. et al.
Ex. 1010, p. 343

Figure 25 illustrates the correlation between aligned memory cells and two levels of a patch in a 16-bit split patch system;

Figures 26 and 27 show how pages of patch data, and superpages of page data are made up in a 16-bit split patch system;

Figures 28 and 29 correspond to Figures 26 and 27 respectively in an 8-bit split patch system;

Figures 30A to C shows modifications of parts of the address translator of Figure 15 used in the split patch system;

Figure 31 is a table to explain the operation of a funnel shifter used in the circuit of Figure 30A;

Figures 32 and 33 illustrate non-aligned split patches in a 16-bit and an 8-bit patch system, respectively;

Figure 34 shows a further modification of part of the address translator of Figure 15 used in the split patch system;
Figures 35A and 35B are tables which illustrate the operation of look-up tables in the circuit of Figure 34;

Figures 36A and 36B shows modifications of a near-page-edge table of Figure 15A used in the split patch system;

Figure 37 illustrates, in part, a modification to the exchange and grid processor of Figure 22 used in the split patch system;

Figures 38 and 39 are tables which illustrate the operation of further tables in a further modification of part of the address translator of Figure 15;

Figure 40 shows the further modification to Figure 15;

Figure 41 shows a modification to Figure 8 which is made in addition

Micron Technology Inc. et al.
Ex. 1010, p. 344

to the modification shown in Figure 40;

Figure 42 is a representation of the VRAM memory space, showing how pages of data are rendered in one section of the memory and then copied to another monitoring section of the memory;

Figure 43 shows a circuit for determining which pages need not be copied from the rendering section to the monitoring section and to the virtual memory;

Figure 44 illustrates the setting and resetting of flags in a table of the circuit of Figure 43;

Figure 45A to 45C are flow diagrams illustrating the copying operations and Figure 45D shows the notation used in Figures 45A to 45C;

Figure 46 is a circuit diagram of a modification to the exchange of Figure 22;

Figure 47A to 47C are simplified forms of the circuit of Figure 46 when operating in three different modes;

Figure 48 shows a modification to part of the flow diagram at Figure 23;

Figure 49 is a schematic diagram of the processors and a microcode memory, with one of the processors shown in detail;

Figures 50A to 50D illustrate three images (Figs 50A to C) which are processed to form a fourth image (Fig. 50D); and

Figure 51 is a system diagram showing in particular a page filing system.

HARDWARE OVERVIEW

Figures 1 to 3 show three different hardware configurations of computer systems embodying the invention. Referring firstly to Figure 1, a host computer 10 has its own backplane in the form of a VME bus 12 which provides general purpose communications between various circuit boards of the computer, such as processor, memory and disk controller boards. To this known configuation, and within a standard housing 14 for the computer 10, there is added a board on which is provided a renderer 16 and a video processor 18, a Futurebus+ 20, and a front-end board 22. The renderer 16 is connected to the VME bus 12 and the Futurebus+ 20, and also communicates with the video processor 18, which in turn drives an external colour monitor 24 having a high-resolution of, for example, 1280 X 1024 pixels. The front-end board 22 is also connected to the Futurebus+ 20 and can communicate with a selection of peripherals, which are illustrated collectively by the block 26. The configuration of Figure 1 is of use when the host computer 10 has a VME backplane 12 and there is sufficient room in the computer housing 14 for the renderer 16, video processor 18, Futurebus+ 20 and front-end board 22, and may be used, for example, with a 'Sun Workstation'.

In the case where the computer housing 14 is physically too small, or where the host computer 10 does not have a VME or Futurebus+ backplane, the configuration of Figure 2 may be employed. In Figure 2, a separate housing 28 is used for the renderer 16, video processor 18, front-end board 22 and Futurebus+ 20, as described above, together with a VME bus 12 and a remote interface 30. In the host computer housing 14, a host interface 32 is connected to the backplane 34 of the host computer 10, which may be of VME, Qbus, Sbus, Multibus II, MCA, PC/AT, etc. format. The host interface 32 and remote interface 30 are connected by an asynchronous differential bus 36 which provides reliable communication despite the physical separation of the host and remote interfaces. The configuration of Figure 2 is appropriate when the host computer 10 is, for example, an 'Apple Mackintosh', 'Sun Sparkstation', 'IBM-PC', or Du Pont Pixel Systems bRISC'.

Micron Technology Inc. et al.
Ex. 1010, p. 346

In the event that a host computer becomes available which has a Futurebus+ backplane and sufficient space in its housing for the additional components, then the configuration of Figure 3 may be employed. In Figure 3, the renderer 16 and the front-end board 22 are directly connected to the Futurebus+ backplane 20.

The general functions of the elements shown in Figures 1 to 3 will now be described in more detail. The host computer 10 supplies data in the form of control information, high level commands and parameters therefor to the renderer 16 via the VME backplane (Figure 1), via the backplane 34, host and remote interfaces 32, 30 and the VME bus 12 (Figure 2), or via the Futurebus+ backplane 20 (Figure 3). Some of this data may be forwarded to the front-end board via the Futurebus+ 20 (Figures 1 and 2), or sent direct via the Futurebus+ backplane 20 (Figure 3) to the front-end board 22.

The Futurebus+ 20 serves to communicate between the renderer 16 and the front-end processor 22 and is used, in preference to a VME bus or the like, in view of its high bit width of 128 bits and its high bandwidth of about 500 to 800 Mbytes/s.

As will be decribed in greater detail below, the renderer 16 includes an image memory, part of which is mapped to the monitor 24 by the video processor 18, and the renderer serves to perform image calculations and rendering, that is the drawing of polygons in the memory, in accordance with the commands and parameter supplied by the host computer 10 or the front-end board 22.

The front-end board 22 serves a number of functions. It includes a large paging RAM, which also interfaces with external disk storage, to provide a massive paging memory, and pages of image data can be swapped between the paging RAM and the image memory of the renderer 16 via the Futurebus+ 20. The front-end board also has a powerful floating-point processing section which can be used for graphics transformation and shading operations. Furthermore, the front-end board

Micron Technology Inc. et al.
Ex. 1010, p. 347

may provide interfacing with peripherals such as a video camera or recorder, monitor, MIDI audio, microphone, SCSI disk and RS 232.

Overall, therefore, the renderer 16, video processor 18 and front-end board 22 can accelerate pixel handling aspects of an application, and also accelerate other computation intensive aspects of an application.

The renderer 16 and video processor 18 will now be described in greater detail with reference to Figure 4, which shows the main elements of the renderer 16 and the main data and address pathways.

The renderer 16 includes a 32-bit internal bus 300, a VME interface 301 which interfaces between the VME bus 12 (Figure 1) or the remote interface 30 (Figure 2) and the internal bus 300, and a Futurebus+ interface 302 which interfaces between the Futurebus+ 20 and the internal bus 300. Also connecting to the internal bus 300 are a control processor 314 implemented by an Intel 80960i, an EPROM 303, 4 or 16 Mbyte of DRAM 304, a real time clock and an I/O block 306 including a SCSI ports. The functions of the control processor 314 and the associated DRAM 304 and EPROM 303 are (a) to boot-up and configure the system; (b) to provide resource allocation for local PRAM 318, 322 of address and grid processors 310, 312 (described in detail below) to ensure that there is no memory space collision; (c) to control the loading of microcode into microcode memories 307, 308 (described below); (d) to run application specific remote procedure calls (RPCs); and (e) to communicate via the I/O block 306 with a diagnostics port of the host computer 10 to enable diagnostics information to be displayed on the monitor 24. The DRAM 304 can also be used as a secondary image page store for the VRAM 700 described below.

The renderer 16 also includes an address processing section 309 comprising an address broadcast bus 311 to which are connected 64 kbyte of global GRAM 316, a data/instruction cache 313 which also connects to the internal bus 300, an internal bus address generator 315 which also connects to the internal bus 300, an address processor 310 with 16 kbyte

Micron Technology Inc. et al.
Ex. 1010, p. 348

of local PRAM 318, and a sequencer 317 for the address processor 310 which receives microcode from a microcode memory 307. The address processor 310 also connects to a virtual address bus 319. The main purpose of the address processing section 309 is to generate virtual addresses which are placed on the virtual address bus under control of microcode from the microcode memory 307.

Also included in the renderer 16 is an address translator 740 (described in further detail below) which receives the virtual addresses on the virtual address bus 319 and translates them into physical addresses of data in the video RAM 700, if the required data is present, or interrupts the address processor 310 to cause the required data to be swapped in from the paging RAM 304 or other page stores on the external buses, if the required data is not present in the VRAM 700.

The renderer 16 furthermore includes a data processing section 321 which is somewhat similar to the address processing section 309 and comprises a data broadcast bus 323, to which are connected 64 kbyte of global GRAM 324, a diagnostics register 325 which also connects to the internal bus 300 and which may be used instead of the I/O block 306 to send diagnostics information to the host computer 10, an internal bus address generator 327 which also connects to the internal bus 300, a grid processor 312 having sixteen processors each with 8 kbyte of local PRAM 322, and a sequencer 329 for the grid processor 312 which receives microcode from a microcode memory 308. The processors of the grid processor 312 also connect to a data bus 331. The main purpose of the data processing section 321 is to receive data on the data bus 331, process the data under control of microcode from the microcode memory 308, and to put the processed data back onto the data bus 331.

The physical VRAM 700 connects with the data bus 331 via an exchange 326 which is described in detail below, but which has the main purposes of shuffling the order of the sixteen pixels read from or written to the VRAM 700 at one time, as desired, to enable any of the sixteen processors in the grid processor 312 to read from or write to any of the

sixteen addressed locations in the VRAM 700 and to enable any of the sixteen processors to transfer pixel data to any other of the sixteen processors.

The last main element of the renderer 16 is a bidirectional FIFO 332 connecting between the broadcast buses 311, 323 of the address and data processing sections 309, 321, which enables virtual addresses to be transferred directly between these two sections.

The front-end board 22 will now be described in greater detail with reference to Figure 5.

The front-end board 22 has an internal bus 502 which communicates with the Futurebus+ 20. A paging memory section 504 is connected to the internal bus 502 and comprises a large paging RAM 506 of, for example, 4 to 256 Mbytes capacity which can be used in conjunction with the DRAM 304 of the renderer, a paging memory control processor 508, and connections to, for example, two external high speed IPI-2 disk drives 510 (one of which is shown) each of which may have a capacity of, for examole, 4 Gbytes, and a data communication speed of 50 Mbytes/s, or two external SCSI drives. The paging RAM 506 enables an extremely large amount of pixel data to be stored and to be available to be paged into the renderer 32 as required, and the fast disk 510 enables even more pixel data to be available ready to be transferred into the paging RAM 506.

Floating point processing is provided by 1 to 4 Intel 80860 processors 516, each rated at 80 MFlops peak. The general purpose processing power can be used on dedicated tasks such as geometric pipeline processing, or to accelerate any part of an application which is compute-intensive, such as floating point fast Fourier transforms. Each of the floating point processors 516 has a 128KByte secondary cache memory 518 in addition to its own internal primary cache memory.

The front-end board 22 may also, if desired, include a broadcast

Micron Technology Inc. et al.
Ex. 1010, p. 350

standard 24-bit frame grabber connected to the internal bus 502 and having a video input 514 and output 516 for connection to video camera or television-type monitor.

The front-end board 22 may also, if desired, include an input/output processor 520 which provides interfacing with MIDI on line 522, SCSI disk on line 524, at least one mouse on line 526, RS232 on line 528, and audio signals on line 530 via a bi-directional digital/analogue convertor 532.

VIDEO RAM AND ADDRESSING THEREOF

Now that an overview of the hardware of the whole system has been set out, the image memory configuration will be described in more detail.

As mentioned above, the VRAM has a of 16 Mbyte capacity. The system is capable of operating with 32-bit pixels, and therefore the image memory has a capacity of 16M x 8/32 = 4 Mpixels. As illustrated in Figures 6A and 6B, pixels are arranged in 4 x 4 groups referred to as 'patches'. Figures 6A and 6B show, respectively, two-and one-dimensional notations for designating a pixel in a patch, as will be used in the following description. In turn, as illustrated in Figure 7A, the patches are arranged in 32 x 32 groups referred to as 'pages'. Furthermore, as illustrated in Figure 7B, the pages are arranged in 4 x 4 groups referred to as 'superpages'. The VRAM therefore has a capacity of 4 Mpixels, or 256k patches, or 256 complete pages, or 16 complete superpages. However, not all pages of a particular superpage need be stored in the memory at any one time, and support is provided for pages from parts of up to 128 different superpages to be stored in the physical memory at the same time.

The VRAM 700 and addressing lines therefor are shown schematically in Figure 8. Each small cube 702 in Figure 8 represents a 32-bit pixel. The pixels are arranged in 512 pixel x 512 pixel banks B(0) -

Micron Technology Inc. et al.
Ex. 1010, p. 351

B(15) lying in the XY plane, and these pixel banks are 16 pixels deep (in the P direction). A line of 16 pixels in the P direction provides an aligned patch 704. The pixels in each bank are addressable as to X address by a respective one of 16 9-bit X address lines AX(0) to AX(15) and are addressable as to Y address by a respective one of 16 9-bit Y address lines AY(0) to AY(15). The Y and X addresses are sequentially supplied on a common set of 16 9-bit address lines A(0) to A(15), with the Y addresses being supplied first and latched in a set of 16 9-bit Y latch groups 706-0 to 706-15 each receiving a row address strobe (RAS) signal on 1-bit line 708, and the X addresses then being supplied and latched in a set of 16 9-bit X latch groups, 707-0 to 707-15 each receiving a respective column address strobe signal CAS(0) to CAS(15) on lines 709(0) to 709(15), respectively.

The memory for each XY bank of pixels (512 pixels x 512 pixels x 1 pixel) is implemented using eight video-RAM (VRAM) chips 710, each 256 K (4-bit) nibbles. Each chip provides a one-eighth thick slice of each pixel bank, whereby 8 x 16 = 128 chips are required. Each Y latch group and X latch group comprises eight latches (shown in detail for Y latch group 706(1) and X latch group 707(1) and a respective one of the X and Y latches is provided on each VRAM chip 710.

In this specification, the banks of memory will sometimes be referred to by the bank number B(0) to B(15) and at other times by a 2-dimensional bank address (bx,by) with the correlation between the two being as follows:

Bank Number	(bx,by)	Bank Number	(bx,by)
B(0)	(0,0)	B(8)	(0,2)
B(1)	(1,0)	B(9)	(1,2)
B(2)	(2,0)	B(10)	(2,2)
B(3)	(3,0)	B(11)	(3,2)
B(4)	(0,1)	B(12)	(0,3)
B(5)	(1,1)	B(13)	(1,3)
B(6)	(2,1)	B(14)	(2,3)
B(7)	(3,1)	B(15)	(3,3)

When a location in the memory 700 is to be accessed, a patch of 16 pixels is made available for reading or writing at one time. If the Y address and X address for all of the VRAMs 710 is the same, then an "aligned" patch of pixels (such as patch 704) will be accessed. However, it is desirable that access can be made to patches of sixteen pixels which are not aligned, but where various pixels in the patch to be accessed are derived from two or four adjacent aligned patches.

It will be appreciated that access to an aligned patch in memory is more straightforward than access to a non-aligned patch, because for an aligned patch the (x,y) address of each pixel in the different XY planes of memory as shown in Figure 8 is the same. Furthermore, the (x,y) address of each pixel in the patch is equal to the bank address (bx,by) of the memory cell from which that pixel is derived. Referring to Figures 9A and 9B, an aligned patch "a" having a patch address (12, 16) in a page "A" having a page address (8, 6) is shown, as an example. The pixels in the aligned patch all have the same address in the sixteen XY banks of the memory, as represented in Figure 9A, and when displayed would produce a 4 X 4 patch of pixels offset from the page boundaries by an integral number of patches, as represented in Figure 9B. In the particular example the absolute address of the aligned patch in the memory would be (8 x 32 + 12, 6 x 32 + 16) = (268, 208).

If, however, a patch "p" is non-aligned, and has a misalignment (mx,my) = (2,1), for example, from the previously considered aligned patch "a" at patch address (12, 16) in page A at page address (8, 6), then some of the pixels of patch "p" will need to be derived from three other aligned patches "b", "c" and "d" having patch addresses (12 + 1, 16), (12, 16 + 1) and (12 + 1, 16 + 1), or (13, 16), (12, 17) and (13, 17) in page A at page address (8, 6). This situation is represented in Figures 10A and 10B. The absolute address of these patches "b", "c" and "d" in the VRAM 700 are (269, 208), (268, 209) and (269, 209); respectively.

A further problem which arises in accessing a non-aligned patch "p" is that the (x,y) address of each pixel in the patch "p" does not correspond to

Micron Technology Inc. et al.
Ex. 1010, p. 353

the bank address (bx,by) in the memory from which that pixel is derived. In the particular example, the following pixel derivations and translations are required.

Address (x,y) of pixel in non-aligned patch "p"	Aligned patch (px,py) from which pixel is derived	Bank address (bx,by) from which pixel is derived	Translation required from bank address (bx,by) to pixel address (x,y) in patch "p"
(0,0)	a (12,16)	(2,1)	(-2,-1)
(1,0)	a (12,16)	(3,1)	(-2,-1)
(2,0)	b (13,16)	(0,1)	(-2,-1) mod 4
(3,0)	b (13,16)	(1,1)	(-2,-1) mod 4
(0,1)	a (12,16)	(2,2)	(-2,-1)
(1,1)	a (12,16)	(3,2)	(-2,-1)
(2,1)	b (13,16)	(0,2)	(-2,-1) mod 4
(3,1)	b (13,16)	(1,2)	(-2,-1) mod 4
(0,2)	a (12,16)	(2,3)	(-2,-1)
(1,2)	a (12,16)	(3,3)	(-2,-1)
(2,2)	b (13,16)	(0,3)	(-2,-1) mod 4
(3,2)	b (13,16)	(1,3)	(-2,-1) mod 4
(0,3)	c (12,17)	(2,0)	(-2,-1) mod 4
(1,3)	c (12,17)	(3,0)	(-2,-1) mod 4
(2,3)	d (13,17)	(0,0)	(-2,-1) mod 4
(3,3)	d (13,17)	(1,0)	(-2,-1) mod 4

From the right hand column above, it will be noted that the translation from the bank address (bx,by) to the corresponding address in the non-aligned patch is constant for a particular non-aligned patch and in particular is equal to the negative of the misalignment (mx,my) of the non-aligned patch "p" from the base aligned patch "a", all translations being in modulus arithmetic with the modulus equal to the patch dimension.

Micron Technology Inc. et al.
Ex. 1010, p. 354

Yet another further complication arises with non-aligned patches, and that is that the patch may extend across the boundary between two or four pages. To provide flexibility, not all pages which make up an image and which are contiguous in the virtual address space need to be stored in the VRAM at one time, and pages are swapped between the paging memory and the VRAM as required. This results in those pages making up an image which are in the VRAM not necessarily being stored adjacent each other in the VRAM, but possibly being scattered in non-contiguous areas of the VRAM.

For example, Figure 11A represents four contiguous pages A, B, C, D in the virtual address space. When these pages are swapped into the physical memory 700, they may be scattered at, for example, page addresses (8,6), (4,8), (12, 12) and (6,10) in the VRAM, as represented in Figure 11B. Now, if it is desired to access a non-aligned patch "p" who base aligned patch "a" in page A has an x or y patch address of 31, then the non-aligned patch "p" may extend into page B, page C or pages B, C and D, depending on the direction of the misalignment. In the example shown specifically in Figure 11, the patch "p" to be accessed has a misalignment (mx,my) = (2,1) relative to base aligned patch "a" having patch address (px,py) = (31,31) in page A having page address (8,6) in the VRAM. It will be appreciated that, in addition to translating accessed pixels between their bank addresses (bx,by) and the addresses (x,y) in the non-aligned patch as described above with reference to Figure 10, it is also necessary to determine the various pages B, C, D which are to be accessed in addition to the basic page A and furthermore to determine the addresses in these other pages B, C, D of the aligned patches to be accessed, it being noted in the example that although the aligned patch "a" in page A has a patch address of (31,31), different patch addresses need to be used in other the pages B, C, D, that is (0,31), (31,0) and (0,0), respectively. The following table sets out, for each of the pixels in the patch "p" to be accessed: the page and patch address of the aligned patch from which that pixel is derived; the translation necessary from the patch address of the basic patch "a" in page A to the patch address of the patch from which the pixel is derived; the bank address from which the pixel is derived; and the translation necessary from this latter address to

Micron Technology Inc. et al.
Ex. 1010, p. 355

the address of the pixel in the patch "p".

Address (x,y) of pixel in non-aligned patch "p"	Page and aligned patch from which pixel is derived P/(px,py)	Translation necessary from address of patch "a" to address (px,py) of patch from which pixel is derived	Bank address (bx,by) of pixel	Translation necessary from bank address (bx,by) to pixel address (x,y) in patch "p"
(0,0)	A/(31,31)a	(0,0)	(2,1)	(-2,-1)
(1,0)	A/(31,31)a	(0,0)	(3,1)	(-2,-1)
(2,0)	B/(0,31)b	(1,0) mod 32	(0,1)	(-2,-1) mod 4
(3,0)	B/(0,31)b	(1,0) mod 32	(1,1)	(-2,-1) mod 4
(0,1)	A/(31,31)a	(0,0)	(2,2)	(-2,-1)
(1,1)	A/(31,31)a	(0,0)	(3,2)	(-2,-1)
(2,1)	B/(0,31)b	(1,0) mod 32	(0,2)	(-2,-1) mod 4
(3,1)	B/(0,31)b	(1,0) mod 32	(1,2)	(-2,-1) mod 4
(0,2)	A/(31,31)a	(0,0)	(2,3)	(-2,-1)
(1,2)	A/(31,31)a	(0,0)	(3,3)	(-2,-1)
(2,2)	B/(0,31)b	(1,0) mod 32	(0,3)	(-2,-1) mod 4
(3,2)	B/(0,31)b	(1,0) mod 32	(1,3)	(-2,-1) mod 4
(0,3)	C/(31,0)c	(0,1) mod 32	(2,0)	(-2,-1) mod 4
(1,3)	C/(31,0)c	(0,1) mod 32	(3,0)	(-2,-1) mod 4
(2,3)	D/(0,0)d	(1,1) mod 32	(0,0)	(-2,-1) mod 4
(3,3)	D/(0,0)d	(1,1) mod 32	(1,0)	(-2,-1) mod 4

A representation of the locations of the pixels in the four aligned patches is shown in Figure 11C.

In the example, the basic patch "a" has a patch address (px,py) of (31,31) and the non-aligned patch "p" to be accessed has a misalignment (mx,

Micron Technology Inc. et al.
Ex. 1010, p. 356

my) of (2,1) relative to the basic patch "a". In the general case of a base patch address (px, py), where 0<=px, py<=31, and a misalignment (mx, my), where 0<=mx, my<=3, the table of Figure 12 sets out which page A, B, C or D should be used when accessing a pixel at bank address (bx,by), where 0<= bx,by<=3, in dependence up bx, by, mx, my, px and py, and the table of Figure 13 sets out the X patch address px, or px + 1 mod 4, and the Y patch address py, or py + 1 mod 4, which should be used in order to obtain the address of the aligned patch a, b, c or d to be accessed, in dependence upon bx, by, mx and my. The increment is calculated using modular arithmetic of base 32. It is also to be noted that for all pixels where (mx, my) < > (0,0), a translation of (-mx, -my) is required between the bank address (bx,by) from which the pixel is derived and the address (x,y) of the pixel in the non-aligned patch "p".

Having described various addressing functions which it is required to be performed, there now follows a description in greater detail of the apparatus for performing these functions.

As described above with reference to Figure 4, the VRAM 700 is addressed by the address processor 310 via the address translator 740, communicates data with the grid processor 312 via the exchange 326 and provides data to the video processor 34. A greater degree of detail of the address translator, VRAM and exchange is shown in Figure 14.

The address translator 740 receives a 48-bit virtual address on bus 319 of a patch origin address. The translator determines whether the required page(s) to access the addressed patch are resident in the VRAM physical memory 700. If not, a page or superpage fault is flagged on line 748, as will be described in detail below. However, if so, the address translator determines the addresses in the sixteen XY banks of the physical memory of the sixteen pixels making up the patch, and addresses the memory 700 firstly with the Y addresses on the sixteen sets of 9-bit lines A(0) to A(15) and then with the X addresses on these lines. The X and Y addresses are generated under control of the X/Y select signal on line 713.

Micron Technology Inc. et al.
Ex. 1010, p. 357

The exchange 326 includes a read surface shifter 742 and a write surface shifter 744. Pixel data is transferred, during a read operation, from the memory 700 to the read surface shifter 742 by a set of sixteen 32-bit data lines D"(0) to D"(15), and, during a write operation, from the write surface shifter 744 to the memory 700 by the same data lines D"(0) to D"(15). The read and write surface shifters 742, 744 receive 4-bit address data from the address translator on line 770, consisting of the least significant two bits of the X and Y address data. This data represents the misalignment (mx, my) of the accessed patch "p" from the basic aligned patch "a". The purpose of the surface shifters is re-order the pixel data in non-aligned patches, that is to apply the translation (-mx, -my) when reading and an opposite translation (mx, my) when writing. Pixel data to be written is supplied by a crossbar 327 forming part of the exchange 326 to the write surface shifter 744, and pixel data which has been read is supplied by the read surface shifter 742 to the crossbar 327, on the 512-bit line 750 made up of a set of 16 32-bit lines. The write surface shifter also receives on line 745 16-bit write enable signals WE(0) - WE(15) from the crossbar 327 one for each pixel, and the write surface shifter 744 re-organises these signals in accordance with the misalignment (mx, my) of the patch "p" to be accessed to provide the sixteen column write enable signals WE"(0) to WE"(15). Each of these signals is then ANDed with a common CAS signal on line 709 to form sixteen CAS signals CAS(0) to CAS(15), one for each of the sixteen banks of memory. This enables masking of pixels within a patch during writing, taking into account any misalignment of the patch.

The address translator 740 will now be described in more detail primarily with reference to Figure 15. The translator 740 includes as shown, a contents addressable memory (CAM) 754, a page address table 756, a near-page-edge table 758, and X and Y incrementers 760X, 760Y. The translator 740 also includes sixteen sections 764(0) to 764(15), one for each output address line A(0) to A(15), and thus for each memory bank B(0) to B(15).

The translator 740 receives a 48-bit virtual address of the origin (0,0) pixel of a patch on the bus 319. It will therefore be appreciated that up to 2^{48} (i.e. 281, 474, 976, 710, 656) different pixels can be addressed. Many

Micron Technology Inc. et al.
Ex. 1010, p. 358

formats of the 48-bit address can be employed, but the following example will be considered in detail.

	BITS	IDENTITY
LSB	0,1	X misalignment (mx) of patch (p) to be accessed relative to basic aligned patch (a)
	2-6	X address (px) of aligned patch (a) in page A
	7,8	X address of page A
	9-15	X portion of superpage address
	16,17	Y misalignment (my) of patch (p) to be accessed relative to basic aligned patch (a)
	18-22	Y address (py) of aligned patch (a) in page A
	23,24	Y address of page A
	25-31	Y portion of superpage address
MSB	32-47	Image ID portion of superpage address

The bits identifying the superpage (i.e. bits 9 to 15, 25 to 31 and 32 to 47) are supplied to the CAM 754. The CAM 754 is an associative memory device which compares the incoming 30-bit word with all of the words held in its memory array, and if a match occurs it outputs the location or address in the memory of the matching value on line 767. The CAM 754 has a capacity of 128 32-bit words. Thirty of these bits are used to store the virtual address of a superpage which is registered in the CAM 754. Thus up to 128 superpages can be registered in the CAM. One of the other bits is used to flag any location in the CAM which is unused. The remaining bit is spare. Figure 16 illustrates how the CAM 754 operates. Upon input of a 30-bit superpage address, e.g. 01234569 (hex), this input value is compared with each of the contents of the CAM. If a match is found and provided the unused flag is not set, the address in the CAM of the match is output, e.g. 1 in the illustration. If no match is found with the contents at any of the 128 addresses of the CAM, then a superpage fault is flagged on line 748S, and the required superpage is then set up in the CAM in the manner described in

Micron Technology Inc. et al.
Ex. 1010, p. 359

detail later.

Referring back to Figure 15, the 7-bit superpage identification output from the CAM 754 on line 767 is used as part of an address for the page address table 756, implemented by a 4k word x 16-bit SRAM. The remaining 5 bits of the address for the page table 756 are made up by: bits 7, 8, 23 and 24 of the virtual address which identify the page within a superpage; and an X/Y select signal on line 713. The page table 756 has registered therein the X and Y page addresses in the VRAM 700 of: a) the basic page A in which the pixel to be accessed is located; b) the page B which is to the right of the page A in the virtual address space; c) the page C which is above the page A in the virtual address space; and d) the page D which is to the right of page C and above page B in the virtual address space, and these addresses are output on lines 771A to 771D, respectively. If these pages A to D are required, but are not stored in the VRAM 700 and thus are not registered in the page table 756, then a page fault is flagged on a line 748p (as described below with reference to Figure 21) and the required page of data is then swapped into the VRAM 700 in the manner described in detail below. However, if all of the pages A to D which may possibly need to be accessed are stored, their addresses are made available on the lines 771A to 771D to all of the sections 764(0) to 764(15), the Y or X addresses being output depending on the state of the X/Y select signal on line 713.

Bits 2 to 6 and 18 to 22 of the virtual address are also supplied to each of the sections 764(0) to 764(15) on lines 772X and 772Y. These denote the patch address (px, py). The X and Y patch addresses together with bits 0,1, 16 and 17 of the virtual address (which indicate the misalignment mx, my of the patch p to be accessed) are also supplied to the near-page-edge table 758, implemented using combinatorial logic, which provides a 2-bit output to the sections 764(0) to 764(15) on line 774, with one bit being high only if the patch X address px is 31 and the X misalignment mx is greater than zero and the other bit being high only if the patch Y address py is 31 and the Y misalignment my is greater than zero.

Furthermore, the X and Y patch addresses (px, py) are also supplied

Micron Technology Inc. et al.
Ex. 1010, p. 360

to the X and Y incrementers 760X, 760Y, and these incrementers supply the incremented values px + 1, mod 32 and py + 1, mod 32, to each of the sections 764(0) to 764(15) on lines 776X, 776Y.

The four bits 0,1, 16 and 17 giving the misalignment mx and my are also supplied to the sections 764(0) to 764(15) on lines 770X, 770Y and are also supplied to the surface shifters 742, 744 on line 770.

Each section 764(0) to 764(15) comprises: a page selection logic circuit 778; X and Y increment select logic circuits 780X 780Y; X and Y 4:1 4-bit page address multiplexers 782X, 782Y; X and Y 2:1 5-bit patch address multiplexers 784X, 784Y; and a 2:1 9-bit address selection multiplexer 786.

The page selection logic circuit 778 implemented using combinatorial logic, provides a 2-bit signal to the page address multiplexers 782X,Y to control which page address A, B, C or D to use. The page selection logic circuit 778 performs this selection by being configured to act as a truth table which corresponds to the table of Figure 12. The circuit 778 receives the 2-bit signal on line 774 from the near-page-edge table 758 and this determines which of the four columns of the table of Figure 12 to use. The circuit 778 also receives the misalignment (mx, my) on lines 770X, 770Y, and this data in combination with which section 764(0) to 764(15) (and thus which bx and by applies) determines which of the four rows in Figure 12 to use. The X and Y page address multiplexers 782X, 782Y therefore supply appropriate page address as four bits to complementary inputs of the X/Y address selection multiplexer 786.

The increment selection logic circuits 780X, 780Y, which are implemented using combinatorial logic, receive the respective X and Y misalignments mx, my and provide respective 1-bit signals to control the patch address multiplexers 784X, 784Y. The increment selection circuits perform this selection by being configured to act as truth tables which correspond to the upper and lower parts, respectively, of the table of Figure 13. It will be noted that selection depends upon the misalignment mx or my in combination with the bx or by position of the memory bank (and thus

which of the sections 764(0) to 764(15) is being considered). The X and Y patch address multiplexers 784X, 784Y therefore output the appropriate 5-bit patch addresses px or px + 1 (mod. 32) and py or py + 1 (mod. 32) which are combined with the X and Y page addresses at the inputs to the X/Y selection multiplexers 786. This latter multiplexer receives as its control signal the X/Y selection signal on line 713 and therefore outputs the 9-bit X or Y address appropriate to the particular section 764(0) to 764(15).

The address translator 740 therefore deals with the problems described above of addressing pixels from different aligned patches a, b, c, d in the memory 700 when a patch "p" to be accessed is misaligned, and of addressing pixels from different pages A, B, C, D in the memory 700 when a patch "p" to be accessed extends across the boundary of the basic page A.

It is still necessary also to perform a translation of the pixel positions in the accessed patch of (-mx,-my) if reading, or (mx,my) if writing. This is performed by the surface shifter 742 for reading and the surface shifter 744 for writing. The read surface shifter 742 will now be described with reference to Figures 17 and 18.

The read surface shifter 742 comprises a pair of 4 x 4 32-bit barrel shifters, 788X, 788Y. The X barrel shifter 788X has four banks 790X(0) to 790X(3) of multiplexers arranged in one direction, and the outputs of the X barrel shifter 788X are connected to the inputs of the Y barrel shifter 788Y, which has four banks 790Y(0) to 790Y(3) of multiplexers arranged in the orthogonal direction. As control signals, the X and Y barrel shifters 788X, Y receive the X and Y misalignments mx, my, respectively.

One of the banks of muliplexers 790X(0) is shown in greater detail in Figure 18, and comprises four 32-bit 4:1 multiplexers 792(0) to 792(3). The data from bank (0,0) is supplied to inputs 0, 3, 2 and 1, respectively, of the multiplexers 792(0) to 792(3). The data from bank (1,0) is supplied to inputs 1, 0, 3 and 2, respectively, of the multiplexers 792(0) to 792(3). The data from bank (2,0) is supplied to the inputs 2, 1, 0 and 3, respectively, of the multiplexers 792(0) and 792(3). The remaining data from bank (3,0) is

Micron Technology Inc. et al.
Ex. 1010, p. 362

supplied to the remaining inputs 3, 2, 1, 0, respectively, of the multiplexers 792(0) to 792(3). The other banks of multiplexers 790X(1) to 790X(3) in the X barrel shifter 788X are similarly connected, and the banks 790Y(0) to 790Y(3) in the Y barrel shifter 788Y are also similarly connected. It will therefore be appreciated that the read surface shifter performs a translation with wrap-around in the -X direction of mx positions and a translation with wrap-around in the -Y direction of my positions as shown in Figure 19.

As shown in the drawings, the write surface shifter 744 may be provided by a separate circuit to the read surface shifter. In this case the write surface shifter is configured similarly to the read surface shifter, except that the inputs 1 and 3 to the multiplexers 792 in the barrel shifter banks are transposed. This results in translations of +mx and +my in the X and Y directions, rather than -mx and -my for the read surface shifter. The part of the write surface shifter which operates on the write enable signals WE(0) to WE(15) is identical to the part which operates on the data signals, except that the signals are 1-bit, rather than 32-bit.

As an alternative to employing separate circuits for the read and write surface shifters 742, 744, a single circuit may be employed, with appropriate data routing switches, and in this case translation provided by the surface shifter may be switched between (-mx, -my) and (+mx, +my), in dependence upon whether the memory is being read or written, as described with reference to Figures 46 and 47.

As mentioned above, if a required superpage is not registered in the CAM 754, then a superpage fault is flagged, on line 748S. This superpage fault is used to interrupt the address processor 310, which is programmed to perform a superpage interrupt routine as follows. Firstly, the address processor checks whether the CAM 754 has any space available for a new superpage to be registered. If not, the address processor selects a registered superpage which is to be abandoned in the manner described below and causes the, or each, page of that superpage which is stored in the VRAM 700 to be copied to its appropriate location in the paging memory.

Micron Technology Inc. et al.
Ex. 1010, p. 363

The registration of that superpage is then cancelled from the CAM 754. Secondly, the new superpage is registered in the CAM 754 at the, or one of the, available locations.

In order to select which superpage to abandon, a determination is made as to which is the least recently used (LRU) superpage which is registered in the CAM 754. To do this, a 128 x 16-bit LRU table 802 is provided, as illustrated in Figure 20. Each of the 128 addresses represents a respective one of the superpages registered in the CAM 754. The 7-bit superpage identification output from the CAM 754 on line 767 is used to address the LRU table 802 each time the superpage identification changes, as detected by the change detector 804. The change detector 804 also serves to increment a 16-bit counter 806, and the content of the counter 806 is written to the addressed location in the LRU table 802.

Accordingly, for all of the registered superpages, the LRU table contains an indication of the order in which those superpages were last used. When registering a new superpage in the CAM 754, the address processor 310 checks the contents of the LRU table 802 to determine which superpage has the lowest count and in that way decides which superpage to abandon.

As also mentioned above, if the required pages of the registered superpage are not stored in the VRAM 700, a page fault is flagged, on line 748P. The page fault generator is shown in Figure 21, and comprises a page fault table 794 constituted by a 2k x 4-bit SRAM, a set of three AND gates 796B, C, D and an OR gate 798. The page fault table 794 is addressed by the 7-bit superpage identity code on line 767, and by the X and Y page addresses on line 768X, Y. At each address, the page fault table 794 contains a 4-bit flag in which the bits denote whether the basic addressed page A and the pages B, C and D, respectively, to the right, above, and to the right and above, page A are stored in the VRAM 700. The page B flag is ANDed by gate 796B with the bit of the near-page-edge signal on line 774 denoting whether the patch "p" to be accessed extends across the boundary between pages A and B. Similarly, the page C flag is ANDed by gate 796C with the bit of the near-page-edge signal on line 774 denoting whether the

Micron Technology Inc. et al.
Ex. 1010, p. 364

patch "p" to be accessed extends across the boundary between pages A and C. Furthermore, the page D flag is ANDed by gate 796D with both bits of the near-page-edge signal, which in combination denote whether the patch "p" to be accessed extends in page D above page B and to the right of page C. The outputs of the three AND gate 796B, C, D and the page A flag are then ORed by the OR gate 798, the output of which provides the page fault flag on line 748P.

From the above, it will be appreciated that a page fault is always generated if the basic page A is not stored in the VRAM, but if page B, C or D is not stored in the memory, a page fault will be generated in response thereto only if the respective page B, C or D will be used, as indicated by the two bits of the near-page-edge signal on line 774.

The page fault signal on line 748P is used to interrupt the address processor 310. The address processor then searches a table in its PRAM 318 for a spare page address in the VRAM 700, causes the required page to be swapped into the VRAM at the spare page address, and update the table in its PRAM 318.

GRID PROCESSOR AND EXCHANGE

As described above with reference to Figure 4, in the operation of the preferred embodiment, the exchange 326 and the VRAM 700 communicate in patches of sixteen pixels of data, each pixel having 32 bits. Furthermore, the grid processor 312 has sixteen processors, each of which processes pixel data and communicates with the exchange 326. Also, the grid processor 312 and the address processor 310 can communicate address data via the FIFO 332.

The exchange 326 includes a crossbar 377, and a logical implementation of the crossbar 377 and of the grid processor 312 is shown in more detail in Figure 22. As shown, the crossbar 377 comprises sixteen 16:1 32-bit data multiplexers 602(0) to 602(15); sixteen 16:1 1-bit write enable multiplexers 603(0) to 603(15); a 512-bit bidirectional FIFO 604 for pixel

Micron Technology Inc. et al.
Ex. 1010, p. 365

data; and a 16-bit bidirectional FIFO 605 for the write enable signals. During a read operation, the 16 pixels of a 4 x 4 patch are supplied from the VRAM 700 (Figure 8) via the read surface shifter 742 and via the FIFO 604 as data D(0) to D(15) to the sixteen inputs of each data multiplexer 602(0) to 602(15). During a write operation, the data multiplexers 602(0) to 602(15) supply data D(0) to D(15) via the FIFO 604 and the write surface shifter 744 to the VRAM and the write enable multiplexers 603(0) to 603(15) supply write enable signals WE(0) to WE(15) via the FIFO 605 to the write surface shifter 744. The FIFOs 604, 605 and also the FIFO 332 are employed so that the grid processor 312 does not need to be stalled to take account of different access speeds of the VRAM 700 in dependence upon whether page-mode of non-page-mode access is taking place.

Each of the data multiplexers 602(0) to 602 (15) is associated with a respective one of sixteen processors 606(0) to 606(15) and communicates therewith respective data signals D'(0) to D'(15), which are logically 32 bits, but which in practice may be implemented physically as 16 bits, with appropriate multiplexing techniques. The data signals D'(0) to D'(15) are also supplied to respective parts of the bus 331. Also, each of the write enable multiplexers 603(0) to 603(15) is associated with a respective one of the sixteen processors 606(0) to 606(15) to 606(15) which supply respective 1-bit write enable signals WE'(0) to WE'(15) to the write enable multiplexers. Each processor 606(0) to 606(15) provides a logical control signal CO(0) to CO(15) to control both its associated data multiplexer 602 and write enable multiplexer 603. Thus, during writing to the memory, any processor may provide any respective one of the data signals by providing the number 0 to 15 of the required data signal as its control signal to its data and write enable multiplexers. Furthermore, during reading from the memory, any processor may read any of the data signals by providing the number 0 to 15 of the required data signal to its data multiplexer. Thus, there is no restriction on data being processable only relative to a particular processor, and each processor can select and control the routing of its own data.

The crossbar 377 shown in Figure 22 is simplified for reasons of clarity, and shows, for example, bi-directional multiplexers, which in

practice are difficult to implement. A modified form of the exchange, incorporating the crossbar and the surface shifters, is shown in Figure 46.

The exchange of Figure 46 comprises sixteen sections, of which one typical section 326(i) is shown for simplicity. The data D"(i) from the memory is supplied via a buffer BA(i) and register RA(i) to one input of a 2:1 multiplexer SA(i) acting as a two-way switch. The output of the switch SA(i) is fed to an input i of the surface shifter 743 which performs surface shifting for read and for write. The corresponding output i of the surface shifter 743 is fed to one input of a multiplexer switch SB(i) and is also fed back to the data D"(i) input via a register RB(i) and a tri-state buffer BB(i). The output of the switch SB(i) is input to a FIFO(i), the output of which forms the other input of switch SA(i) and is also fed to one input of a further switch SC(i). The set of sixteen data lines D(0) to D(15) connect the exchange sections 326(0) to 326(15) and the output of switch SC(i) is connected to data line D(i). In the general case, the output of each switch SC(0) to SC(15) is connected to the data line of the same number.

The sixteen inputs of a 16:1 multiplexer MUX(i) are connected to the data lines D(0) to D(15), and the output of the multiplexer MUX(i) is connected via a register RC(i) and a tri-state buffer BC(i) to the respective processor PROC(i) via the data line D'(i). The output of the multiplexer MUX(i) is also connected to the other input of switch SB(i). Furthermore, the data line D'(i) from the processor PROC(i) is also connected via a buffer BD(i) and a register RD(i) to the other input of the switch SC(i).

The control signal CO(i) for the multiplexer MUX(i) is provided by a switch SD(i) which can select between a hardwired value i or the output of a register RE(i) which receives its input from the output of the register RD(i).

Also, control signals CSB, CSC, CSD and CBC are supplied to the multiplexer switches SB(0) to (15), the multiplexer switches SC(0) to (15), the multiplexer switches SD(0) to SD(15), the tri-state buffers BC(0) to (15) from the microcode memory 308 (Figure 4) of the processing section 321. Furthermore, control signals CSA, CBB and CSS derived from the microcode

Micron Technology Inc. et al.
Ex. 1010, p. 367

memory 307 of the address processing section 309 are supplied to the multiplexer switches SA(0) to (15), the tristate buffers BB(0) to (15) and the surface shifter 743.

The exchange 326 of Figure 46 is operable in three modes. In a read mode, the processors PROC(0) to PROC(15) can read the memory; in a write mode, they can write to the memory; and in a transfer mode, they can transfer pixel data between each other. The values of the control signals for these three modes are as follows:

	CSA	CSB	CSC	CSD	CBB	CBC	CSS
READ	0	0	0	0/1	0	1	0
WRITE	1	1	1	0/1	1	0	1
TRANSFER	X	X	1	1	X	0/1	X

It should be noted that the control signal CSD can select between a "straight-through" mode in which each multiplexer MUX(i) selects its input i and thus data D(i), or a "processor-selection" mode in which it selects an input j and thus data D(j) in accordance with the value j which the processor has loaded into the register RE(i).

The effective configuration of a generalised one of the exchange sections 326(i) of Figure 46 in the read mode is shown in Figure 47A. In this configuration, the data path from the data line D"(i) is via the register RA(i) to the surface shifter 743. In the read mode, the surface shifter applies a shift of (-mx,-my) (mod. 4) to the data paths. From the surface shifter, the data path continues via the FIFO(i) to the data line D(i). The multi-plexer MUX(i) can select if CSD=0 the straight-through path in which its output is D(i), or if CSD=1 the processor selection path in which its output is D(j) were j is the value loaded into the register RE(i). The output data passes via the register RC(i) as data D'(i) to the processor PROC(i).

The effective configuration of the exchange section 326(i) in the write mode is shown in Figure 47B. The data D'(i) from the respective

Micron Technology Inc. et al.
Ex. 1010, p. 368

processor PROC(i) passes via the register RD(i) to the data line D(i). The multiplexer MUX(i) can select, if CSD=0, the straight-through path in which its output is D(i), or if CSD=1 the processor selection path in which its output is D(j) where j is the value loaded into the register RE(i). The output data passes via the FIFO(i) to the surface shifter 743. In the write mode, the surface shifter applies a shift (+mx,+my) (mod. 4) to the data paths. From the surface shifter, the data path continues via the register RB(i) as data D"(i) to the VRAM 700.

It should be noted from Figures 46 and 47B that, in the write mode, the write-enable signal follows the same path WE'(i) to WE(i) to WE"(i) as the data signal path D'(i) to D(i) to D"(i). Thus these paths are logically 33 bits made up from 32 bits for the data signal and 1 bit for the write-enable signal.

In the transfer mode, the effective configuration of the exchange section 326(i) is as shown in Figure 47C. In this configuration the control signal CSD to the switch SD(i) is set to 1 so that the multiplexer MUX(i) receives as its control signal the value j loaded into the register RE(i). There are four phases to a transfer. In the first phase the processors output the values j of the data D(j) which they wish to receive as the lowest four bits of their data lines, and these values j are clocked into the registers RD(i). In the second phase , the processors output the data to be transferred out, and this data is clocked into the registers RD(i), while the values j are clocked out of the registers RD(i) and into the registers RE(i), thus setting the multiplexers MUX(i) to receive the data on the respectively selected lines D(j). In the third phase, the data in the registers RD(i) is clocked out onto the lines D and each multiplexer MUX(i) receives and outputs the data on respectively selected line D(j). In the fourth phase, the outputs of the multiplexers are clocked into the registers RC(i) and the tristate buffers BC(i) are enabled so that the processors can transfer in the data from the buffers BC(i). Thus, each processor PROC(i) receives the data (j) from the processor PROC(j) which was selected by the processor PROC(i) by its output value j in the first phase.

Micron Technology Inc. et al.
Ex. 1010, p. 369

Referring back to Figure 22, the processors 606(0) are connected to the data broadcast bus 323 and to a priority encoder 614 having 16 sections and which is associated with the sequencer 329. The processors 606(0) to 606(15) communicate address data with the data broadcast bus 323 and the FIFO 332 connects the data broadcast bus 323 with the address processor 310. The processors 606(0) to 606(15) can also supply respective "unsatisfied" signals US(0) to US(15) and respective "X waiting" signals XW(0) to XW(15) to the respective sections of the priority encoder and can receive respective "process enable" signals EN(0) to EN(15) from the respective sections of the priority encoder 614. Lastly, the priority encoder 614 has a sequencer enable (SE) output on line 618 to the sequencer 329 which controls the sequence of processing of a series of microcode instructions by the processors 606.

The purpose of the priority encoder 614 is to provide high efficiency in the accessing by the processors 606 of the memory 700. In order to do this, the encoder 614 and processors perform the following process, which is shown in the flow diagram of Figure 23. In Figure 23, the left-hand three columns contain steps which are taken by the processors 606(0)....606(i).....606(15), or PROC(0).....PROC(i).....PROC(15), in parallel with each other, the right-hand column contains steps performed by the priority encoder.

At the beginning of each microcode instruction, there are a series of initialisation steps 620 to 628. In steps 620 to 625, those processors which require access to the memory set (1) their respective unsatified signals US and reset (0) their X waiting signals XW, and those processors which do not require access reset (0) their unsatisfied signals US and their X waiting signals XW. In steps 626, 628, the priority encoder resets (0) the process enable signals EN for all of the processors and also resets (0) the sequencer enable signal SE.

After initialisation, the priority encoder 614 checks through the XW signals, starting with XW(0) in step 630 to find any processor which is X waiting, and if a match is found (step 632) at a processor, designated

Micron Technology Inc. et al.
Ex. 1010, p. 370

PROC(q), then the routine proceeds to step 640. If a match is not found, however, in step 632, then the priority encoder checks through the US signals, starting with US(0) in step 634 to find a processor which is unsatisfied, and if a match is found (step 636) for a processor, designated PROC(q), then the routine proceeds to step 640. If a match is not found, however, in step 636, then this indicates that all processors are satisfied, and accordingly the microcode program can proceed. Therefore, the sequence enable signal SE is set in step 638, and the routine terminates.

In step 640, the process enable signal EN(q) for the selected processor PROC(q) is set. In steps 642, each processor determines whether it is unsatisfied, and if not exits the subroutine of steps 642 to 654. For any processor which is unsatisfied, then in steps 644, that processor determines whether it is the selected processor, and if so supplies, in step 645, to the data broadcast bus 323 as (xq, yq) the virtual address of the base pixel (0,0) of the patch of pixel data which it wishes to process. This address is supplied via the FIFO 332 to the address processor 310, which in response accesses the appropriate locations in the memory 700, swapping in and out pages of pixel data, if required, as described above.

Then, in steps 646, each unsatisfied processor determines whether the y address yi of its required patch of pixel data is equal to the y address yq of the patch which is being accessed. If not, then the processor exits the subroutine of steps 642 to 654. If, however, yi = yq, then in step 648 the processor determines whether the X address xi of its required patch of pixel data is equal to the X address xq of the patch which is being accessed. If so, then the processor resets (0) its unsatisfied signal US(i) and X waiting signal XW(i) in step 650, and accesses the memory for read or write, as appropriate, in step 652. The processor then exits the subroutine of steps 642 to 654. If, in step 648, xi<>xq then in step 654 the X waiting signal XW(i) is set (1), and then the subroutine is exited.

Upon exit from the subroutine of steps 642 to 654 of all processors PROC(0) to PROC(15), the routine proceeds to step 656, where the priority encoder resets (0) the process enable signal EN(q) for the selected processor.

Micron Technology Inc. et al.
Ex. 1010, p. 371

The routine then loops back to step 630.

It will be appreciated from the above that (A) the lowest numbered processor (an arbitrary choice) which is unsatisfied is selected and given access to the memory initially, together with any other processors which require access to the same address as that selected processor. Then, (B) of any remaining unsatisfied processors which require access to the same y address as the selected processor, the lowest numbered processor is given access, together with any other processors requiring the same address. Then, (C) of any remaining unsatisfied processors which require access to the same y address as the last satisfied processor, the lowest numbered processor is given access, together with any other processors which require access to the same address. Step C is repeated, if necessary, and then steps A and B are repeated until all of the processors have been satisfied. The next microcode instruction sequence is then processed.

As example of the operation of the priority encoder and processors in accessing the memory will now be described with reference to the table of Figure 24. In the example, PROC(0) to (3) and (8) to (11) require access to the patches having the base pixel X and Y addresses listed in column 660 of the table, the addresses being in hexadecimal notation. Thus, after the initialisation routine, JS(0) to (3) and (8) to (11) are set to 1 and the other US signals and the XW signals are reset to 0, as shown in column 662.

In the first loop of the main routine, PROC(0) is selected, i.e. q = 0, and thus accesses the memory at (1234, 1234). Because PROC(1) requires the same address, it also becomes satisfied, i.e. US(1) = 0, and accesses the memory at (1234, 1234). Furthermore, because PROC(2) and PROC(10) require the same Y address as PROC(0), they become X waiting, i.e. XW(2) = XW(10) = 1. This is shown in column 664.

In the next loop of the main routine, PROC(2) is found to be X waiting XW(2) = 1), and thus PROC(2) is selected, i.e. q = 2. Therefore PROC(2) becomes satisfied, (US(2) = XW(2) = 0), as shown in column 666, and accesses the memory at (1235, 1234).

Micron Technology Inc. et al.
Ex. 1010, p. 372

In the next loop of the routine, PROC(10) is found to be X waiting XW(10) = 1), and thus PROC(10) is selected, i.e. q = 10. Therefore PROC(10) becomes satisfied, (US(10) = XW(10) = 0), as shown in column 668, and accesses the memory at (1236, 1234).

In the next loop of the routine, no processor is found to be X waiting, and PROC(3) is found to be the first completely unsatisfied processor, i.e. US(3) = 1, YS(3) = 0. Therefore PROC(3) is selected (q = 3), becomes satisfied (US(3) = XW(3) = 0) and accesses the memory at (1235, 1235). Also because PROC(11) has the same Y address as PROC(3), PROC(11) becomes X waiting, i.e. XW(11) = 1, as shown in column 670.

In the next loop of the routine, PROC(11) is found to be the only X waiting processor, (US(11) = XW(11) = 1). Therefore, PROC(11) is selected (q = 11), becomes satisfied (US(11) = XW(11)= 0) and accesses the memory at (1236, 1235), as shown in column 672.

In the next loop of the routine, PROC(8) is found to be the first unsatisfied processor (US(8) = 1). Therefore, PROC(8) is selected (q = 8), becomes satisfied, and accesses the memory at (1234, 1236). Furthermore, because PROC(9) requires the same address, it also becomes satisfied (US(9) = 0) and accesses the memory.

During the next loop of the routine, no processors are found to be unsatisfied, and therefore the sequence enable signal SE is set and the next microcode instruction is processed.

By using the priority encoder as described above, processors which require access to the same patch can access that patch simultaneously. Furthermore, when a plurality of processors require access to different patches having the same Y address, their accesses are made immediately one after the other, in "page mode". Therefore the address translator does not need to re-latch the Y address(es) in the Y address latches 706(0) to (15) (Figures 8 and 14) between such accesses. Thus, a considerable improvement in performance is achieved as compared with a case where the

Micron Technology Inc. et al.
Ex. 1010, p. 373

processors PROC(0) to (15) access their required patches one at a time, sequentially and without reference to any similarity between the addresses to be accessed.

In the system described above, up to sixteen pixels in a patch are processed in parallel by sixteen processors. Preferably, the system is also arranged so that a group of patches, for example, up to 32 patches, are processed in series in order to reduce pipeline start and finish overheads. In this case, the method of operation may be modified, as compared with that shown in Figure 23, in order to increase efficiency, by permitting any processor requiring access to, say, a jth pixel in the group to request that pixel without firstly waiting for all the other processors to complete access to their (j-1)th pixels in the group. To do this, between steps 623 and 630 in Figure 23, for each processor the step "set address of first required pixel in group as (xi, yi)" is include or each processor PROC(i). Furthermore, steps 650 and 652 for each processor as shown in Figure 23 are replaced by the steps shown in Figure 48. In step 682, the memory is accessed at address xi, yi for the particular processor PROC(i). Then, in step 684, it is determined whether or not the processor PROC(i) requires access to a further pixel in the group, and if not in step 686, the unsatisfied flag US(i) and the X waiting flag XW(i) are both reset to 0, similarly to step 650 in Figure 23. However, if so in step 684, then in step 688 the processor PROC(i) sets the address of the next required pixel as (xi, yi). Then, in step 690, it is determined whether or not the new y address yi is equal to the Y address yq of the last accessed pixel. If so, then in step 692, the X waiting flag XW(i) is set to 1, whereas if not, then in step 694, the X waiting flag XW(i) is reset to 0. After steps 692 or 694, the routine proceeds to step 656 as in Figure 23. It will therefore be appreciat ' that, once any processor has accessed a pixel in its series of required t. ty-two pixels, it can immediately make itself ready to access the next pixel in its series, irrespective of how many of their required thirty-two pixels each of the other processors has accessed. This therefore makes good use of the page mode accessing of the VRAM in which a series of pixels with the same Y address are accessed without the need to re-latch the Y address between each access.

Micron Technology Inc. et al.
Ex. 1010, p. 374

A problem which can arise with the modification of Figure 48 is that some of the processors can inordinately race ahead of others of the processors in accessing their thirty-two pixels. For example, in the case where the processors require access to many different Y addresses, it may arise that PROC(0) accesses all of its thirty-two required pixels first, then PROC(1) accesses its thirty-two pixels, and so on. In order to obviate this problem, the following further modification may be made. Basically, access is permitted with the following order of priority: (a) of highest priority, processors which require access in page-mode (i.e. with the same Y address as the last access) are arbitrated for access; (b) of second priority, processors which have progressed least through their series of thirty-two accesses are arbitrated for access; and (c) of lowest priority there is arbitrary selection of any processors still requiring access. This is achieved by maintaining in a register file of each processor a respective local pointer LP(i) indicating which of its 32 accesses it is waiting for, a common low watermark pointer WM for all the processors, and a common high watermark pointer HP for all the processors. Furthermore, the significance of each unsatisfied flag US(i) is Figures 23 and 48 is modified so that US(i)=1 only if LP(i)=WM and the processor PROC(i) is unsatisfied. The process of Figures 23 and 48 is then modified as follows. In the initialisation steps 622 to 625, the additional steps are included of resetting to zero LP(i) and WM in the register files of all processors, and setting HP to the number of accesses in the series, usually 31. The step 642 in Figure 23 is replaced by "LP(i)<>HP?". Furthermore, accompanying step 682 in Figure 48, where a processor accesses the memory, it also increments it local pointer to LP(i)+1. This then has the affect of dealing with priorities "a" and "b" described above. In order to deal with the priority "c" above, an additional decision is included between steps 636 and 638 in Figure 23. If the low watermark pointer WM is less than the high pointer HP, then the low watermark pointer in each of the processor register files is incremented to WM+1, and the process loops back to step 630. However, if WM=31 the process proceeds to step 638. From the above, it will be appreciated that the low watermark pointer is always less than or equal to the lowest local pointer LP(i). When there is no page mode, only those processors whose local pointer LP(i) is equal to the low watermark pointer WM are initially

Micron Technology Inc. et al.
Ex. 1010, p. 375

involved in the access arbitration. If there are none, the watermark pointer is incremented, unless it is equal to HP.

SPLIT-LEVEL PATCHES

It will be noted from the above that the memory is capable of storing pixel data of 32 bits and that the grid processor is capable of processing pixel data logically of at least 32 bits. In some applications, pixel data having a resolution as great as 32 bits is not needed, and all that may be required is 16-bit or 8-bit pixel data. In such cases it is possible to use only 16 or 8 bits of the 32 bits available for each pixel. but this would then result in the VRAM not being used to its full capacity, and pages of pixel data would need to be swapped between the VRAM and the paging memory more often than is necessary.

It may therefore be considered expedient to split the whole image memory into two for 16-bit data, or four for 8-bit data, and thus overlay whole pages of data one on top of another. This would make available the whole capacity of the VRAM, but would suffer from the disadvantage that severe complications would arise when swapping, for example, just one page of 16-bit or 8-bit data between the VRAM and page memory, because it would be necessary to select only half or a quarter of the stored data for transfer from the VRAM to the paging memory, and it would be necessary to mask off half or three-quarters of the VRAM when transferring a page of data from the paging memory to the VRAM.

There now follows a description of an arrangement which avoids these problems associated with transfer of 16-bit or 8-bit data between the VRAM and paging memories.

In essence, the data is overlaid so that at no single address for each of the 128 VRAMs 710 does there exist data for more than one page. This is achieved by overlaying the 8- or 16-bit pixel data in units of a pixel, or more preferably units of a patch, as described below.

Referring to Figure 25, an aligned set of memory cells C(0) to C(127), one from each VRAM chip, and each 4 bits wide, is shown. In the 32-bit arrangement described above, these cells form an aligned patch of 4 x 4 pixels.

In the 16-bit patch-overlay modification, these cells form two layers L(0), L(1) of a 8 x 4 patch. L(0) is provided by C(0) to (3), C(8) to (11), C(16) to (19).... C(120) to (123). L(1) is provided by the remaining cells C(4) to (7), C(12) to (15), C(20) to (27).... C(124) to (127). When the image represented by the two layers of the patch is to be displayed, layer L(0) is displayed immediately to the left of the layer L(1), as shown in Figure 25.

In the 8-bit patch-overlay modification the cells form four layers L(0) to (3) of 16 pixel x 4 pixel patch. The layers are provided by the cells as follows:

Layer L(0) : C(0), C(1), C(8), C(9) .... C(120), C(121)
Layer L(1) : C(2), C(3), C(10), C(11) .... C(122), C(123)
Layer L(2) : C(4), C(5), C(12), C(13) .... C(124), C(125)
Layer L(3) : C(6), C(7), C(14), C(15) .... C(126), C(127)

When the image represented by the four layers of a patch is to be displayed, the layers are displayed left to right in the order L(0), L(1), L(2), L(3).

A different address format needs to be employed when using 16-bit and 8-bit overlayed patches as compared with that used for the more straightforward 32-bit case, and is given in the table below:

	BITS OF VIRTUAL ADDRESS		
	32-bit mode	16-bit mode	8-bit mode
X misalignment	0,1	0,1	0,1
Level	-	2	2,3
X patch address	2-6	3-7	4-8
X page address	7,8	8	-
X portion of superpage address	9-15	9-15	9-15
Y misalignment	16,17	16,17	16,17
Y patch address	18-22	18-22	18-22
Y page address	23,24	23,24	23,24
Y portion of superpage address	25-31	25-31	25-31
Image ID portion of superpage address	32-47	32-47	32-47

It will be noted that, between the different modes, there is no change of identity of the sixteen bits representing the image ID (32-47), the sixteen bits representing the Y address (16-31), the seven bits representing the X portion of the superpage address (9-15), and the two X misalignment bits (0,1). The X patch address is, however, represented by bits 2-6 for 32-bit mode, by bits 3-7 for 16-bit mode, and by bits 4-8 for 8-bit mode. This makes available bit 2 in the 16-bit mode, and the two bits 2 and 3 in the 8-bit mode, to provide the level data, and leaves only one bit 8 in the 16-bit mode, and no bits in the 8-bit mode, for the X page address.

The patch and page arrangements and the address notations used for them are represented in Figures 26-29. Figure 26 shows the arrangement of patches in a single 16-bit page, and Figure 27 shows the arrangement of 8 pages in one complete 16-bit superpage. Figure 28 shows the arrangement of patches in a single 8-bit page, and Figure 29 shows the arrangement of 4 pages in one complete 8-bit superpage.

A number of complications arise when dealing with 16- or 8-bit

Micron Technology Inc. et al.
Ex. 1010, p. 378

data. Firstly, it is necessary to ensure that the X bits of the address are used in the proper manner. In order to do this, the supply of data from the virtual address bus 319 to the page table 756, near page edge table 758 and X patch address incrementor 760X as shown in Figure 15A is modified as shown in Figure 30A. As before, the page Y bits, 23,24 are fed directly to the page table 756 and the patch Y bits, 18-22 are fed directly to the patch Y address multiplexers 784Y, etc. However, the X bits 2-8 (which form the page and patch X addresses in the 32-bit version) are input to a funnel shifter 812. The shift provided by the funnel shifter is controlled by a mode select signal MS on line 814 which is generated by a separate circuit in response to image header information provided prior to an image or graphics processing operation and which indicates whether the pixel data is 32-, 16- or 8-bit. The funnel shifter provides a page X address of up to two bits, a 5-bit patch X address, and the level data L of up to two bits. The relationship between the inputs to and outputs from the funnel shifter 812 is shown in the table of Figure 31, and it will be noted that it corresponds to the required shifting derivable from the table set out above.

The next complication arises due to the need to present the 16- or 8-bit pixels to the grid processor during reading such that the appropriate 16 or 8 bits of each pixel will be processed and not the remaining irrelevant 16 or 24 bits. This complication is overcome by supplying, during a read operation, all 32 bits from a location in the memory to the grid processor, together with shift data ZSFT in response to which the grid processor shifts the read pixel data by an amount corresponding to the ZSFT data, and then processes predetermined bits of the shift data, e.g. bits 0-15 for 16-bit processing, or bits 0-7 for 8-bit processing.

A further complication arises due to the possibility of a read patch of data not being aligned with the patch level boundaries. This complication is overcome in a somewhat similar manner to that described above with respect to 32-bit patches not being aligned with the patch boundaries in the memory. To illustrate the above, reference is made to Figure 32, which shows a 16-bit patch p in which the base pixel is in level

Micron Technology Inc. et al.
Ex. 1010, p. 379

L = 1 of base patch a at (12,16) and is misaligned (mx,my) = (2,1). The address of the patch p in its respective page would therefore be (px,py) = (12,16); L = 1; (mx,my) = (2,1). It will be seen that, because patch p has a non-zero x misalignment, mx>0, part of the patch is at the other level L = 0, and furthermore because both mx>0 and the level of the base pixel is 1, part of the patch p is in another aligned patch b having patch address (13,16). Furthermore, because also the y misalignment my>0, the patch p also extends into aligned patches c and d at patch addresses (12,17) and (13,17) respectively and at levels 1 and 0, respectively. The determination of the further aligned patch addresses b, c, d is performed by the patch x and y address multiplexers 784 X,Y and the patch y address increment select tables 780 Y described above with reference to Figure 15 and by a modified form of the patch X address increment select table 780X which is responsive to the level data L and the mode select signal MS in addition to the X misalignment mx, as shown in Figure 30B. The modified table 780X provides a 1-bit output to the X patch address multiplexer 784X in accordance with the truth table set out in Figure 30C.

The amount of shifting ZSFT required for each pixel in the grid processor so that each pixel occupies bits 0-15 in 16-bit mode and bits 0-7 in 8-bit mode at the grid processor is determined as follows. It will be appreciated from viewing the 16-bit example of Figure 32, that the pixels labelled 6, 7, 10, 11, 14, 15, 2 and 3 will require ZSFT of 16 bits and that the pixels labelled 4, 5, 8, 9, 12, 13, 0 and 1 require zero ZSFT. This is specific to the case where the x misalignment mx = 2 and the base level is 1. It will be appreciated that for the general case of a misalignment mx, where 0< = mx < = 3 and a level L = 0 or 1, the required ZSFT for a pixel at an X location x relative to the base pixel of the patch p will be 0 bits if mx + x < 4 and L = 0, or if mx + x > 3 and L = 1, and will be 16 bits if mx + x > 3 and L = 0, or if mx + x < 4 and L = 1.

As a further illustration, reference is made to Figure 33, which shows an 8-bit patch p which has an address in its respective page of (px,py) = (12,16); L = 1; (mx,my) = (2,1). In this case, the pixels labelled 6,

Micron Technology Inc. et al.
Ex. 1010, p. 380

7, 10, 11, 14, 15, 2 and 3 require a ZSFT of 8 bits, and the pixels labelled 4, 5, 8, 9, 12, 13, 0 and 1 require a ZSFT of 16 bits. In the general case of a misalignment mx, where 0<=mx<=3, and a level L where 0<=L<=3, the required ZSFT for a pixel at an X location x relative to the base pixel of the patch will be zero bits if mx + x <4 and L = 0, or if mx + x >3 and L = 3; will be 8 bits if mx + x <4 and L = 1 or if mx + x >3 and L = 0; will be 16 bits if mx + x <4 and L = 2; or if mx + x >3 and L = 1; and will be 24 bits if x mx + x <4 and L = 3, if if mx + x >3 and L = 2.

In order to provide the required ZSFT value for each pixel, the circuit of Figure 15 includes the addition shown in Figure 34, in addition to being modified as described above with reference to Figures 30 and 31. The level value L and also the bits 0,1 of the vitual address for the misalignment mx are supplied as addresses to four ZSFT tables 818 a to d implemented using combinational logic. The ZSFT tables 818 also receive the mode select signal MS on line 814 and have three sections for 32-, 16- and 8-bit operation which are selected in dependence upon the MS signal. The ZSFT table 818a supplies the ZSFT values ZSFT(0), (4), (8), (12) corresponding to data D(0), (4), (8), (12) supplied from the read surface shifter 742 to the exchange 326; ZSFT table 818b supplies ZSFT (1), (5), (9), (13) for data D(1), (5), (9), (13); ZSFT table 818c supplies ZSFT (2), (6), (10), (14) for data D(2), (6), (10), (14); and ZSFT table 818d supplies ZSFT (3), (7), (11), (15) for data D(3), (7), (11) and (15). It will therefore be appreciated that the four ZSFT tables 818a to d correspond to pixels having X addresses of x = 0, x = 1, x = 2 and x = 3, respectively, in the patch p relative to the base pixel of the patch p.

The table set out in Figure 35A defines the values of ZSFT stored in the ZSFT tables 818a to d for different input misalignments mx, levels L and modes (8-, 16- or 32-bit) and in dependence upon the x value for the particular ZSFT table. As a further example, Figure 35B sets out the values of ZSFT for the particular ZSFT table 818b (x = 1) for all possible values of mx, L and mode. In these tables, the ZSFT values of 0, 1, 2, 3 represent a required shift of 0, 8, 16 and 24 bits respectively.

Micron Technology Inc. et al.
Ex. 1010, p. 381

A further complication which arises when dealing with 8 or 16 bit data is that the X near-page-edge signal no longer needs to be dependent solely upon whether or not 4px + mx >124, but is also dependent upon the mode selected and the level data L. The X near-page-edge signal is set only if the highest X patch address is designated (i.e. px = 31), and if the highest level data is designated (i.e. L = 1 in 16-bit mode, or L = 3 in 8-bit mode), and if the misalignment mx is non-zero. Accordingly, the near-page-edge table 758 shown in Figure 15A is modified as shown in Figure 36A so as to receive the mode select signal MS on line 814 and the level signal L, in addition to the patch address (px,py) and the misalignment (mx,my). The modified table 758 of Figure 36A produces X and Y values NPEx and NPEy of the 2-bit NPE signal as shown by the table set out in Figure 36B.

As described above, during reading, ZSFT data ZSFT(0) to (15) is supplied to the crossbar 327 with the pixel data D(0) to (15). Also, as described earlier with respect to Figure 22, each processor PROC(0) to (15) is capable of reading any of the data D(0) to (15). It is therefore necessary to ensure that the ZSFT data appropriate to the selected pixel data is supplied each processor. Figure 37 shows a modification to the crossbar 377 and part of the grid processor arrangement of Figure 22 for a generalised processor PROC(i) where 0<=i<=15. The modified arrangement is similar to the arrangement of Figure 15 except in the following respects. Firstly, a 16 x 2-bit ZSFT FIFO 678 is provided to receive ZSFT(0) to (15). The output of the ZSFT FIFO 678 is supplied to each of sixteen 16:1 2-bit multiplexers 680(0) to 680(15). The 2-bit outputs of the ZSFT multiplexers 680(0) to 680(15) are supplied to the respective processors PROC(0) to PROC(15) as signals ZSFT'(0) to (15). The ZSFT multiplexers are controlled by the same logical control signals CO(0) to CO(15) as the associated data and write enable multiplexers. It will therefore be appreciated that each processor receives the appropriate ZSFT data for the pixel data which is selects and can then shift the received pixel data by 0, 8, 16 or 24 bits in dependence upon the value 0, 1, 2 or 3 of the received ZSFT data so that the received pixel data then always occupies the first 8 bits of the processor's input register in 8-bit

Micron Technology Inc. et al.
Ex. 1010, p. 382

mode, or the first 16 bits of the input register in 16-bit mode.

It will be appreciated that the arrangement of the multiplexers and FIFOs shown in Figure 37 may be modified in a similar manner to the modification of Figure 22 which is described above with reference to Figures 46 and 47.

A further complication which arises when dealing with 16-bit or 8-bit pixel data is that, during writing to the memory 700, only the appropriate 16 or 8 bits should be written, and the remaining 16 or 24 should not be overwritten. For example, referring to Figure 32, during writing of the patch p as shown, the memory cells which are to store the 16-bit pixels labelled 6, 7, 10, 11, 14, 15, 2 and 3 need to have bits 16 to 31 written, with writing of bits 0 to 15 disabled, and the memory cells which are to store the pixels labelled 4, 5, 8, 9, 12, 13, 0 and 1 need to have bits 0 to 15 written, with bits 16 to 31 being disabled. As a further example, referring to Figure 33, the memory cells which are to store the 8-bit pixels labelled 6, 7, 10, 11, 14, 15, 2 and 3 need to have bits 8 to 15 written, with bits 0 to 7 and 16 to 31 being disabled, and the memory cells which are to store the pixels labelled 4, 5, 8, 9, 12, 13, 0 and 1 need to have bits 16 to 23 written with bits 0 to 15 and 24 to 31 disabled.

In order to deal with this complication, the circuit of Figure 40 is employed, which provides partial write enable signals PWEa to PWEd for the memory banks having x addresses of bx = 0, 1, 2 and 3, respectively. The circuit of Figure 40 comprises four PWE tables 822a to d for the values bx = 0 to 3, respectively. Each PWE table 822 is provided with the bits 0,1 of the virtual address on bus 319 indicating the X misalignment mx, the value L from the circuit 820 of Figure 30, and the mode select MS signal on line 814. The PWE tables contain the data as set out in Figure 38 and therefore a table having a particular value of bx can provide the 4-bit value PWE in dependence upon the input values of mx, L and MS.

In addition to adding the circuit of Figure 40, the connections to the X latch groups 707(0) to (15) (see Figures 8 and 14) are modified as

shown in Figure 41. The column address strobe CAS signal is still ANDed with the write enable signals WE"(0) to (15) to produce the signals CAS(0) to (15) and the addresses A(0) to (15) are also applied to respective groups 707(0) to 707(15) of X latches. The various bits of the partial write enable signals PWE(a) to (d) are connected to write enable inputs of the X latches for the cells 0 to 127 as follows:

PWE bits	Cells
PWE(a) bit 0	0, 1, 32, 33, 64, 65, 96, 97
bit 1	2, 3, 34, 35, 66, 67, 98, 99
bit 2	4, 5, 36, 37, 68, 69, 100, 101
bit 3	6, 7, 38, 39, 70, 71, 102, 103
PWE(b) bit 0	8, 9, 40, 41, 72, 73, 104, 105
bit 1	10, 11, 42, 43, 74, 75, 106, 107
bit 2	12, 13, 44, 45, 76, 77, 108, 109
bit 3	14, 15, 46, 47, 78, 79, 110, 111
PWE(c) bit 0	16, 17, 48, 49, 80, 81, 112, 113
bit 1	18, 19, 50, 51, 82, 83, 114, 115
bit 2	20, 21, 52, 53, 84, 85, 116, 117
bit 3	22, 23, 54, 55, 86, 87, 118, 119
PWE(d) bit 0	24, 25, 56, 57, 88, 89, 120, 121
bit 1	26, 27, 58, 59, 90, 91, 122, 123
bit 2	28, 29, 60, 61, 92, 93, 124, 125
bit 3	30, 31, 62, 63, 94, 95, 126, 127

It will therefore be appreciated that, during writing in the 8-bit or 16-bit mode, only the relevant memory cells are write enabled, and the remaining cells are disabled.

It will be recalled that, in 8-bit mode, the data is processed as the first 8 bits of their 32-bit capacity by the processors, and in 16-bit mode as the first 16 bits. Therefore, in order to ensure that, upon writing, the processors can write to bits 8 to 31 of the memory in 8-bit mode, or bits 16 to 31 of the memory in 16-bit mode, prior to writing, each processor

which is to write duplicates, in 8-bit mode, the pixel data of locations 0 to 7 in its output register at bit locations 8 to 15, 16 to 23 and 24 to 31 of the output register, and duplicates, in 16-bit mode, the pixel data of bit locations 16 to 31. Accordingly, when the enabled bits of the pixel data are written to the memory, the complete data for the pixel is written.

FLAGGING OF MODIFIED PAGES

Referring to Figure 42, it is convenient that a predetermined section 830 of the VRAM 700 is always mapped to the monitor 40, and for simplicity the section will be considered between page addresses (0,0) and (7,7) giving a total mapped area of 8 x 8 x 31 x 31 x 4 x 4 = 1 Mpixel. It is also convenient that images are rendered in another section of the VRAM 700, and for simplicity the section 832 between pages addresses (8,8) and (15,15) will be considered. Then, periodically, the data of the rendering section 832 is copied to the monitoring section, for display on the monitor. It will be appreciated that the data for some pixels may not change between one copying operation and the next, and indeed it can arise that no pixel data changes between two successive copying operations. If these unchanged pixels are unnecessarily copied from the rendering section to the monitoring section, then the performance of system is impaired.

In order to overcome this problem, it may be considered expedient to flag each pixel which is modified during a rendering operation and to copy only those pixels which have been flagged. However, this would require an inordinate amount of memory to store the flags and would require an excessive amount of flag setting, testing and resetting, which would degrade the system performance. In the arrangement described below, therefore, pages which have been changed, or dirtied, in a rendering operation are flagged, and only the flagged dirty pages are copied to the monitoring section of the memory.

It will furthermore be appreciated that, if a page of pixel data is copied from the paging memory to the VRAM, and that if that page is not

modified, or dirtied, in the VRAM, then there is no need to copy that page of data back to the paging memory when the time comes to replace that page in the VRAM with a different page from the paging memory. Accordingly, in the arrangement described in detail below, a flag is set when a page is dirtied in any rendering operation while it is in the VRAM, and when the page is to be replaced, it is copied back to the paging memory only if the flag is set.

It should be noted that pixel data in the VRAM is processed in patches and that a non-aligned patch may extend across a page boundary. Therefore the arrangement described below also includes for each page, dirty flags for the pages B, C and D, as shown in Figure 11, to the right, above and the right and above, of the page A in question. It should be noted that if page A has a virtual page address (PX,PY) then pages B, C and D have virtual page addresses (PX + 1, PY), (PX, PY + 1) and (PX + 1, PY + 1), respectively.

Referring to Figure 43, a dirty-page table 834 is provided by a 2K SRAM which is addressed by the 7-bit superpage identification on line 767 from the CAM 754, and the 2-bit page X address and 2-bit page Y address from the virtual address bus 314 on lines 768X,Y. The eight data bits at each location in the table 834 are assigned as follows:

bit 7	Page A dirty swap	dsA
bit 6	Page B dirty swap	dsB
bit 5	Page C dirty swap	dsC
bit 4	Page D dirty swap	dsD
bit 3	Page A dirty render	drA
bit 2	Page B dirty render	drB
bit 1	Page C dirty render	drC
bit 0	Page D dirty render	drD

Bits 0 to 2 and 4 to 6 of the dirty page data are supplied to respective OR gates 836 (0) to (2) and 836 (4) to (6). At gates 836 (6) and (2), the signals dsB and drB are ORed with the near-page-edge X signal NPEX. At gates 836 (5) and (1), dsC and drC are ORed with the near-page-edge Y signal NPEY, and at gates 836 (4) and (0), the signals dsD and drD are ORed with an ANDed form of the near page edge X and Y signals on line 774. The six bits output from the OR gates, together with a pair of high bits, representing the new signals dsA and drA, are then passed via a register 838 for writing back into the dirty page table 834 under control of a dirty pages write-enable signal DWE on line 840. The 8-bit data line of the dirty page table 834 is also multiplexed onto the 48-bit virtual address bus 319, and the address processor is operable (a) to reset the appropriate dirty swap bits and set the appropriate dirty render bits when a new page is swapped from the paging memory to the VRAM, (b) to set the appropriate dirty swap bits and dirty render bits for a page when rendering operation is carried out on that page, (c) to test the appropriate dirty swap bits for a page when that page is to be replaced by a different page in the VRAM, and (d) to test the appropriate dirty render bits for a page when that page is to be copied from the rendering section to the monitoring section of the VRAM and to reset the dirty render bits.

An example of the operation of the dirty page arrangement will now be described with reference to Figures 42 and 43, the table of Figure 44 and the flow diagrams of Figure 45. Suppose that 4 pages P, Q, R, S of pixel data at (X,Y) page addresses (0,0), (1,0), (2,0), (3,0) in the same superpage are copied into the VRAM at contiguous page addresses (8,8),

(8,9), (8,10), (8,11), and that the superpage has an identification code of 25 in the CAM 754. Suppose also that the rendering section 832 between page addresses (8,8) and (15,15) in the VRAM is copied over to the monitoring section 830 between pages addresses (0,0) and (7,7) in the memory. Suppose also that three rendering operations are carried out in the rendering section, the first rendering operation affecting page Q, the second operation affecting page P and including a misaligned patch which extends into page Q, and the third operation affecting pages Q and S; the pages P to S then being replaced by four other pages.

The dirty page data for pages P to S will be located at addresses 400 (= 25 x 16 + 0 + 0), 401, 402 and 403 in the dirty page table 840. Referring to Figures 44 and 45A, when page P is copied from the paging memory into the VRAM, it is treated as page A for the purposes of Figure 45A. In step 842 bit 7 (dsA) of the dirty flag for page A is reset and bit 3 (drA) of the dirty flag for page A is set. In step 844, the address processor 310 determines whether there is a page B' stored in the physical memory, that is the page to the left of page A. If so, in step 846, bit 6 (dsB) and bit 2 (drB) of the dirty flag for page B' are reset and set respectively. Similar steps 848, 850 and 852, 854 are carried out for pages C' and D', that is pages below and to the left and below of page A in the paging memory. Then, in step 856, page A is copied from the paging memory of the VRAM. The process of Figure 45A is then repeated for pages Q,R & S. It will therefore be appreciated that the dirty flags for pages P to S attain the state as shown in column 902 of Figure 44.

The monitoring section 830 of the VRAM is then to be updated, the address processor 310 carries out the process shown in Figure 45B. In the loop of steps 858 and 860, all of the pages of the rendering section which may possibly need to be copied are selected one-by-one. In step 862, bit 3 (drA) of the selected page (A) is tested, and if set page A is copied to the monitoring section in step 864, and in step 866 bit 3 (drA) for page A, bit 2 (drB) for page B' to the left of page A, bit 1 (drC) for page C' below page A and bit 0 (drD) for page D' to the left and below page A are reset. In step 868, bit 2 (drB) of page A is tested, and if set page B relative to page A is copied to the monitoring section in step 870, and in step 872 bit

Micron Technology Inc. et al.
Ex. 1010, p. 388

2 (drB) of page A and bit 3 (drA) for page B to the right of page A, bit 0 (drD) for page C' below page A, and bit 1 (drC) for page E below and to the right of page A are reset. Somewhat similar steps 874 to 884 are performed for bits 1 and 0 (drC, drD), as shown in Figure 45B, and if set the respective page C or D is copied to the monitoring section and various bits are reset as shown. It will therefore be appreciated that when this process is carried out with the dirty flags in the state as shown in column 902 of Figure 44, all four pages P to S are copied to the monitoring section of the VRAM, and the dirty flags attain the states as shown in column 904.

In the first rendering operation, page Q only is modified, and it will therefore be appreciated that the cicuit of Figure 43 serves to set bit 7 (dsA) and bit 3 (drA) for page Q, as shown in column 906 of Figure 44.

The monitoring section of the VRAM is again updated in accordance with the process of Figure 45B. The only dirty render flag bit set is drA for page Q, and therefore only page Q is copied, and the bit drA for page Q is reset, as shown in column 908.

In the second rendering operation, page P is modified, and also a misaligned patch in page P modifies page Q. As a result, bits 7, 6, 3 and 2 (dsA, dsB, drA, drB) of the page P dirty flag are set, as shown in column 910. Because bits drA and drB for page P are set, pages P and Q are copied to the monitoring section by the process of Figure 45B, and bits 3 and 2 (drA, drB) for page P are then reset, as shown in column 912.

In the third rendering operation, pages Q and S are modified. As a result, bits 7 and 3 (dsA, drA) of the page S flag are set; bit 3 (drA) of the page Q flag is set, and bit 7 (dsA) of the page Q flag remains set, as shown in column 914. Because bits 3 (drA) of pages Q and S are set, pages Q and S are copied to the monitoring section of the VRAM, and these bits are then reset, as shown in column 916.

When the pages P to S are to be replaced, the address processor performs the process of Figure 45C for each selected page to be replaced.

In step 886, a copy flag is reset. Then in step 888, it is determined whether bit 7 (dsA) for page A is set, and if so in step 889 that bit is reset and the copy flag is set. Steps 888 and 889 are then repeated as steps 890 to 895 for bits 6, 5 and 4 (dsB, dsC, dsD) respectively of the dirty page flags for pages B', C' and D' relative to page A. Then in steps 896 and 897, if the copy flag has been set, page A is copied to the paging memory.

Referring back to column 916 of Figure 44, it will be appreciated that as a result of performing the process of Figure 45C for page P, this page is copied to the paging memory because dsA is set for page P (step 888). This is then reset (step 889). Page Q is copied to the paging memory because dsA is set for page Q (step 888). Even if it were not, page Q would be copied because dsB is set for page P (step 890). The flag bits dsA for page Q and dsB for page P are also reset (steps 889 and 891). Page R is not copied because none of dsA for page R (step 888), dsB for page Q (step 890), and dsC and dsD for the pages below, and below and to the left, of page R (step 892 and 894) are set. Page S is copied because dsA is set for page S (step 888). This bit is then reset (step 889). Accordingly, pages P, R and S are copied back to paging memory, and the flags attain the status shown in column 918 of Figure 44.

CONDITIONAL PROCESSING

The processors 606(0) to (15) of the grid processor 312 described above are arranged basically as a SIMD array, SIMD standing for 'Single Instruction - Multiple Data' and meaning that all of the processors receive the same instruction and apply it to their own particular data elements. This can be an efficient and simple way of obtaining good performance from a parallel-processing machine, but it does assume that all of the data elements need exactly the same instruction sequence. However, the processors are preferably arranged, as described below, to be able to deal with conditional instructions. Further detail of such an arrangement is shown in Figure 49.

Figure 49 shows three of the processors PROC 0, PROC i and PROC 15, with PROC i being shown in greater detail, their PRAMs 322(0),

(i), (15), the microcode memory 308 and the processing section broadcast bus 323. The microcode memory 308 supplies microcode instructions of about 90 bits to each respective instruction decode logic (IDL) circuit 100 in each of the processors. The same microcode instruction is supplied to each processor. The instruction decode logic is provided by a gate array which decodes the 90 bit instruction to provide about 140 control bits to various elements in the respective processor including an arithmetic logic unit ALU 102, a 32-bit pixel accumulator (pa) 104, a 1-bit condition accumulator (ca) 106 and a status select circuit 108 which is provided by a gate array. The ALU 102 connects with the data bus D' via the exchange 326 to the VRAM 700, the pa 104 and a stack of pixel registers p0 to pn in the PRAM 322. The main data paths for pixel data are from the data bus D' to the ALU 102 and the pa 104; from the pa 104 to the ALU 102, the data bus D' and selected pixel registers p0 to pn; from the ALU 102 to the data bus D' and the pa 104; and from selected pixel registers p0 to pn to the ALU 102. Various status bits are output from the ALU 102 to the status select circuit 108, such as a "negative" bit, a "zero" bit and an "overflow" bit. Some of these status bits are also fed out externally. Also, external status bits such as the EN flag (see Figures 22, 23) are fed in to the status select circuit 108. Under control of the IDL 100, the status select circuit 108 can select a respective status bit and output it to the ca 106. The ca 106 is associated with a stack of condition registers c0 to cn in the PRAM 322. The ca 106 also connects to the IDL 100 and provides the write enable output WE' of the processor. The main paths for condition and status bits are: from the ALU 102 to the status select circuit 108 and to the external outputs; from the external inputs to the status select circuit 108; from the status select circuit 108 to the ca 106; from the ca 106 to the condition stack registers c0 to cn, the write enable output WE' and the ALU 102; and from the condition stack registers c0 to cn to the ca 106.

The 1-bit input from the ca 106 to the IDL 100 is important. This input condition bit enables the IDL 100 to modify the control outputs from the IDL 100 in dependence upon the value of the condition bit, and accordingly the arrangement provides direct support for microcode instructions from the microcode memory 308 to the IDL 100 which in

high-level language would be represented by, for example, if (condition) then (operation X) else (operation Y). As an example, reference is made to Figures 50A to 50D. Suppose that the VRAM 700 contains three images: image A of Figure 50A which in this simple example is a rectangle of horizontal lines; image B of Figure 50B which is a rectangle of vertical lines; and image C of Figure 50C which is a mask in which the upper-left and lower-right corners are black (say pixel values of 0) and the remainder is white (say pixel values of $(2^{32}-1)$. In the example, it is desired to combine images A and B using image C as a mask to form an output image D such that image A appears where the mask image C is black and image B appears where the mask image C is white. The process performed by the processors under control of the microcode instructions from the microcode memory 308 to perform this operation can be considered, using high-level pseudo-language, to be as follows:

```
1.    For each patch (x,y) in the rectangle:
2.          If pixel in rectangle, ca = 1, else ca = 02.
3.          c0 = ca
4.          pa = A(x,y)
5.          p0 = pa
6.          pa = B(x,y)
7.          p1 = pa
8.          pa = C(x,y)
9.          ca = zero-status (pa)
10.         If ca = 1 then pa = p0 else pa = p1
11.         ca = c0
12.         D(x,y) = pa
13.   Next patch
```

In the above, steps 1 and 13 set up a loop for each patch (x,y) having its origin in the rectangle. For each patch, each processor PROC 0 to PROC 15 will process a different pixel in the patch. In step 2 a test is made to determine whether the particular processor's pixel in the patch is in the rectangle, and if so the ca 106 is set, otherwise it is reset. This value of ca will form the write-enable signal WE'. In step 3, this value which is stored in the ca 106 is put onto the condition stack in c0 and an

associated condition stack pointer is modified accordingly. In step 4, the value of the processor's pixel in the current selected patch in image A is loaded into the pa 104, and in step 5 is transferred to the p0 register. Similarly in step 6, the value of the processor's pixel in the current selected patch in image B is loaded into the pa 104, and in step 7 is transferred to the p1 register. In step 8, the value of the processor's pixel in the current selected patch in the mask image C is loaded into the pa 104, and then in step 9 the zero status bit of the ALU 102 is selected by the status select circuit 108 and is loaded into the ca 106. Thus, if the pixel in the mask image is black, the ca 106 value becomes 1, and if it is white, the ca 106 value becomes 0. The next step 10 is a conditional instruction "If ca = 1 then pa = p0 else pa = p1". The IDL 100 modifies this instruction in dependence upon the value in the ca 106 so that it becomes simply "pa = p0" or "pa = p1" and the modified instruction is used by the processor. In step 11, the signal which was put onto the condition stack at c0 in step 3 is pulled off the stack and placed in the ca 106 in order to constitute the write enable signal WE' and the condition stack pointer is modified accordingly. Lastly, in step 12, the pixel value in the pa 104 is transferred out to the image D at the appropriate pixel position for the processor in the current selected patch.

As a result of the above operations carried out by the processors on the pixels of all of the patches in the rectangle, an image D is formed as shown in Figure 50D.

In the above simple example, the condition stack c0 to cn was used simply to store the initially generated value which will form the write enable signal, and only one register in the stack was employed. By virtue of the provision of more than one register in the condition stack, nesting of the conditional instructions is permitted.

PAGE FILING SYSTEM

As described above, pages of data can be swapped between the VRAM 700, on the one hand, and the paging memory comprising the DRAM 304 (Figure 4), and the paging RAM 504 and fast disk 510 (Figure

Micron Technology Inc. et al.
Ex. 1010, p. 393

5), on the other hand. There now follows a description of how pages are handled as between the VRAM and the paging memory, with reference to the system diagram of Figure 51.

The total system is based on a distributing operating system denoted by the triangle 200. Part of this system constitutes a host page manager module 202 running on the processor 10 of the host computer. Another part constitutes a front-end page manager module 204 running on the i960 control processor 508 of the front-end board 22 and handling the paging RAM 504 and fast disk 510. A further part constitutes a renderer page manager module 206 running on the i960 control processor 314 of the renderer board 16 and handling the VRAM 700 and the DRAM 304. Each of these page manager modules 202, 204, 206 can make a request R to any other module for a page P of image data specified by the virtual page address (VPA) consisting of the following bits of the virtual address:

32-47	Image ID component
25-31	Y superpage component
23, 24	Y page component
9-15	X superpage component
7, 8	X page component

In response, the module to which a request R is made determines whether it is responsible for the requested page, and if so it transfers the page of data P and responsibility therefor to the requesting module, but if not it indicates to the requesting module that it is not responsible for the requested page.

To give two examples of how the filing system would be used, suppose that the page fault table 794 (Figure 21) of the renderer has generated a page fault in respect of a particular page, this page fault is handled by the renderer page manager module 206. Firstly, the module 206 checks with itself whether the required page is stored in the renderer DRAM 304, and if so swaps the page of data into the VRAM 700. If not, the module 206 checks with the front-end page manager module 204 whether it is responsible for the page, and, if so, the page of data is

Micron Technology Inc. et al.
Ex. 1010, p. 394

swapped from the RAM 506 or disk 510, as appropriate, into the VRAM 700. If the front-end module 204 is not responsible, the renderer module 206 asks the host module 202 for the page of data, which is then swapped into the VRAM 700. As another example, suppose that the system is to be closed down and a complete image is to be saved to disk 510. Such saving of an image is handled by the front-end module 204. For each page in the image the module 204 firstly checks with itself whether it is responsible for that page. If it is and the page is already stored on the disk 510, it stays there, and if the page in question is stored in the front-end RAM 506 the data of that page is copied to the disk 510. If the module 204 is not responsible, it checks with the renderer module 206 whether the renderer module has responsibility for the page, and, if so, the page of data is copied from the VRAM 700 or DRAM 304 of the renderer to the disk 510. If not, the front-end module 204 requests the page in question from the host module 202, and the page of data is transferred to the disk 510.

In order to keep track of the pages for which they are responsible, the front-end module 204 and the renderer module 206 each maintain a table 208, 210 containing a list of the virtual page addresses of the pages, and against each address an indication of the location of that page. For example, the location data in the front-end table 208 would comprise an indication of whether the page is in the RAM 506 or on the disk 510. If in the RAM 506, the physical address of that page in the RAM would be included, and if on the disk 510, an indication of the location on the disk would be included. The location data for each virtual page address in the renderer table 210 may contain an indication of whether the page is in the DRAM 304 or the VRAM 700 and the physical address of the page in the respective memory. In the case of a page in the VRAM 700, the physical address of the page need not necessarily be kept in the table 210, because this address can be determined by the module 206 from the CAM 754 and the page table 756 (Figure 15A) of the address translator 740, and indeed it i not necessary for the table 210 to include the virtual page address of the pages in the VRAM 700, because the module can check whether a page is present by referring to the CAM 754 and page table 756 and testing whether or not a page fault is generated.

Micron Technology Inc. et al.
Ex. 1010, p. 395

An important feature of the filing system, in the preferred embodiment, is that the host page manager module 202 is not responsible for the storage of whole pages of data. The host module 202 is used when an image is initially created. The image is specified by the host processor 10 as being of a particular dimension, size, bit width (see Figures 25 to 41) and background colour. In response, the system software 200 allocates to that image the next available image ID (bits 32 to 47 of the virtual address). Until any rendering operations or copying operations are carried out on the image, the colour of every pixel in the new image is the background colour, and the host module 202 therefore merely sets up a table 212 containing the virtual page address of the or each page required in the new image, and against the or each page address the table 212 contains the 32-bit background colo r of the image. There is no need for this 32-bit word of data for the page to be expanded into a full page of data, for example 16k words, until the page is transferred to the control of one of the other modules 204, 206. Accordingly, when one of the other modules requests a page from the host module 202, the host module 202 determines from its table 212 the 32-bit background colour of that page, and then repeatedly sends that 32-bit word to the requesting module, once for each pixel in the page.

In the above description, it is assumed that only one of the modules 202, 204, 206 has responsibility for any given page at any given time and that when a page of data is transferred from one module to another, the sending module cancels the entry for that page from its table 212, 208, 210 and that the receiving module makes an entry in its table for the page. It will be appreciated that the dirty page-swap scheme described with reference to Figures 42 to 45 above will not be effective if the filing system operates in this way, because when, for example, a page is swapped from the disk 510 to the VRAM 700, the entry for that pag is cancelled from the table 208 of the front-end module 204, and so even if the page is not dirtied in the VRAM 700, it would be necessary to swap all of the data-elements of the page back to the disk 510.

The filing system described above may be modified so that it works in conjunction with the dirty page-swap scheme, by including against each

virtual page address in each table 208, 210, 212 a bit indicating whether that page is current. The operation of each module 202, 204, 206 is then modified so that when a module has responsibility for a page, the current bit is set to 1 and when responsibility is transferred to a different module the current bit is reset to zero. Furthermore, when a page which has not be dirtied is to be swapped out of the VRAM 700, the renderer module 206 polls the other modules 202, 204 to check which has an entry in its table for the page with the current bit reset, and instructs that module to set the current bit, obviating the need to copy all of the data-elements for that page from the renderer module to the other module.

In the above arrangement, a single word representing the image background colour is stored for each new image. Rather than storing a single word, a few words may be stored, for example as a patch, and representing, for example, a pattern which is to be repeated in the new image.

MODIFICATIONS AND DEVELOPMENTS

Although preferred embodiments of the invention have been described above, it will be appreciated that many modifications and developments may be made within the scope of the invention. To take a few examples, the non-split-level patches, pages and superpages described above are two-dimensional and have a pixel resolution of 32-bits, a patch size of 4 pixels x 4 pixels, a page size of 32 patches x 32 patches, and a superpage size of 4 pages x 4 pages. It will be appreciated that the system may be configured so as to operate for example with one- or three-dimensional patches, and/or pages and/or superpages, with patches, pages and superpages of different sizes, and with different pixel resolutions. Furthermore, the system may be arranged to operate selectably in different configurations through appropriate use of funnel shifters, switches and the like. In the above description, examples of specific sizes of the memories have been given, but it will be appreciated that other sizes may be used. In the split level patch system, division into two and four in the X direction has been illustrated, but it will be appreciated that other divisors may alternatively or selectably be

Micron Technology Inc. et al.
Ex. 1010, p. 397

employed, that division in other directions may alternatively or selectably be employed, and that division on a pixel basis rather than a patch basis may alternatively or selectably be employed. The dirty page facility described above deals with copying between the rendering section and monitoring section of the VRAM and also with swapping between the VRAM and the paging memory, but it will be appreciated that either of these two features may be employed without the other. In the page filing system, the page manager modules are run on specific processors, but it will be appreciated that each page manager module may be run on different processors, and that the modules may be combined.

RELATED APPLICATIONS

This application is one of a series of eight applications filed the same day under filing numbers PCT/GB90/ , PCT/GB90/ , PCT/GB90/ , PCT/GB90/ , PCT/GB90 , PCT/GB90/ , PCT/GB90 and PCT/GB90/ , and with agent's references PCT 98 1, PCT 99 1, PCT 100 1, PCT 101 1, PCT 102 1, PCT 103 1, PCT 104 1 and PCT 105 1, and the disclosures of the other applications are incorporated herein by reference as if printed below in full.

CLAIMS

1. A memory system for storing an array of data elements which form in a virtual memory space aligned groups (Fig. 7B) of contiguous aligned pages (Fig. 7A) of contiguous data-elements, comprising:

a first memory (304, 506, 510) for storing the data-array;

a second memory (700) for storing part of the data-array;

means (326, 331, 302, 20) for transferring data-elements page-by-page between the first and second memories such that contiguous pages of the data-elements in the virtual memory address space are not necessarily stored contiguously in the address space of the second memory; and

means (740) for translating a group component (bits 9-15, 25-31, 32-47) and a page component (bits 7, 8, 23, 24) of the address of a page in the virtual memory space into the corresponding page address (PAGE A) of that page in the second memory address space, the translating means comprising a contents addressable memory (CAM) (754) which receives the virtual group component and outputs a group code (767) corresponding thereto, and a page table (756) which receives the group code and the virtual page component and outputs the corresponding page address in the second memory address space.

2. A system as claimed in claim 1, wherein the number of group codes which can be stored by the CAM is substantially less than the number of the groups which can be stored in the first memory.

3. A system as claimed in claim 1 or 2, wherein each group code has a number of bits (e.g. 7) which is substantially less than the number of bits (e.g. 30) of the virtual group component.

4. A system as claimed in any preceding claim, wherein the CAM is operable to compare the received virtual group component with the contents at the addresses of the CAM and to output as the group code the address of the CAM whose contents matches the virtual group component.

Micron Technology Inc. et al.
Ex. 1010, p. 399

5. A system as claimed in claim 4, wherein the CAM contents have a bit (e.g. msb) which is not used in the virtual group component, and wherein said bit is set at any location in the CAM at which a virtual group component is not stored.

6. A system as claimed in any preceding claim, wherein the CAM is operable to supply a group fault signal (748S) if a match is not made with the received virtual group component, and further comprising processor means (310) operable in response to the group fault signal to store the received virtual group component at an available location in the CAM.

7. A system as claimed in claim 6, wherein the processor means (310) is operable in response to the group fault signal to determine whether there is any available location in the CAM, and if not to select a virtual group component to delete from the CAM, and to determine whether any pages of data-elements of the selected group are stored in the second memory, and if so to cause the transfer means to transfer those pages of data-elements from the second memory to the first memory.

8. A system as claimed in claim 7, further comprising means (Fig. 20) for determining the least recently used group having a page stored in the second memory, the processor means (310) being operable to select the virtual group component of the determined least recently used group as the virtual group component for deletion from the CAM.

9. A memory system for storing an array of data-elements which form in a virtual memory space contiguous aligned pages (Fig. 7A) of contiguous data-elements, comprising:

a first memory (304, 506, 510) for storing the data-array;

a second memory (700) of smaller capacity than the first memory for storing some of the pages of the data-elements; and

means (310, 304, 508) operable to address locations in the first and second memories and to transfer data-elements between the addressed locations page-by page;

wherein the data-elements provide a plural-dimensional

representation, and wherein the transferred pages include data-elements which are contiguous in each of the plural dimensions.

10. A system as claimed in claim 9, wherein the addressing means is operable to address the second memory with addresses which include components in each of said plural dimensions.

11. A memory system for storing an array of data-elements which form in a virtual memory space contiguous aligned pages (A, B, C, D) of contiguous aligned patches (a, b, c, d) of contiguous data-elements, comprising:

a first memory (304, 506, 510) for storing the data array;

a second memory (700) for storing part of the data-array;

means (326, 331, 302, 20) for transferring data-elements page-by-page between the first and second memories such that contiguous pages of the data-elements in the virtual memory address space (Fig. 11A) are not necessarily stored contiguously in the address space of the second memory (Fig. 11B); and

means (310, 740) for addressing the second memory to access in parallel a patch (p) of data-elements which is not necessarily such an aligned patch, the addressing means including page edge means (758) for determining whether the patch to be accessed includes data-elements in different pages in the second memory, and means (756, 778, 782X, Y) to modify the addresses to the second memory in response to such a determination.

12. A system as claimed in claim 11, wherein the address of a patch (p) to be accessed in the second memory has an aligned page address component (PAGE A), an aligned patch address component (px,py) and a misalignment component indicating a misalignment (mx,my) of the patch (p) to be accessed relative to an aligned patch (a) having an address defined by the aligned page and aligned patch components.

13. A system as claimed in claim 12, wherein the addressing means includes page determining means (756) operable to determine, upon addressing, an aligned page address component for a basic aligned page

(PAGE A) for the patch (p) to be accessed and an aligned page address component of at least one other page (PAGE B, C, D) which is contiguous with the basic aligned page (PAGE A) in the virtual address space.

14. A system as claimed in claim 13, wherein the data-array is a plural-dimensional representation having D-dimensions, and wherein the page determining means (756) is operable to determine the aligned page addresses in the second memory for all aligned pages which differ by one from the basic aligned page address in one direction in each of the D-dimensions in the virtual address space and in any combinations of those one directions.

15. A system as claimed in claim 13 or 14, further comprising page fault determining means (794) for determining, upon addressing, whether the or each aligned page which is to be accessed is not stored in the second memory, the transferring means being operable to transfer the missing page from the first memory to the second memory in response to such a determination.

16. A system as claimed in claim 15, wherein the page fault determining means is operable to indicate a page fault (748P) if the basic page is not stored in the second memory (794 output A), or if any said other aligned page is not stored in the second memory (794 output B, C, D) and (796 B, C, D) that page is required to be accessed.

17. A system as claimed in any of claims 12 to 16, wherein said page edge determining means (758) is operable to determine when the aligned patch address component (px,py) has a maximum value (e.g. 31) for a patch in a page and the misalignment component (mx,my) together with said maximum value indicate that the patch (p) to be accessed extends outside the page (PAGE A) defined by the aligned page address component.

18. A system as claimed in any preceding claim, wherein the data-elements are image pixel data.

19. A system as claimed in any preceding claim, wherein the second memory is a video RAM (700).

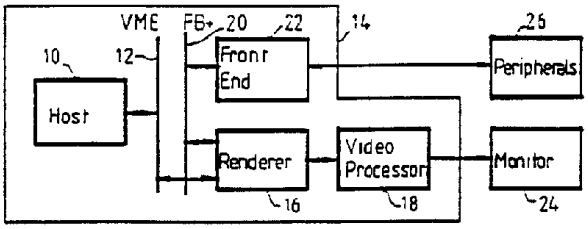


FIG. 1

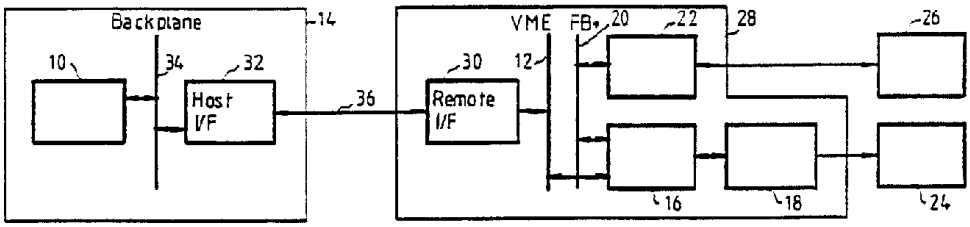


FIG. 2

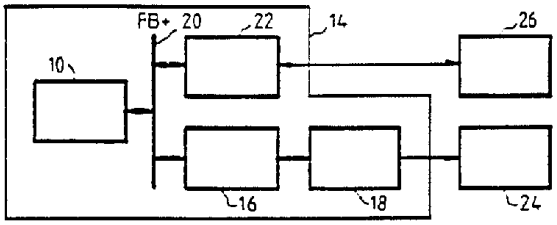


FIG. 3

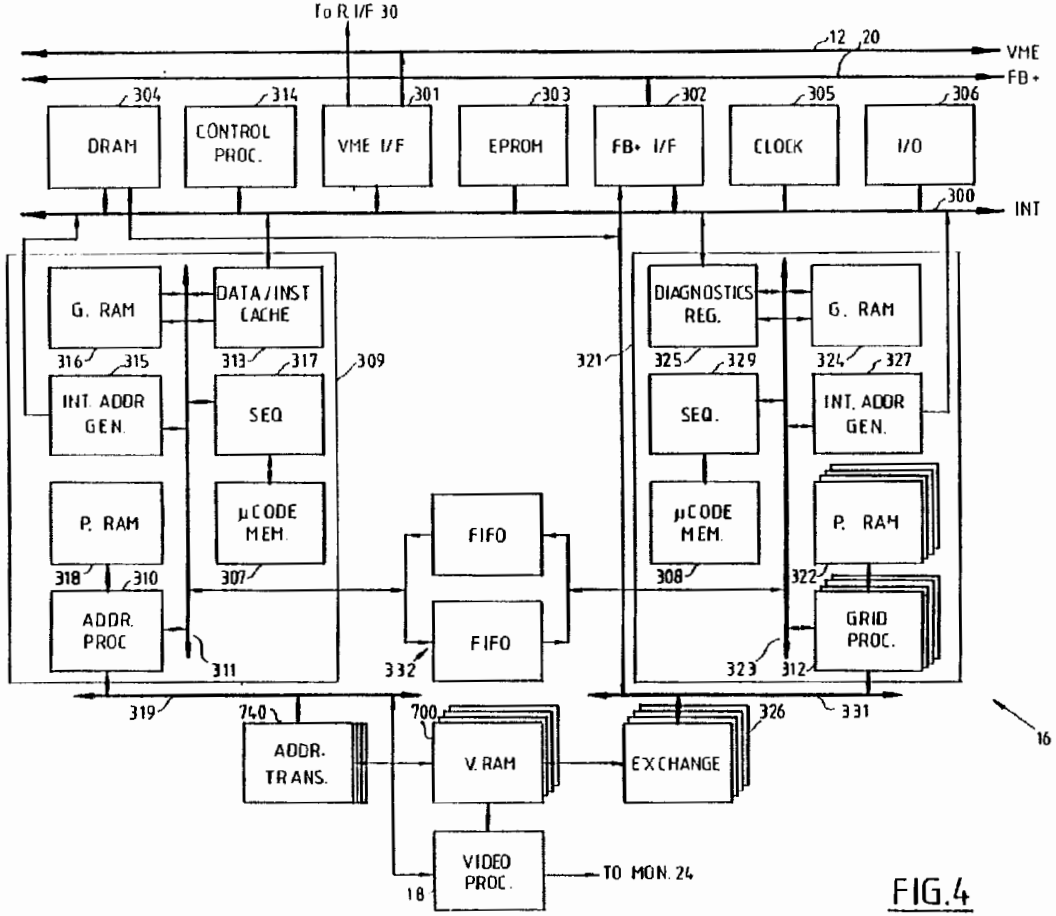


FIG.4

Micron Technology Inc. et al.
Ex. 1010, p. 405

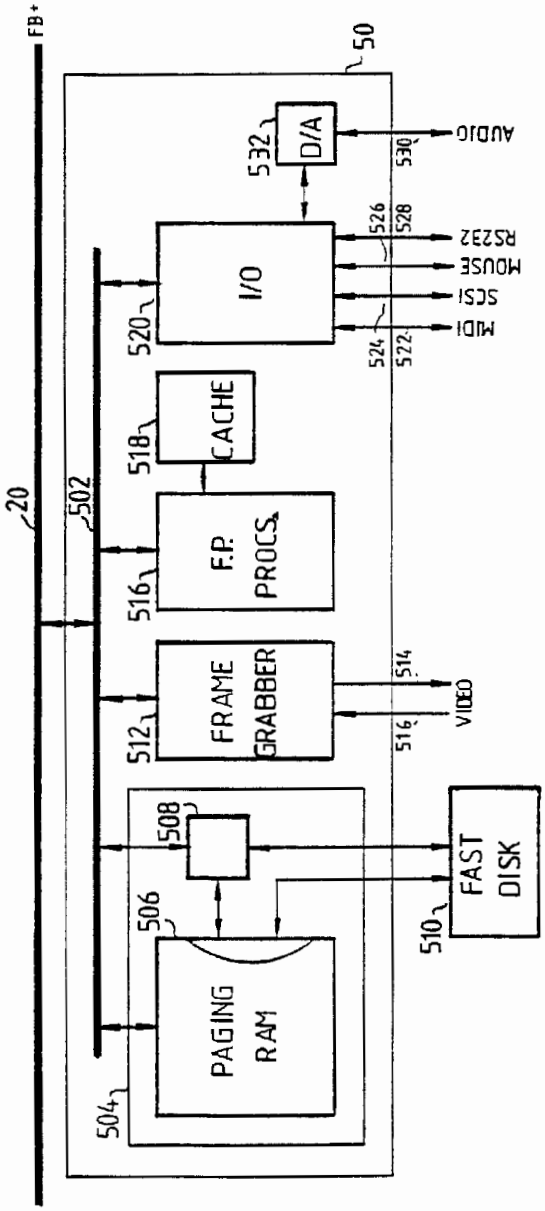


FIG. 5

Micron Technology Inc. et al.
Ex. 1010, p. 406

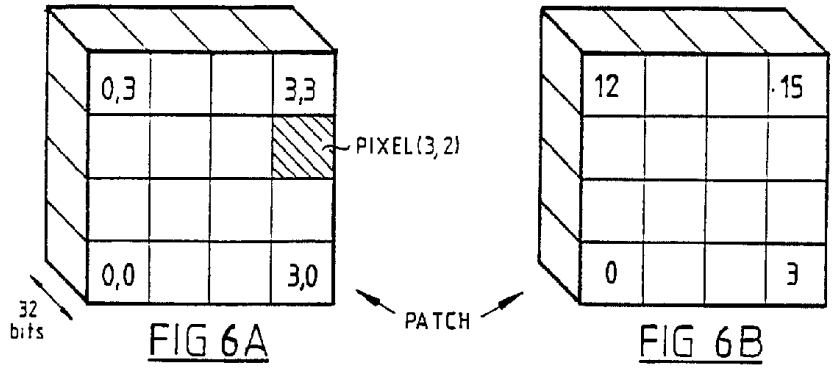


FIG 6A FIG 6B

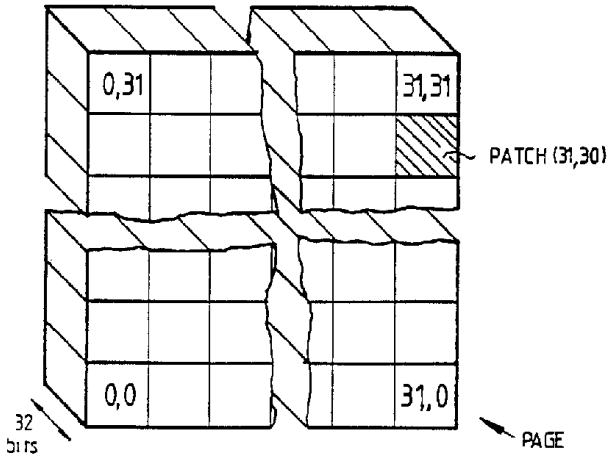


FIG 7A

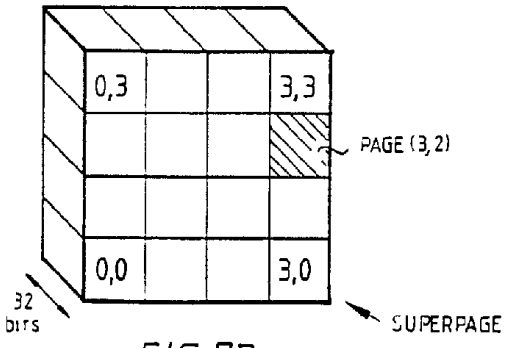


FIG 7B

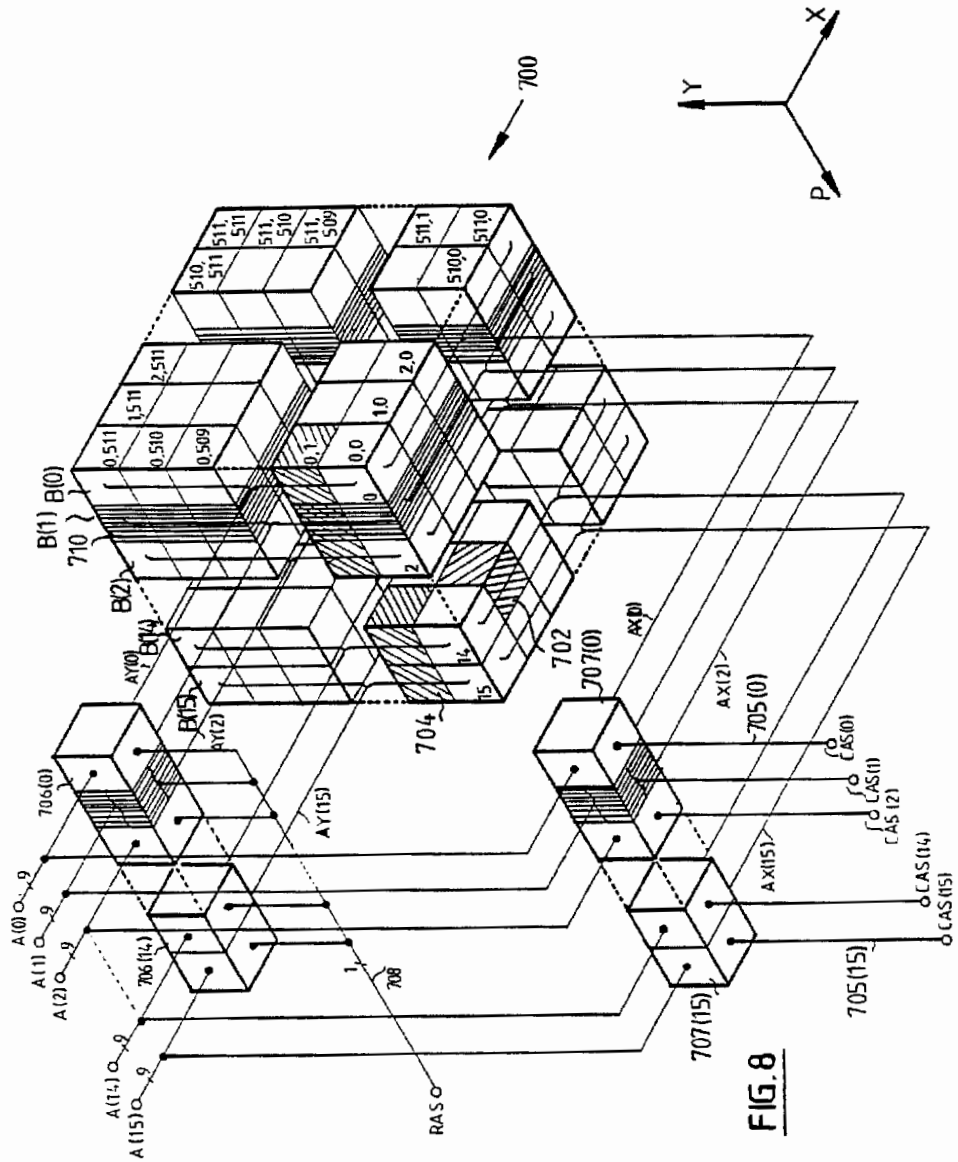


FIG. 8

Micron Technology Inc. et al.
Ex. 1010, p. 408

FIG.9A

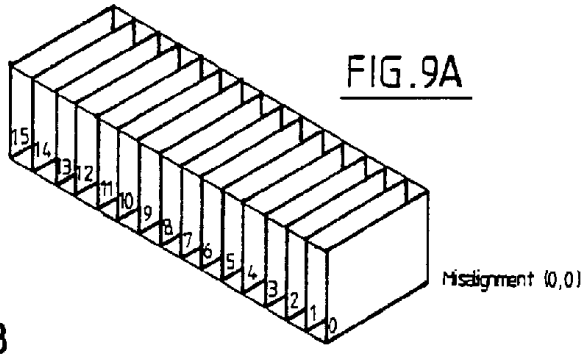
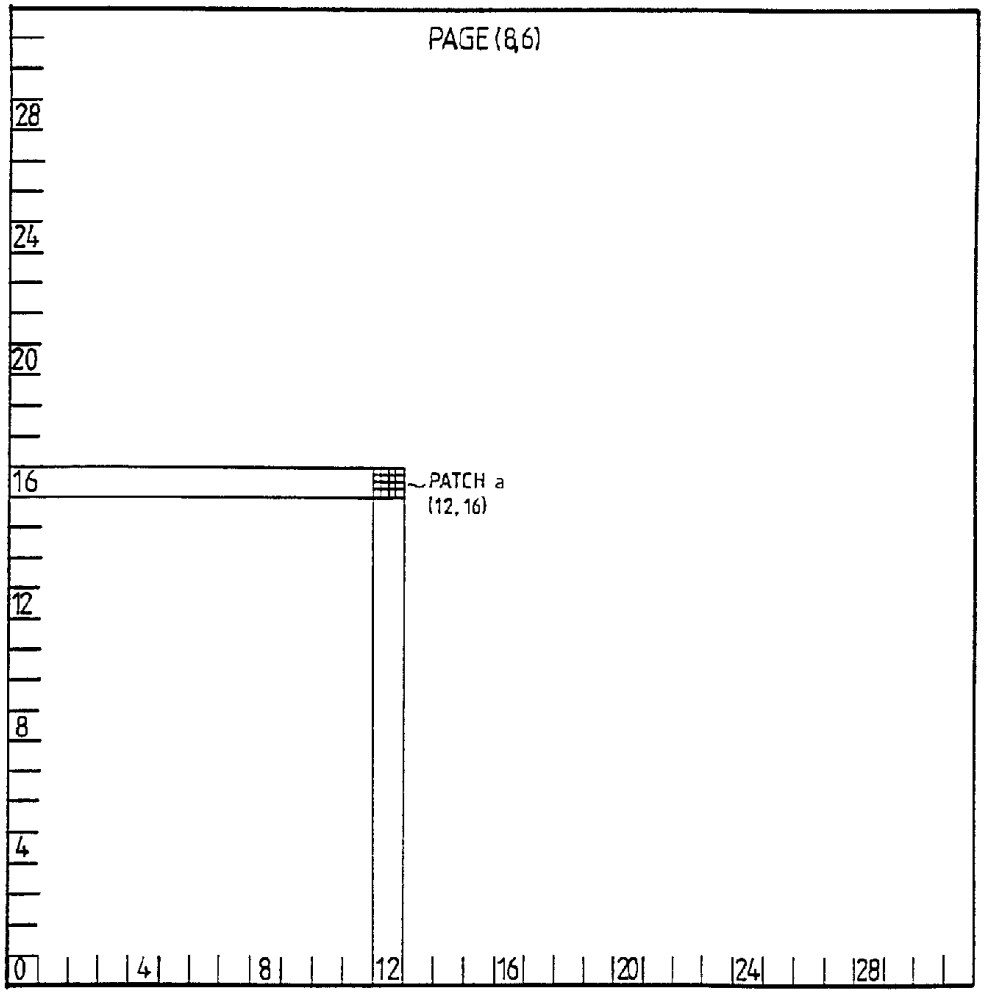

15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
Misalignment (0,0)

FIG.9B


PAGE (8,6)
28
24
20
16
12
8
4
0
4
8
12
16
20
24
28
PATCH a
(12, 16)

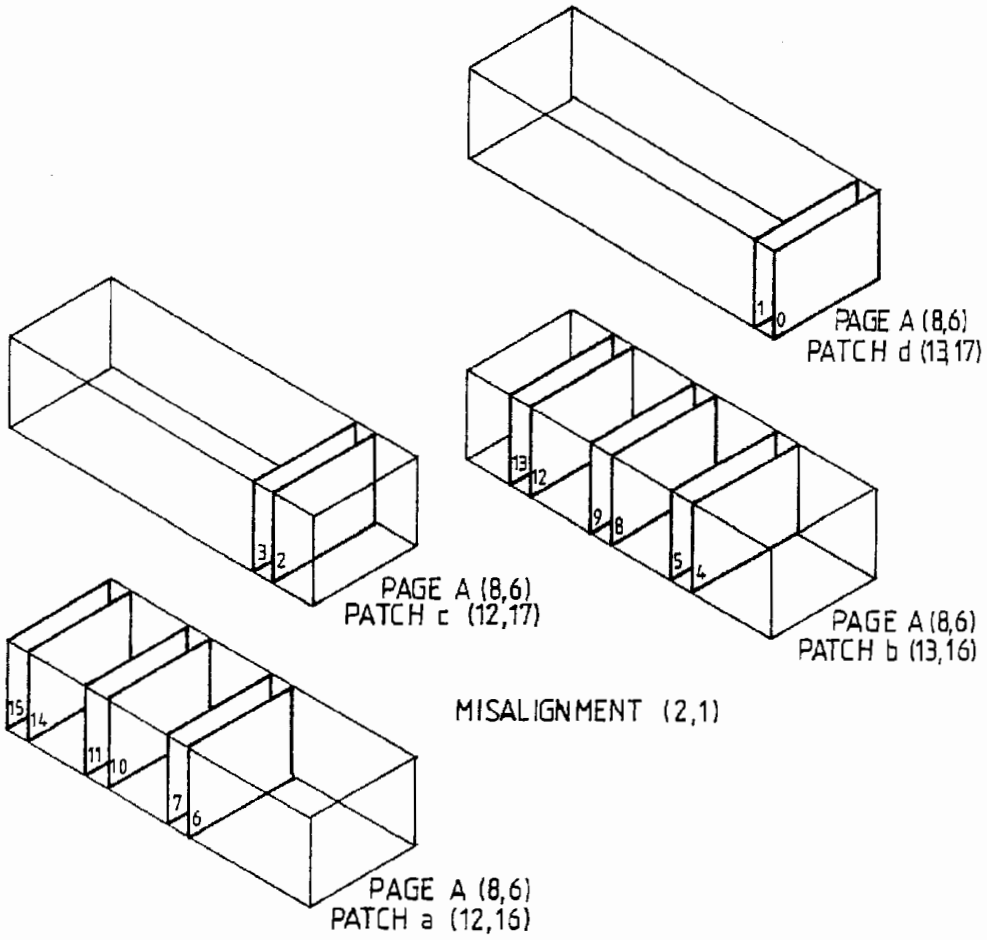

PAGE A (8,6)
PATCH d (13,17)
PAGE A (8,6)
PATCH c (12,17)
PAGE A (8,6)
PATCH b (13,16)
MISALIGNMENT (2,1)
PAGE A (8,6)
PATCH a (12,16)

FIG. 10A

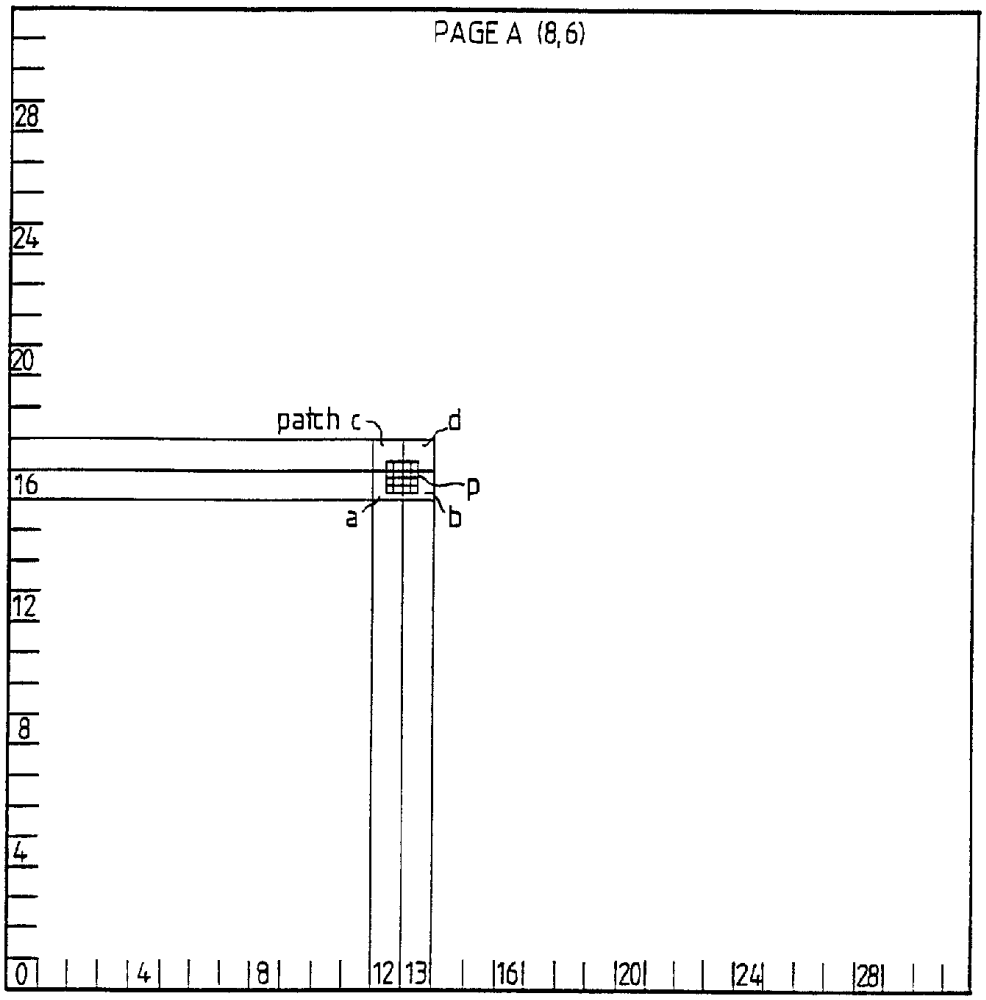


FIG. 10 B

FIG.11A

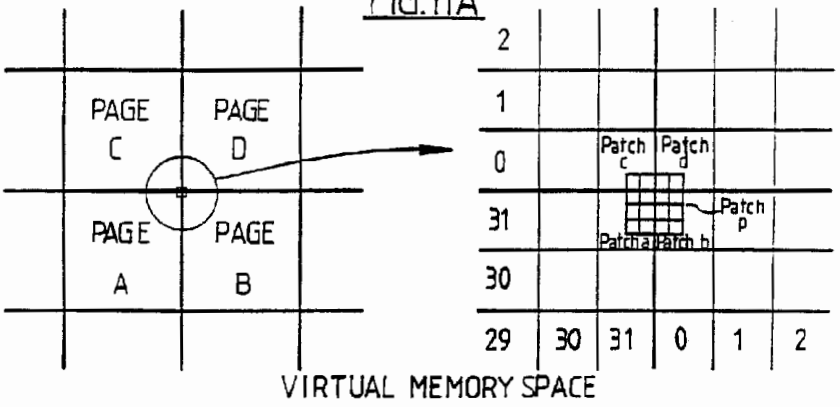
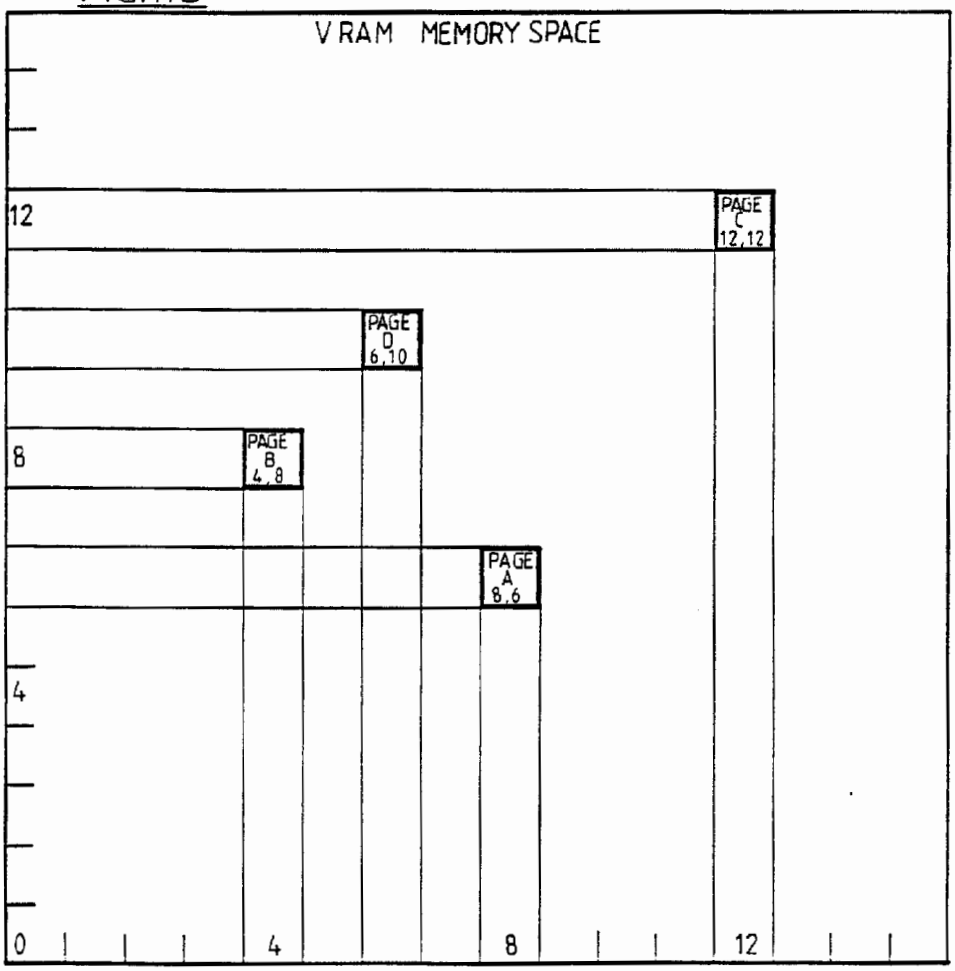


FIG.11B



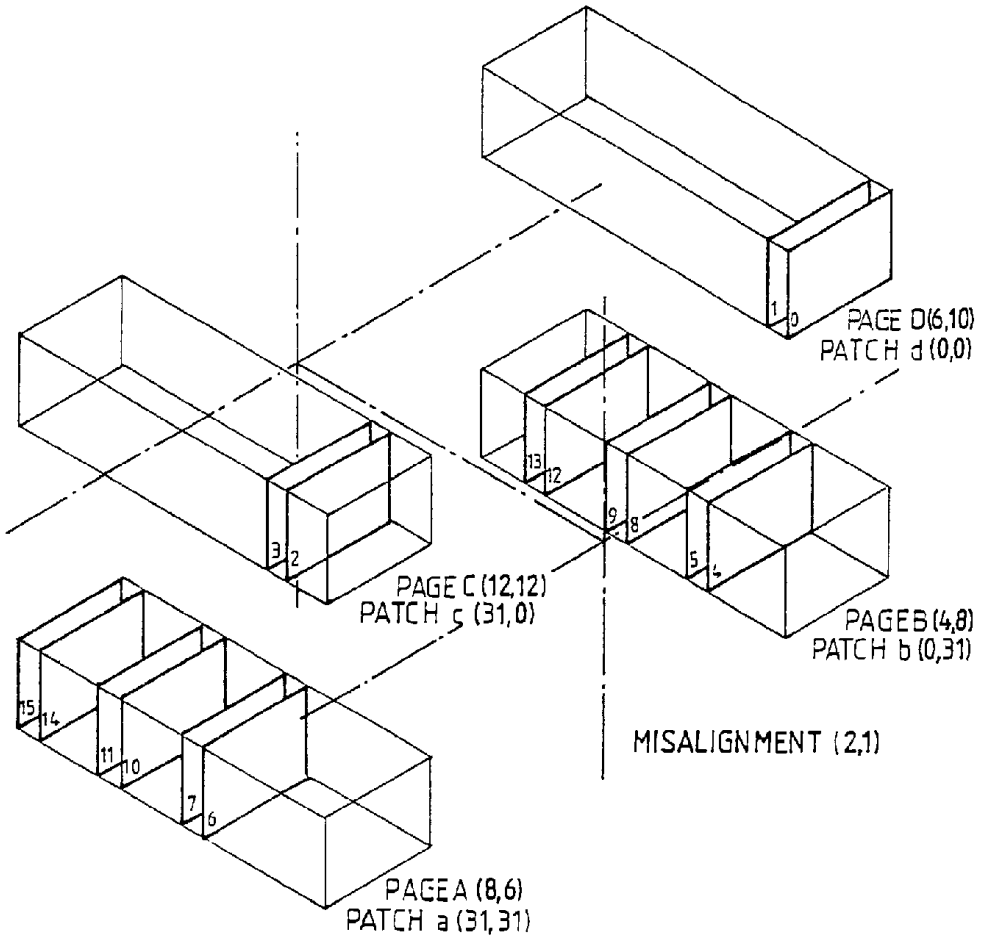

PAGE D(6,10)
PATCH d (0,0)
PAGE C (12,12)
PATCH c (31,0)
PAGE B (4,8)
PATCH b (0,31)
MISALIGNMENT (2,1)
PAGE A (8,6)
PATCH a (31,31)

FIG.11C

PAGE SELECTION

	IF 4px+mx<125 & 4py+my<125	IF 4px+mx>124 & 4py+my<125	IF 4px+mx<125 & 4py+my>124	IF 4px+mx>124 & 4py+my>124
IF bx≥mx & by≥my	A	A	A	A
IF bx<mx & by≥my	A	B	A	B
IF bx≥mx & by<my	A	A	C	C
IF bx<mx & by<my	A	B	C	D

FIG. 12

PATCH ADDRESS INCREMENT SELECTION

X ADDRESS

CONDITION	PATCH X ADDRESS
bx≥mx	px
bx<mx	px+1

Y ADDRESS

CONDITION	PATCH Y ADDRESS
by≥my	py
by<my	py+1

FIG. 13

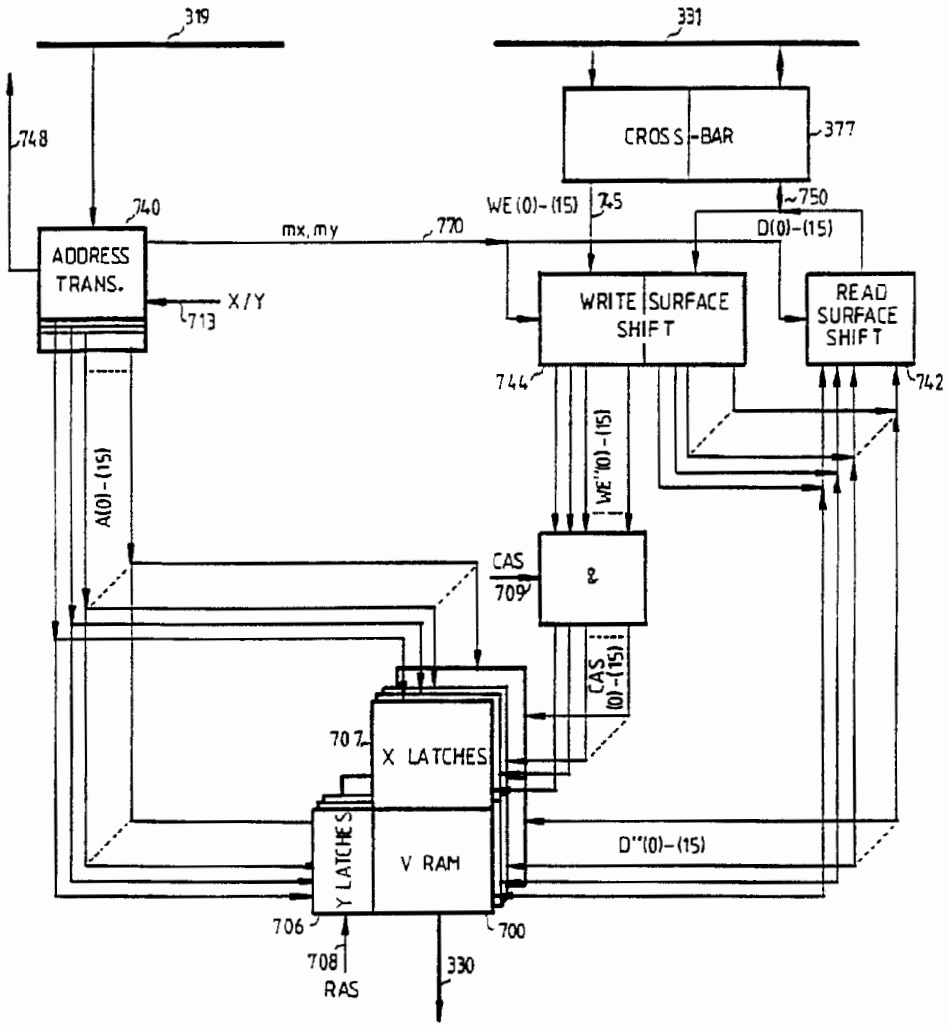

319
331
748
CROSS-BAR
377
740
WE(0)-(15)
745
750
D(0)-(15)
ADDRESS TRANS.
mx, my
770
X/Y
713
WRITE SURFACE SHIFT
READ SURFACE SHIFT
744
742
A(0)-(15)
WE''(0)-(15)
CAS
709
&
CAS (0)-(15)
707
X LATCHES
Y LATCHES
V RAM
D''(0)-(15)
706
700
708
RAS
330

FIG. 14

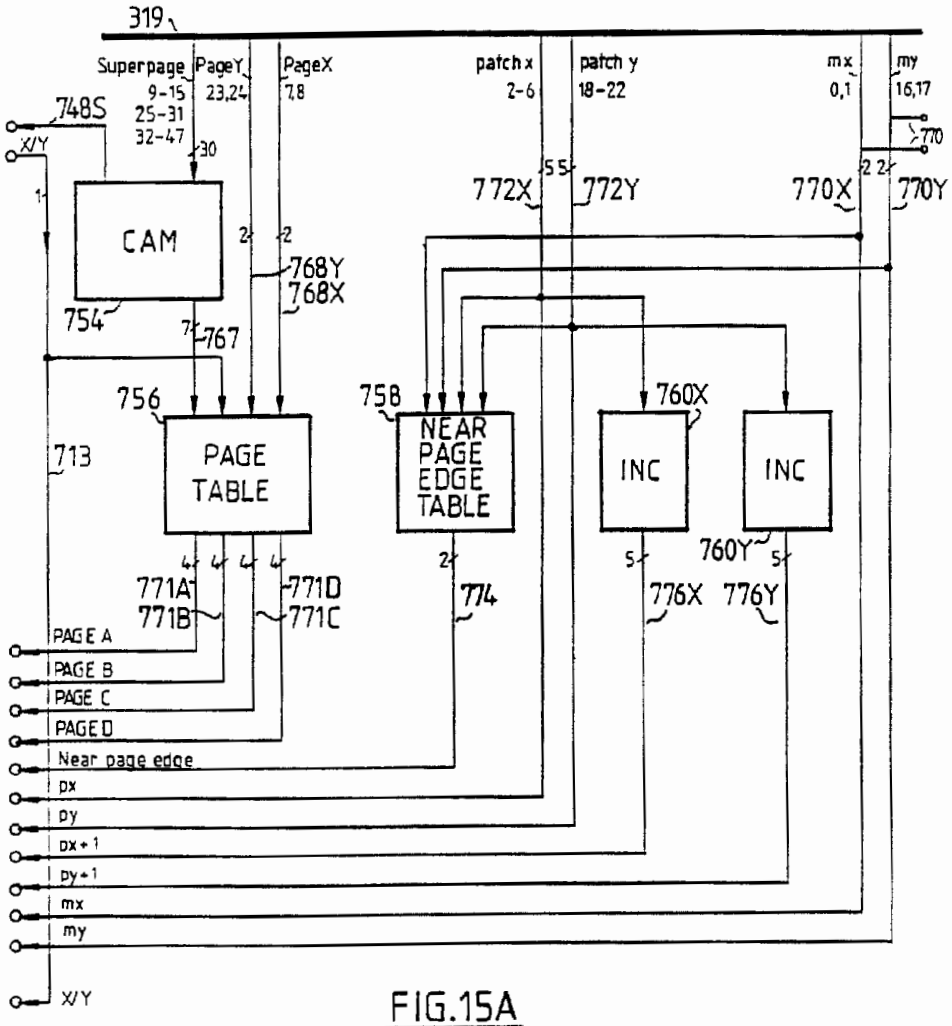


FIG.15A

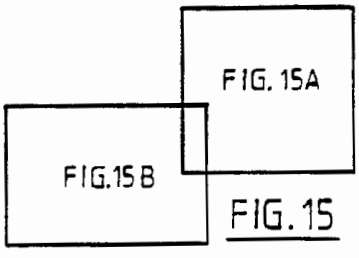
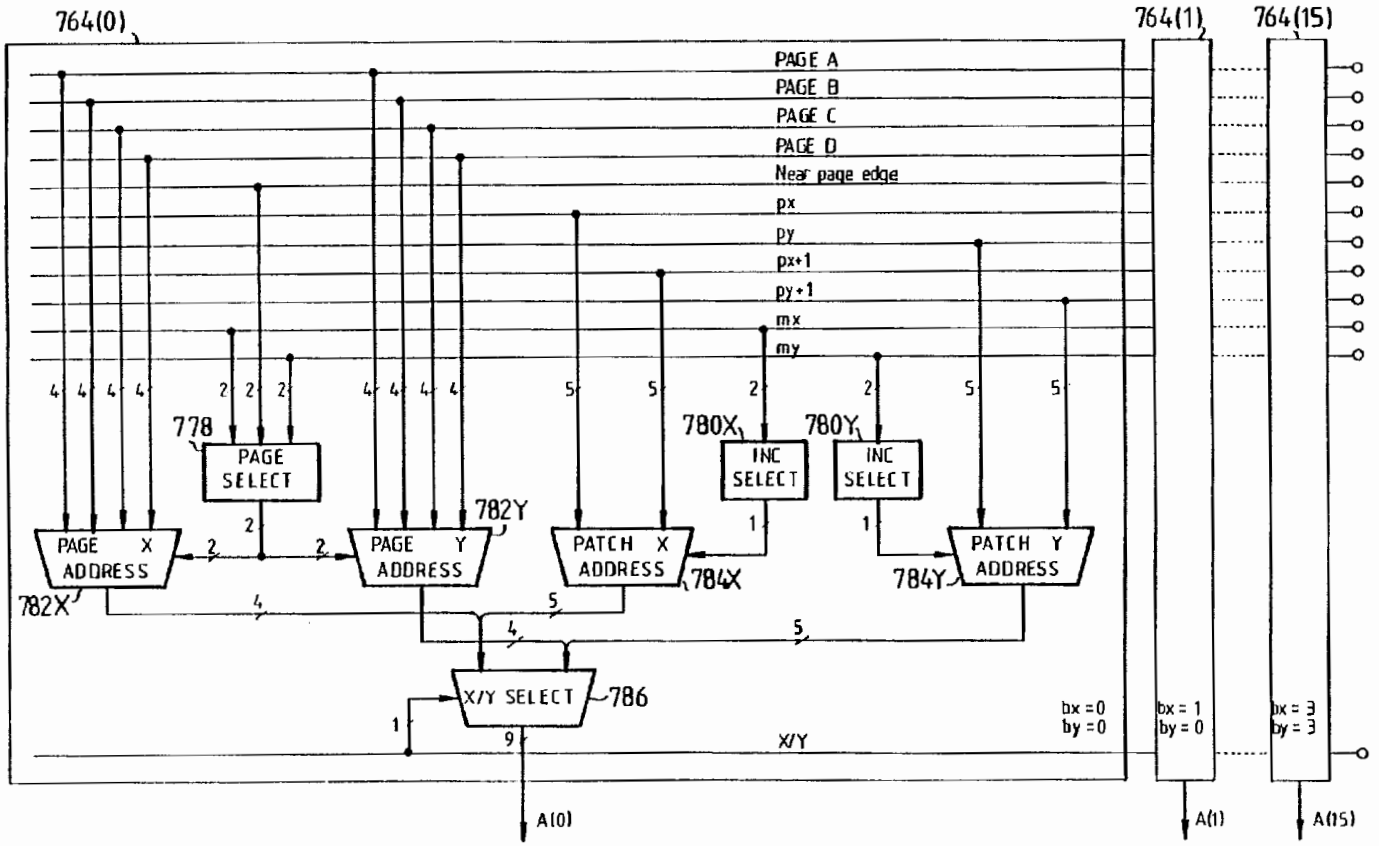


FIG. 15

FIG.15B



Micron Technology Inc. et al.
Ex. 1010, p. 417

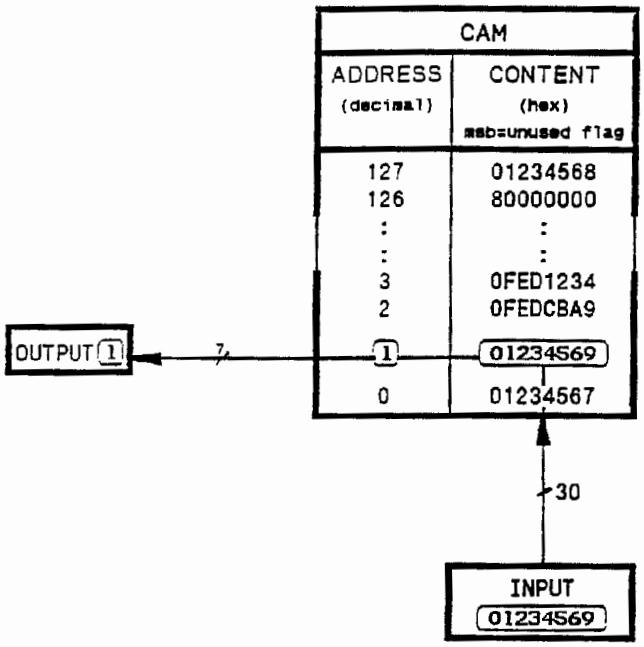

CAM
ADDRESS
(decimal)
CONTENT
(hex)
msb=unused flag
127
126
3
2
1
0
01234568
80000000
0FED1234
0FEDCBA9
01234569
01234567
OUTPUT 1
7
30
INPUT
01234569

FIG. 16

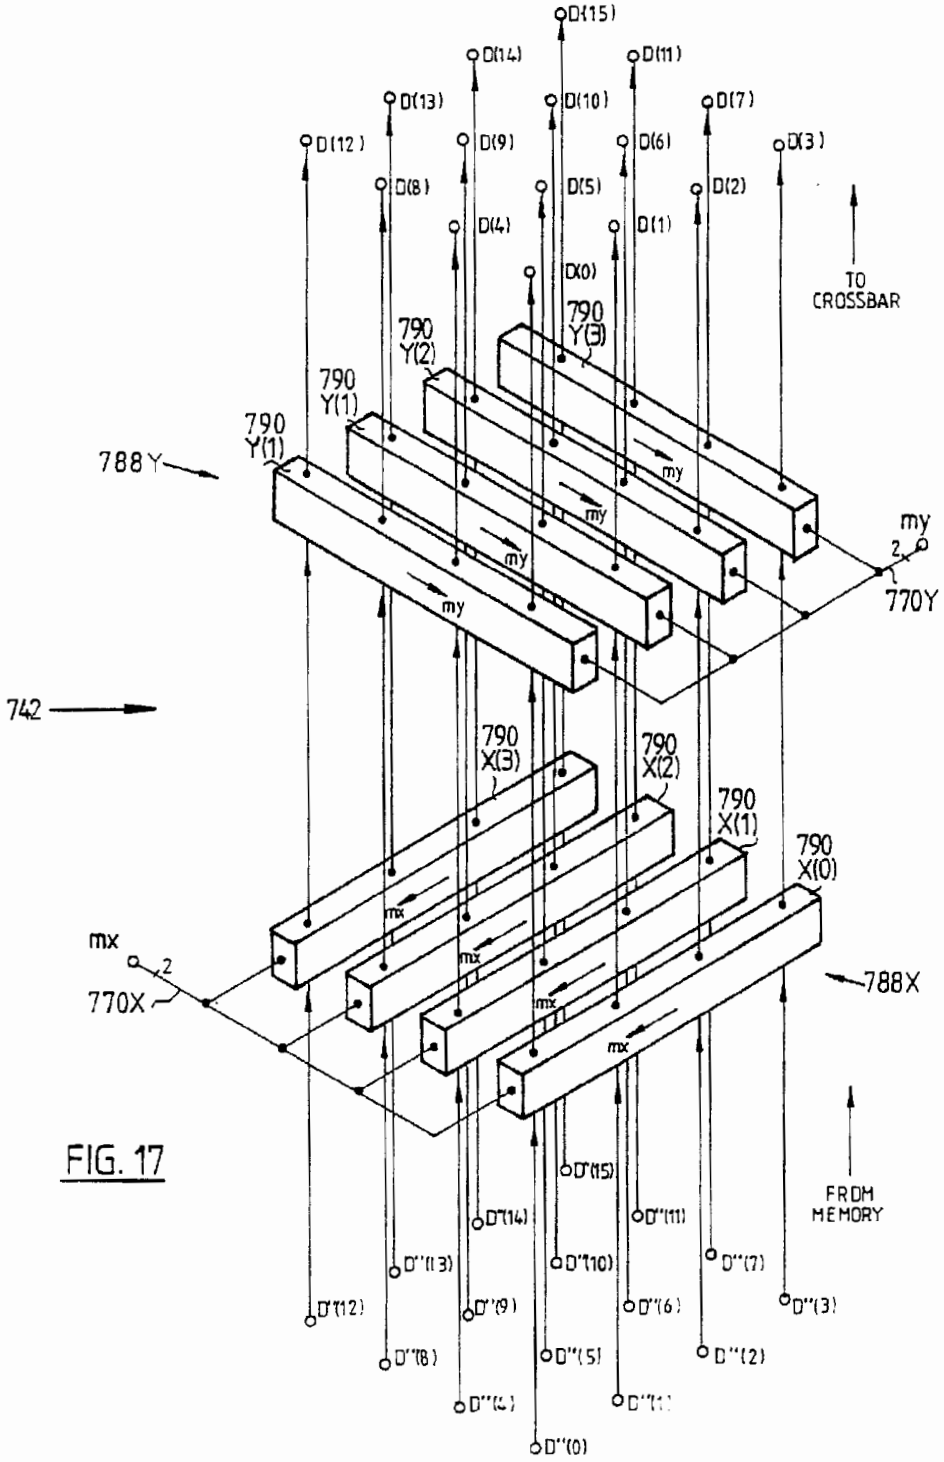

D(15)
D(14)
D(11)
D(13)
D(10)
D(7)
D(12)
D(9)
D(6)
D(3)
D(8)
D(5)
D(2)
D(4)
D(1)
D(0)
TO CROSSBAR
790 Y(3)
790 Y(2)
790 Y(1)
790 Y(1)
788Y
my
2
770Y
742
790 X(3)
790 X(2)
790 X(1)
790 X(0)
mx
2
770X
788X
D''(15)
D''(14)
D''(11)
D''(13)
D''(10)
D''(7)
D''(12)
D''(9)
D''(6)
D''(3)
D''(8)
D''(5)
D''(2)
D''(4)
D''(1)
D''(0)
FROM MEMORY

FIG. 17

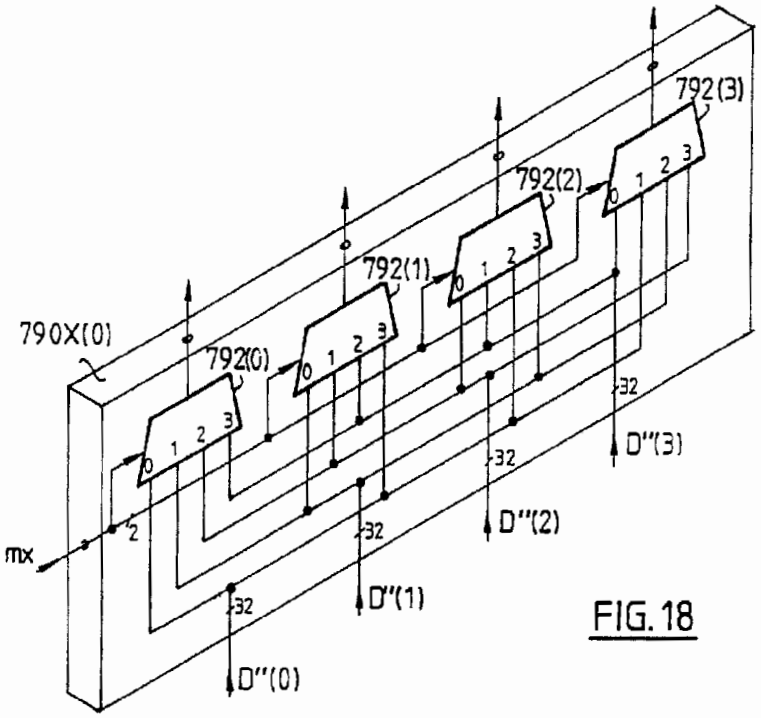


FIG. 18

D"(12)	D"(13)	D"(14)	D"(15)
D"(8)	D"(9)	D"(10)	D"(11)
D"(4)	D"(5)	D"(6)	D"(7)
D"(0)	D"(1)	D"(2)	D"(3)

mx=2

D"(14)	D"(15)	D"(12)	D"(13)
D"(10)	D"(11)	D"(8)	D"(9)
D"(6)	D"(7)	D"(4)	D"(5)
D"(2)	D"(3)	D"(0)	D"(1)

my=1

D"(2)	D"(3)	D"(0)	D"(1)
D"(14)	D"(15)	D"(12)	D"(13)
D"(10)	D"(11)	D"(8)	D"(9)
D"(6)	D"(7)	D"(4)	D"(5)

equal to

D(12)	D(13)	D(14)	D(15)
D(8)	D(9)	D(10)	D(11)
D(4)	D(5)	D(6)	D(7)
D(0)	D(1)	D(2)	D(3)

FIG. 19

FIG. 20

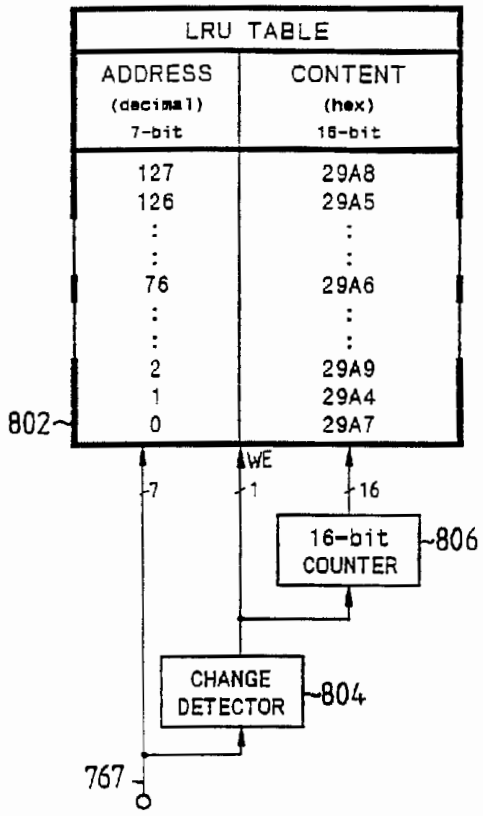

LRU TABLE
ADDRESS (decimal) 7-bit
CONTENT (hex) 16-bit
127 29A8
126 29A5
76 29A6
2 29A9
1 29A4
0 29A7
802
WE
7
1
16
16-bit COUNTER
806
CHANGE DETECTOR
804
767

FIG. 21

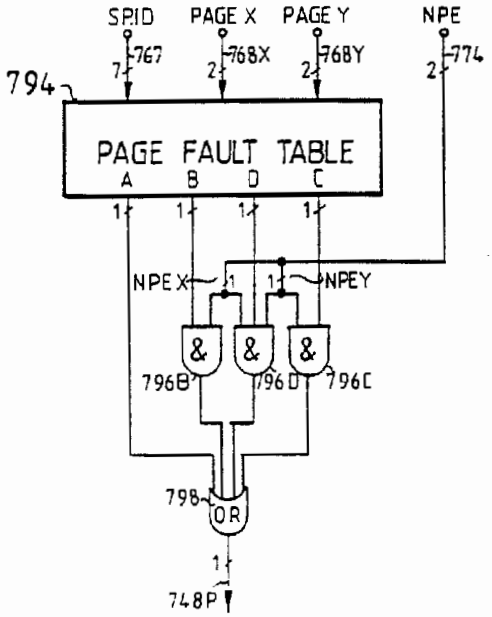
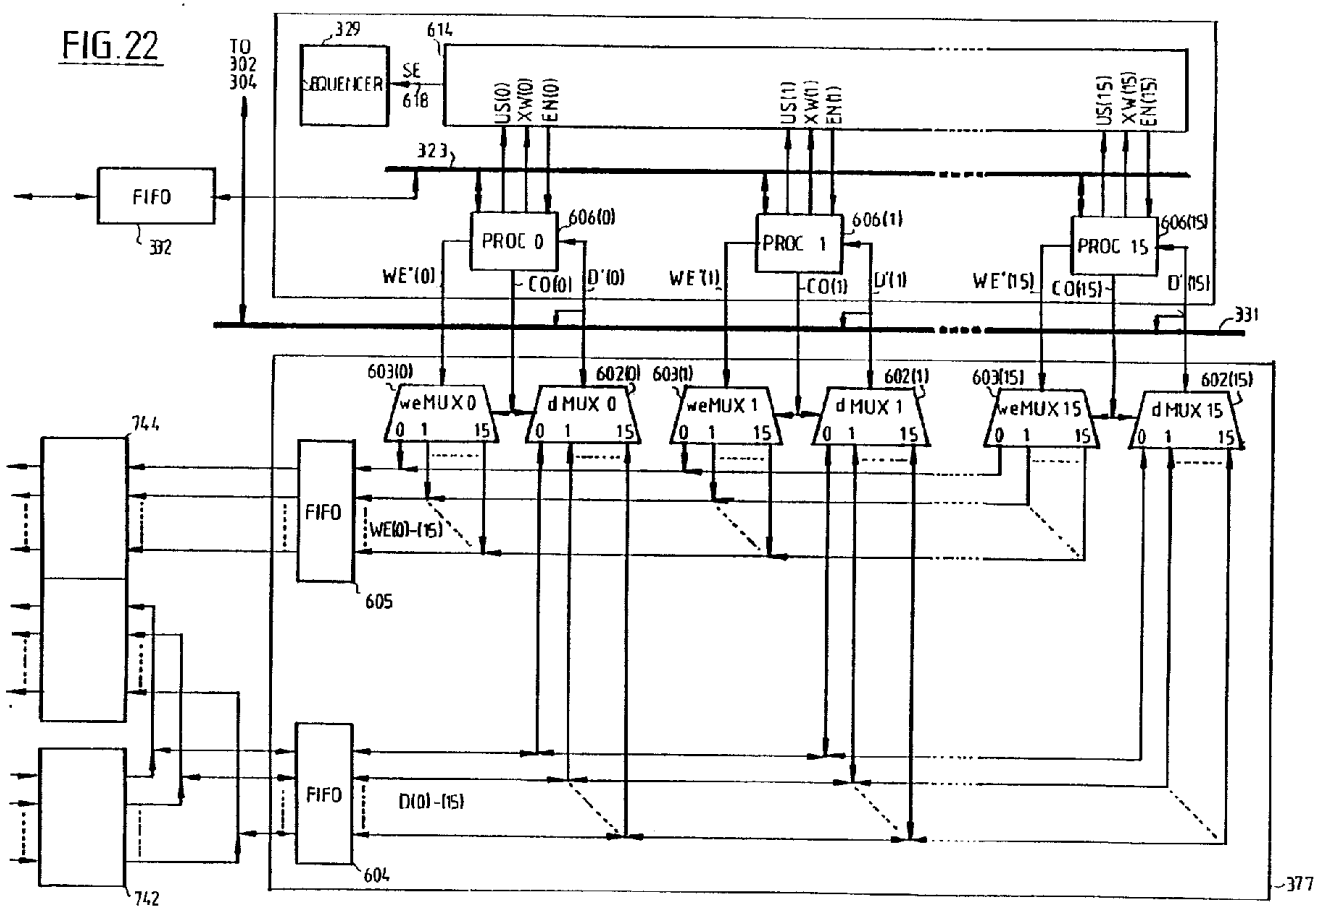

SPID
PAGE X
PAGE Y
NPE
767
768X
768Y
774
7
2
2
2
794
PAGE FAULT TABLE
A B D C
1 1 1 1
NPEX
NPEY
&
&
&
796B
796D
796C
798
OR
1
748P

FIG.22



Micron Technology Inc. et al.
Ex. 1010, p. 422

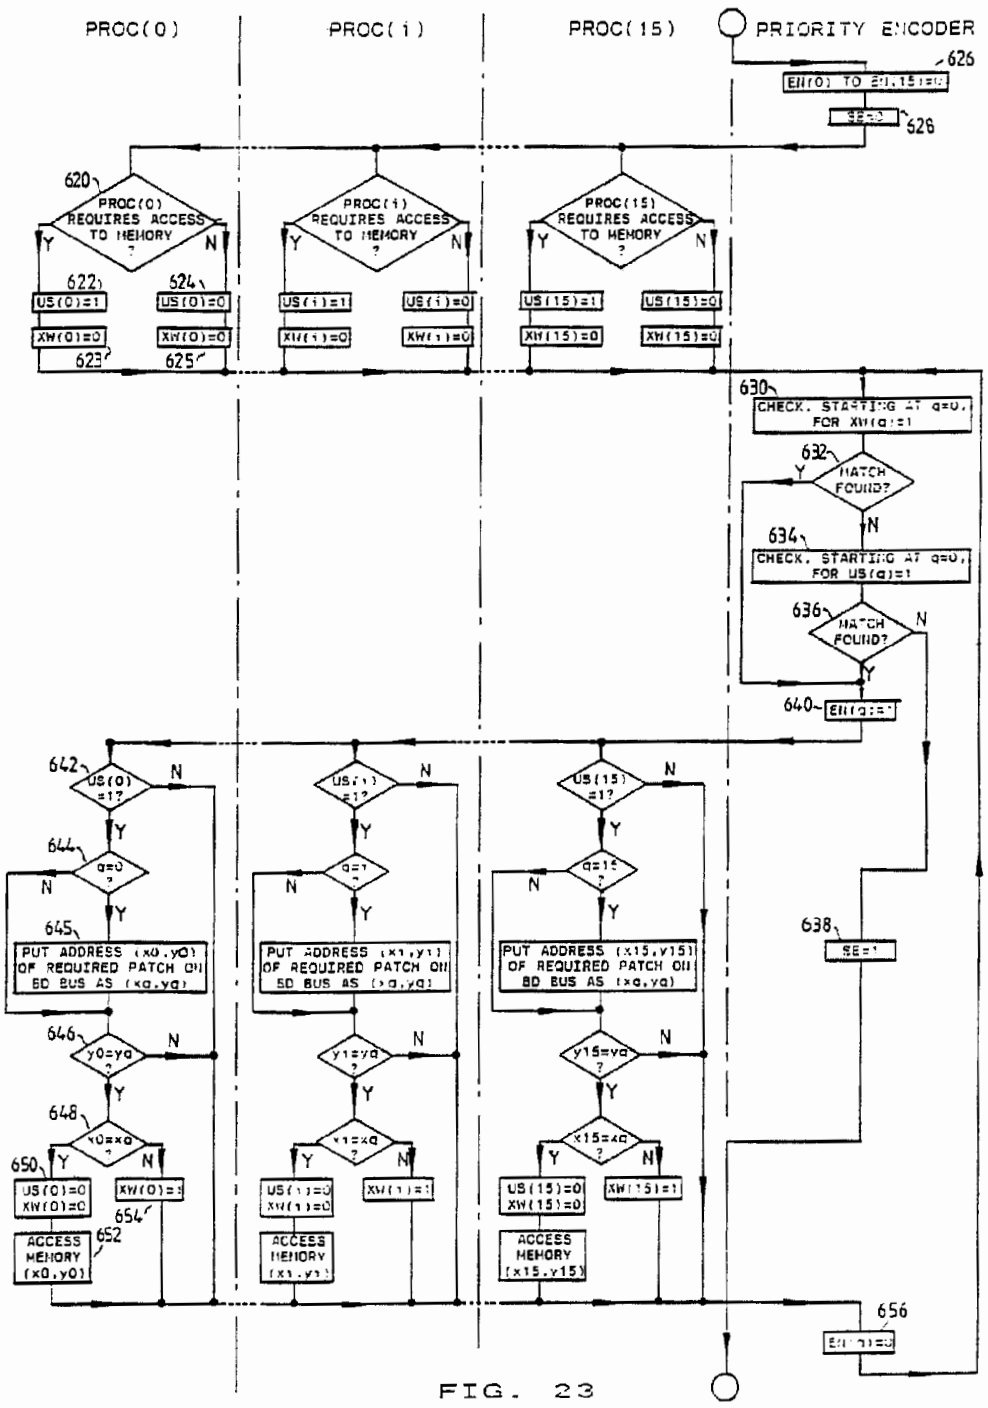

PROC(0)
PROC(i)
PROC(15)
PRIORITY ENCODER
626
EN(0) TO EN(15)=0
SE=0
626
620
PROC(0) REQUIRES ACCESS TO MEMORY ?
PROC(i) REQUIRES ACCESS TO MEMORY ?
PROC(15) REQUIRES ACCESS TO MEMORY ?
Y
N
622
US(0)=1
XW(0)=0
623
624
US(0)=0
XW(0)=0
625
US(i)=1
XW(i)=0
US(i)=0
XW(i)=0
US(15)=1
XW(15)=0
US(15)=0
XW(15)=0
630
CHECK, STARTING AT q=0, FOR XW(q)=1
632
MATCH FOUND?
634
CHECK, STARTING AT q=0, FOR US(q)=1
636
MATCH FOUND?
640
EN(q)=1
642
US(0) =1?
US(i) =1?
US(15) =1?
644
q=0 ?
q=i ?
q=15 ?
645
PUT ADDRESS (x0,y0) OF REQUIRED PATCH ON SD BUS AS (xq,yq)
PUT ADDRESS (xi,yi) OF REQUIRED PATCH ON SD BUS AS (xq,yq)
PUT ADDRESS (x15,y15) OF REQUIRED PATCH ON SD BUS AS (xq,yq)
646
y0=yq ?
yi=yq ?
y15=yq ?
648
x0=xq ?
xi=xq ?
x15=xq ?
650
US(0)=0
XW(0)=0
XW(0)=1
654
US(i)=0
XW(i)=0
XW(i)=1
US(15)=0
XW(15)=0
XW(15)=1
652
ACCESS MEMORY (x0,y0)
ACCESS MEMORY (xi,yi)
ACCESS MEMORY (x15,y15)
638
SE=1
656
EN(q)=0

FIG. 23

	660		662		664		666		668		670		672		674	
i	xi hex	yi hex	Init. US	XW	q=0 US	XW	q=2 US	XW	q=10 US	XW	q=3 US	XW	q=11 US	XW	q=8 US	XW
[0]	1234	1234	[1]	0	0	0	0	0	0	0	0	0	0	0	0	0
[1]	1234	1234	1	0	0	0	0	0	0	0	0	0	0	0	0	0
[2]	1235	1234	1	0	[1]	[1]	0	0	0	0	0	0	0	0	0	0
[3]	1235	1235	1	0	1	0	1	0	[1]	0	0	0	0	0	0	0
[4]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[5]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[6]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[7]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[8]	1234	1236	1	0	1	0	1	0	1	0	1	0	[1]	0	0	0
[9]	1234	1236	1	0	1	0	1	0	1	0	1	0	1	0	0	0
[10]	1236	1234	1	0	1	1	[1]	[1]	0	0	0	0	0	0	0	0
[11]	1236	1235	1	0	1	0	1	0	1	0	[1]	[1]	0	0	0	0
[12]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[13]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[14]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
[15]	-	-	0	0	0	0	0	0	0	0	0	0	0	0	0	0
ADDRESSES READ					x0,y0 x1,y1		x2,y2		x10,y10		x3,y3		x11,y11		x8,y8 x9,y9	

FIG. 24

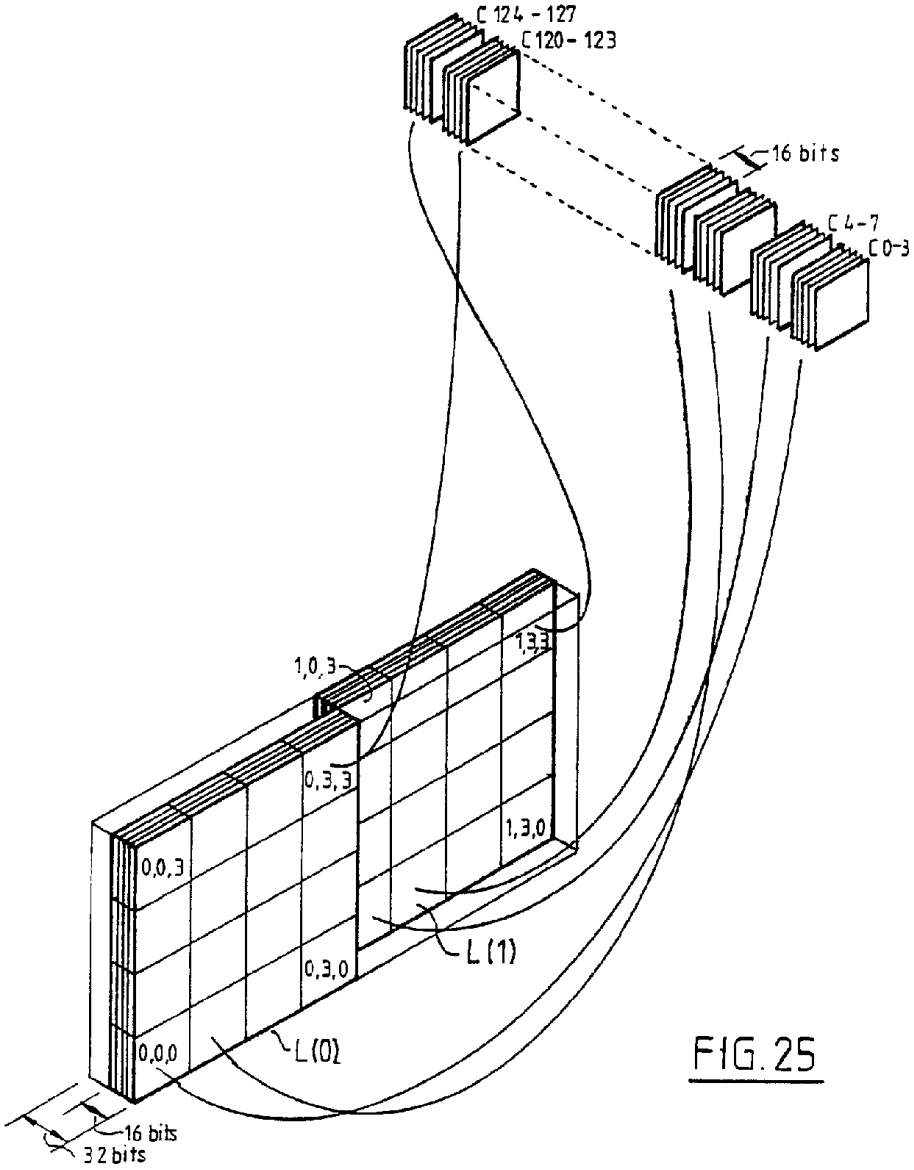

C 124 - 127
C 120 - 123
16 bits
C 4 - 7
C 0-3
1,0,3
1,3,3
0,3,3
1,3,0
0,0,3
L (1)
0,3,0
0,0,0
L (0)
16 bits
32 bits

FIG. 25

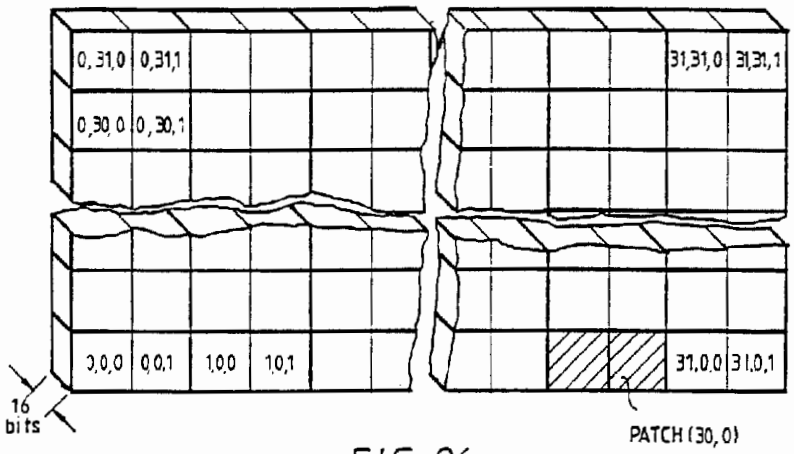

0,31,0
0,31,1
31,31,0
31,31,1
0,30,0
0,30,1
0,0,0
0,0,1
1,0,0
1,0,1
31,0,0
31,0,1
16 bits
PATCH (30,0)

FIG. 26

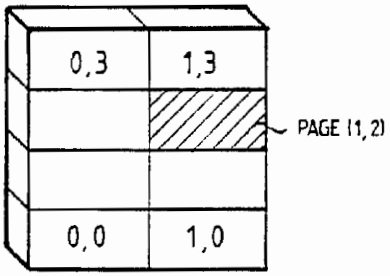

0,3
1,3
PAGE (1,2)
0,0
1,0

FIG. 27

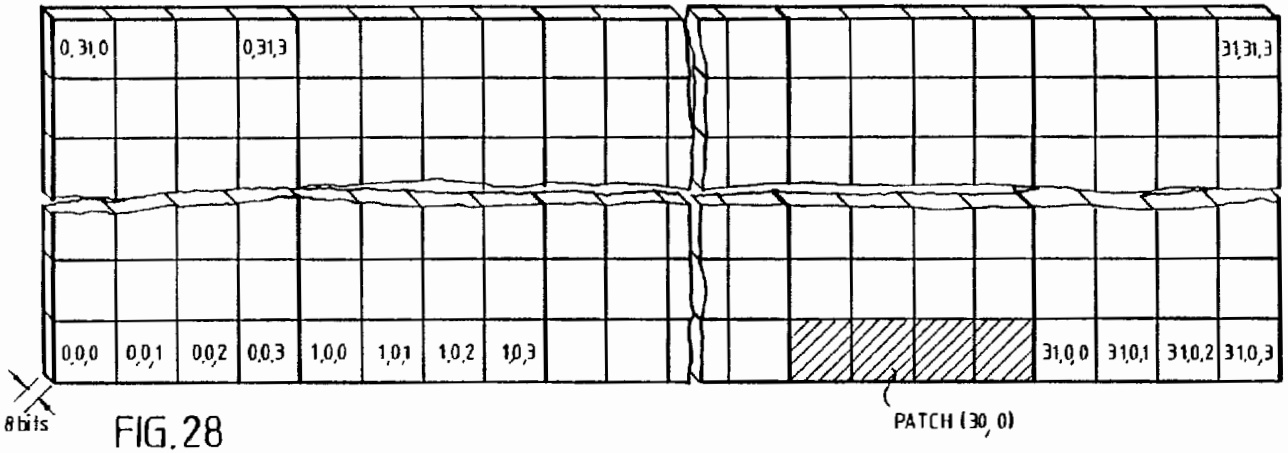

0,31,0
0,31,3
31,31,3
0,0,0
0,0,1
0,0,2
0,0,3
1,0,0
1,0,1
1,0,2
1,0,3
31,0,0
31,0,1
31,0,2
31,0,3
8 bits
PATCH (30, 0)

FIG. 28

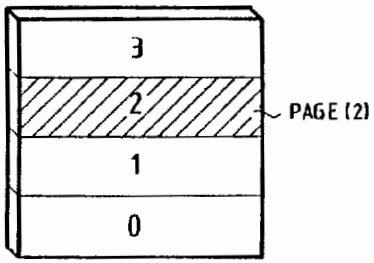

3
2
1
0
PAGE (2)

FIG. 29

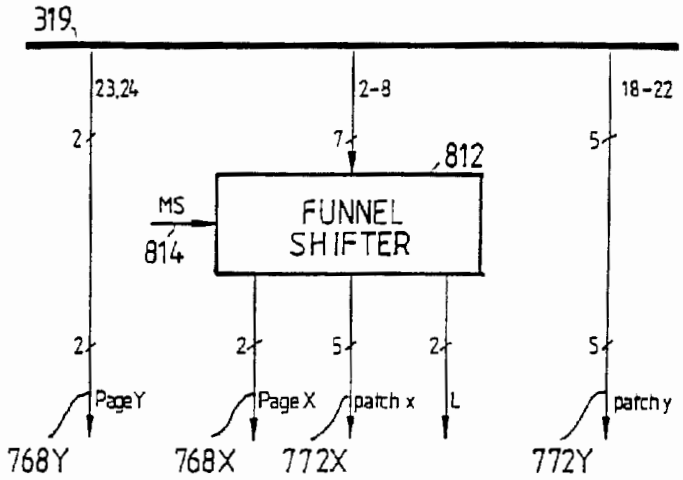


FIG. 30A

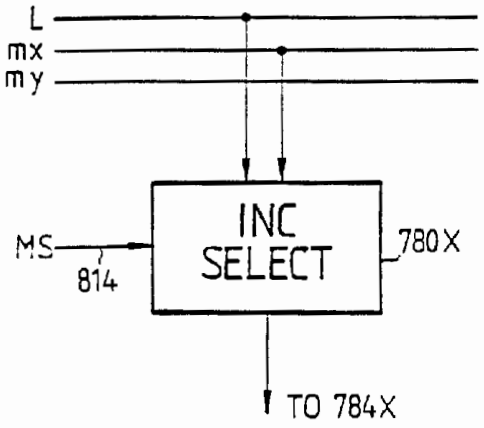


FIG. 30B

CONDITION			
32-bit L ignored	16-bit L=0 or 1	8-bit 0≤L≤3	PATCH X ADDRESS
bx≥mx	L=0 OR bx≥mx	L<>3 OR bx≥mx	px
bx<mx	L=1 AND bx<mx	L=3 AND bx<mx	px+1

FIG. 30C

VIRTUAL ADDRESS BITS

MODE	8	7	6	5	4	3	2
32-bit	PAGE PX (2)		patch px (5)				
16-bit	PX (1)	px (5)					L (1)
8-bit	px (5)						L (2)

FIG. 31

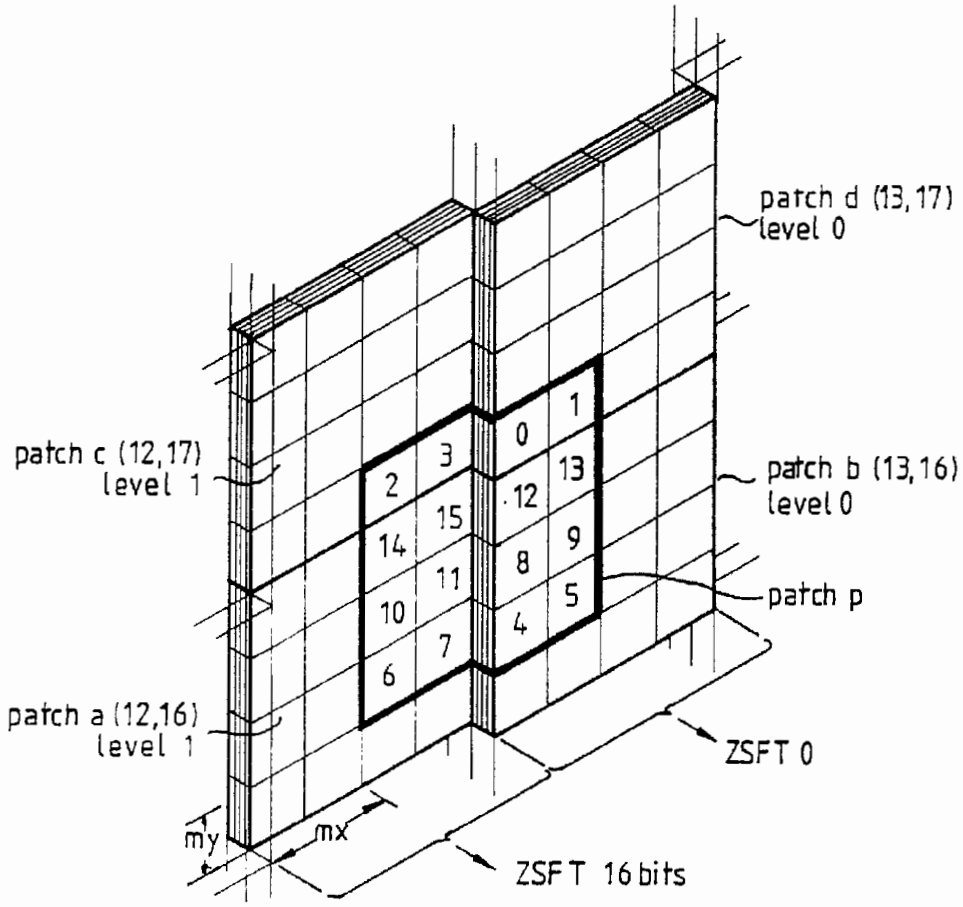

patch d (13,17)
level 0
patch c (12,17)
level 1
patch b (13,16)
level 0
patch p
patch a (12,16)
level 1
ZSFT 0
ZSFT 16 bits
my
mx
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

FIG. 32

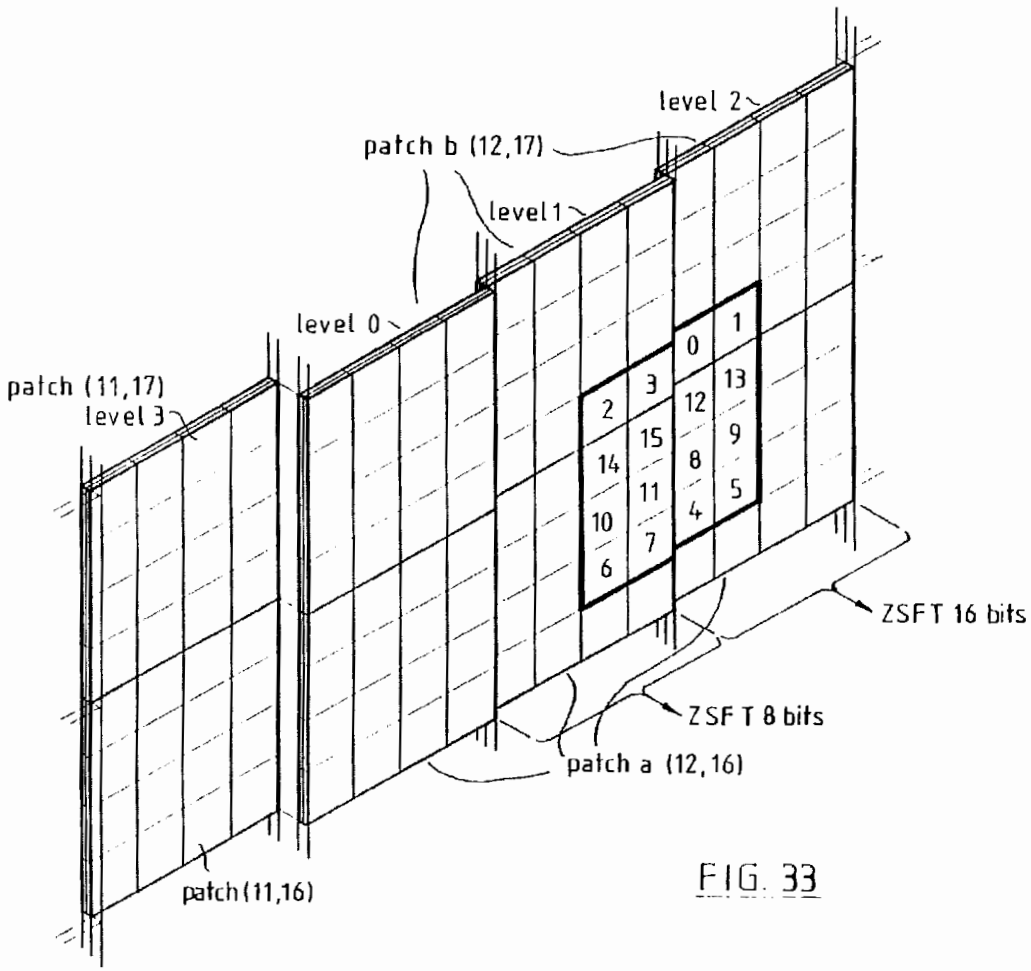


FIG. 33

Micron Technology Inc. et al.
Ex. 1010, p. 431

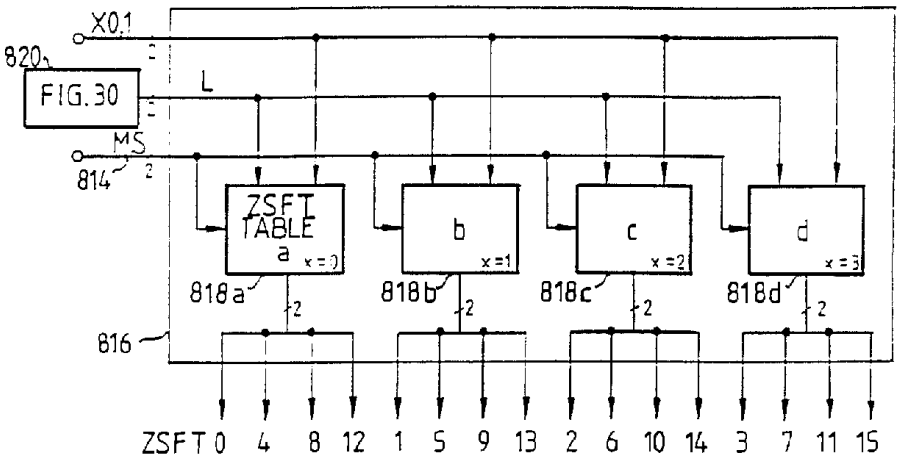


FIG. 34

GENERAL ZSFT TABLE

	8-bit MODE IF L=0	IF L=1	IF L=2	IF L=3	16-bit MODE IF L=0	IF L=1	32-bit
IF mx+x=4	0	1	2	3	0	2	ALWAYS
IF mx+x=3	1	2	3	0	2	0	0

FIG. 35A

ZSFT TABLE 818b (x=1)

ZSFT (1), (5), (9), (13)

mx \ L	8-bit MODE 0	1	2	3	16-bit MODE 0	1	32-bit
0	0	1	2	3	0	2	
1	0	1	2	3	0	2	ALWAYS
2	0	1	2	3	0	2	0
3	1	2	3	0	2	0	

FIG. 35B

Micron Technology Inc. et al.
Ex. 1010, p. 432

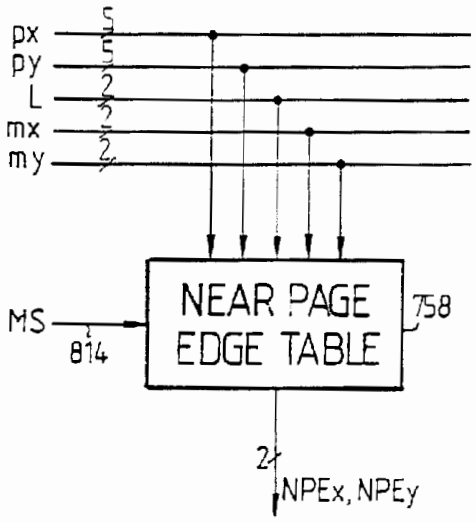
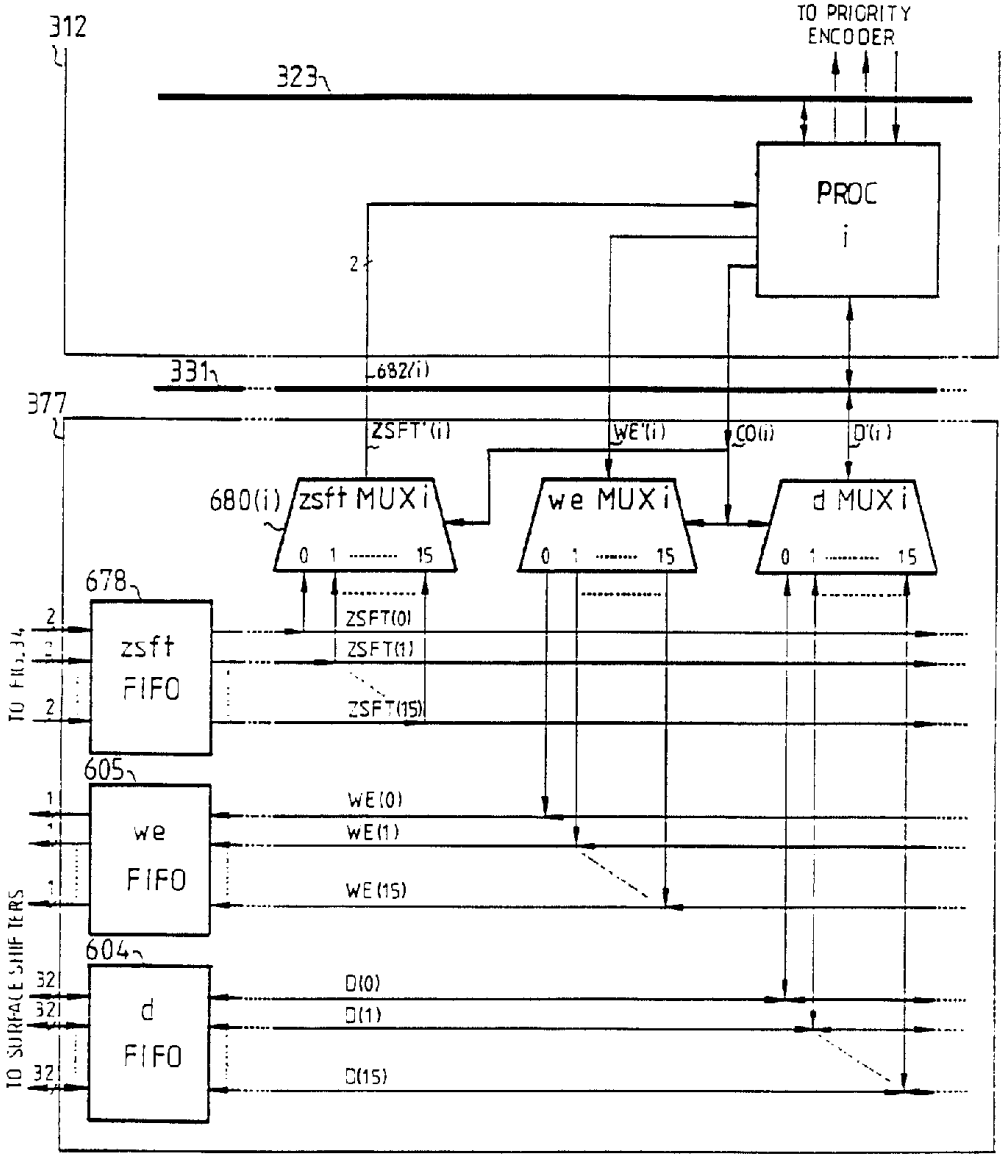


FIG. 36A

CONDITION			
32-bit L=0	16-bit L=0 or 1	8-bit 0≤L≤3	
4px+(4L)+mx<125	8px+4L+mx<253	16px+4L+mx<509	NPEx=0
4px+(4L)+mx>124	8px+4L+mx>252	16px+4L+mx>508	NPEx=1
4py+my<125	4py+my<125	4py+my<125	NPEy=0
4py+my>124	4py+my>124	4py+my>124	NPEy=1

FIG. 36 B

FIG. 37


312
323
TO PRIORITY ENCODER
PROC i
2
331
682(i)
377
ZSFT'(i)
WE'(i)
CO(i)
D'(i)
680(i)
zsft MUX i
we MUX i
d MUX i
0 1 15
678
zsft FIFO
ZSFT(0)
ZSFT(1)
ZSFT(15)
TO FIG.34
605
we FIFO
WE(0)
WE(1)
WE(15)
604
d FIFO
D(0)
D(1)
D(15)
TO SURFACE SHIFTERS
32

GENERAL PWE TABLE

	8-bit MODE				16-bit MODE		32-bit
	IF L=0	IF L=1	IF L=2	IF L=3	IF L=0	IF L=1	
IF bx≥mx	0001	0010	0100	1000	0011	1100	ALWAYS
IF bx<mx	0010	0100	1000	0001	1100	0011	1111

FIG. 38

PWE TABLE (bx=2)

L / mx	8-bit MODE 0	1	2	3	16-bit MODE 0	1	32-bit
0	0001	0010	0100	1000	0011	1100	
1	0001	0010	0100	1000	0011	1100	ALWAYS
2	0001	0010	0100	1000	0011	1100	1111
3	0010	0100	1000	0001	1100	0011	

FIG. 39

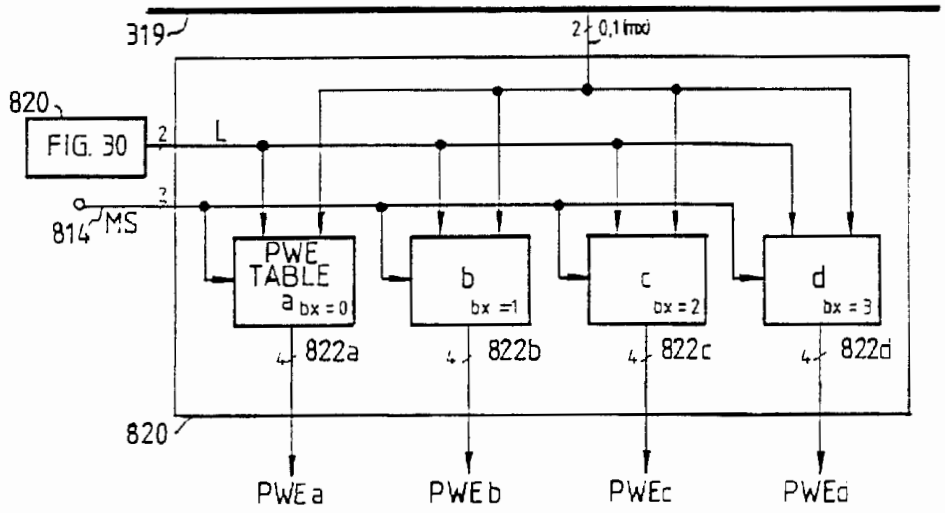


FIG. 40

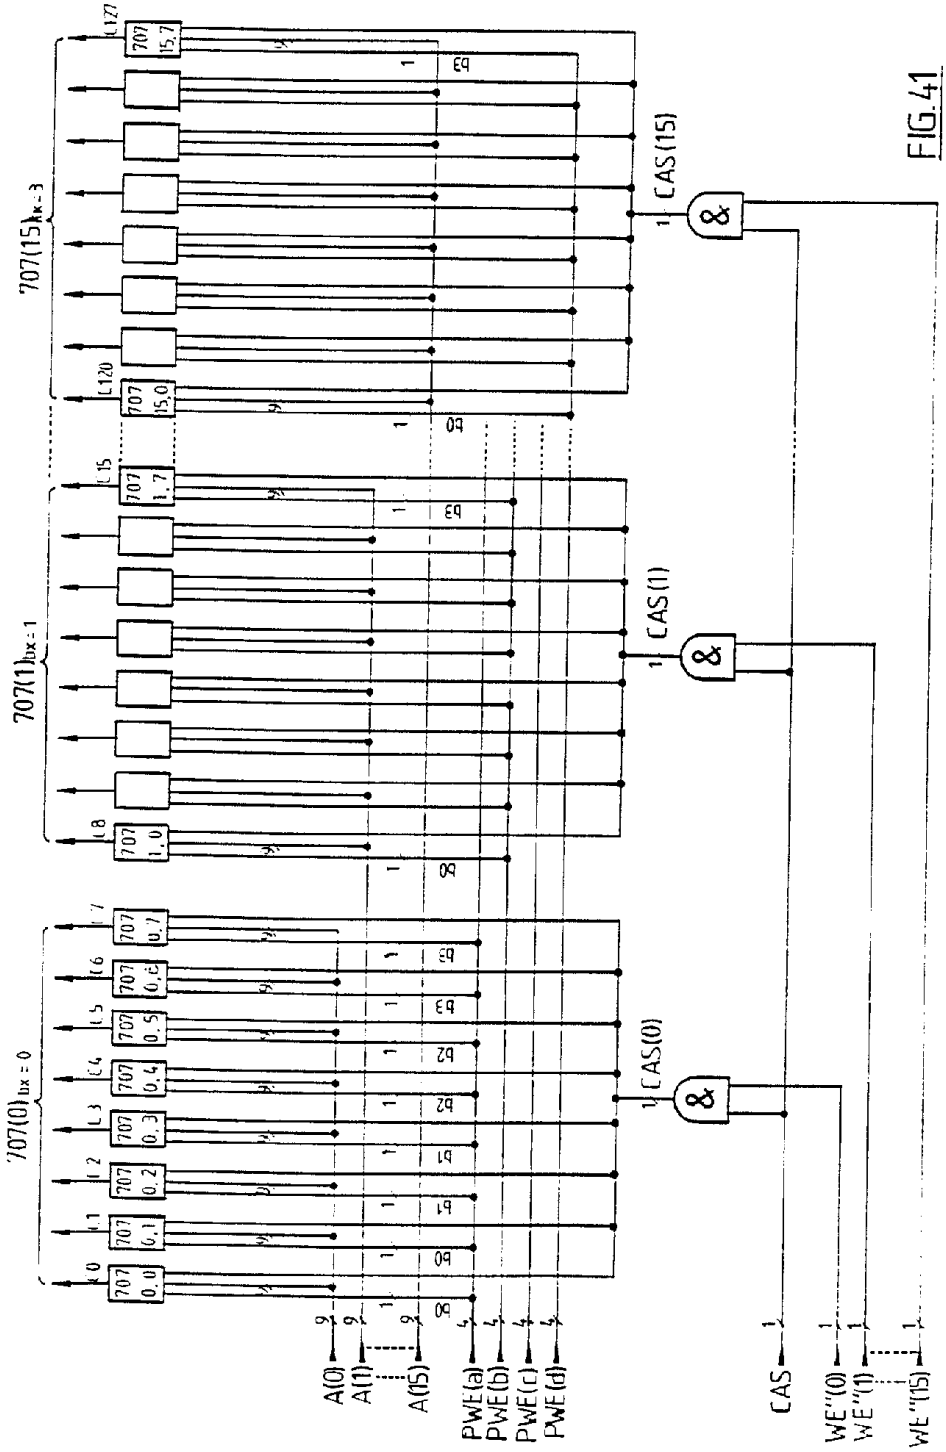

707(15)bx = 3
707(1)bx = 1
707(0)bx = 0
CAS(15)
CAS(1)
CAS(0)
A(0)
A(1)
A(15)
PWE(a)
PWE(b)
PWE(c)
PWE(d)
CAS
WE''(0)
WE''(1)
WE''(15)

FIG.41

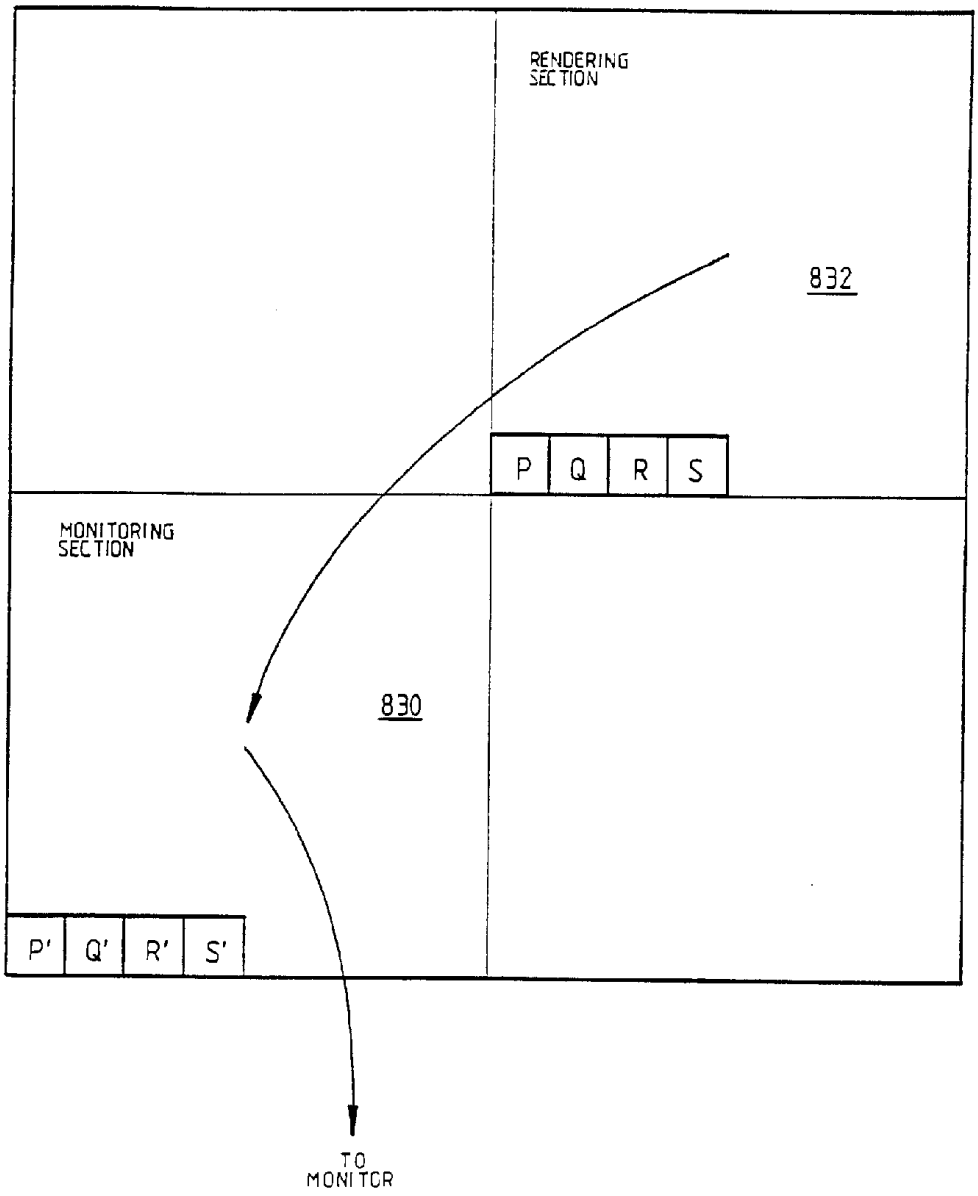

RENDERING SECTION
832
P
Q
R
S
MONITORING SECTION
830
P'
Q'
R'
S'
TO MONITOR

FIG. 42

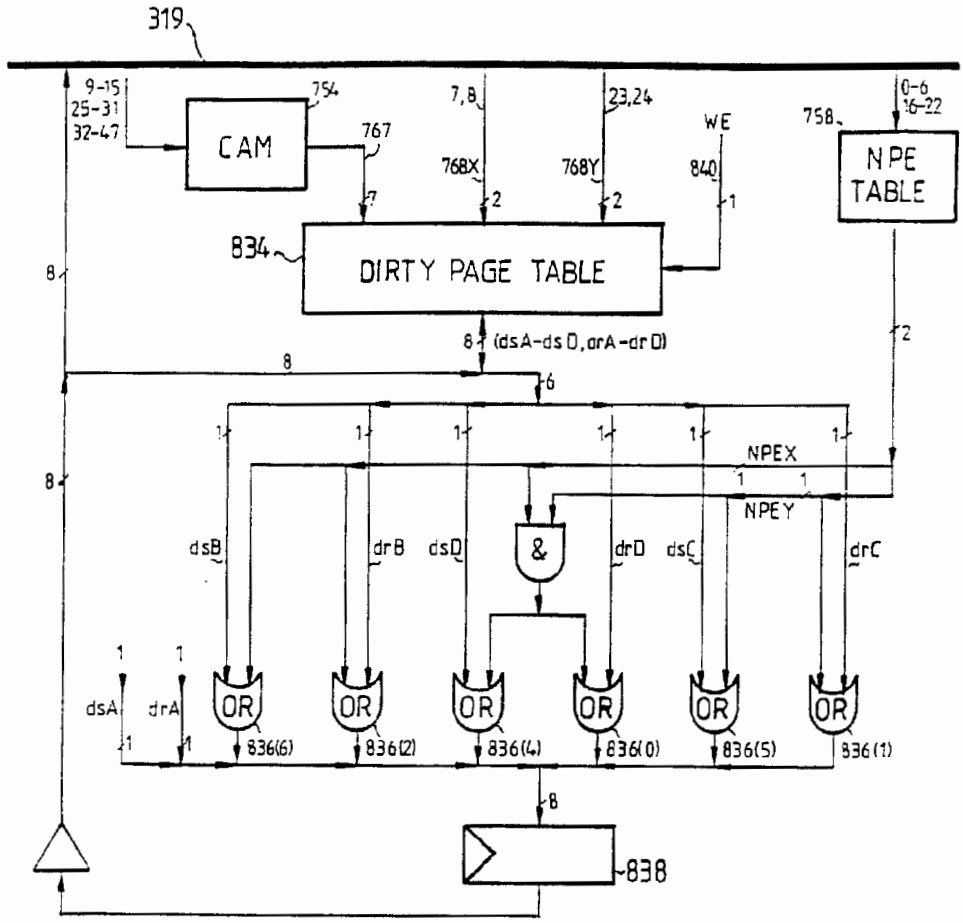

319
9-15
25-31
32-47
CAM
754
767
7,8
768X
23,24
768Y
WE
840
758
0-6
16-22
NPE TABLE
834
DIRTY PAGE TABLE
8
(dsA-dsD, drA-drD)
6
2
NPEX
NPEY
&
dsB
drB
dsD
drD
dsC
drC
dsA
drA
OR
836(6)
836(2)
836(4)
836(0)
836(5)
836(1)
838

FIG. 43

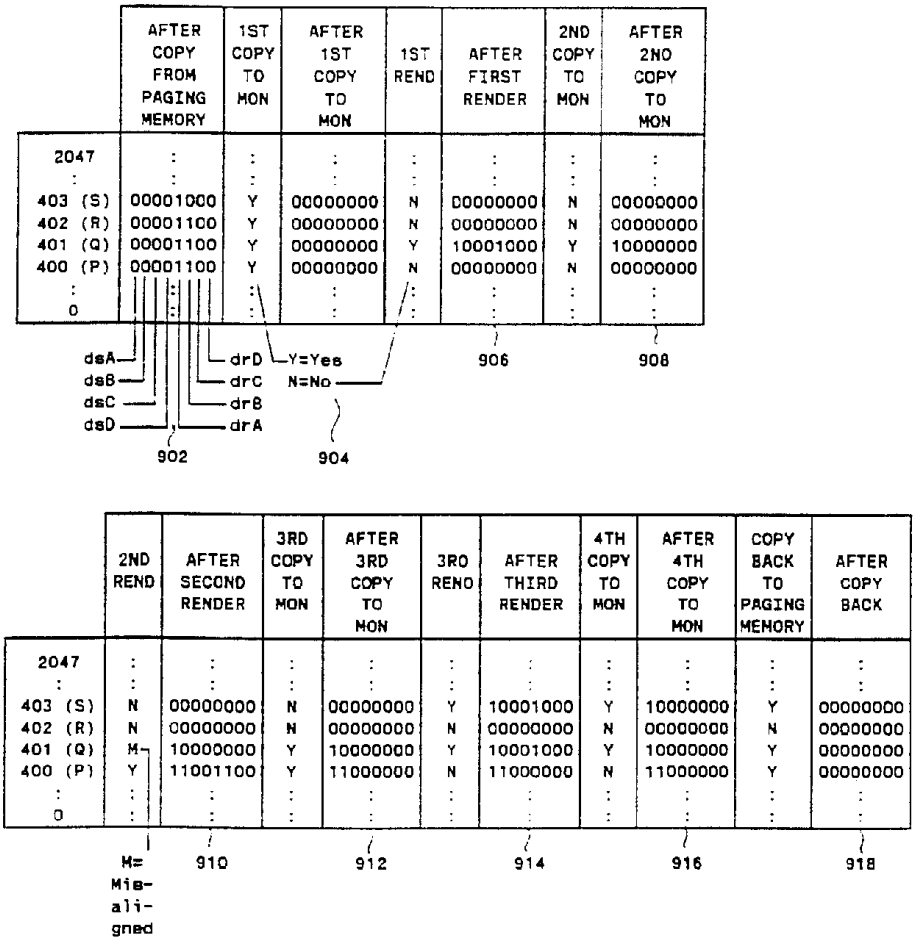

AFTER COPY FROM PAGING MEMORY
1ST COPY TO MON
AFTER 1ST COPY TO MON
1ST REND
AFTER FIRST RENDER
2ND COPY TO MON
AFTER 2NO COPY TO MON
2047
403 (S) 00001000 Y 00000000 N 00000000 N 00000000
402 (R) 00001100 Y 00000000 N 00000000 N 00000000
401 (Q) 00001100 Y 00000000 Y 10001000 Y 10000000
400 (P) 00001100 Y 00000000 N 00000000 N 00000000
0
dsA
dsB
dsC
dsD
drD
drC
drB
drA
Y=Yes
N=No
902
904
906
908
2ND REND
AFTER SECOND RENDER
3RD COPY TO MON
AFTER 3RD COPY TO MON
3RO RENO
AFTER THIRD RENDER
4TH COPY TO MON
AFTER 4TH COPY TO MON
COPY BACK TO PAGING MEMORY
AFTER COPY BACK
403 (S) N 00000000 N 00000000 Y 10001000 Y 10000000 Y 00000000
402 (R) N 00000000 N 00000000 N 00000000 N 00000000 N 00000000
401 (Q) M 10000000 Y 10000000 Y 10001000 Y 10000000 Y 00000000
400 (P) Y 11001100 Y 11000000 N 11000000 N 11000000 Y 00000000
0
M= Mis-ali-gned
910
912
914
916
918

FIG. 44

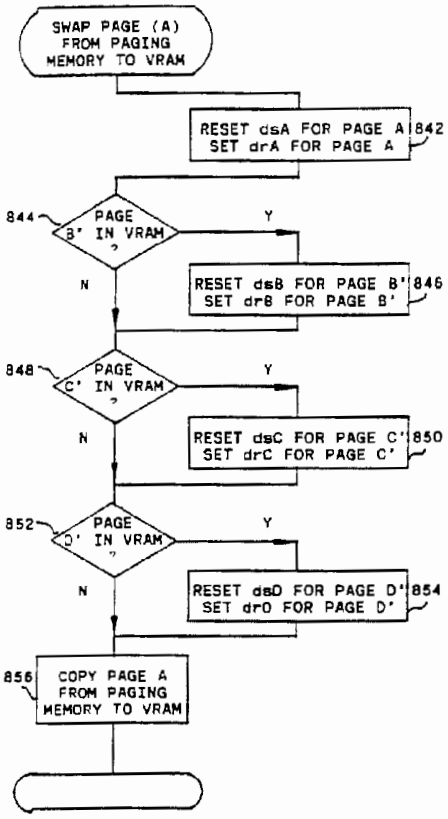


FIG. 45A

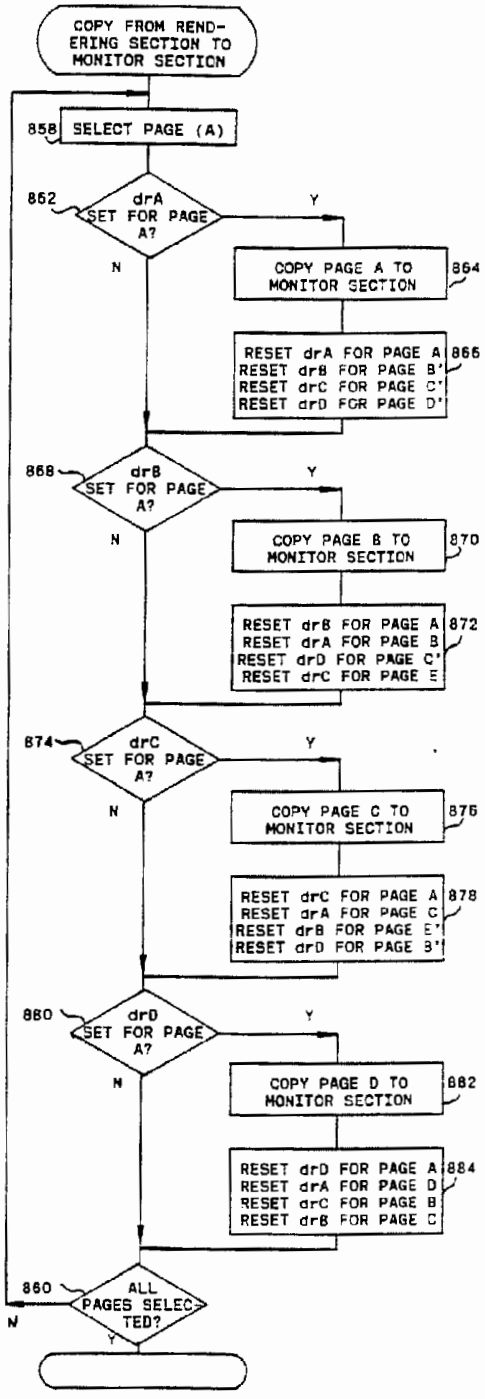


FIG. 45B

PAGE A AT (PX,PY)
PAGE B AT (PX+1,PY); PAGE B' AT (PX-1,PY)
PAGE C AT (PX,PY+1); PAGE C' AT (PX,PY-1)
PAGE D AT (PX+1,PY+1); PAGE D' AT (PX-1,PY-1)
PAGE E AT (PX+1,PY-1); PAGE E' AT (PX-1,PY+1)

PAGE E'	PAGE C	PAGE D
PAGE B'	PAGE A	PAGE B
PAGE D'	PAGE C'	PAGE E

FIG. 45D

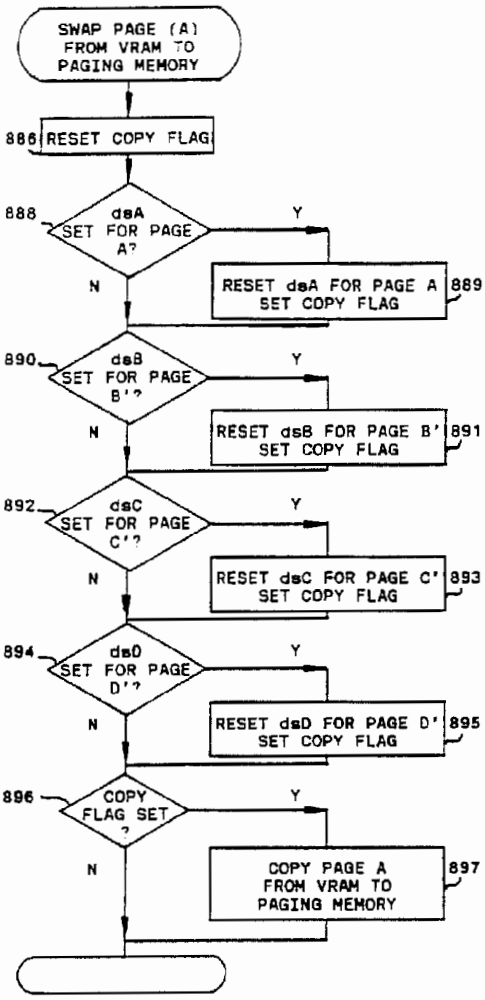

SWAP PAGE (A) FROM VRAM TO PAGING MEMORY
886
RESET COPY FLAG
888
dsA SET FOR PAGE A?
Y
N
RESET dsA FOR PAGE A SET COPY FLAG
889
890
dsB SET FOR PAGE B'?
Y
N
RESET dsB FOR PAGE B' SET COPY FLAG
891
892
dsC SET FOR PAGE C'?
Y
N
RESET dsC FOR PAGE C' SET COPY FLAG
893
894
dsD SET FOR PAGE D'?
Y
N
RESET dsD FOR PAGE D' SET COPY FLAG
895
896
COPY FLAG SET ?
Y
N
COPY PAGE A FROM VRAM TO PAGING MEMORY
897

FIG. 45C

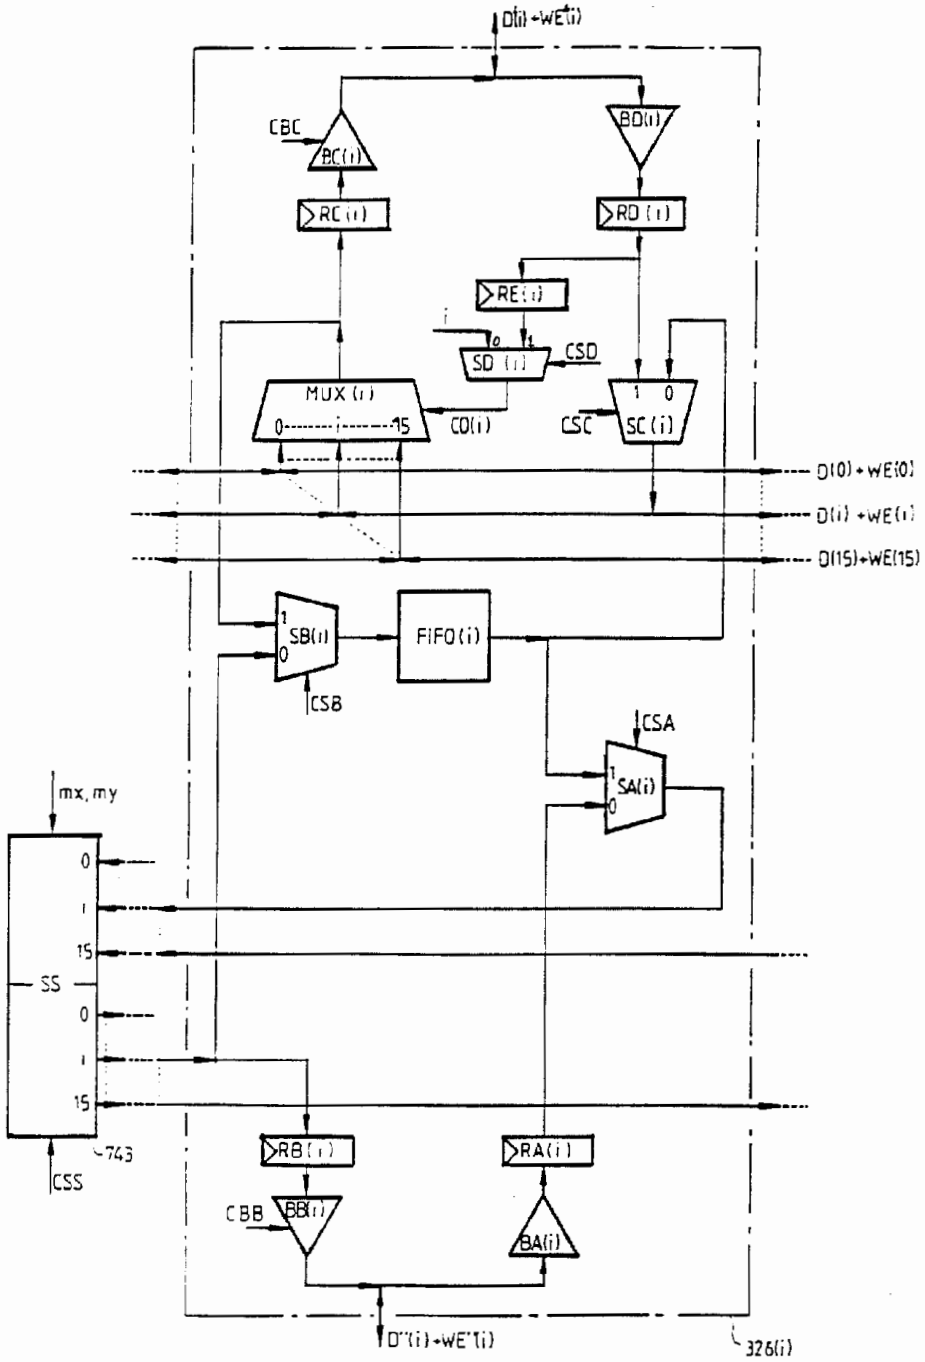

D(i)+WE(i)
CBC
BC(i)
BD(i)
RC(i)
RD(i)
RE(i)
SD(i)
CSD
MUX(i)
0
15
CO(i)
CSC
SC(i)
D(0)+WE(0)
D(i)+WE(i)
D(15)+WE(15)
SB(i)
FIFO(i)
CSB
CSA
SA(i)
mx, my
SS
CSS
743
RB(i)
RA(i)
CBB
BB(i)
BA(i)
D'(i)+WE'(i)
326(i)

FIG. 46

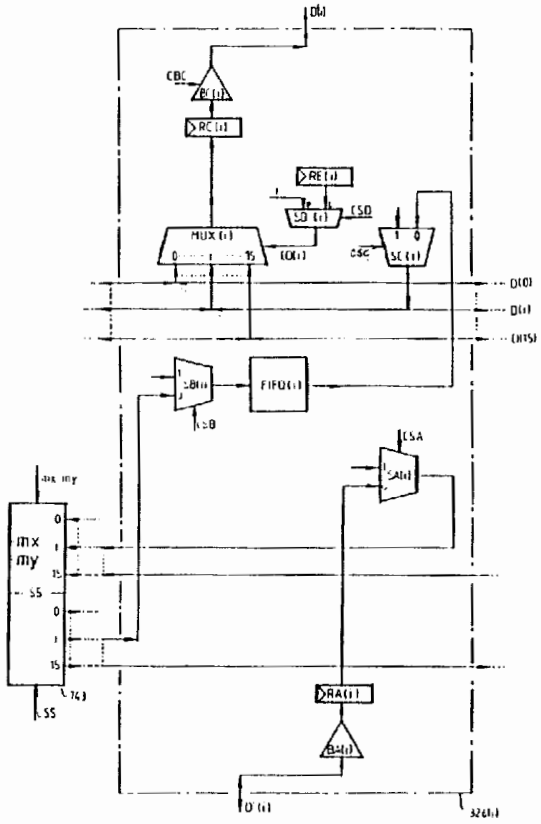


FIG.47A

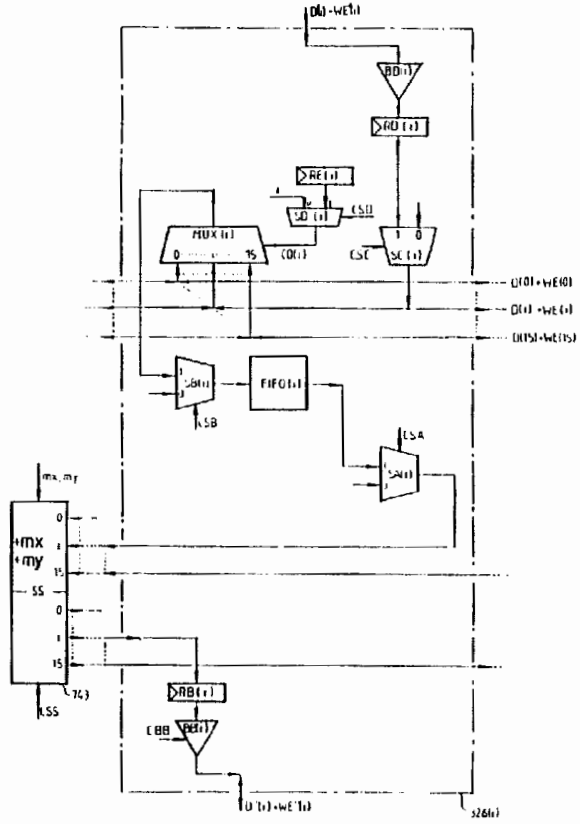


FIG.47B

Micron Technology Inc. et al.
Ex. 1010, p. 443

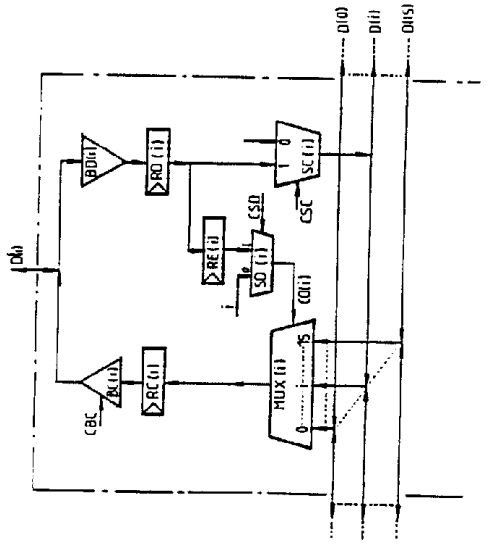


FIG. 47C

Micron Technology Inc. et al.
Ex. 1010, p. 444

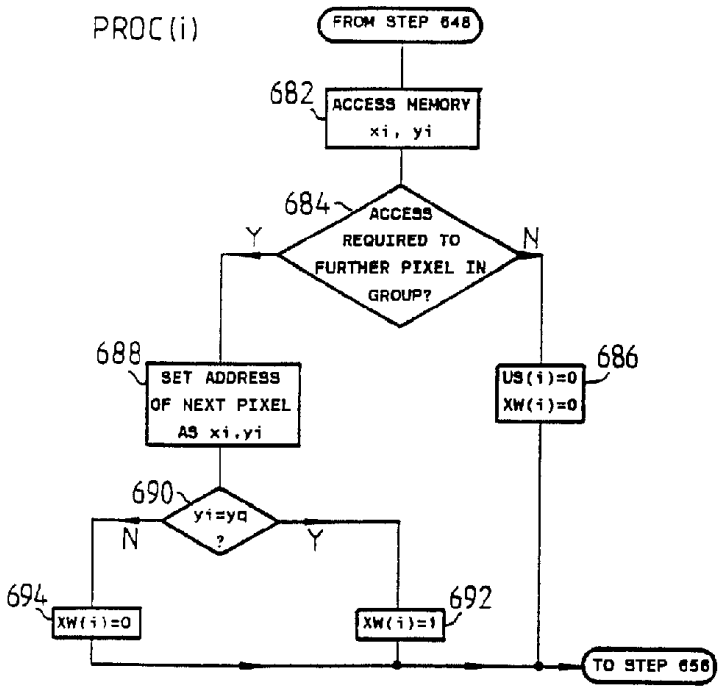

PROC(i)
FROM STEP 648
682
ACCESS MEMORY
xi, yi
684
ACCESS REQUIRED TO FURTHER PIXEL IN GROUP?
Y
N
688
SET ADDRESS OF NEXT PIXEL AS xi,yi
686
US(i)=0
XW(i)=0
690
yi=yq ?
N
Y
694
XW(i)=0
692
XW(i)=1
TO STEP 656

FIG. 48

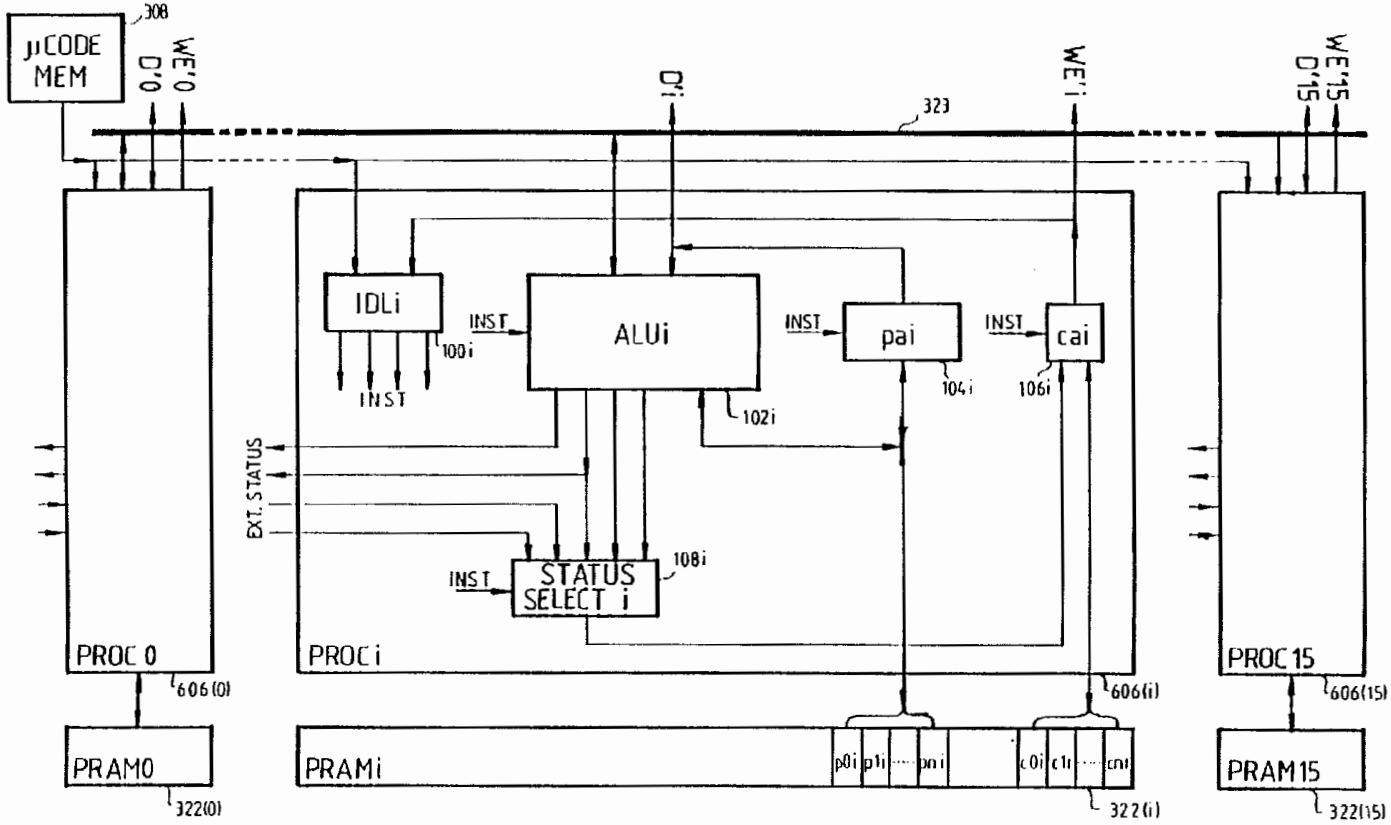

µCODE MEM
308
WE'0
D'0
D'i
323
WE'i
WE'15
D'15
IDLi
INST
100i
ALUi
INST
pai
104i
INST
cai
106i
102i
EXT. STATUS
INST
STATUS SELECT i
108i
PROC 0
606(0)
PROCi
606(i)
PROC 15
606(15)
PRAM0
322(0)
PRAMi
p0i p1i ... pni
c0i c1i ... cni
322(i)
PRAM15
322(15)

FIG. 49

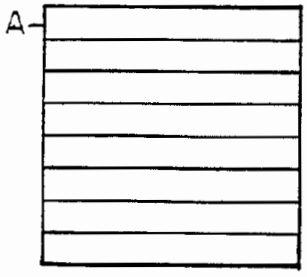


FIG. 50A

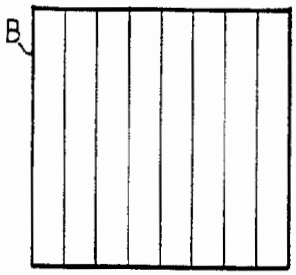


FIG. 50B

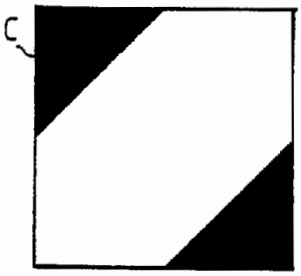


FIG. 50C

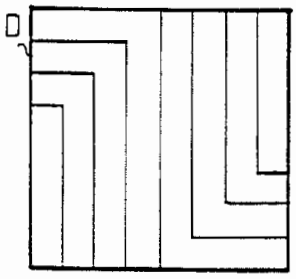


FIG. 50D

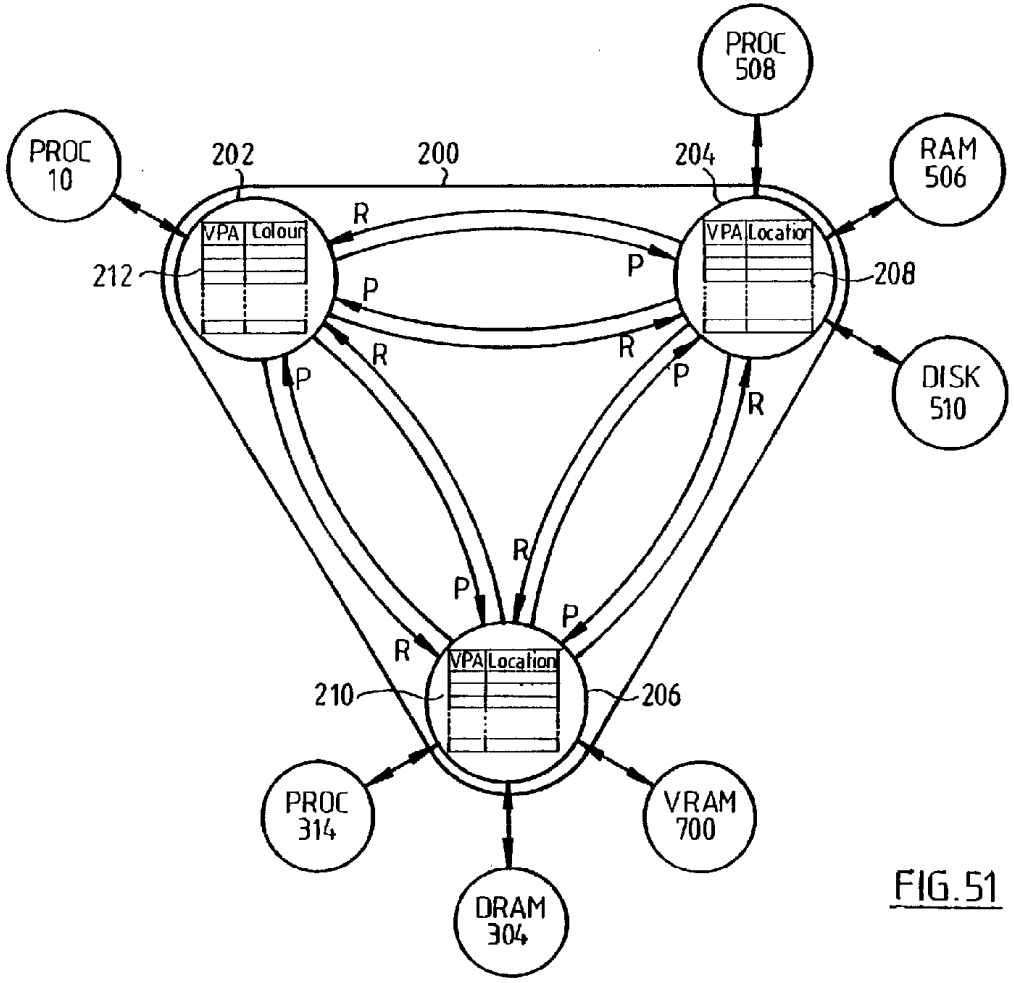


FIG.51

PCT/GB90/01213

Micron Technology Inc. et al.
Ex. 1010, p. 448

INTERNATIONAL SEARCH REPORT

International Application No PCT/GB 90/01213

I. CLASSIFICATION OF SUBJECT MATTER (if several classification symbols apply, indicate all) [6]

According to International Patent Classification (IPC) or to both National Classification and IPC

IPC[5]: G 06 F 12/10, G 09 G 1/16

II. FIELDS SEARCHED

Minimum Documentation Searched [7]

Classification System	Classification Symbols
IPC[5]	G 06 F, G 09 G

Documentation Searched other than Minimum Documentation to the Extent that such Documents are included in the Fields Searched [8]

III. DOCUMENTS CONSIDERED TO BE RELEVANT [9]

Category [10]	Citation of Document, [11] with indication, where appropriate, of the relevant passages [12]	Relevant to Claim No. [13]
A	IEEE Micro, vol. 8, no. 2, April 1988, IEEE, (New York, US), Shinya Kimura et al.: "Implementation of the V60/70 and its FRM function", pages 22-36 see figure 2; page 25, left-hand column, lines 3-1. --	1
A	EP, A, 0254270 (HONEYWELL BULL) 27 January 1988 see figure 1; abstract --	1
A	EP, A, 0338416 (IBM) 25 October 1989 see figure 1; abstract	18,19

---

* Special categories of cited documents: [10]

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

IV. CERTIFICATION

Date of the Actual Completion of the International Search	Date of Mailing of this International Search Report
12th April 1991	27. 05. 91
International Searching Authority	**Signature of Authorized Officer**
EUROPEAN PATENT OFFICE	M. TORIBIO

Form PCT/ISA/210 (second sheet) (January 1985)

ANNEX TO THE INTERNATIONAL SEARCH REPORT ON INTERNATIONAL PATENT APPLICATION NO. GB 9001213 SA 39138

This annex lists the patent family members relating to the patent documents cited in the above-mentioned international search report. The members are as contained in the European Patent Office EDP file on 07/05/91
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A- 0254270	27-01-88	US-A- 4813002	14-03-89
		AU-B- 597697	07-06-90
		AU-A- 7598087	28-01-88
		JP-A- 63044257	25-02-88
EP-A- 0338416	25-10-89	JP-A- 2012523	17-01-90

EPO FORM P0479

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 5 : G11C 7/00 — **A1**

(11) International Publication Number: WO 94/07242

(43) International Publication Date: 31 March 1994 (31.03.94)

(21) International Application Number: PCT/US93/04231

(22) International Filing Date: 5 May 1993 (05.05.93)

(30) Priority data:
949,421 21 September 1992 (21.09.92) US

(71) Applicant: ATMEL CORPORATION [US/US]; 2125 O'Nel Drive, San Jose, CA 95131 (US).

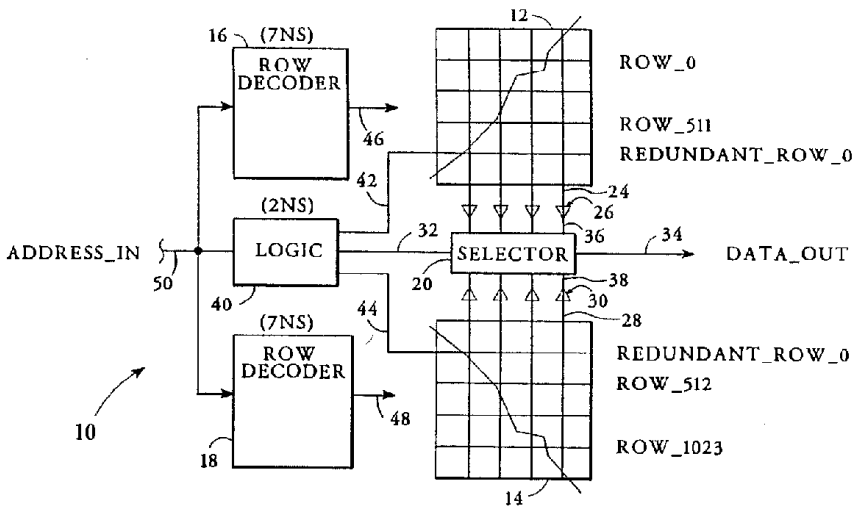
(72) Inventors: PATHAK, Saroj ; 12998 Robleda Road, Los Altos Hills, CA 94022 (US). ROSENDALE, Glen, A. ; 2250 Monroe Street, Apt. #268, Santa Clara, CA 95050 (US). PAYNE, James, E. ; 125 Brook Lane, Boulder Creek, CA 95005 (US).

(74) Agent: SCHNECK, Thomas; Schneck & McHugh, P.O. Box 2-E, San Jose, CA 95109-0005 (US).

(81) Designated States: DE, GB, JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published
With international search report.
With amended claims and statement.

(54) Title: HIGH SPEED REDUNDANT MEMORY



(57) Abstract

A memory circuit (10) in which redundant cell groups are located in a second memory bank (14; 88) to replace any defective cell groups present in a first memory bank (12; 86). Each bank has its own address decoders (16, 18; 90, 92, 94, 96) and read/write circuits. When an address corresponding to a defective cell group is received, the first bank is enabled as usual, while a fast logic circuit (40; 100) enables the second bank, so that both banks are simultaneously enabled. A signal (32; 102) is transmitted by the logic circuit to an output selector (20; 98) where data from a redundant cell group in the second bank is selected for output (34; 84). The time needed to detect an address corresponding to a defective cell group, is masked by the longer time required by the decoders and sense amplifiers of the respective banks.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	**FR**	France	**MR**	Mauritania
AU	Australia	**GA**	Gabon	**MW**	Malawi
BB	Barbados	**GB**	United Kingdom	**NE**	Niger
BE	Belgium	**GN**	Guinea	**NL**	Netherlands
BF	Burkina Faso	**GR**	Greece	**NO**	Norway
BG	Bulgaria	**HU**	Hungary	**NZ**	New Zealand
BJ	Benin	**IE**	Ireland	**PL**	Poland
BR	Brazil	**IT**	Italy	**PT**	Portugal
BY	Belarus	**JP**	Japan	**RO**	Romania
CA	Canada	**KP**	Democratic People's Republic of Korea	**RU**	Russian Federation
CF	Central African Republic			**SD**	Sudan
CG	Congo	**KR**	Republic of Korea	**SE**	Sweden
CH	Switzerland	**KZ**	Kazakhstan	**SI**	Slovenia
CI	Côte d'Ivoire	**LI**	Liechtenstein	**SK**	Slovak Republic
CM	Cameroon	**LK**	Sri Lanka	**SN**	Senegal
CN	China	**LU**	Luxembourg	**TD**	Chad
CS	Czechoslovakia	**LV**	Latvia	**TG**	Togo
CZ	Czech Republic	**MC**	Monaco	**UA**	Ukraine
DE	Germany	**MG**	Madagascar	**US**	United States of America
DK	Denmark	**ML**	Mali	**UZ**	Uzbekistan
ES	Spain	**MN**	Mongolia	**VN**	Viet Nam
FI	Finland				

Description
High Speed Redundant Memory

Technical Field

The invention relates to monolithic integrated memory circuits, particularly those adapted for reliable operation through the use of redundant memory cells.

Background Art

Many of today's large-scale ("LSI") and very-large-scale ("VLSI") integrated circuit memories would be rejected during manufacture if the presence of a few defective memory cells could not be dealt with in an effective manner. One common solution is to provide spare memory cells and some means of selecting these to replace the defective cells. U.S. Pat. No's 4,947,375 and 4,908,798 are typical examples. In such an arrangement, post-production wafer testing is used to identify the addresses of defective cells. At the factory, these addresses are programmed into a cell selection logic circuit, typically independent of the normal address decoders.

During normal operation this independent selection circuit compares a received memory address with the factory programmed addresses of the previously identified groups of defective cells. Typically, when a defective group of cells is addressed, signals are generated by the circuit to deselect the defective cells and to substitute a group of spare replacement cells. In this way a memory having such a redundant architecture can operate reliably despite the presence of a few defective memory cells.

The prior art typically places the spare memory cells in one of two locations, either in a separate region of the same array as the defective cells (U.S. Pat. No's 4,947,375, 5,058,069, 5,058,070 and 5,058,071 are examples of this approach), or in a separate array dedicated for that purpose (U.S. Pat. No's 4,908,798 and 5,033,024 are examples of this latter approach). These devices, while quite valuable in saving an otherwise defective chip, sacrifice speed

because the delay resulting from address decoding, redundant group selection, and defective normal group deselection, is added to other delays typically associated with the speed of memory circuits. The deselection time in particular results in a slowing of memory speed.

Two recent U.S. patents have attacked the speed problem in interesting ways. Kondo, in U.S. Pat. No. 4,918,662, teaches a redundant structure in which the spare memory cells are placed in the same array and use the same sense amplifiers as the normal memory cells. Kondo recognizes the slowing effect of the stray capacitance of an overly long normal cell group selection-inhibit-signal. Kondo partitions this length among many shorter inhibit signal paths, thereby reducing stray capacitance and reducing the "in line" delays. In a completely different approach, Busch et al., in U.S. Pat. No. 4,992,984, disclose a memory module made up of a plurality of identical memory "chips," each having several arrays of storage elements. Factory testing identifies arrays within each chip which are good and those arrays which contain defects. A memory address is applied to all chips in common. A control circuit inside each chip responds to previously programmed addresses corresponding to known defective arrays. The control circuit in the chip having the defective array and the control circuit in the chip having the replacement array both respond to an address corresponding to a known defective array. Both arrays are addressed, but the output data of the replacement array is enabled, while the output data of the defective array is disabled.

Summary of the Invention

The invention is a monolithic integrated circuit memory which contains at least two banks of normal data cells. These cells, which can be SRAMs, DRAMs, PROMs, EPROMs, EEPROMs or flash cells, are arranged in rows and columns, and a few of the normal data cells may be defective. Each bank has its own address decoders and the banks

can be simultaneously enabled. The integrated circuit memory has address lines for receiving a memory address and at least one data output line. The address space is partitioned among the memory banks so that normally only one bank responds to a given memory address.

Each bank has redundant rows or columns which contain one or more memory elements which are programmed with data in a defective element of another bank. In accord with the invention, an entire row or column of one bank is substituted for the row or column containing the defective element of another bank. The redundant row or column is read by the same sense amplifier which reads the rows or columns of the non-redundant elements in the memory bank.

A programmable logic device is programmed to receive all memory address signals and to respond to addresses with a bad memory element. Rather than deselect a bad row or column in the first memory bank, the logic device selects the redundant row or column in the second memory bank and then directs the sense amplifier output from the redundant row or column as the memory output. An output from the sense amplifier with the defective element is available, but is not selected. A selector circuit, operating as a multiplexer, is connected to both memory banks and responds to the logic device for selection of an output from an appropriate memory bank. In the case of an address for a non-defective cell in a first memory bank, the output from a sense amplifier associated with the first bank is selected. In the case of an address for a defective cell in the first memory bank, the output from a sense amplifier associated with the second bank is selected.

The logic circuit operates at a speed faster than memory bank decoders and sense amplifiers combined so that redundant row or column data is available faster than normal data. Since the defective row or column is not deselected, only ignored by the selector circuit, the delay normally present with deselection of a defective row or column is not present and memory output data is delayed by only an amount required by the logic circuit and the selector switch.

Since a switch can be very fast, the total delay of the logic device and the selector can be comparable, or faster, than the speed of a normal decoder. In this way, the use of a redundant memory element is transparent with respect to time delay and no time delay is experienced from deselection of a row or column containing a defective element.

The key ingredients of the present invention include placing a backup row or column for a defective memory bank address into a different memory bank, simultaneously selecting both the defective and backup row or column, and selecting the outputs of the memory bank containing the backup row or column when a defective element is addressed. A fast logic device and a fast selector switch is used, compared to the speed of an address decoder. In a read-only memory application, the correct data must exist in the backup row so that when both defective and redundant rows are selected, the correct data will be available at the sense amplifiers of the bank containing the redundant row. When writing is done, the data is written into the backup rows or columns, and later will be retrieved from the backup row or column.

Brief Description of the Drawings

Fig. 1 is a schematic block diagram showing the present invention in which redundant rows may be substituted.

Fig. 2 is pictorial version of Fig. 1 showing selection of a redundant row.

Fig. 3 is a schematic block diagram showing the present invention in which redundant row or columns may be substituted.

Best Mode for Carrying Out the Invention

With reference to Fig. 1 there is shown a high speed redundant memory 10 capable of reliable operation in the presence of a small number of defective memory cells.

The memory contains first and second memory banks, 12 and 14, each bank includes memory elements which are ar-

ranged in rows and columns. Each memory bank has its own row decoder 16, 18 and each column has its own sense amplifiers, 26, 30. Selector 20 is used to select the output of either first bank sense amplifiers along lines 36 or second bank sense amplifiers along lines 38 depending upon the level of a select signal 32. Selector 20 provides the selected data as memory output data on lines 34.

Each bank includes at least one redundant row of memory elements which share the sense amplifiers for the bank but are not selectable by the row decoder for the bank. Redundant_Row_0 is shown in the first bank 12 and is selectable by logic 40 using a signal carried on line 42. Redundant_Row_1 is shown in the second bank 14 and is selectable by a signal carried on line 44.

The normal rows of each memory bank are shown as Row_0 through Row_511 for first bank 12 and Row_512 through Row_1023 for second bank 14. The normal rows of first bank 12 are selectable by row decoder 16 output signals carried on lines 46, while the normal rows of second bank 14 are selectable by row decoder 18 output signals carried on lines 48. Row select lines 46, 48 are connected to the individual rows of the respective memory banks, generally indicated by the arrows.

Redundant memory 10 receives memory address signals on lines 50. Appropriate subsets of these memory signals are connected as inputs to row decoders 16, 18.

A factory programmable logic device 40 contains a map of bad cells derived from factory probe testing. The map lists the addresses of the bad cells and a corresponding replacement row. The logic circuit is made of non-volatile memory devices, such as PROMs, EPROMs, EEPROMs or flash cells. The speed of logic circuit 40 is about 2 nanoseconds, low compared to the speed of address decoder 16 or 18 which is about 7 nanoseconds. The address decoder 20 has a speed of several nanoseconds. The combined speed of the logic circuit is comparable or less than the speed of an address decoder and associated sense amplifier. When the logic circuit 40 receives memory address signals from lines

50 for bad cells it produces a predetermined signal carried on lines 32, 42 and 44.

In operation, a first predetermined memory address signal corresponding to the address of a defective row cell in second bank 14 is decoded by logic circuit 40 to generate a signal carried on line 42 to select Redundant_Row_0 in first bank 12. This occurs because logic circuit 40 has been programmed with the address of the bad element and responds to the bad address. The logic circuit 40 also generates a signal carried on line 32 to cause selector 20 to select the output of the sense amplifiers 26 of first bank 12 for presentation on memory data output lines 34.

Another operational example is as follows. A second predetermined memory address signal corresponding to the address of a defective row cell in first bank 12 is decoded by logic circuit 40 to generate a signal carried on line 44 to select Redundant_Row_1 in second bank 14. The same predetermined memory signal is decoded by logic circuit 40 to generate a signal carried on line 32 to cause selector 20 to select the output of the sense amplifiers 30 of second bank 14 for presentation on memory data output lines 34.

Fig. 2 depicts the embodiment shown in Fig. 1 in a pictorial fashion to emphasize the basic operation of the invention. In this example, a memory address is received on lines 50 which corresponds to a cell in Defective_Row_K in first bank 12. Defective_Row_K is shown being selected by an output of row decoder 16. Programmable logic circuit 40 receives and responds to the memory address by generating a row selection signal on line 44 connected to and selecting Redundant_Row_1 in second bank 14 at the same time Defective_Row_K is being selected. The contents of each selected row is available at the outputs of the respective sense amplifiers, 26, 30. The outputs of second bank sense amplifiers 30 are selected by selector 20 for output on lines 34 as the memory output data. Programmable logic circuit 40 generates an output select signal on line 32 for this purpose. The row decoder 18 for second bank 14 has no active output signal because the normal rows in second bank 14 do not re-

spond to addresses found the address space of first bank 12.

Fig. 3 depicts a high speed redundant memory 80 capable of reliable operation in the presence of a few bad cells in defective rows or columns. The memory has first and second memory banks, 86 and 88 respectively. Each bank may have bad cells which define a defective row or column and includes at least one redundant row, depicted as Redundant_Row_0 and Redundant_Row_1, and one redundant column, 116 and 118. Each bank has its own row decoder, 90 and 94, and its own column decoder, 92 and 96. A factory programmable logic circuit 100 receives a memory address on lines 82 and generates select signals 102, 104 and 106 in response to programmed addresses. All columns in each bank share a common sense amplifier, 108 and 110. Select signal 102 is used by selector 98 to select the output of one sense amplifier as the single output bit Data_Out on line 84. Select signal 104 is made active in response to the programmed address of a known defective row or column and is applied through programmable switch 112 to select either Redundant_Row_0 or redundant column 116. Similarly, select signal 106 may be programmably connected via switch 114 to select Redundant_Row_1 or redundant column 118. The normal rows and columns or each bank partition the memory address space such that only a single row and column in one bank will be selected at one time.

When the address of a known defective cell group is received on lines 82, both memory banks are enabled and the defective row or column is selected by its address decoders producing a defective output at the bank sense amplifier. Simultaneously, select signal 104 or 106 selects a redundant row or column and a good output is available at the sense amplifier of the bank containing the redundant cell group. Select signal 102 is made active in response to the memory address to select the good sense amplifier output and connect the good output to line 84.

Without loss of generality, the structures depicted in Figs. 1, 2 and 3 can be extended, by means well known to those having ordinary skill in the art, to include

a redundant memory having both read and write capability.

Other embodiments are contemplated. The backup for bank A can be located in bank B, the backup for B located in C, for C in D, and for D in A, so long as both the bank of the defective row and its backup bank are simultaneously enabled and the output of the backup bank selected. Each memory bank may be a portion of a chip, a chip or a laser module, such as a SIMM or SIP.

Claims

1. A high-speed redundant memory, comprising:

a memory having first and second banks of storage cells, each bank having its own address decoders and at least one sense amplifier, each bank having redundant memory cell groups which are not selectable by the address decoders;

logic circuit means for receiving and decoding a memory address which corresponds with a defective cell in one bank and for generating a logic signal in response thereto having an address which corresponds to a redundant cell group in the other bank;

switch means responsive to the logic signal to enable said other bank, said switch means connected to the first and second banks for using the logic signal to select an output signal of a sense amplifier of the other bank as the memory output signal.

2. The redundant memory of claim 1 wherein the redundant memory cell group comprises at least one row.

3. The redundant memory of claim 1 wherein the redundant memory cell group comprises at least one column.

4. The redundant memory of claim 1 wherein the redundant cell group has a redundant sense amplifier.

5. The redundant memory of claim 1 wherein the logic circuit means comprises non-volatile memory cells.

6. The redundant memory of claim 1 wherein the redundant memory cell groups include at least one redundant row and at least one redundant column.

Micron Technology Inc. et al.
Ex. 1010, p. 461

7. The redundant memory of claim 1 wherein said storage cells are selected from the group of SRAM, DRAM, PROM, EPROM, EEPROM or flash memory cells.

8. The redundant memory of claim 1 comprising pairs of memory banks such that the redundant storage cell groups of each bank of a pair serve as substitutes for defective cells in the other bank of the pair, and the decoding of a memory address corresponding to a defective cell group in either bank of the pair results in at least both banks of the pair being simultaneously enabled and wherein the output of the at least one sense amplifier of the bank having the selected replacement group is output by the memory as output data.

9. A high-speed redundant memory for operation in the presence of defective memory elements, comprising:

a plurality of memory banks each with an array of memory elements with regular addresses and each array having its own address decoding means and having redundant elements with addresses whose contents may be read via each bank's associated sense amplifiers,

logic means for recognizing the address of a defective memory element with a regular address in a first memory bank and for selecting a redundant element with an address in a second memory bank, and selecting outputs of the sense amplifiers of the second memory bank when the address of a defective memory element is recognized; and

a selector means responsive to the logic means for outputting data from a redundant memory element in the second memory bank when an address of a defective memory element in the first memory bank is received and otherwise outputting data from a memory element with an address in the first memory bank, said first and second memory banks simultaneously enabled when an address of a defective memory element is received, the combined speed of the logic means being faster than said decoder means and associated sense amplifiers.

10. A method for addressing a redundant memory, comprising the steps of:

providing a memory having first and second banks, each having its own address decoders and sense amplifiers and having redundant memory cells which are not selectable by the address decoders;

receiving and decoding a memory address which corresponds to a defective cell group in a first bank and generating a signal in response thereto;

while the first bank is enabled to read the defective cell group, using the signal (i) to select a redundant cell group in the second bank, (ii) to simultaneously enable the second bank, and (iii) to select an output signal of the sense amplifiers of the second bank; and

outputting the selected output signal of the second bank as memory output data.

AMENDED CLAIMS
[received by the International Bureau on 30 November 1993 (30.11.93) ;
original claims 1, 9 and 10 amended ; new claim 11 added ;
other claims unchanged (3 pages)]

1. A high-speed redundant memory, comprising:

a memory having first and second banks of storage cells, each bank having an address decoder and at least one sense amplifier, each bank having redundant memory cell groups which are not selectable by the address decoder of a respective bank;

logic circuit means for receiving and decoding a memory address which corresponds with a defective cell in one bank and for generating a logic signal in response thereto having an address which corresponds to a redundant cell group in the other bank, the logic circuit means operating at a speed faster than the address decoder and sense amplifier combined of either of said first and second banks,

switch means connected to the first and second banks and to the logic circuit means for using the logic signal of the logic circuit means to select an output signal of a sense amplifier of the other bank as the memory output signal.

2. The redundant memory of claim 1 wherein the redundant memory cell group comprises at least one row.

3. The redundant memory of claim 1 wherein the redundant memory cell group comprises at least one column.

4. The redundant memory of claim 1 wherein the redundant cell group has a redundant sense amplifier.

5. The redundant memory of claim 1 wherein the logic circuit means comprises non-volatile memory cells.

6. The redundant memory of claim 1 wherein the redundant memory cell groups include at least one redundant row and at least one redundant column.

Micron Technology Inc. et al.
Ex. 1010, p. 464

7. The redundant memory of claim 1 wherein said storage cells are selected from the group of SRAM, DRAM, PROM, EPROM, EEPROM or flash memory cells.

8. The redundant memory of claim 1 comprising pairs of memory banks such that the redundant storage cell groups of each bank of a pair serve as substitutes for defective cells in the other bank of the pair, and the decoding of a memory address corresponding to a defective cell group in either bank of the pair results in at least both banks of the pair being simultaneously enabled and wherein the output of the at least one sense amplifier of the bank having the selected replacement group is output by the memory as output data.

9. A high-speed redundant memory for operation in the presence of defective memory elements, comprising:

a plurality of memory banks each with an array of memory elements with regular addresses and each array having an address decoding means and having redundant elements with addresses whose contents may be read via sense amplifiers associated with each bank,

logic means for recognizing the address of a defective memory element with a regular address in a first memory bank and for selecting a redundant element with an address in a second memory bank, and

a selector means responsive to the logic means for outputting data from a redundant memory element in the second memory bank when an address of a defective memory element in the first memory bank is received and otherwise outputting data from a memory element with an address in the first memory bank, said first and second memory banks simultaneously enabled when an address of a defective memory element is received, the speed of the logic means being faster than the combined speed of said address decoding means and associated sense amplifiers of any memory bank.

10. A method for addressing a redundant memory, comprising the steps of:

addressing a memory having first and second banks, each having an address decoder and sense amplifiers and having redundant memory cells which are not selectable by the address decoders of a respective bank,

addressing a logic circuit with a memory address which corresponds to a defective cell group in a first bank and generating a signal from the logic circuit in response thereto in a second amount of time less than the first amount of time;

while the first bank is enabled to read the defective cell group, using the signal from the logic circuit (i) to select a redundant cell group in the second bank, (ii) to simultaneously enable the second bank and (iii) to select an output signal of the sense amplifiers of the second bank, and

outputting the selected output signal of the second bank as memory output data.

11. The memory of claim 1 wherein said logic circuit means is a programmable logic device.

Micron Technology Inc. et al.
Ex. 1010, p. 466

STATEMENT UNDER ARTICLE 19

In response to the International Search Report, claims 1, 9 and 10 were rewritten to point out that the logic means operates faster than the combined speed of the decoder and sense amplifiers, thereby distinguishing applicant's invention from the prior art cited in the International Search Report. In addition, a new claim 11 has been added to point out that the applicant achieves increased speed with a programmable logic device.

Applicant predicates patentability on the basis that the redundant memory element is transparent with respect to time delay in that no time delay is experienced from deselection of a row or column containing a defective element. Applicant achieves this result by having the logic circuit operate at a speed faster than both the memory decoders and sense amplifiers combined. This configuration makes available redundant row or column data faster than normal data.

Micron Technology Inc. et al.
Ex. 1010, p. 467

SUBSTITUTE SHEET

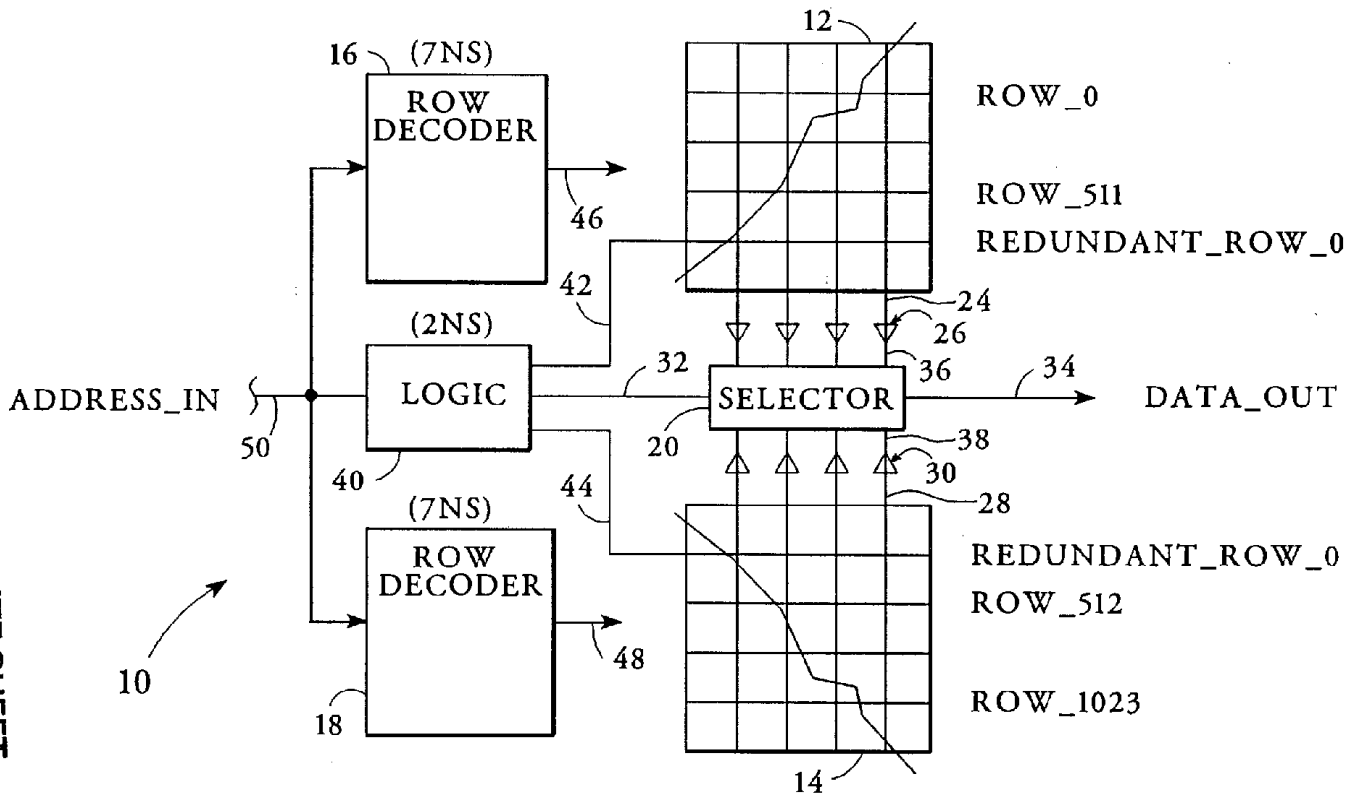


FIG. 1

Micron Technology Inc. et al.
Ex. 1010, p. 468

SUBSTITUTE SHEET

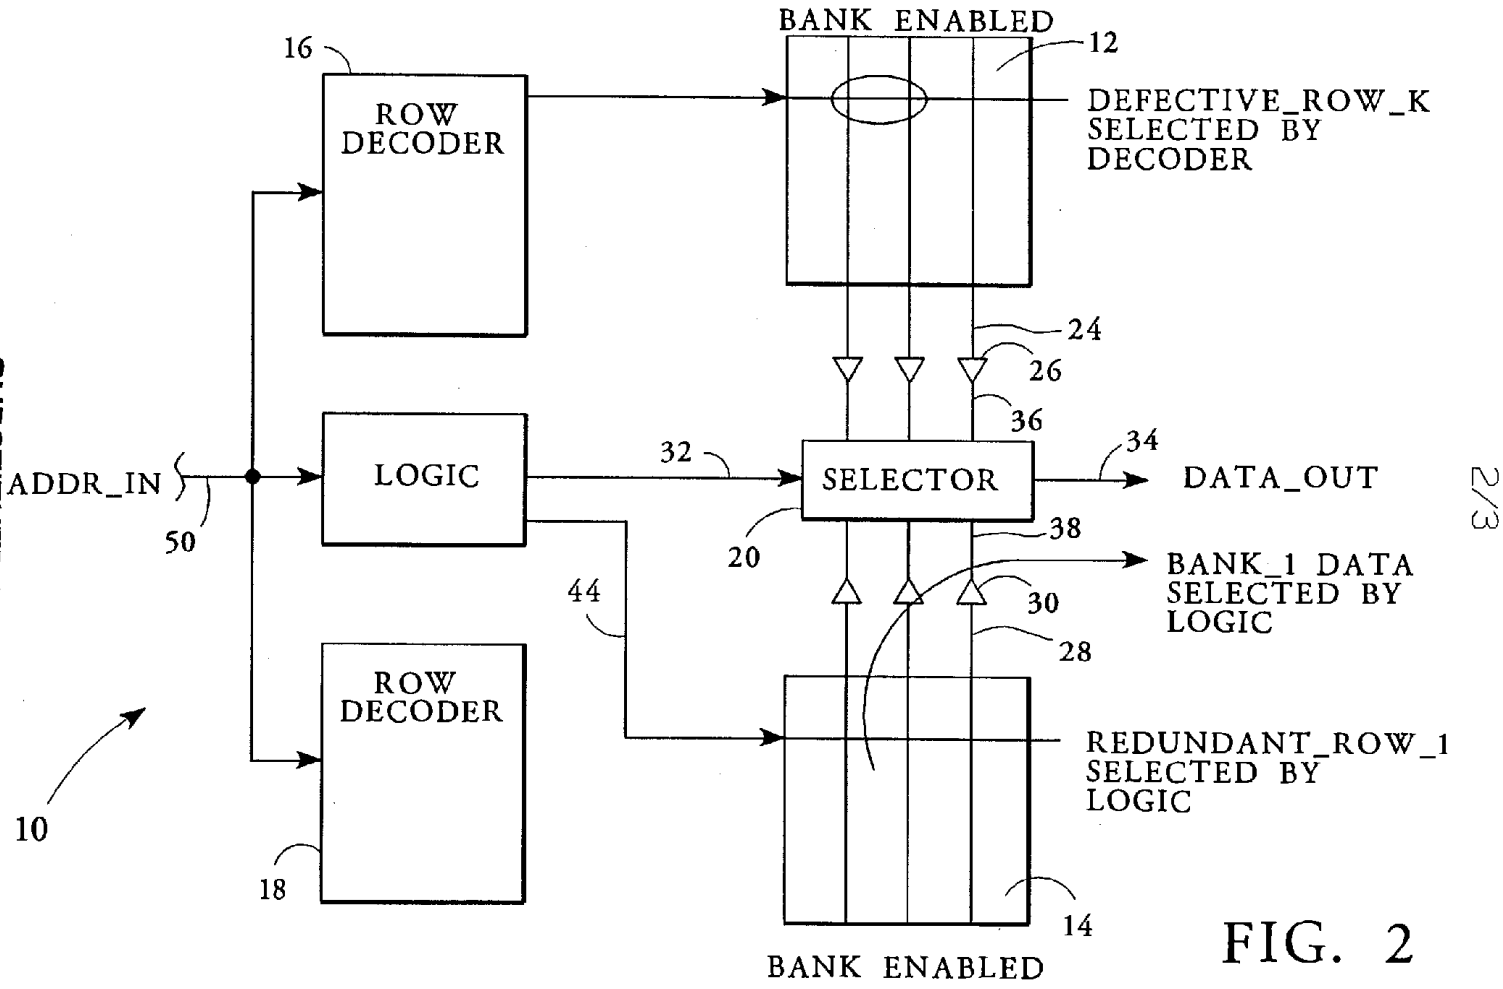


FIG. 2

Micron Technology Inc. et al.
Ex. 1010, p. 469

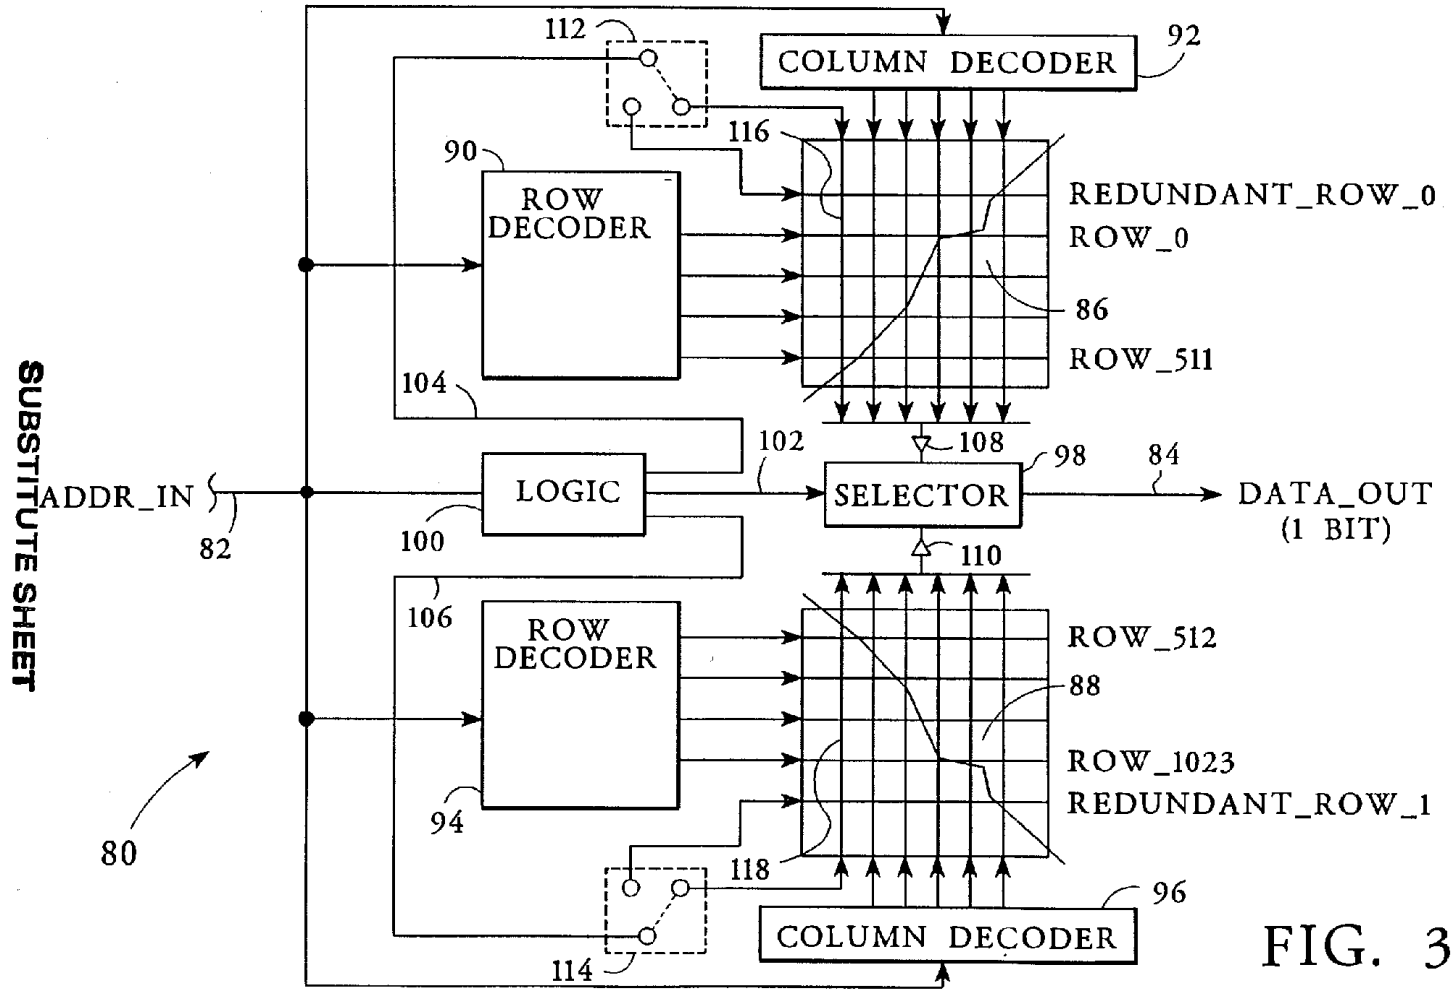


FIG. 3

Micron Technology Inc. et al.
Ex. 1010, p. 470

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US93/04231

A. CLASSIFICATION OF SUBJECT MATTER

IPC(5) :G11C 7/00
US CL :365/200, 201, 230.03; 371/10.2, 10.3
According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : NONE

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS (MEMORY, STORAGE, CELL, ARRAY, MATRIX, BLOCK, REDUNDANT, SPARE, ADDRESS DECODER, SELECTOR, SENSE AMPLIFIER, NONVOLATILE)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	IBM J. RES. DEVELOP., VOL. 24, NO. 3, MAY 1980, FITZGERALD ET AL., "CIRCUIT IMPLEMENTATION OF FUSIBLE REDUNDANT ADDRESSES ON RAMs FOR PRODUCTIVITY ENHANCEMENT", PP. 291-298	1-7, 10 ------- 8, 9
Y	US, A, 5,126,973 (GALLIA ET AL.) 30 JUN 1992. see entire document.	8, 9

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

*	Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A"	document defining the general state of the art which is not considered to be part of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E"	earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&"	document member of the same patent family
"O"	document referring to an oral disclosure, use, exhibition or other means		
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search 22 July 1993	Date of mailing of the international search report 28 JUL 1993
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. NOT APPLICABLE	Authorized officer ANDREW Q. TRAN Telephone No. (703) 305-3495

Form PCT/ISA/210 (second sheet)(July 1992)*

PCT WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06F

A2

(11) International Publication Number: **WO 95/34030**

(43) International Publication Date: 14 December 1995 (14.12.95)

(21) International Application Number: PCT/US95/07062

(22) International Filing Date: 1 June 1995 (01.06.95)

(30) Priority Data:
08/253,499 3 June 1994 (03.06.94) US

(71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).

(72) Inventors: MILLS, Duane, R.; 103 Casselman Street, Folsom, CA 95630 (US). DIPERT, Brian, Lyn; 9840 Burline Street, Sacramento, CA 95827 (US). SAMBANDAN, Sachidanandan; 101 Waterboro Square, Folsom, CA 95630 (US). MCCORMICK, Bruce; 8012 Chris Court, Granite Bay, CA 95746 (US). PASHLEY, Richard, D.; 77710 Haley Drive, Roseville, CA 95661 (US).

(74) Agents: TAYLOR, Edwin, H. et al.; Blakely, Sokoloff, Taylor & Zafman, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025-1026 (US).

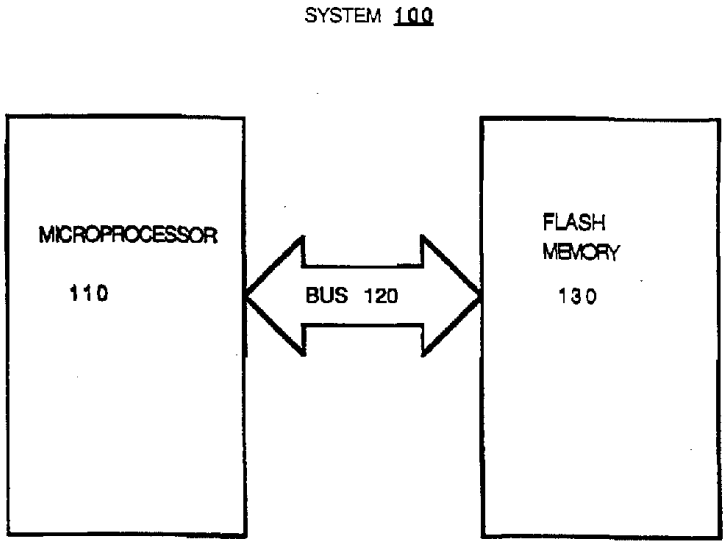
(81) Designated States: AM, AT, AT (Utility model), AU, BB, BG, BR, BY, CA, CH, CN, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, ES, FI, FI (Utility model), GB, GE, HU, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), TJ, TM, TT, UA, UG, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG), ARIPO patent (KE, MW, SD, SZ, UG).

Published
Without international search report and to be republished upon receipt of that report.

(54) Title: FLASH MEMORY BASED MAIN MEMORY

(57) Abstract

A flash memory chip that can be switched into four different read modes is described. Computer systems and hierarchies that exploit these modes are also described. In the first read mode, 'asynchronous flash mode, the flash memory is read as a standard flash memory. In this mode, the reading of the contents of a first address must be completed before a second address to be read can be specified. In the second read mode, synchronous flash mode, a clock signal is provided to the flash chip and a series of addresses belonging to a data burst are specified, one address per clock tick. Then, the contents stored at the addresses specified for the burst are output sequentially during subsequent clock ticks in the order in which the addresses were provided. Alternately, if a single address is provided to the flash chip when it is in the synchronous mode, the subsequent addresses for the burst will be generated within the flash chip and the data burst will then be provided as output from the flash chip. In the third read mode, asynchronous DRAM (dynamic random access memory) mode, the flash memory emulates DRAM. Thus, row and column addresses are strobed into the flash memory using row and column address strobe signals. The flash memory then converts the row and column addresses internally into a single address and provides as output the data stored at that single address. Furthermore, although the flash memory does not need an extended precharge period or to be refreshed, when in the asynchronous DRAM mode, the flash memory responds to precharge periods and refresh cycles as would a DRAM. Therefore, when in the asynchronous DRAM mode, the flash memory can be controlled by a standard DRAM controller. In the fourth read mode, synchronous DRAM mode, the features of the second and third modes are combined to yield a flash memory that emulates a synchronous DRAM.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	**GB**	United Kingdom	**MR**	Mauritania
AU	Australia	**GE**	Georgia	**MW**	Malawi
BB	Barbados	**GN**	Guinea	**NE**	Niger
BE	Belgium	**GR**	Greece	**NL**	Netherlands
BF	Burkina Faso	**HU**	Hungary	**NO**	Norway
BG	Bulgaria	**IE**	Ireland	**NZ**	New Zealand
BJ	Benin	**IT**	Italy	**PL**	Poland
BR	Brazil	**JP**	Japan	**PT**	Portugal
BY	Belarus	**KE**	Kenya	**RO**	Romania
CA	Canada	**KG**	Kyrgystan	**RU**	Russian Federation
CF	Central African Republic	**KP**	Democratic People's Republic of Korea	**SD**	Sudan
CG	Congo			**SE**	Sweden
CH	Switzerland	**KR**	Republic of Korea	**SI**	Slovenia
CI	Côte d'Ivoire	**KZ**	Kazakhstan	**SK**	Slovakia
CM	Cameroon	**LI**	Liechtenstein	**SN**	Senegal
CN	China	**LK**	Sri Lanka	**TD**	Chad
CS	Czechoslovakia	**LU**	Luxembourg	**TG**	Togo
CZ	Czech Republic	**LV**	Latvia	**TJ**	Tajikistan
DE	Germany	**MC**	Monaco	**TT**	Trinidad and Tobago
DK	Denmark	**MD**	Republic of Moldova	**UA**	Ukraine
ES	Spain	**MG**	Madagascar	**US**	United States of America
FI	Finland	**ML**	Mali	**UZ**	Uzbekistan
FR	France	**MN**	Mongolia	**VN**	Viet Nam
GA	Gabon				

FLASH MEMORY BASED MAIN MEMORY

FIELD OF THE INVENTION

The present invention pertains to the field of the architecture of computer systems. More particularly, the present invention relates to computer systems that use a large-block erasable non-volatile semiconductor memory as main memory.

BACKGROUND OF THE INVENTION

As modern computer programs have become increasingly more sophisticated, modern personal computer systems have also had to become more sophisticated in order to accommodate these computer programs. Computer programs are made up of a larger number of code instructions than they once were and on average, require access to larger files of data that are read from, and written to, when executing the programs.

Typically, the heart of a personal computer system is a central processing unit (CPU) that resides on a microprocessor chip. New microprocessor chips that operate at increasingly high operating speeds are constantly being developed in order to permit personal computers to execute the larger programs in a timely manner. Usually, these microprocessor chips are developed using CMOS (complementary metal-oxide semiconductor) technology. The greatest amount of power consumption for CMOS chips occurs on the leading and trailing edges of clock pulses (i.e. when a clock signal transitions from a low voltage state to a higher voltage state and vice versa).

When the operating speed of the microprocessor is increased, the number of clock pulses in a particular time period increases thereby increasing the power consumption of the microprocessor during this time period. Furthermore, more heat is generated by the microprocessor and must be dissipated in order to prevent the damage of components within the computer system.

Both power consumption and heat dissipation pose serious problems when designing a personal computer system. This is

especially true in the case of mobile computers that are typically powered by batteries. The more power that the computer consumes, the less time that the computer can operate off of a given sized battery. Therefore, as the operating speed of the computer is increased, a designer is faced with several unattractive alternatives.

If the same sized batteries are used, then the effective operating time for the computer system must decrease when the operating speed is increased. On the other hand, if the effective operating time is to remain constant then it is necessary to either add additional batteries, thereby increasing the bulk and weight of the computer, or to use an exotic and therefore expensive battery technology (or both).

The trend in mobile computers is towards smaller, faster, less expensive and lighter units. Thus, the need to add additional batteries, or more expensive batteries is a significant disadvantage. This disadvantage is exacerbated by the need to add cooling fans, or to implement other cooling techniques, in order to dissipate the additional heat that is generated by the high speed microprocessors.

Additionally, because the microprocessors are operating at a higher speed, they can execute more instructions in a given amount of time, and therefore can also process a greater amount of data during that period. A bottle neck has developed in computer systems having fast microprocessors that can prevent the higher speed of the microprocessor to be utilized effectively. This bottle neck is the bus (or buses) that provide instructions for the microprocessor to execute and the data that the microprocessor will use when executing the instructions.

If the next instruction to be executed is not available when the microprocessor needs it, then the microprocessor must wait idly (i.e. insert wait cycles) while the required instruction is retrieved and provided to the microprocessor.
Furthermore, if the next instruction to be executed requires data that is not immediately available to the microprocessor, the microprocessor must also idle until the data has been retrieved. During this idle time, the microprocessor clock continues to toggle

thereby needlessly consuming power and generating heat that must be dissipated.

In order to decrease the frequency with which the microprocessor encounters these wait cycles, many modern high performance microprocessors have a small internal cache, called a primary cache. Instructions that are likely to be executed and data that is likely to be needed by the executing instructions are stored in the internal cache so that they may be accessed immediately by the CPU of the microprocessor.

The sequential nature of computer programs is such that when a particular instruction within the program is executed, it is highly probable that the next instruction to be executed will be the instruction that follows the currently executing instruction. Therefore, when an instruction is to be executed, the cache is checked to determine whether a copy of the required instruction is immediately available within the cache. If a copy of the required instruction is stored within the cache (called a cache hit), then the copy of the instruction can be supplied to the CPU immediately from the cache and there is no need for the CPU to wait while the instruction is retrieved to the microprocessor chip from wherever it is stored in the computer system.

On the other hand, if a copy of the required instruction is not stored within the cache (called a cache miss), then the CPU must wait while the instruction is retrieved to the microprocessor chip from wherever it is stored within the computer system. Actually, rather than only retrieving the next instruction to be executed, a cache line is formed by retrieving the next instruction to be executed and a certain number of instructions following the next instruction to be executed. That way, if the subsequent instructions are in fact required to be executed, they will be immediately available to the CPU from within the cache line of the cache. Because of the sequential nature of programs, the benefits of caching also applies to data used by the programs.

Because the internal cache is filled a cache line at a time, many microprocessors can accept data in a burst mode. In a typical burst read, the microprocessor specifies the first address of the data or instructions to be read into a cache line. Then, the data or instructions that are stored at the addresses of the cache line are sent sequentially from where they are stored within the computer system to the microprocessor.

Frequently the internal cache of the microprocessor is formed using static random access memory (SRAM). Because each SRAM cell is formed by six to eight transistors, there is only room on a microprocessor chip for a relatively small SRAM cache. Furthermore, SRAM is volatile meaning that SRAM retains the information stored as long as there is enough power to run the device. If power is removed, the contents of the SRAM cache are lost.

Some microprocessors are dynamic, meaning that if power is removed from them, when power is restored they cannot return directly to the state they were in when the power was removed. When power is restored the microprocessor must be reinitialized, and at least some of the processing progress previously made will probably be lost.

Other microprocessors are static, meaning that they can be placed in an energy saving deep powerdown mode, and then be returned relatively quickly to the state they were in immediately before they entered the deep powerdown mode.

As mentioned earlier, data and instructions are stored within the computer system and provided to the microprocessor over one (or more) bus systems. Because most types of relatively fast random access memory are both volatile and relatively expensive, a typical computer system stores code and data on relatively inexpensive, nonvolatile memory store such as a floppy disk or hard disk.

The typical computer system also has a main memory made of volatile memory because the nonvolatile memory has a relatively slow access speed. When a program is to be executed, the computer system uses a technique known as shadowing to copy the code and

data required to execute the program from the slow nonvolatile memory to the faster volatile memory. The shadow copy in the main memory is then used to execute the program. If any changes are made to the shadow copy during the course of the program execution, the shadow copy can be copied back to the slower nonvolatile memory, when the program finishes execution. Furthermore, because an unexpected power failure will cause the contents of the volatile main memory to be lost, it is common to save intermediate results generated during the course of execution of the program.

The most common form of main memory is dynamic random access memory (DRAM). DRAM is more commonly used than SRAM, even though it is slower than SRAM because DRAM can hold approximately four times as much data as a SRAM of the same complexity.

DRAMs store information in integrated circuits that contain capacitors. Because capacitors lose their charge over time, DRAMs must be controlled by logic that causes the DRAM chips to continuously "refresh" (recharge). When a DRAM is being refreshed, it cannot be read from, or written to, by the microprocessor. Thus, if the microprocessor must access the DRAM while it is being refreshed, one or more wait states occur.

In some computer systems, SRAM is used as main memory in place of DRAM. One advantage of using SRAM as main memory is that SRAM is relatively faster to access than DRAM. Furthermore, because SRAM does not need to be refreshed, it is always available for access by the microprocessor, thereby eliminating the DRAM associated need for the microprocessor to include wait states when accesses are attempted while the DRAM is being refreshed. Moreover, the lack of a refresh requirement simplifies designing a computer system having SRAM based main memory because one does not have to worry about controlling refresh cycles. In fact, a simple battery back-up can be supplied to preserve the contents of the SRAM in the event of a power failure. Of course, if the battery back-up fails, the contents of the SRAM main memory will be lost.

Rather than building a main memory completely from SRAM, it is more common to implement the main memory using DRAM, and then to supplement the DRAM based main memory with a SRAM based external cache memory (i.e. a cache memory that is external to the microprocessor chip). Because the external cache is not contained on the microprocessor chip, it can typically be made to store more data and instructions than can be stored by the internal cache. Because the external cache is not located on the microprocessor chip, however, it must supply the data and instructions to the microprocessor using one of the buses that often form bottlenecks for data and instructions entering and leaving the microprocessor chip.

A high speed microprocessor chip typically interfaces with the rest of the computer system using one or two high speed buses. The first of these buses is a relatively high speed asynchronous bus called a main memory bus. The second of these buses is a relatively high speed synchronous bus called a local bus. The typical operating speed of main memory and local buses is in the range of 16 to 33 MHz and the trend is towards increasingly faster buses.

Although most microprocessors can interface directly with a main memory bus, some microprocessors do not provide an external interface to a local bus. These microprocessors typically interface with a relatively slow speed synchronous bus called an expansion bus. The typical operating speed of an expansion bus is in the range of 8 to 12 MHz.

The main memory (or DRAM) bus is used by the microprocessor chip to access main memory. Usually, rather than interfacing directly to the DRAM chips, the microprocessor is coupled to a DRAM controller chip that, in turn, is coupled to the DRAM chip or chips. The DRAM controller controls accesses to the DRAM chips initiated by the microprocessor. The DRAM controller also controls overhead maintenance such as the refresh cycles for periodically refreshing the DRAM contents. Some microprocessors have the DRAM controller built directly into them. Frequently, the DRAM or SRAM chips are contained in surface-mount packages and several

DRAMs or SRAMs are attached to a small circuit board to form what is called a Single In-line Memory Module (SIMM). One can then relatively easily modify the total amount (or the access speed) of main memory in a computer system by simply swapping one type of SIMM for another. A SRAM based external cache may also be coupled to the microprocessor through the DRAM bus.

If a computer system has a local bus, then the microprocessor can access devices coupled to the local bus at a relatively fast speed. Thus, high bandwidth devices such as graphics adapter cards and fast input/output devices are typically coupled directly to the local bus. Sometimes the external cache is coupled to the local bus rather than to the DRAM bus. It is also possible to supplement (or replace) the main memory on the main memory bus by coupling DRAM to the local bus using a DRAM controller designed to interface with the local bus.

Each device coupled to the local bus has an associated capacitive load. As the load on the local bus is increased, the maximum operating speed for the local bus decreases and the power required to drive the bus increases. Therefore, one device coupled to the local bus can be a peripheral bus bridge from the local bus to another bus called a high speed peripheral bus (e.g. a peripheral component interconnect (PCI) bus). The bus bridge isolates the load of the devices coupled to the high speed peripheral bus from the high speed local bus.

Another device coupled to the local bus is typically an expansion bus bridge that couples the high performance local bus to a lower performance expansion bus. The low bandwidth components of the computer system are then coupled to the lower performance expansion bus. One type of device that is typically coupled to the expansion bus uses flash memory. Flash memory typically is a high-density, nonvolatile, read-write memory. Examples of flash memory based devices include BIOS ROM and hard disk substitutes.

Flash memories differ from conventional EEPROMs (electrically erasable programmable read only memories) with respect to erasure. Conventional EEPROMs use a select transistor for

individual byte erase control. Flash memories, on the other hand, achieve much higher density with single transistor cells.
For a typical flash memory array, a logical "one" means that few if any electrons are stored on a floating gate associated with a bit cell. A logical "zero" means that many electrons are stored on the floating gate associated with the bit cell. Each bit of the flash memory array cannot be overwritten from a logical zero state to a logical one state without a prior erasure. During a flash erase operation, a high voltage is supplied to the sources of every memory cell in a block or in the entire chip simultaneously. This results in a full array or a full block erasure.

After a flash memory array has been erased, a logical one is stored in each bit cell of the flash memory array. Each single bit cell of the flash memory array can then be programmed (overwritten) from a logical one to a logical zero, given that this entails simply adding electrons to a floating gate that contains the intrinsic number of electrons associated with the erased state. Program operations for flash memories are also referred to as write operations.

The read operation associated with a typical flash memory array closely resembles the read operation associated with other read-only memory devices. A read operation for a typical high speed flash memory array takes on the order of 80 nanoseconds (nS). Write and erase operations for a flash memory array are, however, significantly slower. Typically, an erase operation takes on the order of one second. A write operation for a single word of a flash memory array takes on the order of 10 microseconds.

British patent document no. GB 2 251 324 A, published July 1, 1992, describes a computer system that uses flash memory. The patent document discloses various architectures to incorporate a flash memory into a computer system. One architecture referred to therein is a variable file structure. For the variable file structure, computer code is stored contiguously in flash memory, allowing a CPU to execute computer code directly from the flash memory array without the need for RAM. A direct mapped variable file structure is described

that allows direct code execution from all of the flash memory array. A page mapped variable file structure is also described that allows direct code execution from a portion of the flash memory array. Thus, flash memory can serve as the main memory within portable computers, providing user functions similar to those of disk-based systems.

A ROM-executable DOS is available commercially and provides several benefits to both system manufacturers and ultimately end users. First, because most of the operating system is composed of fixed code, the amount of system RAM required to execute DOS is reduced from 50K to 15K, thereby conserving system space and power. Secondly, DOS can now be permanently stored in, and executed from, a single ROM-type of device such as flash memory. This enables systems to be provided that are ready to run right out of the box. Lastly, users enjoy "instant on" performance because the traditional disk-to-DRAM boot function and software downloading steps are eliminated.

For example, by storing application software and operating system code in a Resident Flash Array (RFA), users enjoy virtually instant-on performance and in-place code execution. An RFA also protects against software obsolescence because, unlike ROM, it is in-system updatable. Resident software, stored in flash rather than disk, extends battery life and increases system reliability.

Because erasing and writing data to flash memory is a distinctly different operation than rewriting information to a disk, new software techniques have been developed to allow flash to emulate disk functionality. File management software such as Microsoft's Flash File System (FFS) allows Flash Memory components and flash cards to emulate the file storage capabilities of disk. Microsoft's FFS transparently handles data swaps between flash blocks similar to the way MS-DOS handles swaps between disk sectors. Under FFS, the user can input a MS-DOS or Windows command without regard for whether a flash memory or magnetic disk is installed in the system. Flash filing systems make the management of flash memory devices

completely transparent to the user. Flash filing systems similar to the Microsoft FFS are available or are being developed for other operating systems besides DOS and Windows.

Flash Memory is exceptionally well-suited to serve as a solid-state disk or a cost-effective and highly reliable replacement for DRAMs and battery-backed static RAMs. Its inherent advantages over these technologies make it particularly useful in portable systems that require the utmost in low power, compact size, and ruggedness while maintaining high performance and full functionality.

Flash memory, however, typically has an asynchronous interface wherein an address to be read is specified and then, a set time later, the contents stored at the specified address are output from the flash chip. It is only after the data has been output from the flash chip that the next address to be read can be sent to the flash chip. A high speed bus like the local bus can run at 33 MHz wherein every cycle of the bus takes about 30 nS. A typical high performance flash chip, on the other hand, has a read access time of about 80 nS. Hence, if flash is to be used as main memory, every single memory access to flash involves wait states and zero wait state back to back burst read cycles from flash cannot be supported. This is true for other devices having a read latency similar to that of flash memory. Thus, using prior art technology, it is not practical to use these memories as main memory for a high speed microprocessor.

Micron Technology Inc. et al.
Ex. 1010, p. 483

SUMMARY AND OBJECTS OF THE INVENTION

Therefore, one object of the present invention is to provide an efficient memory hierarchy based on non-volatile memory versus volatile memory wherein both data and applications are stored in random access nonvolatile memory and further wherein applications are executed directly from the random access non-volatile memory.

It is also an object of this invention is to enable flash memory to operate in an optimal synchronous fashion with any synchronous bus.

It is also an object of this invention is to enable flash memory to operate in an optimal synchronous fashion with any synchronous bus to provide a low cost, low power alternative to volatile main memory, and to eliminate the time required to transfer code and data from the hard disk to the main memory.

It is also an object of this invention to enable flash memory to operate in an optimal synchronous fashion with any synchronous bus so that the CPU can execute programs directly out of the flash memory without any degradation in performance when compared to volatile memory based main memory.

It is also an object of this invention to enable flash memory to operate in an optimal synchronous fashion with any synchronous bus and to thereby eliminate the need to incorporate costly memory subsystem designs such as interleaving into the system.

It is also an object of this invention to enable flash memory to operate in an optimal synchronous fashion with any synchronous bus and to thereby support back to back burst cycles and thus ensure that cache line fills are performed in a quick and optimal fashion.

It is also an object of this invention to enable flash memory to operate in an optimal asynchronous fashion with any asynchronous main memory bus.

It is also an object of this invention to enable flash memory to operate in an optimal asynchronous fashion with any asynchronous main memory bus to provide a low cost, low power alternative to volatile memory based main memory and to also eliminate the time

required to transfer code and data from the hard disk to the main memory.

It is also an object of this invention to enable flash memory to operate in an optimal asynchronous fashion with any asynchronous main memory bus such that the CPU can execute programs directly out of the flash memory without any degradation in performance when compared to volatile memory.

It is also an object of this invention to enable flash memory to operate in an optimal asynchronous fashion with any asynchronous main memory bus and to eliminate the need to have custom controllers.

It is also an object of this invention to enable flash memory to operate in an optimal asynchronous fashion with any asynchronous main memory bus to provide a glueless interface to the existing main memory controller and thus reduces cost and loading on the local bus.

A flash memory chip that can be switched into four different read modes is described. Computer systems and hierarchies that exploit these modes are also described. In the first read mode, asynchronous flash mode, the flash memory is read as a standard flash memory. In this mode, the reading of the contents of a first address must be completed before a second address to be read can be specified.

In the second read mode, synchronous flash mode, a clock signal is provided to the flash chip and a series of addresses belonging to a data burst are specified, one address per clock tick. Then, the contents stored at the addresses specified for the burst are output sequentially during subsequent clock ticks in the order in which the addresses were provided. Alternately, if a single address is provided to the flash chip when it is in the synchronous mode, the subsequent addresses for the burst will be generated within the flash chip and the data burst will then be provided as output from the flash chip.

In the third read mode, asynchronous DRAM (dynamic random access memory) mode, the flash memory emulates DRAM. Thus, row and column addresses are strobed into the flash memory

using row and column address strobe signals. The flash memory then converts the row and column addresses internally into a single address and provides as output the data stored at that single address. Furthermore, although the flash memory does not need an extended precharge period or to be refreshed, when in the asynchronous DRAM mode, the flash memory responds to precharge periods and refresh cycles as would a DRAM. Therefore, when in the asynchronous DRAM mode, the flash memory can be controlled by a standard DRAM controller.

In the fourth read mode, synchronous DRAM mode, the features of the second and third modes are combined to yield a flash memory that emulates a synchronous DRAM. Thus, addresses to be read as a data burst are specified by strobing row and column addresses into the flash memory using RAS and CAS signals. The data of the data burst is then provided sequentially as output from the flash memory on subsequent clock ticks.

Other objects, features, and advantages of the present invention will be apparent from the accompanying drawings and from the detailed description which follows below.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements, and in which:

Figure 1 illustrates a computer system having a microprocessor that is coupled to a flash main memory by a high speed bus;

Figure 2 illustrates computer system that combines a static microprocessor with a flash main memory to form a low power consumption, but high performance, computer system;

Figure 3 illustrates a flash memory subsystem that interlaces several asynchronous flash memory units to support synchronous back to back data burst read cycles;

Figure 4 is a timing diagram that illustrates a read cycle for an asynchronous flash memory;

Figure 5 illustrates a state transition diagram for the flash memory subsystem of Figure 3;

Figure 6 illustrates a block diagram of a flash memory integrated circuit having a synchronous flash interface to support synchronous back to back data burst read cycles from internal banks of flash arrays;

Figure 7 is a timing diagram that illustrates a back to back data burst read cycle for the flash memory having the synchronous flash interface;

Figure 8 is a block diagram of a computer system illustrating a generic synchronous flash interface flash memory unit coupled to a microprocessor using a high speed synchronous bus and bus specific glue logic;

Figure 9 is a block diagram of a flash memory integrated circuit having an asynchronous main memory interface;

Figure 10 is a timing diagram that illustrates a column and row address multiplexed asynchronous read cycle for a flash memory having an asynchronous main memory interface;

Figure 11 is a timing diagram that illustrates a refresh cycle for a flash memory having an asynchronous main memory interface;

Figure 12 is a block diagram of an asynchronous main memory interface single in-line memory module for the flash memory integrated circuit having the asynchronous main memory interface;

Figure 13 is a block diagram of a computer system that uses a dynamic random access memory controller to access a flash memory based asynchronous main memory interface single in-line memory module;

Figure 14 is a block diagram of a flash memory integrated circuit having a synchronous main memory interface; and

Figure 15 is a block diagram of a computer system that uses a synchronous dynamic random access memory controller to access a flash memory based synchronous main memory interface single in-line memory module.

DETAILED DESCRIPTION

Figure 1 illustrates a computer system 100 wherein microprocessor 110 is coupled to flash memory 130 by high speed bus 120. In computer system 100, flash memory 130 is used to replace some, if not all, of volatile main memory (not shown). Thus, flash memory 120 is used in lieu of volatile main memory such as dynamic random access memory (DRAM) or static random access memory (SRAM) to provide a high speed non-volatile main memory for computer system 100.

A computer hierarchy based upon volatile main memory loses all information in main memory when power is turned off. A flash-based non-volatile main memory, however, reduces or eliminates the lengthy process of obtaining information from disk when power is turned on. Therefore flash main memory based computer system 100 has higher system performance when a program is initially executed than would a volatile main memory based computer system.

Furthermore, by using flash memory as a main store, the duplication of shadowing information on both disk and RAM is no longer necessary thereby reducing memory cost by eliminating memory duplication.

Moreover, by having information stored in non-volatile flash memory, power consumption is reduced because battery backup of volatile memory is eliminated and because disk accesses are minimized or eliminated.

Figure 2 illustrates a computer system 200 wherein static microprocessor 210 is coupled to flash main memory 230 and battery-backed SRAM 240 by high speed bus 220. Bus bridge 280 couples high speed bus 280 to low speed bus 260. Slow non-volatile memory 250 is coupled to low speed bus 260.

In a typical modern volatile main memory based computer system, there will be four to eight megabytes of volatile main memory, of which, approximately half will be used to store an operating system, executable code for applications and non-changing data files. The rest of the main memory typically contains data files

(and executable code for infrequently used applications) that are read from and written to frequently. The former type of information is stored in what is commonly referred to as "read mostly" files and the latter type of information is stored in what is commonly referred to as "read/write" files. In computer system 200, flash memory 230 is used to store the read mostly files and SRAM 240 is used to store the read/write files that are currently in use.

The benefits obtainable by using a computer system having a flash memory based main memory architecture of system 200 is best illustrated by example. Suppose that a user had a computer and frequently used the computer in a windows operating system environment to run a word processor application (to write letters). Further suppose that the user occasionally used the computer to play video games.

Because most computer operating systems (and applications) are designed to be executed from random access memory (i.e. memory that can be read from and written to), no effort is made to divide their memory address space into read only and read/write areas. In actuality, very little of the memory address space will be written to. Most of the memory address space will be read from to provide code instructions and data that infrequently changes.

ROM-able programs are written to provide code that can be executed from read-only memory (ROM). In ROM-able programs, the code is carefully written to ensure that most of the memory address space is read-only. If there are any parts of the program memory address space that must be written to, these parts are contained in a separate read/write memory address space (that can be saved to a read/write memory device). ROM-able versions of Windows, the underlying Disk Operating System (DOS), and popular applications such as word processing programs are currently commercially available.

Thus, the user of computer system 200 can benefit by storing the read-only parts of ROM-able DOS, ROM-able Windows and a ROM-able word processing program in flash memory 230. The read/write

Micron Technology Inc. et al.
Ex. 1010, p. 490

parts of these programs and the rest of the applications and files can be stored in slow non-volatile memory 250.

Then, when the user starts computer system 200 to write a letter, the read-only parts of the DOS, Windows and the word processor will be resident in flash memory 230. If these programs need any read/write files (e.g. system configuration information), these (typically small) files will quickly be loaded from slow non-volatile memory 250 to battery backed SRAM 240 over bus bridge 280 and low speed bus 260. In short, the major parts of the DOS, Windows and word processing programs will be resident in non-volatile flash memory 230 and can be executed in place directly from non-volatile main memory 230 as soon as the system powers up. Therefore, it will appear to the user as though computer system 200 starts up instantly with the word processor loaded and ready to write the letter.

The user can then retrieve into SRAM 240 a file containing a previously written letter to modify (from where it was stored in slow non-volatile memory 250). Alternately, the user can open a new file in the memory address space of SRAM 240 and can use the new file to start an entirely new letter.

After working on the letter for a while, suppose the user decides to play one of the computer games stored in slow non-volatile memory 250. In such a case, the program associated with the selected game will be loaded into SRAM 250 and the game will begin.

If the user wishes to take a break, the user can suspend the game and leave computer system 200 running. While the user is gone, static microprocessor 210 will sense user inactivity and send a powerdown signal to flash memory 230 over powerdown signal line 270. Then static microprocessor 210 and flash memory 230 will enter a state of low power consumption.

Flash memory 230 and slow non-volatile memory 250 are both non-volatile. Additionally, flash memory 230 will draw little power when in its energy saving powerdown mode and non-volatile memory 250 will not have to be accessed while the user is away. Furthermore, because microprocessor 210 is a static microprocessor, it

Micron Technology Inc. et al.
Ex. 1010, p. 491

will draw very little power but will retain its state when it enters its deep powerdown mode. Moreover, the read/write files (including the executable code for the game) stored in battery backed SRAM 240 will not be lost.

Later, when the user returns, the game can be continued almost instantly from the place where it was suspended. This is because, even though very little power will be consumed while the user is away, the state of computer system 200 will be retained.

Furthermore, if the user desires to upgrade the DOS or Windows operating systems or to upgrade the program stored in flash main memory 230, this will be easy to do because of the ability to write to flash memory. The erasing and then programming operations required to write to flash memory will mean that it will take longer to store programs in flash memory 230 than would be the case where the main memory was volatile memory such as DRAM. Installation of a program, however, is performed infrequently and is usually a relatively prolonged process anyway. Therefore, the extra time required to install a program into non-volatile flash main memory 230 will usually be insignificant when compared to the benefit of having the program instantly available for execution from non-volatile flash memory 230 whenever system 200 is started.

The ability to write to flash main memory 230 further means that the user can install a different program in place of (or in addition to) the word processing program currently stored in flash main memory 230. Therefore, the user can switch to a rival word processing program by installing the rival word processing program in flash main memory 230. Alternately, the user can, for example, install a spreadsheet or graphics design program in flash memory 230 in place of the word processing program that is currently stored there.

Moreover, provided there is sufficient room in the flash memory 230, the user can install additional applications there. Thus, for example, if the user has a particular computer game that is the current favorite, it can be installed to reside in the flash memory 230.

Note that a program can be installed in flash memory 230, even if the program is not ROM-able. As described previously, because flash memory can be written to, a flash file system will permit standard applications to execute in place from flash memory 230. The asymmetric read/write nature of flash will, however, cause a performance penalty to be incurred when standard (i.e. non-ROM-able) programs are executed in place from flash main memory 230.

There are approaches that can be used, however, to mitigate these performance problems. One approach is for the operating system to mark the memory address space of the application as being read-only while the application is executing. Then, if the application attempts to write to flash main memory 230, an interrupt will occur and an error handler can cause the write to be redirected to SRAM memory 240. Later, when there is a lull in processing, the information can be written from SRAM 240 to flash 230 as a background operation. In effect, a portion of SRAM 240 can be used as a write cache for flash main memory 240.

One benefit of using a portion of SRAM 240 as a write cache is that applications stored in flash 230 do not need to be ROM-able. Another benefit is that the write areas of the application memory address space will probably be relatively well localized. In the fortunate case where this is so, any subsequent read or write operations to the write cache will benefit from the faster access of SRAM 240. Also, because it is frequently the case that the information to be written back to the application memory space is ephemeral (i.e. not to be saved from one execution of the application to the next), any ephemeral information in the write cache when the application terminates will not have to be written back to the flash memory 230.

From the above example, one can see that, when compared to a computer system that has the traditional volatile main memory based architecture, there are many beneficial and synergistic effects that stem from having the non-volatile main memory computer system architecture of computer system 200.

First, computer 200 can start instantly when power is first applied. Thus, there is no need to spend time and power loading the operating system and frequently used applications such as the word processor into main memory; they are already there and virtually instantly available.

Second, there is no need to store a copy of an operating system (e.g. DOS and Windows) or to store frequently used application programs (e.g. the word processing program) on slow non-volatile memory 250. Thus, slow non-volatile memory 250 can be smaller than would be required if shadowing of these files was required.

Third, using techniques that will be described below in greater detail, unlike asynchronous (volatile) DRAM, a burst read can be made to flash memory 230. Thus, a cache line for internal cache 215 of microprocessor 210 can be filled from flash memory 220 faster than would be the case if asynchronous DRAM based main memory was used.

Fourth, SRAM 240 can be written to, and read from, faster than can volatile DRAM. Thus, a cache line of internal cache 215 of microprocessor 210 can be filled from SRAM 240 faster than would be the case if a DRAM based main memory was used. Also, programs executing from SRAM 240 can be accessed, and hence executed, faster than would be the case if a DRAM based memory was used.

Fifth, a portion of SRAM 240 can be used to provide an external cache for flash memory 220 thereby further enhancing the execution speed.

Sixth, a lower performance microprocessor can be used to provide the same level of performance as would be the case if a DRAM based main memory was used. This is because the low bandwidth constraint of a volatile DRAM based main memory caused by the need to load the operating system and programs will be eliminated. The operating system and programs are instead stored in non-volatile flash main memory 230. This is also because in system 200 there will be a better match between the data and program instruction needs of microprocessor 210 and the ability of the flash

230/SRAM 240 based main memory to supply them than would be the case with a DRAM main memory system. This will result in fewer idle wait states for microprocessor 210. Frequently, microprocessor 210 will be able to operate with zero wait states. Therefore, even though there will be fewer cycles in which microprocessor 210 can do work because it is a lower performance microprocessor, fewer of the cycles that are available to the microprocessor will be wasted while the microprocessor is waiting for code or data.

Seventh, because microprocessor 210 can have a lower clock rate for a given level of performance, microprocessor 210 will be less expensive, have fewer heat dissipation problems and use less power than if a DRAM based main memory that required that a faster processor to be used.

Eighth, because microprocessor 210 and flash 230 can enter a deep powerdown mode until they are needed, further energy savings will result.

Thus, a non-volatile read/write general purpose computer hierarchy can be achieved with the architecture described above when the dominant portion of main (execute) memory is flash memory.

Depending upon the requirements of system 200, archival storage (i.e. slow nonvolatile memory 250) may or may not be present in the system. If it is present, the needs of system 200 can govern the technology used to provide it (i.e. mechanical, optical or solid state). Also, depending upon the particular needs of the system 200, battery-backed SRAM 240 can alternately be replaced by non-volatile RAM (NVRAM), electrically erasable programmable ROM (EEPROM), DRAM, SRAM, or battery-backed DRAM.

The underlying theme of each of the above configurations for system 200 is that the CPU of microprocessor 210 is fetching code and/or data from random access non-volatile (flash) memory 230 and writing back different data and/or programs on an as-needed basis. As described previously, the new information may or may not be written to the same place in main memory. A RAM based system is

inherently volatile. A ROM based main memory system, although non-volatile, does not provide the capability to write back to the ROM.

Flash is not a symmetrical technology as is the DRAM., i.e. it does not allow for byte alterability at read speeds. Therefore, to modify a file or program in one embodiment, the portion to be altered is moved into byte alterable memory (e.g. NVRAM, EEPROM, SRAM or DRAM), modified, and then written back to flash main memory. This alterable memory can be separate memory chips and/or integrated onto a CPU or microprocessor. Data that is not being altered is read directly by the CPU from Flash. This configuration essentially eliminates disk down load for both read and unaltered data to provide faster read performance than a disk based system. For writes, the write back is comparable in performance to disk. Because the read cycle is faster, however, a higher overall system performance is provided.

The key to each of these configurations is that a flash memory is used as a main memory and is read at a speed that approximates, or surpasses, the read access speed of DRAM. There are basically four possible ways that this can be accomplished. The first way is to provide a hardware controller that interleaves read operations to more than one conventional asynchronous flash main memory chip.

Multiple Chip Interleave for Asynchronous Flash Main Memory

Figure 3 illustrates a flash memory sub-system wherein four flash memory integrated circuits 320 (#1 through #4) are interlaced under the control of a controller unit to provide a burst read capability. Although Figure 3 shows four flash integrated circuits 320, it is to be understood that if the data width of each individual integrated circuit 320 is not that of a double-word, an appropriate number of flash integrated circuits 320 would be provided in parallel with each of the circuits 320 of the figure to provide a double word tdata width. In this example, the controller is implemented using fully programmable gate array (FPGA) 310. In alternate embodiments, the interlace control is provided using alternate technologies.

Figure 4 illustrates the read cycle for a typical asynchronous flash memory such as flash memory integrated circuits 320 of Figure 3.

From Figure 4 it can be seen that before and after a read cycle the data lines (D_{OUT}) of the flash memory are tri-stated (indicated in the figure as OPEN). This means that the data lines have been placed into a high impedance state.

The read cycle of Figure 4 begins with the transition of output enable (OE#) and chip enable (CE#) signals from low voltage (V_L) to high (V_H). Next, if it is not already high, a write enable (WE#) signal is transitioned high to signify a read operation and the address to be read is placed on the address pins (A_0 through A_{19}) of the flash chip. The data stored at the address specified is then retrieved internally from the flash array. The read process takes a known amount of time that is constant. Once that amount of time has elapsed, the OE# signal is brought low and the data that was read from the specified address is driven by the flash chip out onto the data lines (D_{OUT}). The read cycle ends when the CE# and OE# signals are again brought high (perhaps to begin the next read cycle).

Note in Figure 4 that a new address to be read cannot be specified until the previous read cycle has completed. Therefore, a single standard asynchronous flash memory unit cannot be interlaced to support rapid, back to back read cycles. One must instead wait for the first read cycle to complete before a next read cycle can be begun.

The total access time of a standard asynchronous flash component (T_{ACC}) can be divided into two parts. First, there is the time to decode addresses and chip enable, select the correct bits in the flash array to be read and then sense their stored data (T_{ACC}-T_{OE}). Second, there is the time required to drive this information onto the data bus through the output buffers (T_{OE}). Once a single address for a data burst read is known, the rest of the addresses of the data burst read are also known. Therefore, referring again to Figure 3, interlacing exploits this property of a data burst to amortize the first access time (T_{ACC}-T_{OE}) across two or more flash chips to provide a burst that has an improved average read access time.

Consider the case of a 20 MHz synchronous local bus (i.e. each clock cycle takes approximately 50 nS). At that speed, a typical

Micron Technology Inc. et al.
Ex. 1010, p. 497

microprocessor will have a maximum burst read capability of 2-1-1-1 bursts. That means that when the microprocessor reads a burst of data, the burst will be made of four double-words of data that are each part of the same cache line. Furthermore, once the microprocessor specifies the first address of the burst, it will be unable to accept the contents stored at that address for two cycles (approximately 100 nS in this example). Then, the microprocessor will be able to accept the respective contents of the second, third and fourth addresses of the cache line during the next three clock cycles.

If flash memory chips 320 each have a (T_{ACC}-T_{OE}) less than 2 clock cycles (approximately 85 nS) and a (T_{OE}) less than 40 nS, then flash memory sub-system 300 will be able to provide a burst read that matches the maximum burst read attainable by the microprocessor. This is because, as soon as FPGA 310 receives the first address of the data burst to be read, it will enable all four flash memory chips 320 and order each chip 320 to read one of the addresses of the data burst in parallel with the other chips 320. After the first two cycles, the data stored at the first address of the data burst will be available from flash 320 #1. Then, on each successive cycle thereafter, the data stored at the next address of the data burst will be available (from flash 320 #2, then #3 and finally #4).

This means that in this example flash memory sub-system 300 will be able to provide a data read burst of four double-words of data during a five clock cycle period (approximately 250 nS) thereby providing an average read time of approximately 50 nS. If interleaving were not used (e.g. if only one flash memory chip 320 was used), the read for each address would take three cycles, (i.e. two cycles for (T_{ACC}-T_{OE}) and one for (T_{OE})) and each read would have to wait for any earlier read to complete. Using the numbers presented above, each address of the data "burst" to a single asynchronous flash 320 would take 150 nS. Therefore, without interlacing, the fastest burst speed attainable would be a 3-3-3-3 data burst that would take 600 nS to complete (4 times 150).

Micron Technology Inc. et al.
Ex. 1010, p. 498

Clock signal CLK coordinates the operation of flash memory sub-system 300. In sub-system 300, four byte flash memory pages are selected by common addresses A21 through A2 of the AX through A0 memory space. The higher order addresses (AX through A22) decode the 4 Mbyte flash array in the system memory map. Addresses A1 and A0, along with the RD# (read) signal input to FPGA 310, cause FPGA 310 to generate output enable signals (OE#) for flash memory components 320 (#1 through #3). Therefore, in one embodiment, system addresses of flash sub-system 300 are associated with flash memory components 320 as follows:

System Address	Component	Component Address
0	#1	0
1	#2	0
2	#3	0
3	#4	0
4	#1	1
5	#2	1
etc. . . .		

A data burst series of read accesses from memory sub-system 300 will proceed as follows:

First, the processor will signal that it is to read data from system address 0. Signal CE# will enable flash memories 320 #1 through #4 and signal ADDR will provide them with address 0. FPGA 310 will decode system address bits A1 and A0 and, because they are both zero, will enable the OE# signal for flash component 320 #1. Because this is the first access to the four byte page, the access will take 150 nS (i.e., three 50 nS clock cycles).

Second, the processor, executing sequential code, will read from system address 1. Components 320 #1 through # 4 will remain enabled, and address (ADDR) inputs to them will remain as zeros. FPGA 310 will then decode the "01" on system address bits A1 and A0, and enable the OE# line for component 320 #2. Because the data for component 320 #2 has already been selected and sensed, the access

time for component 320 #2, therefore, will only have a 50 nS delay (i.e., one clock cycle).

Then, similar 50 nS accesses will follow for system addresses 2 and 3, reading from components #3 and #4, respectively.

If the next data burst begins with address 4, the processor will then read from system address 4 on its next cycle. This will change system address A2 from zero to one thereby incurring another full 150 nS delay for this access from component 320 #1. This is because a new 4-byte page, and therefore, new data from each of the flash memories 320 is to be read. Once again, however, subsequent accesses to system addresses 5, 6, and 7 will each only have a 50 nS duration.

FPGA 310 decodes lower addresses and, correspondingly, generates OE# signals to flash memory devices 320 within an interleaving page. The earlier example was for a four-byte page, so addresses A1 and A0 are decoded. For a two-byte page, only address A0 would be examined, whereas an eight-byte page would use addresses A2 through A0.

FPGA 310 will examine all upper address bits (i.e. AX to A22 in this example), to determine first if the memory subsystem 300 is being accessed, and then if access to the same page is occurring, and adjust wait states back to the processor accordingly (via the READY signal output). This function is accomplished by the comparison logic sub-block 340 of Figure 3.

Furthermore, FPGA 310 will distinguish between a read or write to flash memories 320, and adjust wait states accordingly. Figure 5 illustrates a state transition diagram for hardware interleaving that is implemented in FPGA logic 310.

State transition diagram 500 begins in initial state 510 wherein the FPGA is waiting for an access to a flash memory 320 of sub-system 300. When the FPGA detects the beginning of an access to the flash memory 320 of sub-system 300, a test is made in state 520 to determine whether the access is to the same four-byte page as was the previous access. If the access is to the same four-byte page as was the previous

access, then transition 524 is taken to state 530 and a test is made to determine whether the access is a read operation.

If it is determined in state 530 that the access is a read operation, then the information to be read is already available and a transition is made on path 534 to ready state 560 and the read is performed. After the read is performed, a transition is made from ready state 560 back to initial state 510 and the next access is awaited.

Returning now to state 530, if it is determined in state 530 that the access is not a read operation, then a transition is made on path 532 to wait state(s) generation state 550 and an appropriate number of wait states are performed while an erase is performed (if needed) and the information is programmed.

Returning now to state 520, if it is determined in state 520 that the access is not to the same four-byte page as the previous access, then a transition is made on path 522 to state 540 and a new address is latched for the page compare. Then, in state 550, an appropriate number of wait states are performed while the new four-byte page is accessed.

Note that the flash memory writes cannot take advantage of interleaving as can reads because there are no TACC or TOE equivalents for write operations. Although memory sub-system 300 supports back-to-back read bursts and reduces the average read access time for sequential accesses, it increases the complexity of the system hardware required over that required for access to a conventional asynchronous flash memory chip.

As described previously, however, a write cache or software interleaving techniques can be used to maximize program and erase performance of memory sub-system 300. Furthermore, in the case where a ROM-able program is being executed in place from flash memory sub-system 300, there will not be any writes to the flash memory during the execution of the program. In such a case, the slower write performance of the flash memory will only be a factor on the rare occasion that a new ROM-able version of the program is installed.

Micron Technology Inc. et al.
Ex. 1010, p. 501

Synchronous Flash Interface (SFI) for Flash Main Memory

Besides requiring the additional system complexity of FPGA 310, a memory subsystem such as subsystem 300 will have to have at least two flash memory chips 320 to interleave between (four in the example of Figure 3). Furthermore, if additional flash memory chips are to be added to sub-system 300, they too must (at least) be added in pairs. Moreover, each flash memory chip added to sub-system 300 is enabled simultaneously and, hence sub-system 300 consumes a large amount of power. Alternately, an entire synchronous flash interface can be incorporated onto a single flash memory chip.

Figure 6 illustrates a block diagram of a synchronous flash interface (SFI) flash memory integrated circuit 600 that incorporates a complete synchronous flash interface in a single flash memory chip. The synchronous flash interface provides a means of converting the synchronous cycles that appear from a controller to the asynchronous protocol that the flash core understands. The device is internally split into two banks and the least significant address is used to distinguish between the two banks. In short the device is interleaved internally. The user can direct back to back cycles to the alternate banks and, after an initial delay equal to an asynchronous flash access time, data will output onto the output bus at a rate that is approximately equivalent to half the asynchronous flash access time. Hence, the device using this interface creates an average access time for sequential read accesses that is significantly less than the access time of an asynchronous flash device.

A clock input is a part of the interface. An address latch enable pin is present to indicate that a valid address is present on the address bus. All the external operations of the device are synchronized to the rising edge of the clock. ALE# is sampled on the rising edge of the clock and, if it is sampled valid, the address that is present on the address bus is latched into the part and an access commences. The user can alternately strobe addresses into the device every clock cycle. After an initial delay the data corresponding to these addresses is output on the data bus in synchronism with the rising edge of the

Micron Technology Inc. et al.
Ex. 1010, p. 502

clock. The user can cycle the device at frequencies as high as 33 MHz and can expect back to back burst cycles to be performed without any wait states in between. Also part of the design is a programmable latency that enables the user to define the clock cycle latency between address strobe and valid data.

In Figure 6, the flash array of SFI 600 is divided into two banks, flash bank A 610, and flash bank B 620. Mode register 680 can be set to enable or disable the SFI mode. Thus, when mode register 680 is set, signal interlace enable (IE) enables interlace control logic and address transition detection (ATD) logic 570, transistor-to-transistor logic (TTL) stages 672, bank select logic 674 and bus logic 640. When SFI is enabled, interlace control 670 and bank select logic 674 operate to interlace read (and write) operations between flash bank A 610 and flash bank B 620 in a manner that is analogous to the way that FPGA 310 interlaced accesses to flash chips 320 in sub-memory system 300 of Figure 3. Thus, a single SFI flash chip 600 can support back-to-back read bursts when mode register 680 has been set to SFI mode.

In the SFI mode, an address to be accessed is presented to TTL stages 672 on address buses A19 through A0 along with a clock pulse (CLK) and an address latch enable signal (ALE#). Output enable and write enable signals are also provided to TTL stages 672. Addresses A19 through A1 are then provided from TTL stages 672 to bank A address latch 630 and bank B address latch 632. Signals CLK, ALE# and A0 are provided to interlace control logic and ATD logic 670.

Based on the low order bit of the address (A0), interface control logic 670 controls bank A address latch 630 and bank B address latch 632 using signals address latch enable A (ALEA) and address latch enable B (ALEB), respectively. Thus, interface control logic 670 thereby steers the remaining higher order bits of the address to be accessed to flash bank A 510 or flash bank B 520. Signals address transition detected A (ATDA) and address transition detected B (ATDB) output from interlace control logic and address transition detection logic 670 prepare flash array banks A 610 and B 620, respectively, to be read.

Data multiplexer (MUX) 650 receives as input signal S/A OUT A from the sense amplifiers of flash bank A 610 when bank A 610 is read. Data MUX 650 also receives as input signal S/A OUT B from the sense amplifiers of flash bank B 620 when bank B is read. Bank select logic 674 receives signals CLK, ALE# and A0 as input and provides output signal bank select (BS). Signal BS is used to control data MUX 650 so that it interlaces the data read from banks A 610 and B 620 and steers the data into output buffer 660. Output buffer 660 then drives the data read out onto data input/output (I/O) pads 690 of SFI flash chip 600.

Thus, if a read access is being performed when SFI flash 600 is in the SFI mode, the next address to be accessed can be begun while the preceding access is being performed. If the next address to be read belongs to the bank that is not presently being read, then the next address can be steered to the appropriate bank (A 610 or B 620) while the output from the previous access is being latched into output buffer 660 and is being driven onto data I/O pads 690. Also part of the design is a programmable latency that enables the user to define the clock cycle latency between address strobe and the presentation of valid data on data I/O pads 690.

Alternately, more than two flash banks can be used in SFI chip 600 provided that an appropriate number of address bits are used to select the proper bank. Thus, if four banks are interlaced, two address bits would be used to select among the banks.

Figure 7 is a timing diagram that illustrates a sustained read burst cycle for SFI flash 600 of Figure 6. In Figure 7, synchronizing clock signal CLK is operating at 33 MHz, therefore each pulse (T1 through T8) of signal CLK is 30 nS long. The read burst begins on the rising edge of clock pulse T1 when the first address to be read (ADDR 0) is placed on the address pins and address latch enable signal ALE# is brought low to indicate that a valid address is ready to be read. Note that the low order bit A0 of address ADDR 0 is low on the rising edge of clock pulse T1. Therefore, address ADDR 0 belongs to bank A and bank A will be read to retrieve the data stored at address ADDR 0.

The read burst continues on the rising edge of clock pulse T2 when the second address to be read (ADDR 1) is placed on the address pins. In one embodiment, address latch enable signal ALE# is brought high and then low to indicate that a valid address is ready to be read. Alternately, address latch enable signal ALE# is simply held low to indicate that another valid address is ready to be read. Note that the low order bit A0 of address ADDR 1 is high on the rising edge of clock pulse T2. Therefore, address ADDR 1 belongs to bank B and bank B will be read to retrieve the data stored at address ADDR 1. Because addresses ADDR 0 and ADDR 1 belong to different flash banks, both address locations can be read concurrently.

The read burst continues on the rising edge of clock pulse T3 when the third address to be read (ADDR 2) is placed on the address pins and address latch enable signal ALE# is held low to indicate that another valid address is ready to be read. Note that during period T3 the read from ADDR 0 of bank A (begun during period T1) completes. Thus, during period T3, bank select signal BS is brought low to steer the data read from ADDR 0 to the output buffer where it is latched. The data of ADDR 0 is then driven out onto the data I/O pads of the SFI chip at the beginning of pulse T4.

Meanwhile, because the low order bit A0 of address ADDR 2 is low on the rising edge of clock pulse T3, address ADDR 2 belongs to bank A and bank A will be read to retrieve the data stored at address ADDR 2. The bank A output from address ADDR 0 has been latched into the output buffer, so the reading of ADDR 2 of bank A can begin. Furthermore, addresses ADDR 2 and ADDR 1 belong to different banks, therefore both address locations can be read concurrently.

The burst read can continue indefinitely (or at least until the address space of SFI flash chip 600 is exhausted) with a new address being input on each clock pulse and the data for that address being output two pulses later.

Referring again to Figure 6, in order for SFI chip 600 to be truly compatible with the high speed synchronous bus with which it communicating, bus logic block 640 receives and provides bus specific

signals for SFI 600. Thus, the particular high speed synchronous bus may require, for example, that a parity bit be received with every write operation and then sent again as a handshake. In one embodiment, interface logic 600 will latch and examine all bits written to flash memory 600 during write operations and provide this parity data during the post-write "handshake". Alternately, by examining data during write attempts to flash memory 600, interface logic 640 can use the same parity algorithm used by the high speed synchronous bus to generate the required parity information for the post-write "hand-shake".

Furthermore, the bus protocol may be such that the parity bit may need to be sent from SFI 600 when data is read. Because flash memory does not require that parity bits be stored with data, bus logic 640 would use the same parity algorithm used by the high speed synchronous bus to generate the necessary parity bit when data is read. Thus, during flash memory 600 read operations, logic 640 will derive parity bits from the output of flash memory 600 for transfer to the high speed synchronous bus. Therefore, by examining data both during write attempts to flash memory 600 and read attempts from flash memory 600, interface logic 640 will generate the required parity information.

When mode register 680 is not set to indicate SFI mode, the synchronized bank interlacing control functionality of interlace control 670 is disabled and SFI flash chip 600 will function as would a standard asynchronous flash memory chip. In such a case, the CLK and ALE# signals will be ignored as will the bus specific input signals to bus logic 640. Instead signals CE# and OE# are used to access SFI flash 600 using an asynchronous read signal of the type depicted in Figure 4.

In one embodiment, mode register 680 is implemented using content addressable memory. Alternately, in another embodiment there is no mode register 680 and the SFI functionality of chip 600 cannot be disabled and SFI chip 600 will not function as a standard asynchronous flash memory.

In one embodiment, the SFI flash memory 600 requires one address per word/byte read. Some burst processor buses, on the other hand, initiate a multiple-read burst access by providing a start address, and then repeatedly signaling for reads with no additional address information provided. Alternately, to accommodate these processors, SFI flash chip 600 will latch the start address of a burst, and using on-chip counters, will generate subsequent addresses throughout the remainder of the burst read sequence. In such a case, SFI chip 600 will contain a full-address latch, and a multi-bit counter that matches the addressable space in SFI flash memory array 600. After being provided the start address, the counter logic will auto-increment (in linear or non-linear order, depending on the bus) and feed addresses to the SFI flash memory banks. This allows a multiple-burst access of indeterminate length to occur at highest system performance.

Figure 8 illustrates a computer system 800 that uses an alternate generic SFI chip approach. In Figure 8, SFI flash chip 860 has a generic synchronous flash memory interface. Glue logic 830 is a separate integrated circuit chip that is placed between microprocessor 810 and SFI flash chip 860 to provide the circuitry required to interface SFI chip 860 to high speed synchronous bus 820.

One benefit of the approach of Figure 8 is that providing a generic interface control 870 within SFI flash chip 860 permits a burst read of banks A 880 and B 890. Thus, rather than requiring pairs of asynchronous flash chips in order to perform interlacing, a flash main memory sub-system can be expanded in units of a single chip. At the same time, when compared to FPGA 310 of Figure 3, the design of SFI glue logic 830 is simpler because glue logic 830 does not have to include the generic synchronous interface control logic 870 that would be required for any SFI chip 860.

Furthermore, a single SFI chip 860 can be used to interface with a large number of different synchronous buses because the bus specific control logic 850 is implemented within glue logic 830 instead of within SFI chip 860.

Micron Technology Inc. et al.
Ex. 1010, p. 507

An additional benefit to this approach is that, when compared to a standard asynchronous flash chip, SFI chip 860 need only have two additional pins (i.e. CLK and ALE#). Any other signals that are required for synchronous bus 820 are provided by glue logic 830 in a manner that is transparent to SFI chip 860. Typically, a flash chip is mounted in a small package such as a thin small outline package (TSOP). Because of the small size of the package, a standard asynchronous flash chip will use almost all of the pins available to the package. When providing a generic synchronous flash interface, the two additional pins required are available even within constraints imposed by mounting the SFI flash chip 860 in a small package.

Note that glue logic 830 includes burst logic 840. This is because, as mentioned earlier, different microprocessors 810 use different burst ordering sequences.

A common sequence is the linear addressing sequence. In the linear addressing sequence, a data burst increments sequentially from the beginning address of the burst. Typically, in this sequence, the data burst will begin with the lowest address in the cache line being read and will increment through the addresses of the cache line.

The other burst sequence is a non-linear addressing sequence. In this sequence, the burst read does not necessarily have to begin with the lowest address of a cache line. If the burst does begin with the lowest address of the cache line, then the non-linear addressing sequence behaves as does the linear sequence. If, on the other hand, the burst begins on an address other than the lowest address of the cache line, the subsequent addresses of the burst are generated in a manner that guarantees that every address of the cache line will be accessed and that they will be accessed in a way that interleaves between odd and even address banks.

Burst logic 840 takes a first data address for a data burst and generates the rest of the addresses of the burst in accordance to whatever addressing sequence is required by microprocessor 810.

Note that because both the linear and non-linear burst sequences interleave between odd and even banks, SFI chip 860 will

interleave between banks regardless of the sequence generated by burst logic 840 (be it linear or non-linear).

Asynchronous Main Memory Interface (AMMI) for Flash Main Memory

The synchronous flash interface (SFI) supports back to back data bursts thereby permitting flash memory to be used as main memory. There may be occasions, however, where system constraints prevent the use of the synchronous flash interface to a high speed synchronous bus. One such case is where there are already several components that are coupled to the high speed synchronous bus. In such a case, the additional coupling of the SFI chip (or its glue logic controller chip) will increase the capacitive load of the high speed synchronous bus to the point where it can no longer support a desired level of performance. Another such case is where the SFI chip requires a glue logic controller chip and the cost constraints imposed when designing the system prohibit the additional expense of this chip.

Furthermore, many portable systems do not have a local bus (i.e. do not have a high speed synchronous bus). Typically, in these systems, the CPU is integrated with the peripheral device controllers onto a single microprocessor chip. In such systems, only an asynchronous main memory bus and a relatively slow expansion bus are available within the system.

Asynchronous flash memory as it exists today must interface within these systems to the expansion bus through a custom controller. Bus cycles on the expansion bus are typically slow. Therefore, because each expansion bus cycle is typically longer than the asynchronous access time of a flash memory, the synchronous flash interface provides no real benefit when interfacing with the expansion bus.

An Asynchronous Main Memory Interface (AMMI) permits flash memory to interface to the asynchronous high speed main memory bus in any computer system. In effect, the AMMI permits a flash memory chip to emulate an asynchronous DRAM chip. In

Micron Technology Inc. et al.
Ex. 1010, p. 509

addition to providing a low cost, low power flash main memory, the AMMI also eliminates the time required to transfer programs or data from the hard disk to the DRAM by permitting nonvolatile flash memory to be used as main memory in place of volatile DRAM. Thus, as was the case with SFI, the CPU can execute programs directly out of a flash main memory without any degradation in performance when compared to volatile main memory.

Because the AMMI permits a flash chip to emulate a DRAM chip, the AMMI eliminates the need to have custom controllers. It provides a glueless interface to the existing DRAM controller of a system and thus reduces cost and loading on the local bus.

The AMMI improves system performance and provides an alternative to hard disks in portable computer systems. It also provides for faster execution out of flash with minimal change to the system architecture. Moreover, the system performance can be increased without requiring any additional peripheral logic.

Figure 9 illustrates a block diagram of one embodiment of a flash memory chip having an asynchronous main memory interface. In the embodiment illustrated in Figure 9, flash chip 900 can operate as a standard asynchronous flash memory chip or can emulate a DRAM. The particular addressing mode to be used is determined by whether mode register 910 has been set or not. Thus, if mode register 910 has been set to signify that an asynchronous main memory interface is desired, AMMI flash chip 900 will emulate a DRAM chip. On the other hand, if mode register 910 has not been set to signify that an asynchronous main memory interface is desired, AMMI flash chip 900 will not emulate a DRAM chip. In the latter case, the DRAM emulation circuitry of AMMI chip 900 would be disabled and AMMI flash chip 900 will operate as would a standard asynchronous flash memory chip. In one embodiment, mode register 910 is implemented using content addressable memory.

Figure 10 illustrates a read cycle of an AMMI flash chip 900 when it is emulating a DRAM. Because the AMMI flash chip 900 is

Micron Technology Inc. et al.
Ex. 1010, p. 510

emulating a DRAM chip, Figure 10 also illustrates a read cycle for a typical DRAM chip.

The asynchronous main memory interface provides a means of interpreting DRAM cycles and converting these cycles to a protocol that can be used to access a flash memory array. To reduce the pin count on a typical DRAM device, rather than using a single 20 bit address, as does flash, a DRAM multiplexes a ten bit row address and a ten bit column addresses. Address signal multiplexing for a DRAM is accomplished by providing two additional pins not found on a prior art asynchronous flash device.

Thus, in a DRAM, a Row Address Strobe (RAS#) pin and a Column Address Strobe (CAS#) pin are used to control the strobing in and latching of respective row and column addresses. Typically, these addresses are latched on the falling edges of these strobes and are used to access specific bits inside the DRAM array. Because AMMI flash memory chip 900 can emulate a DRAM chip, AMMI flash chip 900 has row and address pins for accepting row and column address strobe signals, respectively.

Some DRAM devices use row and column addresses that are not ten bits. For example, the row and column addresses of a DRAM device may only require eight bits to be expressed. Throughout the following discussion ten bit row and column addresses will be used. It is to be understood, however, that AMMI flash 900 can be practiced using other sized row and column addresses.

In Figure 10 it can be seen that the read cycle for an AMMI flash chip 900 that is emulating a DRAM begins with the RAS# and CAS# signals both high and the multi-bit data output signal (D_{OUT}) open (tri-stated). The ten bit row address is then placed on the A_0 through A_9 address pins of the flash chip 900 and the RAS# signal is brought low to indicate that a valid address is to be strobed into the AMMI flash. (Note that in the case of an eight bit row address, the eight bit row address would be placed on the A_0 through A_7 address pins of the flash chip.)

After the row address has been strobed into AMMI flash chip 900, the RAS# signal continues to be held low for a predetermined period of time. In a standard DRAM, the RAS# signal is held high at the end of the read cycle prior to the next read cycle to precharge. Of course, because flash memory is not based on the capacitor-based storage technology of the DRAM, the flash array does not need the lengthy precharge time required by a DRAM. On the other hand, because AMMI flash 900 is emulating DRAM, it must react as would a DRAM.

Once the row address has been latched within AMMI flash chip 900, a read operation is signified by transitioning a WRITE# signal from low to high. Then, the ten bit column address is asserted on the A_0 through A9 address pins of AMMI flash chip 900. Next, the CAS# signal is asserted and the column address is latched within AMMI flash chip 900. The ten bit row address and the ten bit column address are then combined within AMMI flash chip 900 to form a twenty bit address used to specify a location within the flash array to be read. Once read, the contents of the twenty bit address location are driven onto the data pins of AMMI flash chip 900 as multi-bit signal data out (D_{OUT}). Then, after a CAS# goes high, the multi-bit signal data out (D_{OUT}) is tri-stated again.

Meanwhile, after the row address has been strobed into AMMI flash chip 900 and the RAS# has been held low long enough for a row precharge to occur (if a DRAM was being accessed), the RAS# signal transitions back to high. The CAS# signal also transitions back to high after the column address has been strobed into AMMI flash chip 900.

Referring again to Figure 9, the mode value stored within mode register 910 is provided to control logic 920. One output of control logic 920 is an address multiplexer select (AMS) signal. Signal AMS controls address multiplexer (mux) 930. Address mux 930, in turn, steers the ten bit address that is asserted on address pins A_0 through A9 of AMMI flash 900 into the proper portion of address latch 940.

Address latch 940 is divided into a ten bit wide high portion 942 and a ten bit wide low portion 944 for storing, respectively, the ten high and low order bits of flash address 950. Demultiplexed address 950 is a twenty bit address used to address flash array 960. Another output from control logic 920 is a multi-bit address latch control (ALC) signal that controls the latching of addresses within the high portion 942 and the low portion 944 of address latch 940.

The mode value stored within mode register 910 is also provided to control bypass multiplexer (MUX) 955. The inputs to bypass multiplexer 955 are demultiplexed address 950 from address latch 940 and the twenty bit external flash address asserted on address pins A_0 through A_{19} of the flash chip 900 when flash chip 900 is operating as a standard asynchronous flash chip. The output from bypass multiplexer 900 is flash address 957 that is provided to flash array 960.

Thus, when the AMMI flash chip 900 is operating as a standard asynchronous flash chip, a twenty bit flash address will be asserted on address pins A_0 through A_{19} of the flash chip 900. Signal 915 from mode register 910 will set bypass multiplexer 955 to steer the twenty bit flash address asserted on address pins A_0 through A_{19} of the flash chip 900 to flash array 960 as flash address 957.

Alternately, when AMMI chip 900 is emulating a DRAM chip, address pins A_0 through A_9 of flash chip 900 will be multiplexed. First, a ten bit row address will be placed on address pins A_0 through A_9 of flash chip 900 and the RAS# signal will indicate that the row address is ready to be strobed into the flash chip. Then, a ten bit column address will be placed on address pins A_0 through A_9 of flash chip 900 and the CAS# signal will indicate that the column address is ready to be strobed into the flash chip.

Note that in the DRAM emulation mode, no address signals will be placed on address pins A_{10} through A_{19} of the flash chip. Therefore, when the row address is to be strobed into AMMI flash device 900, control logic 920 will set address mux 920 to steer the ten bit row address that is asserted on address pins A_0 through A_9 of the

flash chip into the low portion 944 of address latch 940. Signal ALC will then cause the ten bit row address that is asserted on address pins A_0 through A_9 of flash chip 900 to be latched by low portion 944 of address latch 940.

Later, when the CAS# signal indicates that the column address is to be strobed into AMMI flash device 900, control logic 920 will set address mux 920 to steer the ten bit column address that is asserted on address pins A_0 through A_9 of the flash chip into the high portion 942 of address latch 940. Signal ALC will then cause the ten bit column address that is asserted on address pins A_0 through A_9 of the flash chip to be latched by high portion 942 of address latch 940. In this way, the ten bit row address will be combined with the ten bit column address internally within AMMI flash memory 900 to form a twenty bit demultiplexed flash address 950 that is stored in latch 940 and can be used to address the flash array 960. Signal 915 from mode register 910 will set bypass multiplexer 955 to steer the twenty bit demultiplexed flash address 950 stored in address latch 940 to flash array 960 as flash address 957.

Regardless of whether the AMMI chip 900 is operating as a standard flash chip or is emulating a DRAM chip, once a complete twenty bit flash address 950 has been provided to flash array 960, the contents stored within flash array 960 at address 950 are then retrieved from flash array 960 and sent over bus 970 to data latch 980. Data latch control (DLC) signal from control logic 920 controls data latch 980. At the proper time, data latch 980 drives the data that has been read from flash array 960 out onto the data bus 990 (and data pins of AMMI flash 900). Then, at the completion of the read cycle, control logic 920 causes data latch 980 to tri-state. Other DRAM read accesses such as fast page mode are handled similarly.

In an alternate embodiment, AMMI flash chip 900 is a dedicated DRAM emulation device and cannot function as a standard flash chip. In this alternate embodiment, mode register 910 and address pins A_{10} through A_{19} of the flash chip are not required. This

embodiment enables a lower pin count interface to AMMI chip 900 and is especially attractive to space constrained design situations.

In addition, control logic 920 contains logic that is used to interpret DRAM cycles that are not needed for flash 900 when it is emulating a DRAM. As mentioned before, unlike DRAM, flash memory does not need to have an extensive precharge or the various refresh cycles commonly required by DRAM. Figure 11 illustrates a typical refresh cycle for an AMMI flash chip that is emulating a DRAM. Because an AMMI flash chip that is emulating a DRAM will behave as would a DRAM chip, Figure 11 illustrates a typical refresh cycle for a DRAM.

Note that the refresh cycle of Figure 11 begins as does the read cycle of Figure 10. Thus, at the beginning of the refresh cycle, the column address strobe signal (CAS#) is brought high and a row address is placed on address pins A_0 through A_9 of AMMI flash chip 900 while the row address strobe signal (RAS#) is brought low.

Unlike the read cycle of Figure 10, however, during the refresh cycle of Figure 11, after the row address (of the row to be refreshed) has been strobed into AMMI flash memory chip 900, no column address is strobed into the flash chip. Instead, the RAS# signal is held low for a period sufficiently long to permit the row specified to be refreshed (if a DRAM was being refreshed) and the CAS# signal is held high throughout this period. During the refresh cycle, the data pins of AMMI flash chip 900 remain open (tri-stated).

Referring again to Figure 9, when in the DRAM emulation mode, control logic 920 will identify a refresh cycle pattern and hold data latch 980 tri-stated throughout the refresh cycle. Other standard refresh cycles (e.g. CAS before RAS (CBR) refresh, self-refresh and hidden refresh cycles) are well known in the art and handled similarly. In this manner, a glueless interface from the AMMI flash chip 900 to a DRAM controller is achieved.

Alternately, some of the inefficiency associated with emulation of a DRAM can be eliminated while the benefits of DRAM emulation are retained by controlling access to AMMI flash chip 900 using a

DRAM controller that has been modified to compensate for the fact that some, if not all, of the "DRAM" that it is controlling is actually AMMI flash 900 that is emulating DRAM. In such a case, the multiplexing of row and column addresses to the AMMI flash chip 900 would be retained. The DRAM controller would be modified, however, so that the DRAM controller will not force AMMI flash chip 900 to wait for a lengthy (and unnecessary) row precharge period to complete. The DRAM controller would also be modified to eliminate refresh cycles for the AMMI flash memory 900. That way the DRAM controller will not force AMMI flash chip 900 to be tied up during lengthy (and unnecessary) refresh cycles. Thus, in the case of a programmable DRAM controller, the DRAM controller can be programmed to generate DRAM control signals that have been optimized to access an AMMI flash chip 900. An example of a suitable programmable DRAM controller is described in U. S. Patent Number 4, 785,428, Programmable Memory Array Control Signals, Bajwa et al., issued November 15, 1988.

Some programmable DRAM controllers permit more than one type of DRAM to be controlled simultaneously by the same controller even when each DRAM type being controlled has different timing cycles than the others. This is accomplished by having more than one bank of DRAM and permitting the controller to have different access cycle timings programmed for each bank. The DRAM are then grouped so that all of the DRAM in a particular bank are of the same type.

Thus, in an embodiment where a programmable controller can be programmed on a bank-by-bank basis, the AMMI flash memory 900 is placed in one bank and standard DRAM is placed in another bank. Then, the programmable DRAM controller can be programmed to control the DRAM in the DRAM bank using standard DRAM timing cycles. The programmable controller can be also be programmed to control the bank of AMMI flash memory 900 that is emulating DRAM in a way that has been optimized to meet the needs of AMMI flash memory 900.

Micron Technology Inc. et al.
Ex. 1010, p. 516

Alternately, rather than using a programmable controller, a special non-programmable DRAM controller can be provided that has been optimized for access to AMMI flash memory 900. This special non-programmable controller could either control AMMI flash memory 900 exclusively, or could also control one or more banks of conventional DRAM.

In yet another alternate embodiment, a DRAM controller that has been designed to interface to a high speed synchronous bus (i.e. a local bus) can be used to control access to AMMI flash memory 900.

Programming an AMMI flash chip 900 that is emulating DRAM is performed in a fashion similar to writing to DRAM. As is the case with a read operation, the control logic 920 interprets the commands on the control pins of the flash chip and issues them in accordance to the protocol required by asynchronous main memory flash device 900.

Because programming (and erasing) flash memory takes longer than writing to DRAM, approaches similar to those mentioned earlier can be used to minimize the impact of the time required to program an AMMI flash device 900. For example, flash file system software can insert wait states whenever a write to AMMI flash 900 is required. The number of writes to the AMMI flash 900 memory can be significantly reduced by using the AMMI flash device 900 to store ROM-able programs and data that are infrequently changed and by supplementing the AMMI flash memory 900 with a memory technology that can be written to quickly. Alternately, a write cache can be provided to store information that is to be programmed into the AMMI flash 900. The programming of the AMMI flash 900 memory can then be performed in the background during what would otherwise be idle periods.

Figure 12 illustrates a scheme for creating a DRAM-compatible SIMM using one or more AMMI flash memory chips. This scheme allows AMMI SIMM 1200 containing an AMMI flash memory chip 900 to plug directly into a SIMM socket intended for DRAM. Thus, a standard DRAM controller can control AMMI SIMM 1200. The

AMMI SIMM enables the system to access data and direct-execute code stored in the flash memory 900 as if it were stored in DRAM while requiring no system hardware modification.

Figure 12 illustrates an AMMI SIMM 1200 that is populated with a single AMMI flash chip 900. In an alternate embodiment, AMMI SIMM 1200 is populated with two or more AMMI flash memory chips 900.

An advantage of the AMMI flash chip 900 is that if a DRAM controller exists in the system, the hardware interface to the AMMI flash memory chip 900 is dramatically simplified by the ability of the AMMI chip to emulate a DRAM. AMMI SIMM 1200 further simplifies the addition of flash main memory into a computer system. This is because AMMI SIMM 1200 allows a DRAM controller to interface, not only to individual flash memory components, but also to an industry-standard SIMM containing one or more AMMI flash devices 900. All additional circuitry required for flash memory compatibility is included directly on the AMMI SIMM 1200.

When compared to a similar density DRAM chip, an AMMI flash memory chip 900 can have several extra pins. For example, the AMMI flash chip 900 may have a RY/BY# output pin to indicate when the chip is ready to be accessed and when the chip cannot be accessed because it is busy (for example, when the flash array is being programmed or when a flash block is being erased). Therefore, RY/BY# output pin of AMMI flash chip 900 on SIMM 1200 is simply not connected (i.e. is "floated"). Instead of using the RY/BY# pin to determine when AMMI flash chip 900 is ready or busy, system software is used to poll component status registers within chip 900 to determine the program/erase status of the chip.

When it is emulating a DRAM, AMMI flash chip 900 will multiplex row and addresses onto address pins A_0 through A_9 when they are strobed into chip 900. Therefore, address pins A_{10} through A_{19} are not used by AMMI flash chip 900 when it is part of AMMI SIMM 1100. Thus, address pins A_{10} through A_{19} of AMMI flash chip

Micron Technology Inc. et al.
Ex. 1010, p. 518

900 are not connected when AMMI flash 900 is populating AMMI SIMM 1200.

The RAS# and CAS# signals are both provided directly to AMMI flash chip 900. They are also provided to read/write logic 1210 along with the DRAM WRITE# signal. The WRITE# signal indicates whether a read operation or a write operation is to be performed. Logic 1210 of AMMI SIMM 1200 translates the RAS#, CAS# and WRITE# signals provided to AMMI SIMM 1200 into appropriately timed output enable (OE#) and write enable (WE#) signals for AMMI flash chip 900.

Although not all flash memory chips require an external programming voltage (Vpp) supply, in the embodiment of Figure 12, AMMI flash chip 900 requires an external Vpp (e.g. 12 Volts). Vpp is not provided to a SIMM as part of the standard SIMM interface. System operating voltage Vcc (e.g. 3.3 or 5 Volts), however, is provided to a SIMM as part of a standard SIMM interface. Therefore, in the embodiment of Figure 12, Vpp converter 1220 is provided as part of AMMI SIMM 1200 to take the Vcc system operating voltage and convert it to a Vpp for AMMI flash chip 900. Suitable Vpp voltage converters 1220 are well known in the art and are commercially available. These voltage converters 1220 typically operate as DC-to-DC converters or as voltage pumps.

The WP# input of flash chip 900 is tied to the flash memory Vcc input. Vcc is also provided directly to AMMI chip 900.

Finally, the PWD# input for AMMI flash memory 900 is coupled to an on-SIMM programming voltage monitoring device 1230 such as the Maxim MAX705, manufactured by Maxim Integrated Products, Sunnyvale, California. The Vpp monitor circuit 1230 will hold AMMI flash memory 900 in Deep Powerdown mode unless the supply voltage is within tolerance, thereby protecting AMMI flash memory 900 from unintended writing or erasure.

This scheme allows AMMI SIMM 1200 containing AMMI flash memory 900 to use a Vpp program/erase voltage that is not part of a standard SIMM interface and yet also to plug directly into a SIMM

socket intended for DRAM. Of course, in an alternate embodiment wherein AMMI flash memory chip 900 does not require the external supply of Vpp, AMMI SIMM 1200 will not require Vpp converter 1220. Vpp monitor 1230 would then be used to monitor Vcc.

For an embodiment wherein the SIMM pinout permits the addition of generate purpose input/output signals, one can switch off Vpp when not programming or erasing. One can also control signals RP# to provide power management and WP# to control block lock and unlock.

Figure 13 is a block diagram that illustrates the main memory organization of a computer system that uses AMMI flash SIMMs as main memory. In computer system 1300, microprocessor 1310 includes on-board DRAM controller 1330. DRAM controller 1330, in turn, is coupled to DRAM SIMM 1360 and AMMI flash SIMM 1350 by main memory bus 1320. DRAM SIMM 1360 and AMMI flash SIMM 1350 together form part of the main memory for microprocessor 1310.

When microprocessor 1310 is to access code or data stored in DRAM SIMM 1360 or AMMI flash SIMM 1350, DRAM controller 1330 of microprocessor 1310 will place an access request on main memory bus 1320. The access request will be made to DRAM SIMM 1360 or AMMI flash SIMM 1350 according to the asynchronous protocol of main memory bus 1320. Because AMMI flash SIMM 1350 contains one or more AMMI flash devices that are emulating an asynchronous DRAM device, DRAM controller 1330 will be able to access AMMI flash SIMM 1350 using a DRAM access protocol (i.e. by strobing row and column addresses). DRAM controller 1330 will also generate DRAM refresh cycles and provide precharge delays for AMMI flash SIMM 1350 even though a flash device in AMMI flash SIMM 1350 would not need them.

Furthermore, in computer system 1300, microprocessor 1310 is coupled to DRAM controller 1335 by high speed synchronous local bus 1325. DRAM controller 1335, in turn, is coupled to DRAM SIMM 1365 and AMMI flash SIMM 1355 by main memory bus 1345. DRAM

Micron Technology Inc. et al.
Ex. 1010, p. 520

SIMM 1365 and AMMI flash SIMM 1355 together form part of the main memory for microprocessor 1310.

When microprocessor 1310 is to access code or data stored in DRAM SIMM 1365 or AMMI flash SIMM 1355, microprocessor 1310 will place an access request on local bus 1325. The access request will be made according to the synchronous protocol of local bus 1325. DRAM controller 1335 will then interpret the access request and then make the access request to DRAM SIMM 1365 or AMMI flash SIMM 1355 using the protocol of main memory bus 1345. Because AMMI flash SIMM 1355 contains one or more AMMI flash devices that are emulating an asynchronous DRAM device, DRAM controller 1330 will be able to access AMMI flash SIMM 1355 using a DRAM access protocol (i.e. by strobing row and column addresses). DRAM controller 1335 will also generate DRAM refresh cycles and provide a precharge delay for AMMI flash SIMM 1350 even though a flash device in AMMI flash SIMM 1355 would not need them.

Secondary cache 1375 is also coupled to local bus 1325. If microprocessor 1310 attempts to access data or instructions from DRAM SIMMs 1360 or 1365 or AMMI flash SIMMs 1350 or 1355 that are available in cache 1375, the access will be made to cache 1375 instead.

Note that alternate embodiments of computer system 1300 exist wherein computer system 1300 uses a subset of the main memory capability depicted in Figure 13. Therefore, in one embodiment, microprocessor 1300 will address all of its main memory through main memory bus 1320. Alternately, microprocessor 1300 will address all of its main memory through local bus 1325. In such a case, microprocessor 1310 need not include on-board DRAM controller 1330.

In yet another embodiment, secondary cache 1375 is not available in system 1300. Furthermore, alternate embodiments exist wherein all of the main memory controlled by DRAM controller 1330 or DRAM controller 1535, is comprised of AMMI flash SIMMs 1350 or 1355, respectively. Moreover, in yet another embodiment, rather than

being coupled directly to microprocessor 1310, local bus 1325 is coupled to main memory bus 1320 by a bus bridge.

Synchronous Main Memory Interface (SMMI) for Flash Main Memory

As processors have become faster, asynchronous DRAM has become increasingly unable to supply data and instructions to these processors at a speed fast enough to meet the processing speed sustainable by the processors. One recent solution has been the introduction of synchronous DRAM. Synchronous DRAM uses an external clock signal to synchronize interleaving within a DRAM integrated circuit and thereby provide a faster burst access than could be provided asynchronously.

Figure 14 illustrates a block diagram of a flash memory integrated circuit that is able to emulate a synchronous DRAM integrated circuit chip. In the embodiment depicted in Figure 14, synchronous main memory interface (SMMI) flash memory unit 1400 combines the internal flash bank interlacing features of SFI flash 600 (Figure 6) with the row and column address multiplexing of AMMI flash memory 900 (Figure 9).

In the embodiment of SMMI flash 1400 depicted in Figure 14, SMMI flash 1400 functions in one of four modes as determined by the setting of mode register 1480. In the first mode, asynchronous flash mode, the flash memory 1400 is read as a standard flash memory. In this mode, the reading of the contents of a first address must be completed before a second address to be read can be specified.

In the second mode, synchronous flash mode, a clock signal CLK is provided to SMMI flash chip 1400 and a series of addresses belonging to a data burst are specified, one address per clock tick. Then, the contents stored at the addresses specified for the burst are output sequentially during subsequent clock ticks in the order in which the addresses were provided. Alternately, if a single address is provided to SMMI flash chip 1400 when it is in the synchronous mode, the subsequent addresses for the burst will be generated within

Micron Technology Inc. et al.
Ex. 1010, p. 522

SMMI flash chip 1400 and the data burst will then be provided as output from the flash chip.

In the third mode, asynchronous DRAM (dynamic random access memory) mode, SMMI flash memory 1400 emulates asynchronous DRAM. Thus, row and column addresses are strobed into flash memory 1400 using row (RAS#) and column (CAS#) address strobe signals. SMMI flash memory 1400 then converts the row and column addresses internally into a single address and provides as output the data stored at that single address. Furthermore, although the SMMI flash memory 1400 does not need an extended precharge period or to be refreshed, when in the asynchronous DRAM mode, the flash memory 1400 responds to precharge periods and refresh cycles as would an asynchronous DRAM. Therefore, when in the asynchronous DRAM mode, SMMI flash memory 1400 can be controlled by a standard DRAM controller.

Finally, in the fourth mode, synchronous DRAM mode, the features of the second and third modes are combined to yield a flash memory that emulates a synchronous DRAM. Thus, addresses to be read as a data burst are specified by strobing row and column addresses into the flash memory using RAS# and CAS# signals. The data of the data burst is then provided sequentially as output from the flash memory on subsequent clock ticks.

In an alternate embodiment, SMMI flash 1400 only has one mode, synchronous DRAM emulation mode, or has synchronous DRAM mode and only one or two of the other modes.

In the four mode embodiment depicted in Figure 14, mode register 1480 has a multi-bit output mode enable signal ME that reflects the current mode of SMMI flash 1400. Signal ME is provided to bank select logic 1474, bus logic 1440, TTL stages 1472 and logic and ATD control block 1470. These circuits control the function of SMMI flash 1400 differently according to the current mode of SMMI flash device 1400.

Thus, bus logic 1440 will function as does bus logic 640 of Figure 6 when SMMI flash device 1400 is operating in synchronous flash

mode and will be disabled in the other modes. Note that in an alternate embodiment wherein SMMI flash device 1400 is operating as a generic synchronous flash device, bus logic 1440 will not be part of chip 1400, but instead will be provided off chip.

When SMMI flash device 1400 is operating synchronously, in either the synchronous flash or synchronous DRAM emulation modes, TTL stages 1472, bank select logic 1474 and logic and ATD circuit 1470 will control the interlacing of accesses to flash banks A 1410 and B 1420. In an alternate embodiment, interlacing will be performed among more than two flash banks within SMMI flash device 1400 when it is placed in these synchronous modes.

Similarly, when SMMI flash device 1400 is emulating a DRAM, in either the asynchronous or synchronous DRAM emulation modes, TTL stages 1472, bank select logic 1474 and logic and ATD circuit 1470 will control the multiplexing of row and column addresses by address multiplexers 1492 and 1494 and the tri-stating of output buffer 1460.

A SMMI flash based synchronous SIMM can be formed using the techniques described above in connection with AMMI SIMM 1200 of Figure 12. This would be accomplished by populating AMMI SIMM 1200 with one or more SMMI flash devices 1400 (rather than AMMI flash devices 900) and by providing a clock signal to each SMMI flash device.

Figure 15 is a block diagram that illustrates the main memory organization of a computer system that uses SMMI flash SIMMs as main memory. In computer system 1500, microprocessor 1510 includes on-board synchronous DRAM controller 1530. Synchronous DRAM controller 1530, in turn, is coupled to synchronous DRAM SIMM 1560 and SMMI flash SIMM 1550 by synchronous main memory bus 1520. Synchronous DRAM SIMM 1560 and SMMI flash SIMM 1550 together form part of the main memory for microprocessor 1510.

When microprocessor 1510 is to access code or data stored in synchronous DRAM SIMM 1560 or SMMI flash SIMM 1550, synchronous DRAM controller 1530 of microprocessor 1510 will place

Micron Technology Inc. et al.
Ex. 1010, p. 524

an access request on synchronous main memory bus 1520. The access request will be made to synchronous DRAM SIMM 1560 or SMMI flash SIMM 1550 according to the synchronous protocol of synchronous main memory bus 1520. Because SMMI flash SIMM 1550 contains one or more SMMI flash devices that are emulating a synchronous DRAM device, synchronous DRAM controller 1530 will be able to access SMMI flash SIMM 1550 using a synchronous DRAM access protocol (i.e. by strobing row and column addresses in conjunction with providing a clock signal). DRAM controller 1530 will also generate DRAM refresh cycles and provide precharge delays for SMMI flash SIMM 1550 even though a flash device in SMMI flash SIMM 1550 would not need them.

Furthermore, in computer system 1500, microprocessor 1510 is coupled to synchronous DRAM controller 1535 by high speed synchronous local bus 1525. Synchronous DRAM controller 1535, in turn, is coupled to synchronous DRAM SIMM 1565 and SMMI flash SIMM 1555 by synchronous main memory bus 1545. Synchronous DRAM SIMM 1565 and SMMI flash SIMM 1555 together form part of the main memory for microprocessor 1510.

When microprocessor 1510 is to access code or data stored in synchronous DRAM SIMM 1565 or SMMI flash SIMM 1555, microprocessor 1510 will place an access request on local bus 1525. The access request will be made according to the synchronous protocol of local bus 1525. Synchronous DRAM controller 1535 will then interpret the access request and then make the access request to synchronous DRAM SIMM 1565 or SMMI flash SIMM 1555 using the synchronous protocol of synchronous main memory bus 1545. Because SMMI flash SIMM 1555 contains one or more SMMI flash devices that are emulating a synchronous DRAM device, synchronous DRAM controller 1530 will be able to access SMMI flash SIMM 1555 using a synchronous DRAM access protocol (i.e. by strobing row and column addresses and by providing a clock signal). Synchronous DRAM controller 1535 will also generate synchronous DRAM refresh cycles and provide a precharge delay for SMMI flash

SIMM 1550 even though a flash device in SMMI flash SIMM 1555 would not need them.

Secondary cache 1575 is also coupled to local bus 1525. If microprocessor 1510 attempts to access data or instructions from synchronous DRAM SIMMs 1560 or 1565 or SMMI flash SIMMs 1550 or 1555 that are available in cache 1575, the access will be made to cache 1575 instead.

Note that alternate embodiments of computer system 1500 exist wherein computer system 1500 uses a subset of the main memory capability depicted in Figure 15. Therefore, in one embodiment, microprocessor 1500 will address all of its main memory through main memory bus 1520. Alternately, microprocessor 1500 will address all of its main memory through local bus 1525. In such a case, microprocessor 1510 need not include on-board synchronous DRAM controller 1530.

In yet another embodiment, secondary cache 1575 is not available in system 1500. Furthermore, alternate embodiments exist wherein all of the main memory controlled by synchronous DRAM controller 1530 or synchronous DRAM controller 1535, is comprised of SMMI flash SIMMs 1550 or 1555, respectively. Moreover, in yet another embodiment, rather than being coupled directly to microprocessor 1510, local bus 1525 is coupled to main memory bus 1520 by a bus bridge.

In the foregoing specification the invention has been described with reference to specific exemplary embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative rather than restrictive sense.

CLAIMS

What is claimed is:

1. A computer system comprising:

a microprocessor;

a high speed bus;

a nonvolatile main memory, the non-volatile main memory providing first data to the microprocessor through the high speed bus.

2. The computer system as set forth in claim 1, wherein the non-volatile main memory provides the first data to the microprocessor through the high speed bus with zero microprocessor wait states.

3. The computer system as set forth in claim 1, wherein the nonvolatile main memory is flash memory.

4. The computer system as set forth in claim 1, wherein the high speed bus is a synchronous bus.

5. The computer system as set forth in claim 1, wherein the high speed bus is an asynchronous bus.

6. The computer system as set forth in claim 1, wherein the high speed bus is a local bus.

7. The computer system as set forth in claim 1, wherein the high speed bus is a main memory bus.

8. The computer system as set forth in claim 1, wherein the first data includes program instructions to be executed by the microprocessor.

9. The computer system as set forth in claim 1, wherein the microprocessor has an internal cache, the internal cache storing data for use by the microprocessor.

10. The computer system as set forth in claim 1, wherein the microprocessor is a static microprocessor.

11. The computer system as set forth in claim 1, wherein the nonvolatile main memory can be placed into an energy saving deep powerdown state.

Micron Technology Inc. et al.
Ex. 1010, p. 527

12. The computer system as set forth in claim 11, wherein the microprocessor can order the nonvolatile memory to enter and to leave the deep powerdown state.

13. The computer system as set forth in claim 1, further comprising a volatile main memory, the volatile main memory providing second data to the microprocessor through the high speed bus.

14. The computer system as set forth in claim 13, wherein the second data includes program instructions to be executed by the microprocessor.

15. The computer system as set forth in claim 13, wherein the volatile main memory is battery backed.

16. The computer system as set forth in claim 13, wherein the volatile main memory is static random access memory.

17. The computer system as set forth in claim 13, wherein the volatile main memory is dynamic random access memory.

18. The computer system as set forth in claim 13, wherein the volatile main memory operates as a write cache for third data to be stored in the nonvolatile main memory.

19. The computer system as set forth in claim 13, wherein the volatile main memory operates as an external cache for the first data stored in the nonvolatile main memory.

20. The computer system as set forth in claim 1, wherein the computer system further comprises nonvolatile store memory coupled to the nonvolatile main memory.

21. The computer system as set forth in claim 1, wherein the computer system further comprises:

a low speed bus;

a bus bridge coupling the high and low speed buses; and

a nonvolatile store memory coupled to the low speed bus.

22. The computer system as set forth in claim 1, wherein the first data has a first part and a second part, the nonvolatile main memory further comprising:

a plurality of nonvolatile memory units, the plurality of nonvolatile memory units including first and second memory units, the first part of the first data being stored in the first nonvolatile memory unit and the second part of the first data being stored in the second nonvolatile memory unit;

an interlace control unit, the interlace control unit controlling access by the microprocessor to the plurality of nonvolatile memory units, if the first data is to be provided to the microprocessor, the interlace controller causing the first nonvolatile memory unit to provide the first part of the first data to the microprocessor and the interlace controller causing the second nonvolatile memory unit to begin retrieving the second part of the first data while the first nonvolatile memory unit is providing the first part of the first data to the microprocessor, the interlace controller then causing the second nonvolatile memory unit to provide the second part of the first data to the microprocessor after the first nonvolatile memory unit has provided the first part of the first data to the microprocessor.

23. The computer system as set forth in claim 22, further comprising a clock, the clock providing a clock signal to the interlace controller, the interlace controller causing the first nonvolatile memory unit to provide the first part of the first data to the microprocessor on a first pulse of the clock and the interlace controller causing the second nonvolatile memory unit to provide the second part of the first data to the microprocessor on a second pulse of the clock.

24. The computer system as set forth in claim 23, wherein the first part of the first data has an associated first address and the microprocessor requests that the first data be provided to the microprocessor by providing the first address to the interlace controller on a third pulse of the clock, the third pulse occurring before the first and second pulses of the clock.

25. The computer system as set forth in claim 24, wherein an address latch enable signal is provided to the interlace controller on

the third pulse to indicate that the first address is being provided to the interlace controller.

26. The computer system as set forth in claim 24 wherein the second part of the first data has an associated second address and the interlace controller calculates the second address based upon the first address.

27. The computer system as set forth in claim 26, wherein the interlace controller calculates the second address by incrementing the first address by a predetermined amount.

28. The computer system as set forth in claim 27, wherein the predetermined amount is one.

29. The computer system as set forth in claim 26, wherein the interlace controller calculates the second address by decrementing the first address by a predetermined amount.

30. The computer system as set forth in claim 29, wherein the predetermined amount is one.

31. The computer system as set forth in claim 24, wherein the second part of the first data has an associated second address and a request that the second part of the first data be provided to the microprocessor is made by providing the second address to the interlace controller on a fourth pulse of the clock, the fourth pulse occurring after the third pulse of the clock.

32. The computer system as set forth in claim 31, wherein an address latch enable signal is provided to the interlace controller on the fourth pulse to indicate that the second address is being provided to the interlace controller.

33. The computer system as set forth in claim 22, wherein the first and second parts of the first data belong to a single cache line.

34. The computer system as set forth in claim 22, wherein the first nonvolatile memory unit is a first integrated circuit and the second nonvolatile memory units is a second integrated circuit.

35. The computer system as set forth in claim 34, wherein the interlace controller is a third integrated circuit.

36. The computer system as set forth in claim 22, wherein the first and second nonvolatile memory units are separate banks of a first integrated circuit.

37. The computer system as set forth in claim 36, wherein the interlace controller is a second integrated circuit.

38. The computer system as set forth in claim 36, wherein the interlace controller is part of the first integrated circuit.

39. The computer system as set forth in claim 22, wherein the first part of the first data has an associated first row address and first column address and wherein the microprocessor initiates access to the first part of the first data by providing the first row and column addresses to the interlace controller.

40. The computer system as set forth in claim 39, wherein the first row address is provided during a row address strobe pulse and the first column address is provided during a column address strobe pulse.

41. The computer system as set forth in claim 39, wherein the second part of the first data has an associated second row address and second column address and wherein the microprocessor initiates access to the second part of the first data by providing the second row and column addresses to the interlace controller.

42. The computer system as set forth in claim 41, wherein the second row address is provided during a row address strobe pulse and the second column address is provided during a column address strobe pulse.

43. The computer system as set forth in claim 22, wherein the first data has an associated row address and column address and wherein the microprocessor initiates access to the first data by providing the row and column addresses to the nonvolatile main memory.

44. The computer system as set forth in claim 43, wherein the row address is provided during a row address strobe pulse and the column address is provided during a column address strobe pulse.

45. A memory device comprising:

a plurality of nonvolatile memory banks including first and second banks;

an interlace controller, the interlace controller controlling access to the plurality of nonvolatile memory banks, the interlace controller causing the first bank to provide first data and causing the second bank to begin retrieving second data while the first bank is providing the first data, the interlace controller then causing the second bank to provide the second data after the first bank has provided the first data.

46. The memory device as set forth in claim 45, wherein a clock signal is provided to the interlace controller, the interlace controller causing the first bank to provide the first data on a first pulse of the clock and the interlace controller causing the second bank to provide the second data on a second pulse of the clock.

47. The memory device as set forth in claim 46, wherein the first data has an associated first address specified to the interlace controller on a third pulse of the clock, the third pulse occurring before the first and second pulses of the clock.

48. The memory device as set forth in claim 47, wherein an address latch enable signal is provided to the interlace controller on the third pulse to indicate that the first address is being provided to the interlace controller.

49. The memory device as set forth in claim 47 wherein the second data has an associated second address and the interlace controller calculates the second address based upon the first address.

50. The memory device as set forth in claim 49, wherein the interlace controller calculates the second address by incrementing the first address by a predetermined amount.

51. The memory device as set forth in claim 50, wherein the predetermined amount is one.

52. The memory device as set forth in claim 49, wherein the interlace controller calculates the second address by decrementing the first address by a predetermined amount.

53. The memory device as set forth in claim 52, wherein the predetermined amount is one.

54. The memory device as set forth in claim 47, wherein the second data has an associated second address specified to the interlace controller on a fourth pulse of the clock, the fourth pulse occurring after the third pulse of the clock.

55. The memory device as set forth in claim 54, wherein an address latch enable signal is provided to the interlace controller on the fourth pulse to indicate that the second address is being provided to the interlace controller.

56. The memory device as set forth in claim 45, wherein the first and second data belong to a single cache line.

57. The memory device as set forth in claim 45, wherein the interlace controller can enter a noninterlace mode, when in the noninterlace mode, the interlace controller will not permit a second access to the plurality of nonvolatile memory banks to begin until a first access to the plurality of nonvolatile memory banks has completed.

58. The memory device as set forth in claim 57, wherein the interlace controller can enter an address multiplexing mode, when in the address multiplexing mode, a particular portion of one of the plurality of nonvolatile memory banks is specified to the interlace controller by a row and column address pair.

59. The memory device as set forth in claim 58, wherein the row address is provided during a row address strobe pulse and the column address is provided during a column address strobe pulse.

60. The memory device as set forth in claim 58, wherein the interlace controller combines the row and column address pair to form a bank address, the interlace controller then uses the bank address to specify the particular portion of one of the plurality of nonvolatile memory banks to access.

61. The memory device as set forth in claim 58, wherein the first data has an associated first row address and first column address

pair and wherein the interlace controller initiates access to the first data upon receiving the first row and column addresses.

62. The memory device as set forth in claim 61, wherein the first row address is received during a row address strobe pulse and the first column address is received during a column address strobe pulse.

63. The memory device as set forth in claim 61, wherein the interlace controller combines the first row and column address pair to form a first bank address, the interlace controller then uses the bank address to specify that the first data of the first bank is to be accessed.

64. The memory device as set forth in claim 61, wherein the second data has an associated second row address and second column address pair and wherein the interlace controller initiates access to the second data by upon receiving the second row and column addresses.

65. The memory device as set forth in claim 64, wherein the second row address is received during a row address strobe pulse and the second column address is received during a column address strobe pulse.

66. The memory device as set forth in claim 65, wherein the interlace controller combines the second row and column address pair to form a second bank address, the interlace controller then uses the second bank address to specify the second data of the second bank is to be accessed.

67. The memory device as set forth in claim 45, wherein the interlace controller can enter an address multiplexing mode, when in the address multiplexing mode, a particular portion of one of the plurality of nonvolatile memory banks is specified to the interlace controller by a row and column address pair.

68. The memory device as set forth in claim 67, wherein the row address is provided during a row address strobe pulse and the column address is provided during a column address strobe pulse.

69. The memory device as set forth in claim 68, wherein the interlace controller combines the row and column address pair to form a bank address, the interlace controller then uses the bank

address to specify the particular portion of one of the plurality of nonvolatile memory banks to access.

70. The memory device as set forth in claim 45, wherein the first data is identified by a first address and wherein the memory device further comprises a first address latch, the interlace controller controlling the first address latch and causing the first address latch to latch the first address prior to the interlace controller causing the first bank to provide the first data.

71. The memory device as set forth in claim 70, wherein the second data is identified by a second address and wherein the memory device further comprises a second address latch, the interlace controller controlling the second address latch and causing the second address latch to latch the second address prior to the interlace controller causing the second bank to provide the second data.

72. The memory device as set forth in claim 45, wherein the first data is identified by a first address and wherein the interlace controller further comprises an address transition detector, when the address transition detector detects an address transition to an address of the first bank, the interlace controller causes the first bank to prepare to provide data.

73. The memory device as set forth in claim 72, wherein the second data is identified by a second address and wherein when the address transition detector detects an address transition to an address of the second bank, the interlace controller causes the second bank to prepare to provide data.

74. The memory device as set forth in claim 45, further comprising:

an output buffer, the output buffer storing output data; and

a bank selector, the bank selector selecting the first bank to provide to the output buffer the first data as the output data when the first bank is providing the first data and the bank selector selecting the second bank to provide to the output buffer the second data as the output data when the second bank is providing the second data.

75. The memory device as set forth in claim 45, further comprising bus specific logic, the bus specific logic receiving bus specific input signals and providing bus specific output signals in response.

76. The memory device as set forth in claim 45, further comprising a parity generator, the parity generator generating a parity signal based upon the first data when the first data is provided.

77. The memory device as set forth in claim 45, wherein the parity generator generates a parity signal based upon the second data when the second data is provided.

78. The memory device as set forth in claim 45, wherein at least one of the plurality of nonvolatile memory banks is comprised of flash memory.

79. A memory device, comprising:

a nonvolatile memory bank;

an address demultiplexer, the address demultiplexer receiving a multiplexed row and column address pair and combining the row and column address to form a bank address, the bank address being used to access the memory bank.

80. The memory device as set forth in claim 79, wherein the nonvolatile memory bank is a flash memory bank.

81. The memory device as set forth in claim 79, wherein the nonvolatile memory bank is a flash memory and the memory device emulates a dynamic random access memory (DRAM) device.

82. The memory device as set forth in claim 79, wherein the address demultiplexer comprises:

an address latch, the address latch having a row address latch and a column address latch;

an address multiplexer, when the address multiplexer receives the row address, the address multiplexer provides the row address to the row address latch to be latched therein, when the address multiplexer receives the column address, the address multiplexer provides the column address to the column address latch to be latched therein;

once the row and column addresses have been latched in the address latch, the address latch provides the row and column addresses to the memory bank as the bank address used to access the memory bank.

83. The memory device as set forth in claim 79, wherein the row address is provided during a row address strobe pulse and the column address is provided during a column address strobe pulse.

84. The memory device as set forth in claim 83, wherein the address demultiplexer comprises:

an address latch, the address latch having a row address latch and a column address latch;

an address multiplexer, during the row address strobe pulse, the address multiplexer receives the row address and provides the row address to the row address latch to be latched therein, and during the column address strobe pulse, the address multiplexer receives the column address and provides the column address to the column address latch to be latched therein;

once the row and column addresses have been latched in the address latch, the address latch provides the row and column addresses to the memory bank as the bank address used to access the memory bank.

85. The memory device as set forth in claim 79, wherein the address demultiplexer can enter a nonmultiplexed addressing mode, when in the nonmultiplexed addressing mode, the address demultiplexer receives a bank address and uses the bank address to access the memory bank.

86. The memory device as set forth in claim 79, wherein the nonvolatile memory bank does not require refreshing and wherein the address demultiplexer further comprises a refresh cycle detector to detect a refresh cycle.

87. The memory device as set forth in claim 86, wherein the memory device further comprises a tri-statable latch coupled to the memory bank and to the refresh cycle detector, the latch adapted to hold data for an access to the memory bank, and further wherein if

the refresh cycle detector detects a refresh cycle, the refresh cycle detector causes the latch to tri-state during the refresh cycle.

88. A memory sub-system comprising:

at least two nonvolatile memory devices, each memory device for storing a corresponding portion of a page of data;

an interlace controller coupled to the nonvolatile memory devices and controlling an interlaced read of the page of data, if the interlace controller detects a read operation to the page, the interlace controller enables the memory devices and causes each memory device to read the corresponding portion of the page, the interlace controller then sequentially enabling each memory device to output the corresponding portion of the page.

89. The memory sub-system as set forth in claim 88 wherein each memory device is a flash memory device.

90. The memory sub-system as set forth in claim 88, wherein there are a plurality of pages of data including a first and second page and wherein the interlace controller further comprises a page transition detector to detect a transition from a read operation for the first page to a read operation for the second page and if the transition is detected, the page transition detector causes the interlace controller to control the reading of the second page from the memory devices.

91. The memory sub-system as set forth in claim 90, wherein the page transition detector detects every page transition from one page of the plurality of pages to another page of the plurality of pages and if the transition is detected, the page transition detector causes the interlace controller to control the reading from the memory devices of the page to which the transition was made.

92. The memory sub-system as set forth in claim 88, wherein the interlace controller receives a clock signal comprising a series of clock pulses and wherein the interlace controller sequentially enables each memory device to output the corresponding portion of the page during a different one of the pulses of the clock signal.

93. The memory sub-system as set forth in claim 92, wherein each of the portions of the page has a corresponding address, wherein

the interlace controller receives a different address of the corresponding addresses during a corresponding different one of the clock cycles and wherein the interlace controller then causes the each of the portions of the page to be output in the order in which the addresses were received.

94. The memory sub-system as set forth in claim 93, wherein the interlace controller receives an address latch enable pulse when the different address of the corresponding addresses is to be received.

95. The memory sub-system as set forth in claim 93, wherein if the interlace controller does not receive addresses for every portion of the page, the interlace controller will only cause the portions of the page to be read that correspond to the addresses received.

96. A system comprising:

a high speed synchronous bus having a data burst access bus protocol;

a nonvolatile memory device having at least two banks of nonvolatile memory and a generic interlace controller to control interlaced accesses to the banks according to a generic data burst access memory protocol;

a glue logic device coupled between the high speed synchronous bus and the nonvolatile memory device, the glue logic receiving the burst access bus protocol from the high speed bus of a burst access to the nonvolatile memory, converting the burst access bus protocol from the high speed bus to the generic data burst access memory protocol, and then providing the converted generic data burst access memory protocol to the memory device.

97. The system as set forth in claim 96, wherein at least one of the banks of nonvolatile memory is flash memory.

98. The system as set forth in claim 96, wherein the data burst access bus protocol uses a linear addressing sequence.

99. The system as set forth in claim 96, wherein the data burst access bus protocol uses a non-linear addressing sequence.

100. The system as set forth in claim 96, wherein the glue logic device further comprising bus specific logic, the bus specific logic

receiving bus specific input signals from the bus and providing bus specific output signals to the bus in response.

101. The system as set forth in claim 96, further comprising a parity generator, the parity generator generating a parity signal based upon the data burst being accessed.

102. A memory module comprising:

a nonvolatile memory;

a memory module board upon which the nonvolatile memory is mounted, the memory module board, if the memory module board provides the nonvolatile memory with a multiplexed row and column address pair indicating a memory location within the nonvolatile memory to be accessed, the nonvolatile memory accessing the memory location indicated.

103. The memory module set forth in claim 102, wherein the nonvolatile memory is comprised of at least one flash memory device.

104. The memory module as set forth in claim 102, wherein the nonvolatile memory emulates dynamic random access memory.

105. The memory module as set forth in claim 102, wherein the nonvolatile memory emulates asynchronous dynamic random access memory.

106. The memory module as set forth in claim 102, wherein the nonvolatile memory emulates synchronous dynamic random access memory.

107. The memory module as set forth in claim 102, further comprising a voltage converter coupled to the nonvolatile memory and to the memory module board, the voltage converter converting an operating voltage from the memory module board into a programming voltage and providing the programming voltage to the nonvolatile memory.

108. The memory module as set forth in claim 107, wherein the nonvolatile memory can be disabled and wherein the memory module further comprises a programming voltage monitor coupled to the nonvolatile memory and to the voltage converter, the voltage

monitor monitoring the programming voltage from the voltage converter programming voltage and if the programming voltage provided to the nonvolatile memory is not within a predetermined voltage range, the voltage monitor disabling the nonvolatile memory.

109. The memory module as set forth in claim 108, wherein the nonvolatile memory is disabled by entering a deep powerdown mode.

110. The memory module as set forth in claim 102, wherein the memory module is a single in-line memory module.

111. A system comprising:

a nonvolatile memory module having a nonvolatile memory device that emulates a dynamic random access memory (DRAM) device; and

a DRAM controller, coupled to the nonvolatile memory module and controlling access to the nonvolatile memory module.

112. The system as set forth in claim 111, wherein the nonvolatile memory device is a flash memory device.

113. The system as set forth in claim 111, wherein the nonvolatile memory module comprises a plurality of nonvolatile memory devices.

114. The system as set forth in claim 111, wherein the nonvolatile memory device emulates an asynchronous DRAM device.

115. The system as set forth in claim 111, wherein the nonvolatile memory device emulates a synchronous DRAM device.

116. The system of claim 111, further comprising a volatile memory module having a volatile memory device, the volatile memory module coupled to, and controlled by, the DRAM controller.

117. The system as set forth in claim 111, further comprising a high speed bus, coupled to the DRAM controller and indicating to the DRAM controller the accesses to the nonvolatile memory to be controlled by the DRAM controller.

118. The system as set forth in claim 117, wherein the high speed bus is an asynchronous bus.

119. The system as set forth in claim 117, wherein the high speed bus is a synchronous bus.

120. The system as set forth in claim 117, wherein the high speed bus is a main memory bus.

121. The system as set forth in claim 117, wherein the high speed bus is a local bus.

122. The system as set forth in claim 117, further comprising a memory cache coupled to the high speed local bus.

123. The system as set forth in claim 117, further comprising a microprocessor coupled to the high speed local bus.

124. A method of accessing a nonvolatile memory device, the method comprising the steps of:

providing a first address to the memory device, the first address specifying a first location within the nonvolatile memory device to be read;

receiving a series of output data from the memory device, the output data from the series of output data including first data stored at the first address.

125. The method as set forth in claim 124, further comprising the step of:

providing a clock signal comprising a series of output clock pulses to the memory device, each output clock pulse signifying to the memory device when a different one of the series of output data is to be received from the memory device.

126. The method as set forth in claim 125, further comprising the step of:

providing an address latch signal to the memory device when the first address is provided to the memory device to indicate to the memory device that the first address is being provided.

127. The method as set forth in claim 124, further comprising the steps of: providing a clock signal to the memory device, the clock signal comprising a series of address input clock pulses followed by a series of output clock pulses, each input clock pulse of the series of input clock pulses being provided with a

Micron Technology Inc. et al.
Ex. 1010, p. 542

corresponding address specifying a location in the memory device to be read, each output clock pulse signifying to the memory device when a different one of the series of output data is to be received from the memory device.

128. The method as set forth in claim 127, further comprising the steps of:

providing an address latch signal to the memory device with each address input clock pulse of the series of address input clock pulses to indicate to the memory device that the address corresponding to the input clock pulse is being provided.

129. The method as set forth in claim 124, wherein the step of providing the first address to the memory device further comprises the steps of:

providing a row address to the memory device;

providing a column address to the memory device, the row and column address together indicating the first address to the memory device.

130. The method as set forth in claim 129, further comprising the step of:

providing a row address strobe to the memory device when providing the row address to the memory device, to indicate to the memory device that the row address is being provided.

131. The method as set forth in claim 130, further comprising the step of:

providing a column address strobe to the memory device when providing the column address to the memory device, to indicate to the memory device that the column address is being provided.

132. The method as set forth in claim 129, further comprising the step of:

prior to providing the row address strobe to the memory device, indicating to the memory device that the memory device is to enter a multiplexed address mode.

133. The method as set forth in claim 124, further comprising the step of:

prior to providing the first address to the memory device, indicating to the memory device that the memory device is to enter a synchronous address mode.

134. A method of accessing a nonvolatile memory device, comprising the steps of:

providing a row address to the memory device;

providing a column address to the memory device;

receiving output data from the memory device, the output data being contents stored by the memory device at a location corresponding to the row and column address provided to the memory device.

135. The method as set forth in claim 134, further comprising the step of providing a row address strobe to the memory device when the row address is provided to the memory device to indicate that the row address is being provided.

136. The method as set forth in claim 134, further comprising the step of providing a column address strobe to the memory device when the column address is provided to the memory device to indicate that the column address is being provided.

137. The method as set forth in claim 134, further comprising the step of:

receiving a series of additional output data from the memory device, the additional output data from the series of output data being received after the output data corresponding to the row and column address provided to the device.

138. The method as set forth in claim 137, further comprising the step of:

providing a clock signal comprising a series of output clock pulses to the memory device, each output clock pulse signifying to the memory device when a different one of the series of output data is to be received from the memory device.

139. The method as set forth in claim 138, further comprising the step of:

prior to providing the row address to the memory device, indicating to the memory device that the memory device is to enter a synchronous address mode.

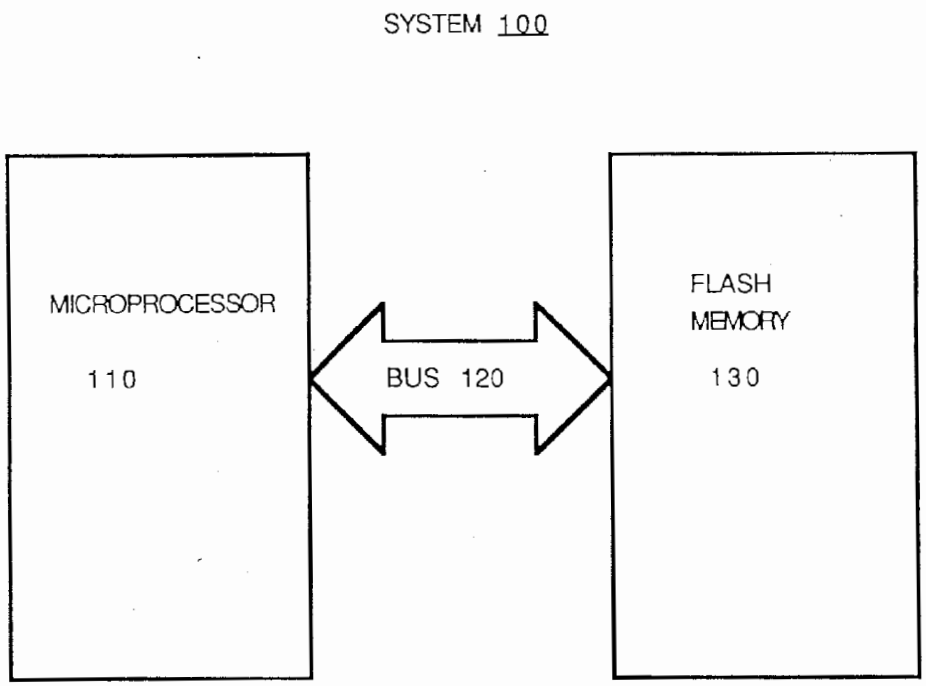


FIGURE 1

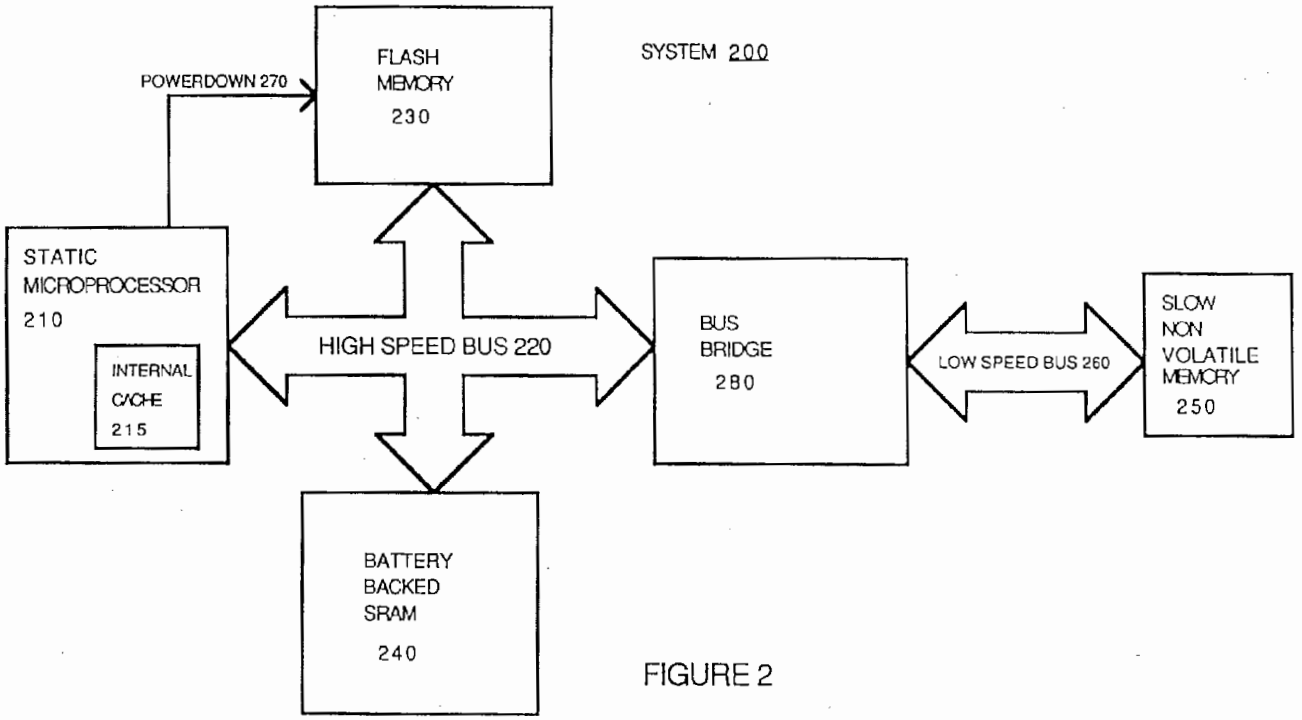

SYSTEM 200
FLASH MEMORY 230
POWERDOWN 270
STATIC MICROPROCESSOR 210
INTERNAL CACHE 215
HIGH SPEED BUS 220
BUS BRIDGE 280
LOW SPEED BUS 260
SLOW NON VOLATILE MEMORY 250
BATTERY BACKED SRAM 240

FIGURE 2

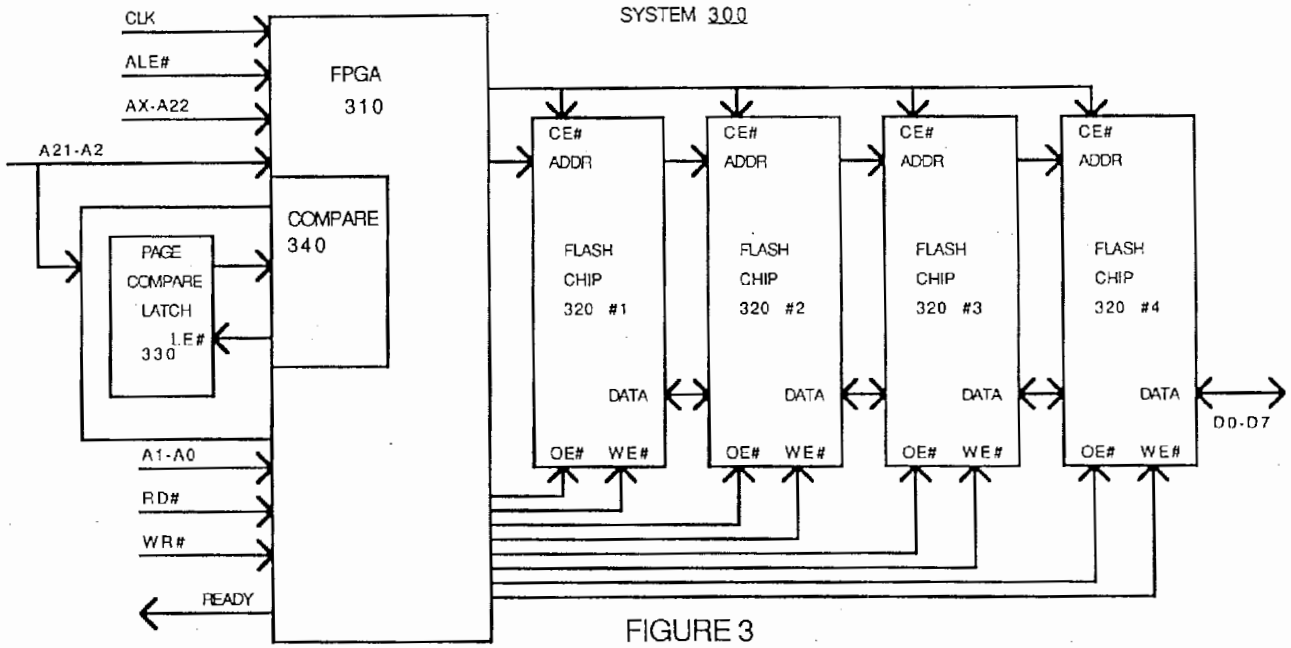

SYSTEM 300
CLK
ALE#
AX-A22
A21-A2
FPGA
310
COMPARE
340
PAGE
COMPARE
LATCH
330
L.E#
A1-A0
RD#
WR#
READY
CE#
ADDR
FLASH
CHIP
320 #1
320 #2
320 #3
320 #4
DATA
OE#
WE#
D0-D7

FIGURE 3

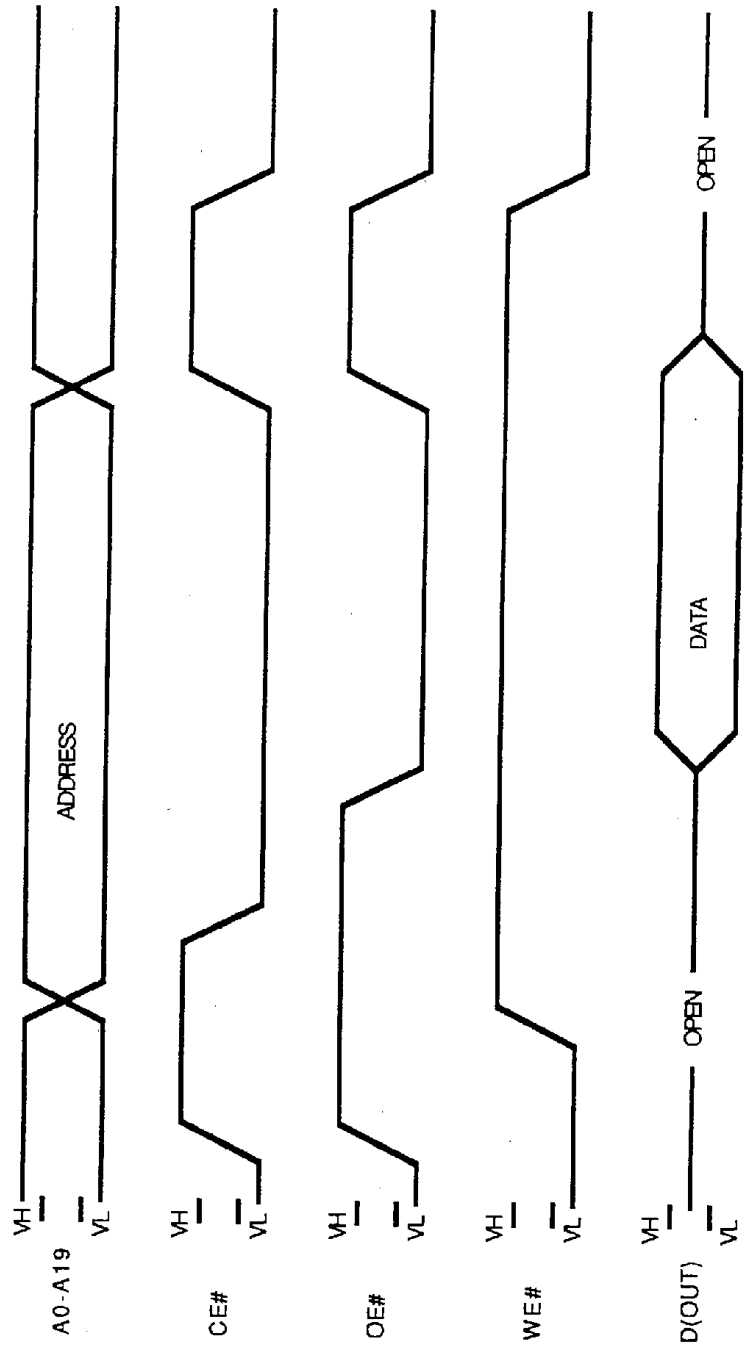


FIGURE 4

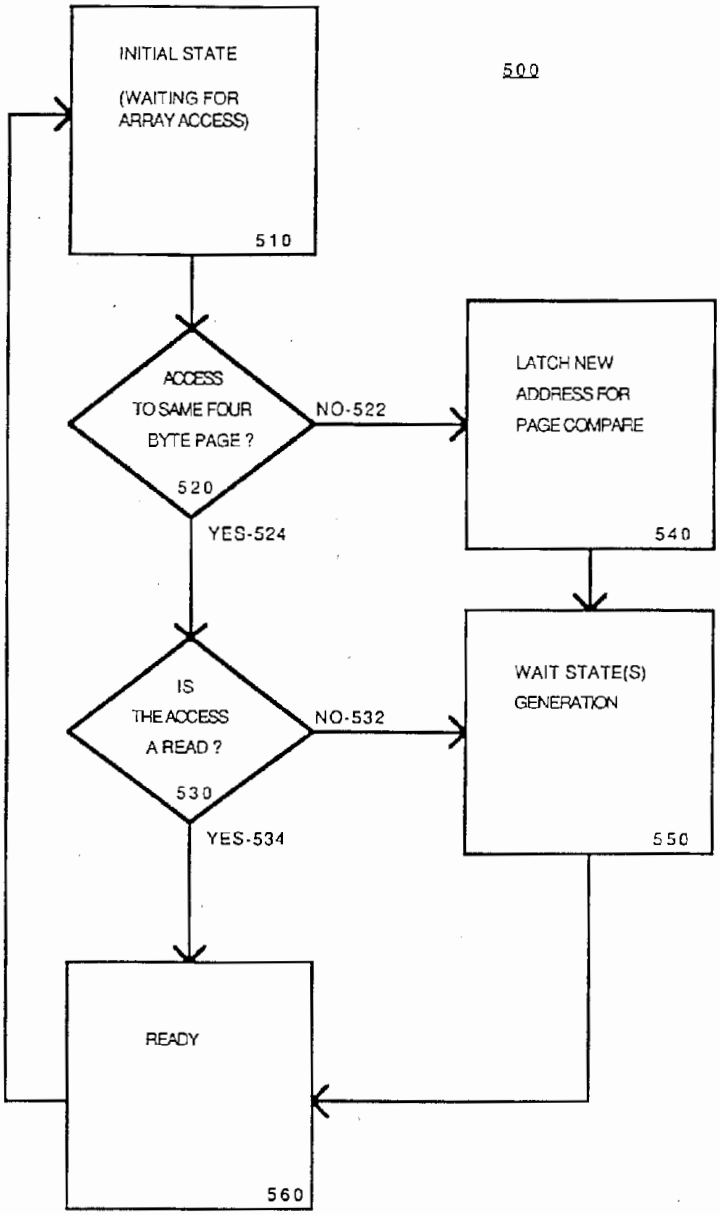

INITIAL STATE
(WAITING FOR ARRAY ACCESS)
510
500
ACCESS TO SAME FOUR BYTE PAGE ?
520
NO-522
LATCH NEW ADDRESS FOR PAGE COMPARE
540
YES-524
IS THE ACCESS A READ ?
530
NO-532
WAIT STATE(S) GENERATION
550
YES-534
READY
560

FIGURE 5

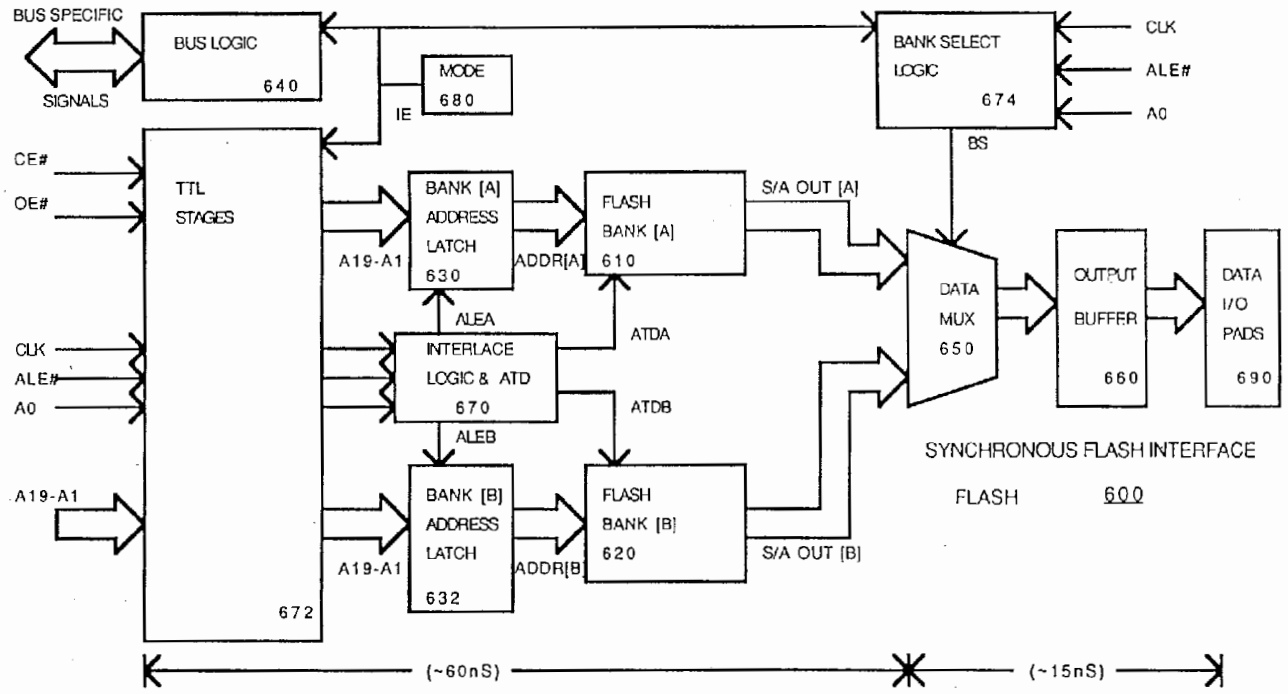

BUS SPECIFIC
SIGNALS
BUS LOGIC
640
MODE
680
IE
BANK SELECT
LOGIC
674
CLK
ALE#
A0
BS
CE#
OE#
TTL
STAGES
BANK [A]
ADDRESS
LATCH
630
A19-A1
ADDR[A]
FLASH
BANK [A]
610
S/A OUT [A]
ALEA
INTERLACE
LOGIC & ATD
670
ATDA
ATDB
ALEB
CLK
ALE#
A0
A19-A1
BANK [B]
ADDRESS
LATCH
632
A19-A1
ADDR[B]
FLASH
BANK [B]
620
S/A OUT [B]
672
DATA
MUX
650
OUTPUT
BUFFER
660
DATA
I/O
PADS
690
SYNCHRONOUS FLASH INTERFACE
FLASH 600
(~60nS)
(~15nS)

FIGURE 6

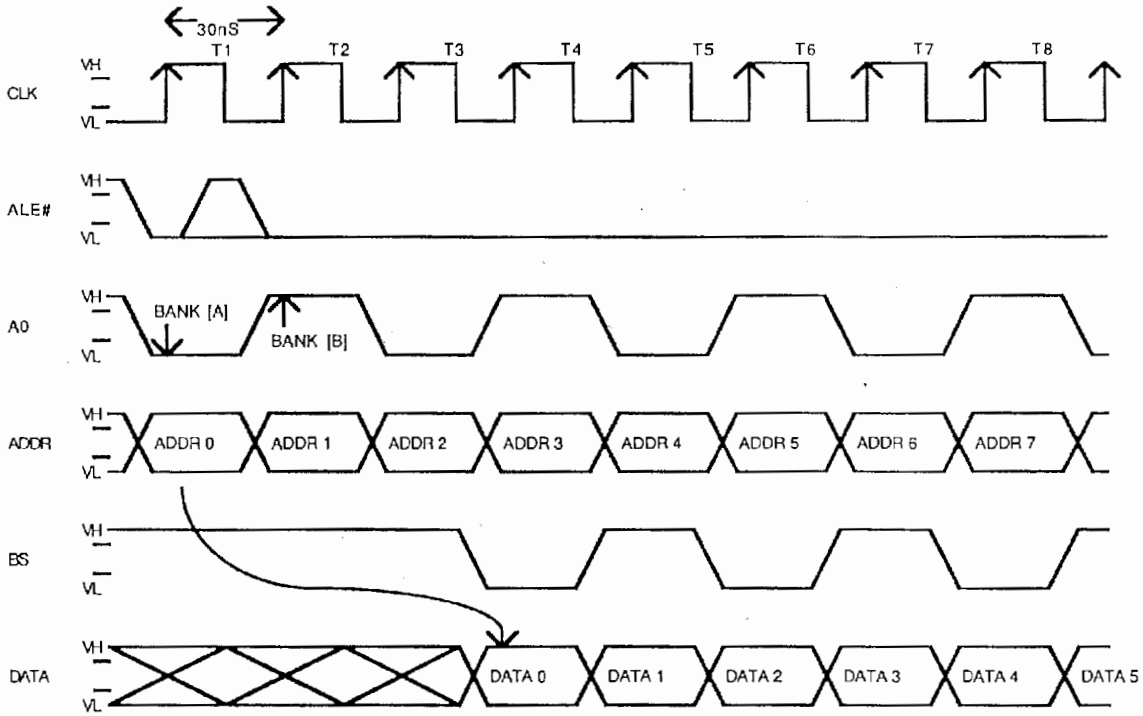


FIGURE 7

Micron Technology Inc. et al.
Ex. 1010, p. 552

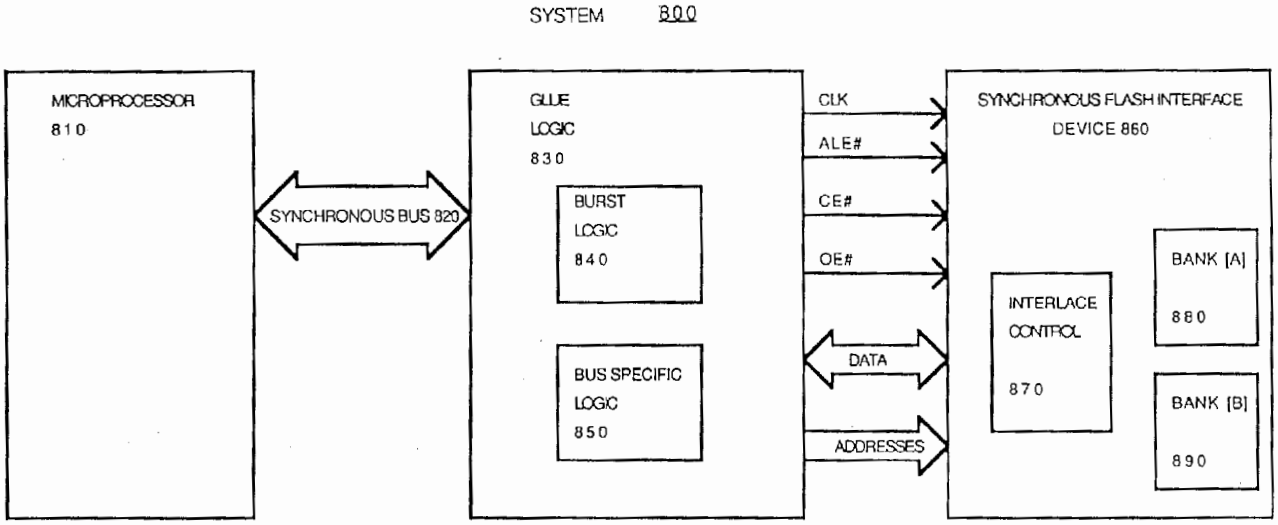

SYSTEM 800
MICROPROCESSOR
810
SYNCHRONOUS BUS 820
GLUE
LOGIC
830
BURST
LOGIC
840
BUS SPECIFIC
LOGIC
850
CLK
ALE#
CE#
OE#
DATA
ADDRESSES
SYNCHRONOUS FLASH INTERFACE
DEVICE 860
INTERLACE
CONTROL
870
BANK [A]
880
BANK [B]
890

FIGURE 8

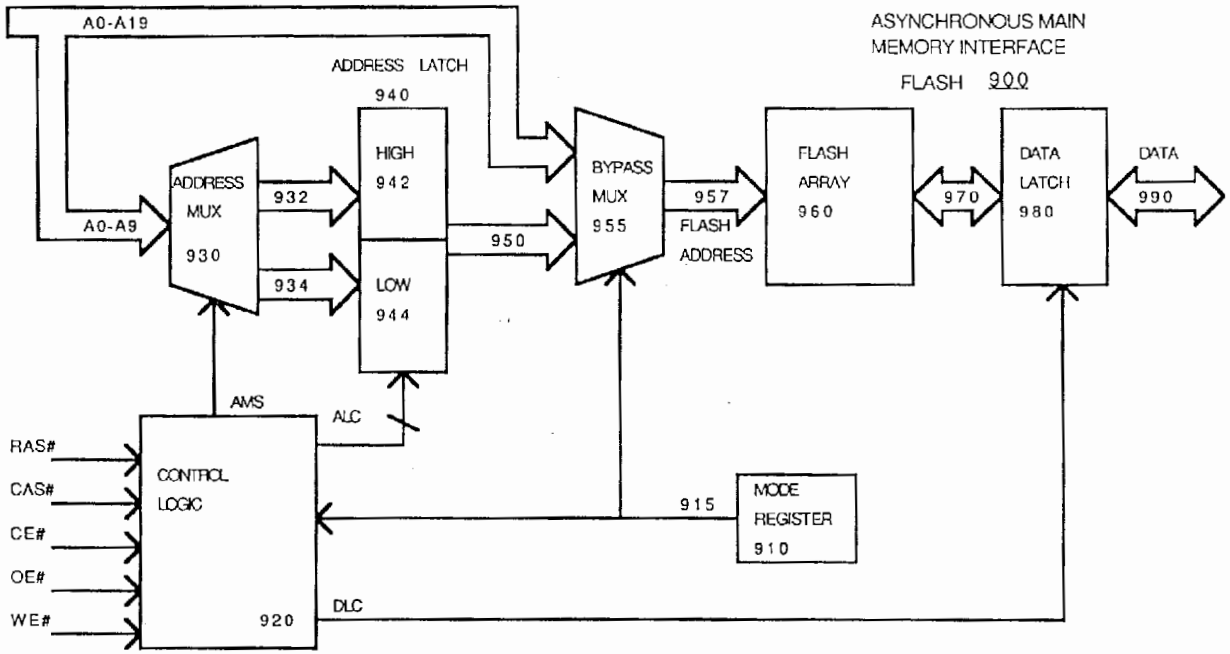

A0-A19
ADDRESS LATCH
940
ASYNCHRONOUS MAIN
MEMORY INTERFACE
FLASH 900
ADDRESS
MUX
930
A0-A9
932
934
HIGH
942
LOW
944
950
BYPASS
MUX
955
957
FLASH
ADDRESS
FLASH
ARRAY
960
970
DATA
LATCH
980
DATA
990
AMS
ALC
RAS#
CAS#
CE#
OE#
WE#
CONTROL
LOGIC
920
915
MODE
REGISTER
910
DLC

FIGURE 9

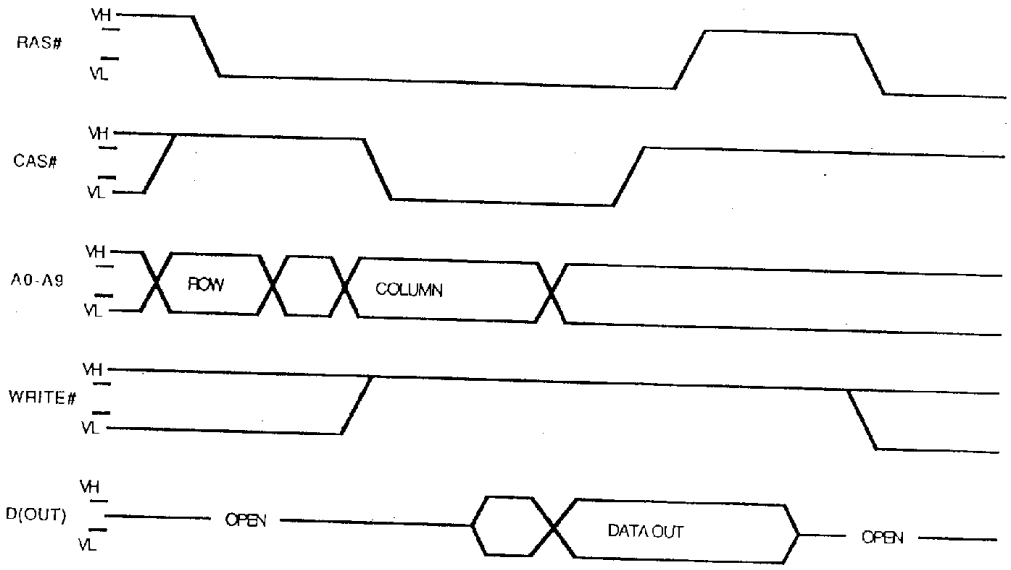


FIGURE 10

Micron Technology Inc. et al.
Ex. 1010, p. 555

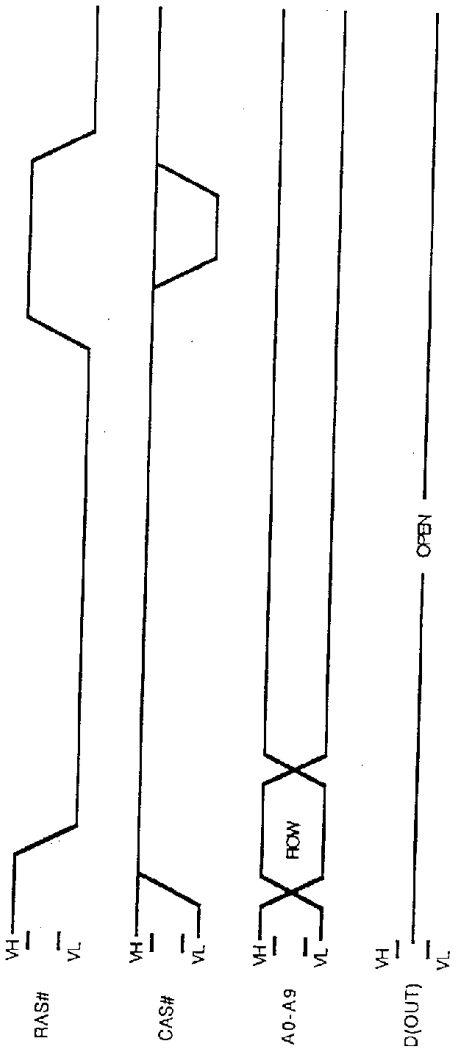
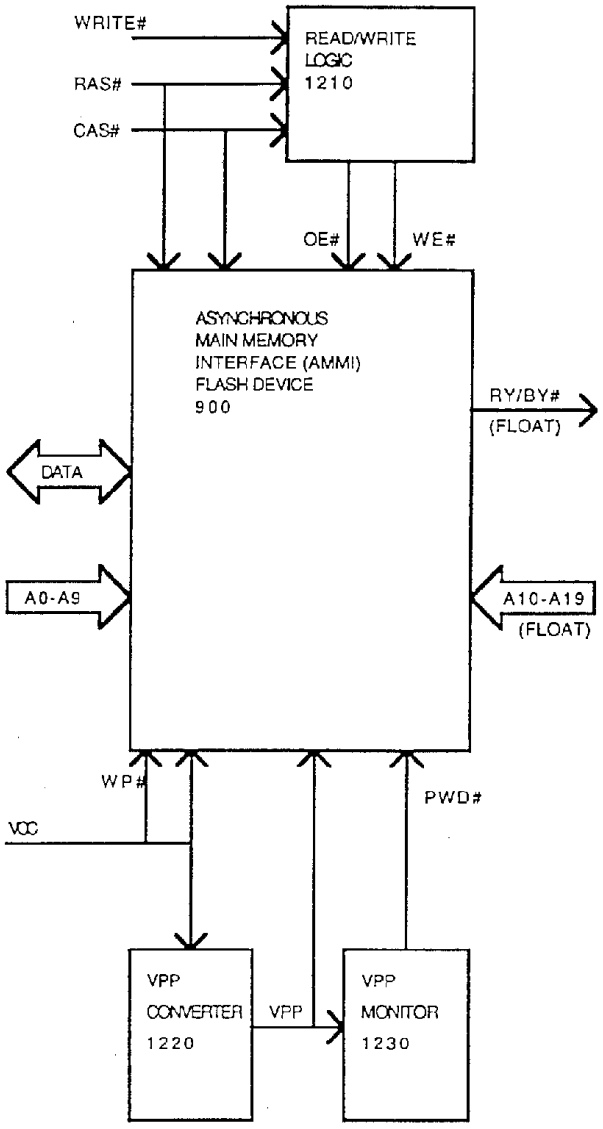


FIGURE 11



AMMI SINGLE IN-LINE
MEMORY MODULE (SIMM)
1200

FIGURE 12

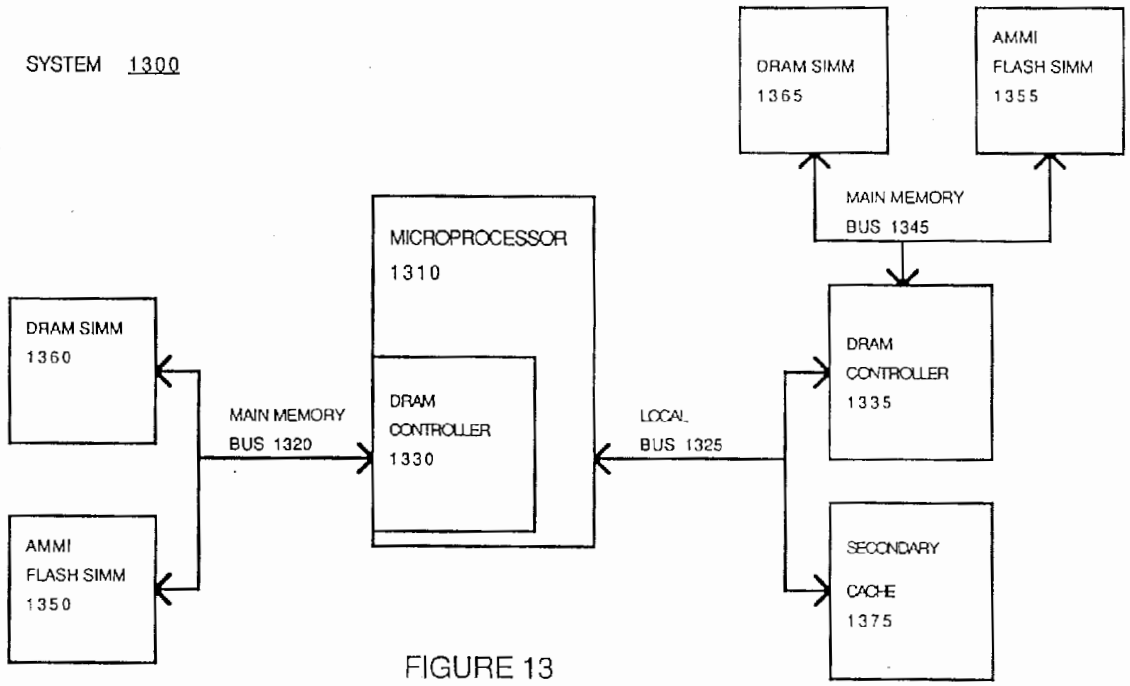

SYSTEM 1300
DRAM SIMM 1365
AMMI FLASH SIMM 1355
MAIN MEMORY BUS 1345
MICROPROCESSOR 1310
DRAM SIMM 1360
MAIN MEMORY BUS 1320
DRAM CONTROLLER 1330
LOCAL BUS 1325
DRAM CONTROLLER 1335
AMMI FLASH SIMM 1350
SECONDARY CACHE 1375

FIGURE 13

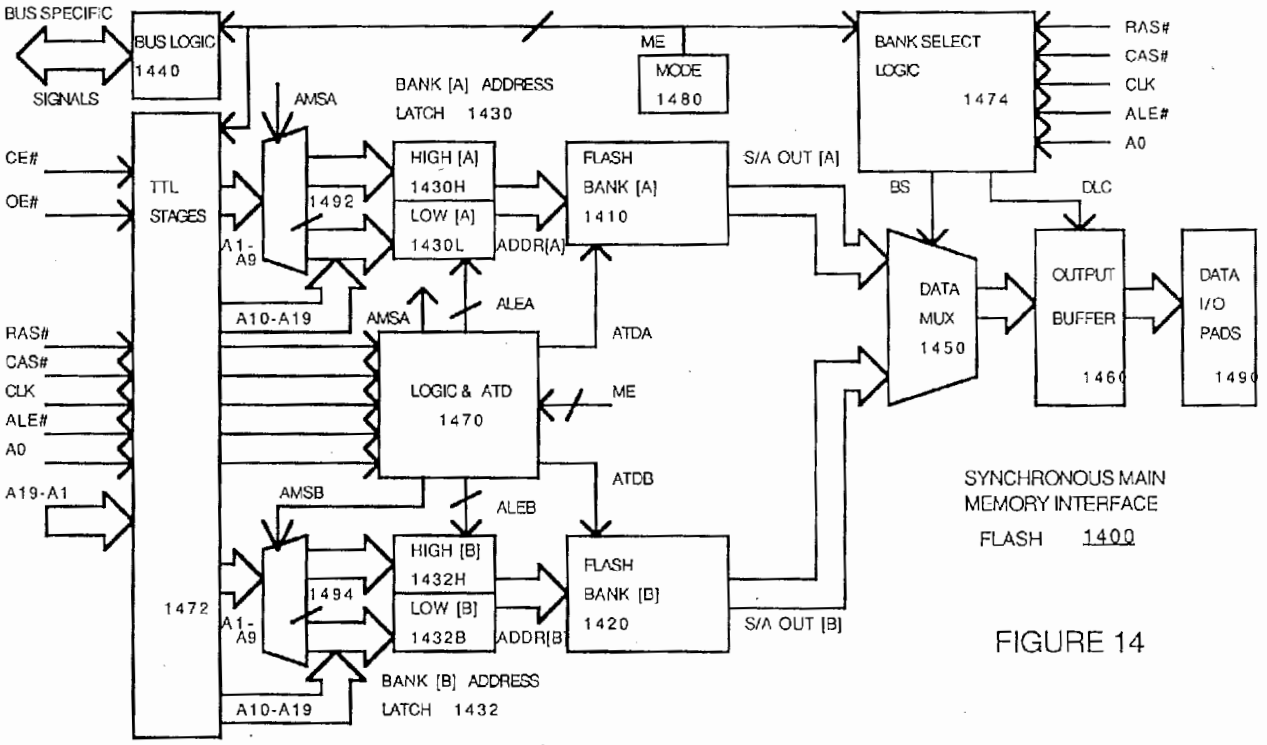


FIGURE 14

Micron Technology Inc. et al.
Ex. 1010, p. 559

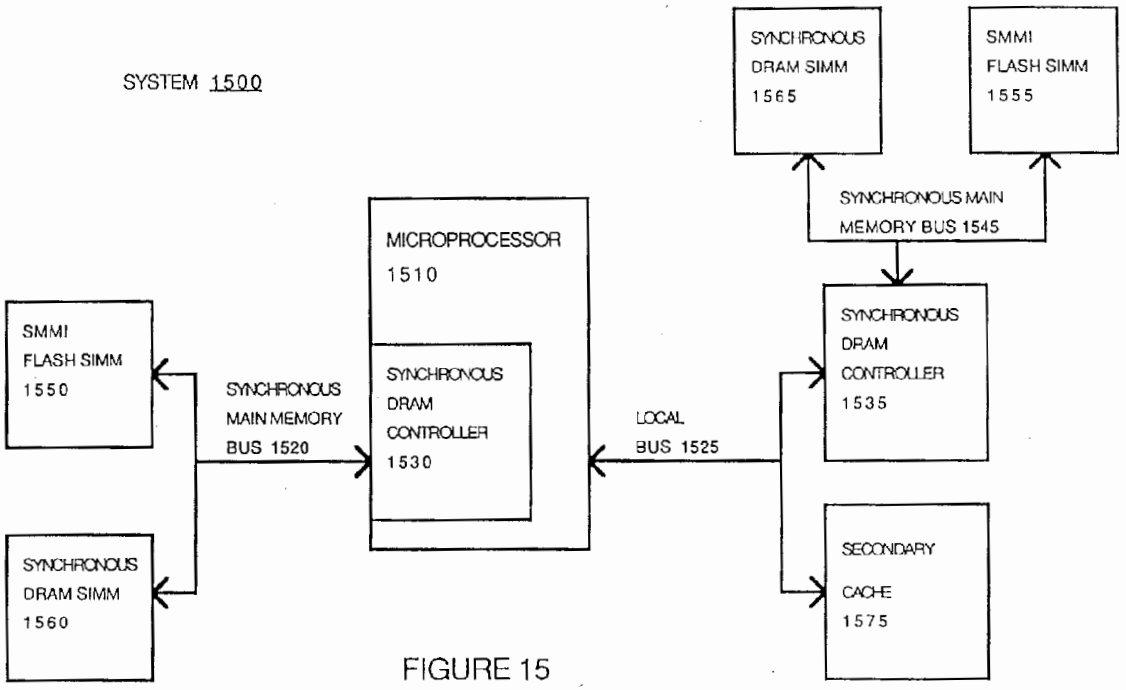


FIGURE 15

Micron Technology Inc. et al.
Ex. 1010, p. 560

世界知的所有権機関
国際事務局

PCT

特許協力条約に基づいて公開された国際出願

(51) 国際特許分類6 **G06F 12/06**	**A1**	(11) 国際公開番号	**WO99/00734**
		(43) 国際公開日	1999年1月7日(07.01.99)

(21) 国際出願番号 PCT/JP98/02866

(22) 国際出願日 1998年6月26日(26.06.98)

(30) 優先権データ
特願平9/171571 1997年6月27日(27.06.97) JP

(71) 出願人（米国を除くすべての指定国について）
株式会社 日立製作所(HITACHI, LTD.)[JP/JP]
〒101-8010 東京都千代田区神田駿河台四丁目6番地
Tokyo, (JP)
(72) 発明者；および
(75) 発明者／出願人（米国についてのみ）
中込儀延(NAKAGOME, Yoshinobu)[JP/JP]
西尾洋二(NISHIO, Yoji)[JP/JP]
佐藤高史(SATO, Takashi)[JP/JP]
〒187-8588 東京都小平市上水本町五丁目20番1号
株式会社 日立製作所 半導体事業部内 Tokyo, (JP)
(74) 代理人
弁理士 玉村静世(TAMAMURA, Shizuyo)
〒271-0092 千葉県松戸市松戸1333番地1
コスモ松戸ステーションスクエア209号 Chiba, (JP)

(81) 指定国 CN, JP, KR, SG, US, 欧州特許 (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

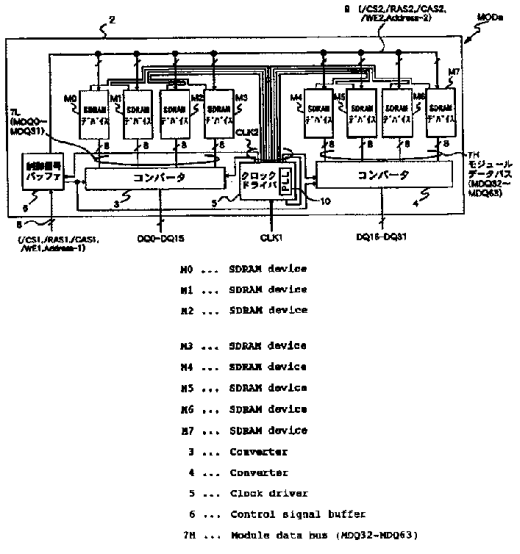
添付公開書類
国際調査報告書

(54)Title: MEMORY MODULE AND DATA PROCESSING SYSTEM

(54)発明の名称 メモリモジュール及びデータ処理システム

(57) Abstract

A memory module enabling a high speed access. A memory module (MODa) has converters (3, 4) including register buffers (RBUFa0-RBUFa15) that are provided between data buses (7L, 7H) connected to a plurality of parallelly operated semiconductor memory devices (M0-M7) and external data input/output terminals (DQ0-DQ31). Because the bus connection between the inside and the outside of the memory module is buffered by the register buffers, the data buses on the mother board is separated from the data buses on the memory module. Thus, the data buses on the mother board no longer see the load of wiring after the buffer registers, so that the load that should be driven by the memory controller can be reduced, thus facilitating improvement in the data transfer frequency on the memory buses.



Micron Technology Inc. et al.
Ex. 1010, p. 561

(57)要約

本発明はアクセス動作の高速化を実現できるメモリモジュールを提供する。すなわち、並列動作される複数個の半導体記憶装置（Ｍ０～Ｍ７）が接続されたデータバス（７Ｌ，７Ｈ）と外部データ入出力端子（ＤＱ０～ＤＱ３１）との間にレジスタバッファ（ＲＢＵＦａ０～ＲＢＵＦａ１５））を含むコンバータ（３、４）設けてメモリモジュール（ＭＯＤａ）を構成する。メモリモジュール内部と外部とのバス接続は前記レジスタバッファでバッファリングされるから、マザーボード上のデータバスとメモリモジュール上のデータバスとは分離される。したがって、マザーボード上のデータバスからは前記バッファレジスタ以降の配線負荷が見えなくなり、メモリコントローラが駆動すべき負荷を低減でき、メモリバス上のデータ転送周波数を向上させることが容易になる。

PCTに基づいて公開される国際出願のパンフレット第一頁に掲載されたPCT加盟国を同定するために使用されるコード(参考情報)

AL	アルバニア	FI	フィンランド	LK	スリ・ランカ	SI	スロヴェニア
AM	アルメニア	FR	フランス	LR	リベリア	SK	スロヴァキア
AT	オーストリア	GA	ガボン	LS	レソト	SL	シエラ・レオネ
AU	オーストラリア	GB	英国	LT	リトアニア	SN	セネガル
AZ	アゼルバイジャン	GD	グレナダ	LU	ルクセンブルグ	SZ	スワジランド
BA	ボスニア・ヘルツェゴビナ	GE	グルジア	LV	ラトヴィア	TD	チャード
BB	バルバドス	GH	ガーナ	MC	モナコ	TG	トーゴー
BE	ベルギー	GM	ガンビア	MD	モルドヴァ	TJ	タジキスタン
BF	ブルギナ・ファソ	GN	ギニア	MG	マダガスカル	TM	トルクメニスタン
BG	ブルガリア	GW	ギニア・ビサオ	MK	マケドニア旧ユーゴスラヴィア共和国	TR	トルコ
BJ	ベナン	GR	ギリシャ	ML	マリ	TT	トリニダッド・トバゴ
BR	ブラジル	HR	クロアチア	MN	モンゴル	UA	ウクライナ
BY	ベラルーシ	HU	ハンガリー	MR	モーリタニア	UG	ウガンダ
CA	カナダ	ID	インドネシア	MW	マラウイ	US	米国
CF	中央アフリカ	IE	アイルランド	MX	メキシコ	UZ	ウズベキスタン
CG	コンゴー	IL	イスラエル	NE	ニジェール	VN	ヴィエトナム
CH	スイス	IN	インド	NL	オランダ	YU	ユーゴースラビア
CI	コートジボアール	IS	アイスランド	NO	ノールウェー	ZW	ジンバブエ
CM	カメルーン	IT	イタリア	NZ	ニュー・ジーランド		
CN	中国	JP	日本	PL	ポーランド		
CU	キューバ	KE	ケニア	PT	ポルトガル		
CY	キプロス	KG	キルギスタン	RO	ルーマニア		
CZ	チェッコ	KP	北朝鮮	RU	ロシア		
DE	ドイツ	KR	韓国	SD	スーダン		
DK	デンマーク	KZ	カザフスタン	SE	スウェーデン		
EE	エストニア	LC	セントルシア	SG	シンガポール		
ES	スペイン	LI	リヒテンシュタイン				

明 細 書

メモリモジュール及びデータ処理システム

技術分野

本発明は、メモリモジュール及びそれを用いたデータ処理システムに係り、殊に、メモリモジュールの高速インタフェース技術に関し、例えばＳＤＲＡＭ（Synchronous Dynamic Random Access Memory：シンクロナス・ダイナミック・ランダム・アクセス・メモリ）を用いた高速メモリモジュールに適用して有効な技術に関するものである。

背景技術

近年、マイクロプロセッサの高性能化にともなって、主記憶についても高速化の要求が高まってきている。特にマルチメデイアパソコンなど大量の画像データの入出力を高速に行うシステムにおいて特に高速の主記憶に対する要求は高まる一方である。これに伴い、主記憶に用いられているＤＲＡＭ（Dynamic Random Access Memory：ダイナミック・ランダム・アクセス・メモリ）に関しては、ファーストページモードから、より高速なデータ転送が可能なＥＤＯ(Extended Data-OUT)ページモード品への移行が進み、最近ではシステムクロックに同期してＤＲＡＭ内部をパイプライン動作させる事により 100MHz 程度のデータ転送率を実現したＳＤＲＡＭも主記憶に使われ始めている。

一方、ＤＲＡＭの応用分野は多岐にわたっており、ゲーム機などのように数個程度を使うものから、ワークステーションなどのように数１００個使うシステムなど、様々なニーズに対応する必要が生じている。大容量を実装する場合にはメモリバスを複数設け、並列化する事により

Micron Technology Inc. et al.
Ex. 1010, p. 563

システムとしてデータ転送速度を上げる事が可能である。しかし、その場合、メモリシステムを構成する増設単位は大きくなり、小容量のメモリしか必要としないシステムには適さないという問題があった。

小さな増設単位でも高いデータ転送速度を得るために、Rambus（ランバス）ＤＲＡＭ（以下ＲＤＲＡＭと略す）と呼ばれる新しい高速ＤＲＡＭも提案され、一部小容量システムに使われている。これは、入出力回路とメモリバスを高速動作向けに新しくしたＤＲＡＭである。すなわち、ＲＤＲＡＭは、コマンドによってアクセス制御されるプロトコル制御方式を採用するもので、予め決められた長さのパケットに読出しリクエストや書き込みリクエスト及びアドレス情報などを入れておき、それをＲＤＲＡＭが解読してリード動作やライト動作を行う。ＲＤＲＡＭは、例えば第２０図に示されるように、プロセッサボードなどのマザーボード上のメモリバスＩＯ０～ＩＯ１５に個々に搭載され、ＲＤＲＡＭ単位でその動作が選択される。前記パケットに入れられた制御情報とリード又はライトのデータ情報は同一バス上に伝達される。このようなプロトコル制御方式によるＲＤＲＡＭは、５００Ｍバイト／秒のような高速でデータの入出力を行うことができる。

尚、ＲＤＲＡＭについて記載された文献の例としては、500Mbyte/sec Data-Rate 512Kbits x9 DRAM Using a Novel I/O Interface (1992 Symposium on VLSI Circuits Digest of Technical Papers, pp.66-67)、日経エレクトロニクス（日経ＢＰ社 1992.3.30 発行　No.550）第１９７～２０９頁がある。

このような背景のもと、種々のシステムによってＤＲＡＭへの要求仕様は異なったものになりつつある。したがって、ある世代のＤＲＡＭに対して、極めて多くの品種を開発する必要が生じている。これによって、開発コストの上昇、品種構成に応じた製造面での対応、性能の異なる

Micron Technology Inc. et al.
Ex. 1010, p. 564

種々の製品をテステイングするためのオーバヘッドなど、色々な問題が生じると予想される。

こうした状況下、パソコンやワークステーションなどＤＲＡＭが用いられるシステムの大半はメモリチップを複数個搭載したモジュール（メモリモジュール）の形で実装されており、メモリ単体の性能よりもメモリモジュールの性能が重要である。したがって、全てのメモリの中に高速動作のための複雑な入出力回路などを取り込まなくとも、メモリモジュールで必要な機能、性能が実現できれば大半のシステムの要求に応じることができる、ということが本発明者によって明らかにされた。本発明者は、上記の観点から、高性能化若しくは高速化を実現するメモリモジュールについて検討した。

前記ＳＤＲＡＭを用いて高速メモリモジュールを構成する場合には次の問題点のあることが本発明者によって明らかにされた。すなわち、現状のＳＤＲＡＭを用いたジュールは、例えば第１９図に示されるように、プロセッサボードのようなマザーボード上の６４ビットのようなメモリバスＩＯ０～ＩＯ６３の所定の複数本の信号線毎に１個のＳＤＲＡＭのデータ入出力端子を結合して６４ビットのようなメモリバスとの間で並列的にデータの入出力が可能にされている。このように、複数のＳＤＲＡＭを用いて所要バス幅のメモリバスとデータの入出力を行う構成では、ＳＤＲＡＭ間におけるアクセスタイムなどの特性ばらつきを補償する回路手段がないと、高データ転送レートを実現することは困難である。また、データ入出力ビット数が×４構成のＳＤＲＡＭを用いて６４ビットバスに対応するには少なくとも１６個のＳＤＲＡＭが必要になり、メモリの増設単位が大きくなる。

一方、ＲＤＲＡＭを用いて高速メモリモジュールを構成する場合には次の問題のあることが本発明者によって見出された。第１に、ＲＤＲＡ

Micron Technology Inc. et al.
Ex. 1010, p. 565

Mは、それ固有のプロトコル制御回路を個々のメモリチップが保有するから、全てのＲＤＲＡＭデバイスが一つのバスに共通接続される。このため、デバイスの入出力容量を相当小さくしないと、バスの負荷容量成分が大きくなり過ぎ、例えば３２個のＲＤＲＡＭを実装するのも難しくなる。

第２に、ＲＤＲＡＭでは各デバイス毎に高精度のクロック同期機構、例えば高精度のＰＬＬ(Phase Locked Loop)又はＤＬＬ(Delayed Locked Loop)が必要であり、また、外部端子数の低減のために信号のビット幅を内部で６倍又は８倍にするように直並列変換を行う必要があり、内部バスのビット幅が大きく（例えば１２８ビット）、デバイスのチップサイズが大型化してします。例えば同じ記憶容量のＳＤＲＡＭに対して１０～１５％程度チップが大型化する。

第３に、バス幅に対するデータが全て一つのＲＤＲＡＭデバイスから出力されるため、ＲＤＲＡＭでは一つのデバイスに消費電力が集中し、デバイスの冷却に特別な考慮を払わなければならないと予想される。

第４には、ＲＤＲＡＭではアクテイブ電力を減らさないと周波数向上は難しい。システムに要求される性能向上に見合った低電力（低電圧化）を半導体集積回路の微細化技術でカバーするのは難しい。

第５には、高速になればなる程、実装系の電気特性の制約が大きく、実使用状態とテスト系の電気特性を整合させる工夫が必要になり、デバイステストのオーバヘッドが大きくなると予想される。

本発明は、上記事情に鑑みてなされたものであり、その目的は、アクセス動作の高速化を実現できるメモリモジュールを提供することにある。

本発明の別の目的は実装するメモリに高速動作のための複雑な入出力回路や制御回路を取り込まなくても、高速メモリアクセスの為のプロ

Micron Technology Inc. et al.
Ex. 1010, p. 566

トコルなど必要な機能や性能を容易に実現することができるできるメモリモジュールを提供することにある。

本発明の更に別の目的は、マイクロプロセッサの高速化に追従してデータ処理速度を容易に向上させることができるデータ処理装置を提供することにある。

本発明の上記並びにその他の目的と新規な特徴は本明細書の以下の記述と添付図面から明らかにされるであろう。

発明の開示

本願において開示される発明のうち代表的なものの概要を簡単に説明すれば下記の通りである。すなわち、並列動作される複数個の半導体記憶装置が接続されたデータバスと外部データ入出力端子との間にレジスタバッファを設けてメモリモジュールを構成する。メモリモジュールはプロセッサボードなどのマザーボードに実装され、当該マザーボード上のメモリコントローラとメモリモジュール内の半導体記憶装置との間でのデータのやりとりはレジスタバッファを介して行われる。

これによれば、メモリモジュール内部と外部とのバス接続は前記レジスタバッファでバッファリングされるから、前記マザーボード上のデータバスとメモリモジュール上のデータバスとは分離される。したがって、マザーボード上のデータバスからは前記レジスタ以降の配線負荷が見えなくなり、メモリコントローラが駆動すべき負荷を低減できる。これにより、メモリバス上のデータ転送周波数を向上させることが容易になる。また、メモリモジュール上についても同様に信号線負荷が低減される。従来のメモリモジュールをマザーボードに実装すると、マザーボード上のメモリバスをメモリモジュール内への負荷の大きな分岐を途中に多数有することになる。上記によれば、メモリバス上におけるそのよ

Micron Technology Inc. et al.
Ex. 1010, p. 567

うな負荷の分岐は実質的になくなる。したがって、メモリモジュール内のデータバス上では、分岐のない2地点間（Point-to-Point）のデータ伝送が可能になるため、メモリモジュール内の動作周波数の向上も比較的容易である。

前記レジスタバッファを配線負荷の分離手段として位置付ける観点とは別の観点による本発明は、並列動作される複数個の半導体記憶装置が接続されたデータバスと外部端子との間に速度変換手段を設けてメモリモジュールを構成する。速度変換手段は、前記データバスから前記外部端子に与えるデータを所定ビット数単位で並列から直列に変換し、前記外部端子から前記データバスに与えるデータを所定ビット数単位で直列から並列に変換する。前記マザーボード上のメモリコントローラと半導体記憶装置との間でのデータのやりとりは前記速度変換手段を介して行われる。この速度変換手段によれば、マザーボード上のメモリバスの信号線数は、メモリモジュール内のデータバスの信号線数よりも少なくなる。したがって、マザーボードのメモリバス上の信号周波数よりも低い動作周波数の半導体記憶装置を用いることができる。また、例えばマザーボードのメモリバスをＲＤＲＡＭに対応した仕様とするとき、一つのバスあたりに実装できる記憶容量を一桁以上向上させることができる。

また、メモリモジュールのインターフェースとは異なる制御手法でメモリモジュール内の半導体記憶装置を制御可能にする為の、プロトコル制御手段を追加してメモリモジュールを構成する。これにより、メモリモジュール内のバスのアクセス制御仕様（メモリアクセスプロトコル）を、前記メモリバスのアクセス制御仕様と相異させることができる。したがって、速度変換手段とプロトコル制御手段の具体的な処理内容次第では、ＥＤＯＤＲＡＭ、ＳＤＲＡＭなど汎用の半導体記憶装置を用いな

Micron Technology Inc. et al.
Ex. 1010, p. 568

がら、メモリルモジュールではＲＤＲＡＭやSyncLink ＤＲＡＭなどに適用されるプロトコルで動作させることができる。これにより、メモリモジュールに実装する半導体記憶装置毎に、プロトコルデコーダや内外の位相同期の為のＰＬＬ、ＤＬＬなどの複雑な制御手段を搭載する必要がなくなる。したがって、ＥＤＯＤＲＡＭ，ＳＤＲＡＭ，ＲＤＲＡＭ，Syclink ＤＲＡＭなど多種の半導体記憶装置を品種展開しなくとも、外部とのインタフェース仕様が異なる種々のメモリモジュールを容易に提供することができる。これは、各種仕様のメモリモジュールの為の半導体メモリを開発するコストの削減にもなる。

また、今後、データ転送の高速化のために位相変調、振幅変調、更にはＱＰＳＫ（Quadrature Phase Shift Keying：直交位相変調）に代表されるような多値変調などにより多重化された信号伝送が行なわれる場合を想定すると、メモリモジュールにそのためのエンコーダ／デコーダ手段（変復調手段）を搭載することができる。これにより、各半導体記憶装置毎に、変復調の為の複雑な信号処理回路を設ける必要がなくなる。

更に、前記速度変換手段やプロトコル制御手段は、半導体記憶装置とは別の半導体集積回路で形成することができる。したがって、速度変換手段やプロトコル制御手段は、半導体記憶装置とは異なるプロセスで形成できるため、より高速のＣＭＯＳ論理プロセス、バイポーラプロセス、化合物半導体プロセスなど、マザーボードのメモリバス上でのデータ伝送に最適な半導体プロセスを使用することができる。

また、メモリコントローラとの間のデータ伝送に光伝送を用いれば、超高速のデータ伝送が可能である。これを考慮する時、外部からの光信号を電気信号に変換し、メモリモジュール内部の電気信号を光信号に変換して外部に出力する光電変換手段をメモリモジュールに搭載するこ

Micron Technology Inc. et al.
Ex. 1010, p. 569

とができる。メモリモジュール上にハーフミラーを設ければ光ファイバーによるバス結合が可能である。光の伝走路は大気中であってもよい。光の場合にはモメモリジュールの挿入による信号劣化が少ないため、一つの伝走路に接続できるメモリモジュールの数を飛躍的に増大させることが可能である。

上記で開示された発明のうち代表的なものによって得られる効果を整理すれば下記の通りである。

すなわち、マザーボード上のデータバスとメモリモジュール上のデータバスがレジスタバッファで分離されるため、マザーボード上のデータバスの負荷が最小になり、マザーボード上、メモリモジュール上ともに、より高いデータ転送速度を達成することが可能となる。

メモリコントローラとメモリのデータバス幅を一致させなくともデータのやりとりが可能になり、メモリコントローラのデータバス幅を小さくして高速で動作させ、一方メモリモジュール上のデータバス幅を大きくして低速で動作させることが可能となる。これにより、例えばメモリコントローラのピン数を削減することができる。

モリチップの読み出し、書き込み、アドレス指定などの制御法をコントローラからの制御法と変えることができるため、見掛け上異なる仕様のメモリモジュールを構成することができる。これにより、各種メモリモジュール用の半導体記憶装置の品種の数を少なくすることができる。また、メモリチップ設計やテストに最適な仕様を選ぶことができるため、メモリモジュールのコストを低減することができる。

上記より、メモリモジュールに搭載すべき半導体記憶装置の入出力仕様に制約されずに、高速化に特化したインターフェース仕様でメモリコントローラあるいはマザーボードを設計することができ、これにより、設計の自由度が向上してデータ処理システムの性能向上を図り易くな

Micron Technology Inc. et al.
Ex. 1010, p. 570

る。

図面の簡単な説明

第１図は本発明に係るメモリモジュールの一例を示すブロック図である。

第２図は第１図に示されるメモリモジュールのチップ配列と配線の概略を示す説明図である。

第３図は第１図に示されるメモリモジュールに含まれるコンバータの一例を示す論理回路図である。

第４図は第１図に示されるメモリモジュールを用いたデータ処理システムの部分的なシステム構成図である。

第５図は第１図に示されるメモリモジュールのアクセス動作タイミングの一例を示すタイミングチャートである。

第６図は第１図に示されるメモリモジュールに対して記憶容量を倍増したメモリモジュールのブロック図である。

第７図はＲＤＲＡＭの代わりにＳＤＲＡＭを用いて構成されたプロトコル制御方式のメモリモジュールのブロック図である。

第８図は第７図に示されるメモリモジュールのアクセス動作タイミングの一例を示すタイミングチャートである。

第９図は速度変換アダプタを採用したマザーボードの一例を示すブロック図である。

第１０図は速度変換アダプタを採用した更に別のマザーボードの一例を示すブロック図である。

第１１図は外部インタフェースを光で行うようにしたメモリモジュールの一例を示すブロック図である。

第１２図は第１１図に示されたメモリモジュールを用いたマザーボ

Micron Technology Inc. et al.
Ex. 1010, p. 571

ードの一例を示すブロック図である。

第１３図は外部インタフェースを多値変調信号で行うようにしたメモリモジュールの一例を示すブロック図である。

第１４図は第１３図のメモリモジュールを用いたマザーボードの一例を示すブロック図である。

第１５図は冗長メモリデバイスを有するメモリモジュールの一例を示すブロック図である。

第１６図は第１５図のメモリモジュールの正規のメモリデバイスに欠陥が無い場合におけるメモリデバイスと外部データ入出力端子との接続態様を示すブロック図である。

第１７図は第１５図のメモリモジュールの正規のメモリデバイスに欠陥がある場合におけるメモリデバイスと外部データ入出力端子との接続態様の一例を示すブロック図である。

第１８図はＥＣＣ回路を有するメモリモジュールの一例を示すブロック図である。

第１９図はＳＤＲＡＭデバイスのデータ入出力端子を対応するメモリバスに直接接続する形式のメモリモジュールを概念的に示した説明図である。

第２０図はＲＤＲＡＭデバイスをメモリバスに接続した状態を示す説明図である。

発明を実施するための最良の形態

《レジスタバッファと速度変換》

先ず最初にレジスタバッファを用いて速度変換を行うようにしたメモリモジュールを説明する。

第１図には本発明に係るメモリモジュールの一例ブロック図が示さ

Micron Technology Inc. et al.
Ex. 1010, p. 572

れる。メモリモジュールＭＯＤａは、ガラスエポキシ樹脂基板の表面に所要の配線パターンが形成されて成るような配線基板２に、例えば８個のＳＤＲＡＭデバイス（半導体記憶装置）Ｍ０～Ｍ７、コンバータ３，４、クロックドライバ５、及び制御信号バッファ（タイミング制御回路）６が実装されて、構成される。

前記ＳＤＲＡＭデバイスＭ０～Ｍ７は並列データ入出力ビット数が夫々８ビット（×８）とされる。７Ｌで示されるものは下位３２ビットのモジュールデータバス、７Ｈで示されるものは上位３２ビットのモジュールデータバスである。４個のＳＤＲＡＭデバイスＭ０～Ｍ３の各々のデータ入出力端子はモジュールデータバス７Ｌの信号線に１対１対応で結合され、４個のＳＤＲＡＭデバイスＭ４～Ｍ７の各々のデータ入出力端子は前記モジュールデータバス７Ｈの信号線に１対１対応で結合される。第１図においてＭＤＱ０～ＭＤＱ６３は８個のＳＤＲＡＭデバイスＭ０～Ｍ７全部のデータ入出力端子を意味する。

メモリモジュールＭＯＤａのデータ入出力端子ＤＱ０～ＤＱ１５，ＤＱ１６～ＤＱ３１は合計３２ビットとされる。前記コンバータ３は下位側１６ビットのデータ入出力端子ＤＱ０～ＤＱ１５と前記下位３２ビットのモジュールデータバス７Ｌとの間でデータの直列・並列変換を行う。同様に、前記コンバータ４は上位１６ビットのデータ入出力端子ＤＱ１６～ＤＱ３１と前記上位３２ビットのモジュールデータバス７Ｈとの間でデータの直列・並列変換を行う。第１図の例では、コンバータ３，４は、モジュールデータバス７Ｌ，７Ｈ（ＭＤＱ０～ＭＤＱ６３）とデータ入出力端子ＤＱ０～ＤＱ３１との間の速度変換手段を実現する。

前記制御信号バッファ６は、外部から供給されるアクセス制御情報８を入力し、内部動作に合わせて内部アクセス制御情報９を各ＳＤＲＡＭ

Micron Technology Inc. et al.
Ex. 1010, p. 573

デバイスＭ０～Ｍ７に並列に供給する。外部アクセス制御情報８は、チップセレクト信号／ＣＳ１、ロウアドレスストローブ信号／ＲＡＳ１、カラムアドレスストローブ信号／ＣＡＳ１、及びライトイネーブル信号／ＷＥ１によって代表されるアクセス制御信号と、アドレス信号Ａｄｄｒｅｓｓ－１とを含む。内部アクセス制御情報９は、チップセレクト信号／ＣＳ２、ロウアドレスストローブ信号／ＲＡＳ２、カラムアドレスストローブ信号／ＣＡＳ２、及びライトイネーブル信号／ＷＥ２よって代表されるアクセス制御信号と、アドレス信号Ａｄｄｒｅｓｓ－２とを含む。

前記クロックドライバ５は外部から供給されるクロック信号ＣＬＫ１を入力して内部クロック信号ＣＬＫ２を生成する。クロックドライバ５から出力されるクロック信号ＣＬＫ２の一つは帰還され、ＰＬＬ（又はＤＬＬ）１０を介して内部クロック信号ＣＬＫ２を外部クロック信号ＣＬＫ１に位相同期させる。内部クロック信号ＣＬＫ２は外部クロック信号ＣＬＫ１と実質的に同一のクロック信号と見なすことができる。前記内部クロック信号ＣＬＫ２は、前記ＳＤＲＡＭデバイスＭ０～Ｍ７、コンバータ３，４及び制御信号バッファ６に供給され、それらの動作基準クロック信号とされる。

ここで前記ＳＤＲＡＭデバイスＭ０～Ｍ７それ自体の機能について簡単に説明する。前記アクセス制御信号／ＣＳ２，／ＲＡＳ２，／ＣＡＳ２，／ＷＥ２及びアドレス信号Ａｄｄｒｅｓｓ－２はクロック信号ＣＬＫ２の立ち上がりエッジで参照される。信号／ＣＳ２はそのローレベルによってコマンド入力サイクルを開始する。コマンド入力サイクルによって始めてその他のアクセス制御信号及びアドレス信号が参照可能にされる。コマンドサイクルでは、信号／ＲＡＳ２，／ＣＡＳ２，／ＷＥ２及びアドレス信号Ａｄｄｒｅｓｓ－２の一部のレベルの組み合わ

Micron Technology Inc. et al.
Ex. 1010, p. 574

せが解読され、その解読結果に従ってＳＤＲＡＭの動作内容が決定される。例えば、／ＲＡＳ２＝“Ｌ”，／ＣＡＳ２＝“Ｈ”，／ＷＥ２＝“Ｈ”の組み合わせは、ロウアドレスストローブ・バンクアクティブコマンドとされ、ロウアドレス系に対する動作が行われる。／ＲＡＳ２＝“Ｈ”，／ＣＡＳ２＝“Ｌ”，／ＷＥ２＝“Ｈ”の組み合わせは、カラムアドレスリードコマンドとされ、ロウアドレス系動作の後にカラムアドレス系を介するリード動作を指示する。／ＲＡＳ２＝“Ｈ”，／ＣＡＳ２＝“Ｌ”，／ＷＥ２＝“Ｌ”の組み合わせは、カラムアドレスライトコマンドとされ、ロウアドレス系動作の後にカラムアドレス系を介するライト動作を指示する。

この例では、メモリモジュールＭＯＤａに対する外部からのアクセス制御形態は、ＳＤＲＡＭデバイスに対するアクセス制御形態と基本的に同じとされる。但し、外部とメモリモジュール１との間のデータ転送速度を上げるため、メモリモジュール１は、クロック信号ＣＬＫ１（＝ＣＬＫ２）の立ち上がりと立ち下がりの両方のクロックエッジに同期してデータの入出力（書き込みと読み出し）を行うＤＤＲ（Double Data-Rate：ダブルデータレート）形式のＳＤＲＡＭの制御仕様を有する。一方、ＳＤＲＡＭデバイスＭ０～Ｍ７は、前述の通り、クロック信号ＣＬＫ２の立ち上がりエッジのみに同期して動作するＳＤＲ（Single Data-Rate：シングルデータレート）形式のＳＤＲＡＭの制御仕様で動作される。

前記メモリモジュールＭＯＤａの外部データ入出力端子ＤＱ０～ＤＱ３１は３２ビット（メモリモジュール１が搭載されるマザーボード上のメモリバスのデータバス幅が３２ビット）であるのに対し、メモリモジュールＭＯＤａ上のモジュールデータバス７Ｌ，７Ｈのバス幅は全体で６４ビットである。マザーボードのメモリバス上でのデータ転送周波

Micron Technology Inc. et al.
Ex. 1010, p. 575

数は、メモリモジュールＭＯＤａのモジュールデータバス上でのデータ転送周波数の２倍とされる。これにより、データバス全体のデータ転送レート（データ転送周波数×データバスのビット幅）は、マザーボード上のメモリバスとメモリモジュールＭＯＤａ上のモジュールデータバスとの間で一致する。

第２図には前記メモリモジュールＭＯＤａのチップ配列と配線の概略を示す。第２図に示される配線はその一部が省略されているが、実際には第１図と同じように形成されている。第２図の例では、ＳＤＲＡＭデバイスＭ０～Ｍ７、コンバータ３，４、クロックドライバ５、及び制御信号バッファ６は夫々別々に半導体集積回路化されている。第２図においてＶＤＤ及びＶＤＤＱは電源端子、ＶＳＳ及びＶＳＳＱは接地端子である。特に電源端子ＶＤＤＱ及び接地端子ＶＳＳＱはデータ出力バッファへの給電に専用化された端子である。

第３図には前記コンバータ３の一例が示される。例えば１ビットの外部データ入出力端子ＤＱ０はＳＤＲＡＭデバイスＭ０の２ビットのデータ入出力端子ＭＤＱ０，ＭＤＱ１に対応され、それらの間にはレジスタバッファＲＢＵＦａ０が配置されている。その他のデータ入出力端ＤＱ１～ＤＱ１５に関しても同様のレジスタバッファＲＢＵＦａ１～ＲＢＵＦａ１５が配置されている。各レジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５の構成は相互に同一である。

前記レジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５は並直変換機能を有する。詳細な構成が例示されたレジスタバッファＲＢＵＦａ０において、Ｂ１で示されるものは、データ入出力端子ＤＱ０に結合された入力バッファである。前記入力バッファＢ１の出力には２個の入力レジスタＲＥＧ１，ＲＥＧ３が結合され、入力レジスタＲＥＧ１の出力は出力ゲートＢＣ１を介してデータ入出力端子ＭＤＱ０に接続され、同様に、

Micron Technology Inc. et al.
Ex. 1010, p. 576

入力レジスタＲＥＧ３の出力は出力ゲートＢＣ３を介してデータ入出力端子ＭＤＱ０に接続される。

第３図においてＢＣ５で示されるものは前記データ入出力端子ＤＱ０に結合された出力バッファである。出力バッファＢＣ５とデータ入出力端子ＭＤＱ０との間には、入力バッファＢ２、出力レジスタＲＥＧ２及び出力ゲートＢＣ２が直列配置され、出力バッファＢＣ５とデータ入出力端子ＭＤＱ１との間には、入力バッファＢ３、出力レジスタＲＥＧ４及び出力ゲートＢＣ４が直列配置されている。前記レジスタＲＥＧ１～ＲＥＧ４及び出力バッファＢＣ１～ＢＣ５は、タイミングジェネレータＴＧＥＮから出力される制御信号φ１Ｒ，φ２Ｗ，ｆａｉ２Ｒ，φ１－１Ｗ，φ１－１Ｒ，φ１－２Ｗ，φ１－２Ｒによってラッチ動作と出力動作が制御される。前記タイミングジェネレータＴＧＥＮは制御信号バッファ６から供給される制御信号とクロック信号ＣＬＫ２に基づいてそれら制御信号を生成する。タイミングジェネレータＴＧＥＮによる制御内容は後で詳述する。特に図示はしないが、コンバータ４も上記と同様に構成されている。

第４図には前記メモリモジュールＭＯＤａを用いたプロセッサボードＰＣＢの部分的なシステム構成図が示される。パーソナルコンピュータではプロセッサボードをＰＣボードとも称する。プロセッサボードやＰＣボードのような回路をマザーボードとも称する。

マザーボードＰＣＢはマイクロプロセッサＭＰＵを中心に構成され、マイクロプロセッサＭＰＵが結合されたＣＰＵバス（システムバス）１１には代表的に示されたメモリコントローラＢＣＯＮＴが結合されている。メモリコントローラＢＣＯＮＴは、メモリバス１２、周辺バス１３、グラフィックバス１４などの動作速度の異なるバスをＣＰＵバス１にインタフェースするためのバスアクセス制御を行うコントローラで

Micron Technology Inc. et al.
Ex. 1010, p. 577

ある。前記メモリコントローラＢＣＯＮＴはバスコントローラ或いはインタフェースコントローラとも称される。グラフィックバス１４にはグラフィックアクセラレータのようなグラフィックコントローラＧＣＯＮＴが結合されている。周辺バス１３には図示を省略するＩＤＥ（Integrated Device Electronics）コントローラなどが接続されている。メモリバス１２には複数個の前記メモリモジュールＭＯＤａが結合されている。マイクロプロセッサＭＰＵは前記複数個のメモリモジュールＭＯＤａをメインメモリとしてワーク領域やデータの一時記憶領域に利用する。この例に従えば、マイクロプロセッサＭＰＵがメインメモリに割り当てられたアドレス空間をアクセスする時、メモリコントローラＢＣＯＮＴは、前記アクセス制御情報をメモリバス１２に出力すると共に、書き込み動作では書き込みデータをメモリバス１２に出力し、読出し動作ではメモリバス１２に読出されたデータを取り込んでマイクロプロセッサＭＰＵに与える。複数個のメモリモジュールＭＯＤａに対するチップ選択信号／ＣＳ１はメモリモジュール毎に固有の信号とされ、前記メモリコントローラＢＣＯＮＴから出力される。

第５図にはメモリモジュールＭＯＤａのアクセス動作タイミングの一例が示される。第５図において、メモリコントローラＢＣＯＮＴは、最初、時刻ｔ０に同期させて、／ＣＳ１＝“Ｌ”、／ＲＡＳ１＝“Ｌ”、／ＣＡＳ１＝“Ｈ”、／ＷＥ＝“Ｈ”によってモジュールアクティブコマンドを発行すると共に、ロウアドレス信号ＲＡｉを出力する。メモリモジュールＭＯＤａの制御信号バッファ６は、時刻ｔ０において、前記／ＣＳ１＝“Ｌ”、／ＲＡＳ１＝“Ｌ”、／ＣＡＳ１＝“Ｈ”、／ＷＥ＝“Ｈ”を、クロック信号ＣＬＫ２（ＣＬＫ１）の立ち上がりエッジで参照する。制御信号バッファ６はこれに応答して、／ＣＳ２＝“Ｌ”、／ＲＡＳ２＝“Ｌ”、／ＣＡＳ２＝“Ｈ”、／ＷＥ＝“Ｈ”によってメ

Micron Technology Inc. et al.
Ex. 1010, p. 578

モリチップアクティブコマンドとロウアドレス信号ＲＡｉをアクセス制御情報９としてＳＤＲＡＭデバイスＭ０～Ｍ７に並列的に出力する。ＳＤＲＡＭデバイスＭ０～Ｍ７は、時刻ｔ１に、メモリチップアクティブコマンドを認識して、ロウアドレス系の動作を開始する。

次にメモリコントローラＢＣＯＮＴは、時刻ｔ２に同期させて、／ＣＳ１＝“Ｌ”、／ＲＡＳ１＝“Ｈ”、／ＣＡＳ１＝“Ｌ”、／ＷＥ＝“Ｌ”によりモジュールライトコマンドを発行すると共に、カラムアドレス信号ＣＡｉを出力する。更にメモリコントローラＢＣＯＮＴは、書き込みデータD(i)，D(i+1)を出力する。このとき、メモリコントローラＢＣＯＮＴによるデータD(i)，D(i+1)の転送は、前記ダブルデータレート形式で、クロック信号ＣＬＫ１の立ち上がり（時刻ｔ２）及び立ち下がり(時刻ｔ３)に同期して行われる。メモリモジュールＭＯＤａの制御信号バッファ６は、時刻ｔ４において、前記メモリモジュールライトコマンドをクロック信号ＣＬＫ２（ＣＬＫ１）の立ち上がりエッジで参照し、これに応答して、／ＣＳ２＝“Ｌ”、／ＲＡＳ２＝“Ｈ”、／ＣＡＳ２＝“Ｌ”、／ＷＥ＝“Ｌ”によりメモリチップライトコマンドとカラムアドレス信号ＣＡｉをアクセス制御情報９としてＳＤＲＡＭデバイスＭ０～Ｍ７に並列的に出力する。また、制御信号バッファ６は、前記メモリチップライトコマンドに応答して、コンバータ３，４にライトデータD(i)，D(i+1)の直列・並列変換動作を指示する。これにより、データ入出力端子ＤＱ０～ＤＱ３１に直列的に与えられたデータD(i)，D(i+1)は、６４ビットの並列データとして時刻ｔ４にモジュールデータバス７Ｌ，７Ｈに供給され、ＳＤＲＡＭデバイスＭ０～Ｍ７のデータ入出力端子ＭＤＱ０～ＭＤＱ６３を介して並列的に８個のＳＤＲＡＭデバイスＭ０～Ｍ７に書き込まれる。

ここで前記直列・並列変換動作を更に説明する。例えば端子ＤＱ０に

Micron Technology Inc. et al.
Ex. 1010, p. 579

最初に供給されるデータD(i)0 はクロック信号ＣＬＫ２（ＣＬＫ１）の立ち上がりに同期してレジスタＲＥＧ１に格納され、後続のデータD(i+1)0 はクロック信号ＣＬＫ２（ＣＬＫ１）の立ち下がりに同期してレジスタＲＥＧ３に取り込まれる。そしてクロック信号ＣＬＫ２の次の立ち上がりに同期して双方のレジスタＲＥＧ１，ＲＥＧ３の出力が開かれて、データD(i)0，D(i+1)0 が端子ＭＤＱ０，ＭＤＱ１に向けて並列出力される。その他の端子ＤＱ１～ＤＱ３１に与えられるデータに関しても同じように直列・並列変換される。

書き込みの後、第５図の例では同一ロウアドレスに対するリード動作が行われる。即ち、メモリコントローラＢＣＯＮＴは、時刻ｔ５に同期させて、／ＣＳ１＝“Ｌ”、／ＲＡＳ１＝“Ｈ”、／ＣＡＳ１＝“Ｌ”、／ＷＥ＝“Ｈ”によりモジュールリードコマンドを発行すると共に、カラムアドレス信号ＣＡｊを出力する。メモリモジュールＭＯＤａの制御信号バッファ６は、クロック信号ＣＬＫ２（ＣＬＫ１）の立ち上がりエッジに同期する時刻ｔ５に、前記モジュールリードコマンドを参照する。制御信号バッファ６はこれに応答して、／ＣＳ２＝“Ｌ”、／ＲＡＳ２＝“Ｈ”、／ＣＡＳ２＝“Ｌ”、／ＷＥ＝“Ｈ”によってメモリチップリードコマンドとカラムアドレス信号ＣＡｊをアクセス制御情報９としてＳＤＲＡＭデバイスＭ０～Ｍ７に並列的に出力する。ＳＤＲＡＭデバイスＭ０～Ｍ７は、時刻ｔ６に、メモリチップリードコマンドを認識して、カラムアドレス系の動作を開始する。カラムアドレス系の動作の結果、時刻ｔ７にＳＤＲＡＭデバイスＭ０～Ｍ７のデータ入出力端子ＭＤＱ０～ＭＤＱ６３からモジュールデータバス７Ｌ，７Ｈにリードデータ D(j)，D(j+1)が並列的に出力される。この出力動作に同期して、制御信号バッファ６は、前記コンバータ３，４にリードデータD(j)，D(j+1)の並列・直列変換動作を指示する。これにより、モジュールデー

Micron Technology Inc. et al.
Ex. 1010, p. 580

タバス７Ｌ，７Ｈ上の並列データＤ(j)，Ｄ(j+1)は３２ビットづつデータ入出力端子ＤＱ０～ＤＱ３１から直列的に出力される。この出力動作は、前記ダブルデータレート形式により、クロック信号ＣＬＫ１の立ち上がりエッジ(時刻ｔ８)及び立ち下がりエッジ（時刻ｔ９）の双方に同期して行われる。

ここで前記並列・直列変換動作を更に説明する。ＳＤＲＡＭデバイスＭ０～Ｍ７のデータ入出力端子ＭＤＱ０～ＭＤＱ６３から並列的に６４ビットのデータがモジュールデータバス７Ｌ，７Ｈに読出されたとき、例えば、データ入出力端子ＭＤＱ０，ＭＤＱ１から並列出力されたデータＤ(j)0，Ｄ(j+1)0 は、クロック信号ＣＬＫ２（ＣＬＫ１）の立ち上がりに同期してレジスタＲＥＧ２，ＲＥＧ４に並列的にラッチされる。そしてクロック信号ＣＬＫ２の次の立ち上がりに同期してレジスタＲＥＧ２のデータＤ(j)0 がデータ入出力端子ＤＱ０から出力され、これに続くクロック信号ＣＬＫ２の立ち下がりに同期してレジスタＲＥＧ４のデータＤ(j+1)0 がデータ入出力端子ＤＱ０から出力される。その他の端子ＭＤＱ１～ＭＤＱ６３から供給されるデータに関しても同じように並列・直列変換される。

以上説明した図１のメモリモジュールＭＯＤａによれば、コンバータ３，４は外部からの書き込みデータに対しては直列・並列変換を行い、外部への読出しデータに対しては並列・直列変換を行い、内外のデータ転送レートに対する速度変換手段としての機能を有する。したがって、メモリモジュールＭＯＤａの入出力は高速だがメモリバス１２のバス幅は小さく、一方メモリモジュールＭＯＤａのモジュールデータバス７Ｌ，７Ｈに対してはそれよりも低速でバス幅が広い構成を実現することができる。したがって、マザーボードＰＣＢのメモリバス上の信号周波数よりも低い動作周波数のＳＤＲＡＭデバイスＭ０～Ｍ７を用いるこ

Micron Technology Inc. et al.
Ex. 1010, p. 581

とができる。更に、メモリコントローラＢＣＯＮＴのピン数（パッケージの外部端子数）の削減にも役立つ。

前記マザーボードＰＣＢ上のメモリコントローラＢＣＯＮＴとメモリモジュールＭＯＤａ内のＳＤＲＡＭデバイスＭ０～Ｍ７との間でのデータのやりとりはレジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５を介して行われる。これによれば、メモリモジュールＭＯＤａの内部と外部とのバス接続は前記レジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５でバッファリングされるから、前記マザーボードＰＣＢ上のデータバスとメモリモジュールＭＯＤａ上のモジュールデータバス７Ｌ，７Ｈとは分離される。したがって、マザーボードＰＣＢ上のデータバスからは前記レジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５以降の配線負荷が見えなくなり、メモリコントローラＢＣＯＮＴが駆動すべき負荷を低減できる。これにより、メモリバス１２上のデータ転送周波数を向上させることが容易になる。すなわち、メモリモジュールＭＯＤａの入出力容量成分を最小化することができる。

また、メモリモジュールＭＯＤａのモジュールデータバス７Ｌ，７Ｈ上についても同様に信号線負荷が低減される。例えば第１９図に示されるような従来のメモリモジュールをプロセッサボードに実装すると、プロセッサボード上のメモリバスにメモリモジュール内への負荷の大きな分岐を途中に多数有することになる。レジスタバッファＲＢＵＦａ０～ＲＢＵＦａ１５を有する上記メモリモジュールＭＯＤａの場合には、メモリバス１２上におけるそのような負荷の分岐は実質的になくなる。したがって、メモリモジュールＭＯＤａ内のモジュールデータバス７Ｌ，７Ｈ上では、分岐のない２地点間のデータ伝送が可能になるため、メモリモジュールＭＯＤａ内の動作周波数の向上も比較的容易になる。

第６図には第１図に対して記憶容量を倍増したメモリモジュールＭ

Micron Technology Inc. et al.
Ex. 1010, p. 582

ＯＤｂの例が示される。このメモリモジュールＭＯＤｂは、ＳＤＲＡＭデバイスＭ０ａ～Ｍ７ａとＭ０ｂ～Ｍ７ｂを有し、８個のＳＤＲＡＭデバイスＭ０ａ～Ｍ７ａはチップ選択信号／ＣＳ２ａによってコマンドサイクルが指示され、ＳＤＲＡＭデバイスＭ０ｂ～Ｍ７ｂはチップ選択信号／ＣＳ２ｂによってコマンドサイクルが指示される。チップ選択信号／ＣＳ２ａ，／ＣＳ２ｂは、外部から供給されるチップ選択信号／ＣＳ１ａ，／ＣＳ１ｂに対応され、チップ選択信号／ＣＳ１ａ，／ＣＳ１ｂは選択的に何れか一方がイネーブルにされる。これにより、第１図の例ではメモリモジュール上に８個のＳＤＲＡＭデバイスが搭載されていたのに対して、第６図のメモリモジュールは２倍の１６個のＳＤＲＡＭデバイスを搭載して利用することができる。尚、第６図では配線などの図示を部分的に省略してあるが、上記相違点以外の構成は第１図と同じである。

この構成においてモジュールデータバス７Ｌ，７Ｈの負荷成分（配線抵抗及び寄生容量成分）は増えるが、外部データ入出力端子ＤＱ０～ＤＱ３１とモジュールデータバス７Ｌ，７Ｈとの間には前記コンバータ３，４が配置されているので、コンバータ３，４に含まれる前記レジスタバッファの作用により、ＳＤＲＡＭデバイスを２倍搭載した第６図の構成でも、外部データ入出力端子ＤＱ０～ＤＱ３１から見たときの電気特性は第１図のメモリモジュールＭＯＤａと実質的に同じである。この例ではメモリモジュールＭＯＤｂ上のモジュールデータバスに２組みのＳＤＲＡＭデバイスを接続しているが、これに限らず、さらに多数組のＳＤＲＡＭデバイスを接続しても構わない。

《レジスタバッファとプロトコル変換》

次に、レジスタバッファを用いた速度変換と共にプロトコル変換を行うようにしたメモリモジュールを説明する。

Micron Technology Inc. et al.
Ex. 1010, p. 583

第７図にはＲＤＲＡＭの代わりにＳＤＲＡＭを用いて Rambus（ランバス）などのプロトコル制御方式に準拠させたメモリモジュールの一例が示される。同図に示されるメモリモジュールＭＯＤｃは、ガラスエポキシ樹脂基板の表面に所要の配線パターンが形成されて成るような配線基板２２に、例えば１６個のＳＤＲＡＭデバイス（半導体記憶装置）Ｍ０～Ｍ１５及びプロトコルコンバータ２３が実装されて成る。プロトコルコンバータ２３は、特に制限されないが、１個の半導体集積回路として構成される。前記プロトコルコンバータ２３は、クロックドライバ２５、プロトコル制御回路２６、並直変換機能を有するレジスタバッファＲＦＵＦ０～ＲＢＵＦ１５、入出力ゲートＩＯＧ０～ＩＯＧ１５及びタイミングジェネレータ２８などを有する。

前記ＳＤＲＡＭデバイスＭ０～Ｍ１５は並列データ入出力ビット数が夫々４ビット（×４）とされる。２７で示されるものは６４ビットのモジュールデータバスである。前記１６個のＳＤＲＡＭデバイスＭ０～Ｍ１５のデータ入出力端子ＭＤＱ０～ＭＤＱ６３はモジュールデータバス２７の信号線に１対１対応で結合されている。

ＩＯＰ０～ＩＯＰ１５はメモリモジュールＭＯＤｃの１６個の入出力端子である。前記レジスタバッファＲＢＵＦｃ０及び入出力ゲートＩＯＧ０は、入出力端子ＩＯＰ０と前記ＳＤＲＡＭデバイスＭ０の４ビットのデータ入出力端子ＭＤＱ０～ＭＤＱ３との間に配置されている。その他のレジスタバッファＲＢＵＦ１～ＲＢＵＦ１５と入出力ゲートＩＯＧ１～ＩＯＧ１５も同様に、入出力端子ＩＯＰ１～ＩＯＰ１５と前記ＳＤＲＡＭデバイスＭ１～Ｍ１５とに対応して配置されている。更にレジスタバッファＲＢＵＦ０～ＲＢＵＦ１５は６４ビットのコマンドバス２９を介してプロトコル制御回路２６に接続されている。

第７図において代表的にその詳細が示された入出力ゲートＩＯＧ０

Micron Technology Inc. et al.
Ex. 1010, p. 584

において、入出力端子ＩＯＰ０に直列的に与えられたデータは、入力バッファＢＣ１９を介して、順次入力レジスタＲＥＧ１１，ＲＥＧ１３，ＲＥＧ１５，ＲＥＧ１７にラッチされ、４ビット単位で出力ゲートＢＣ１１，ＢＣ１３，ＢＣ１５，ＢＣ１７を介して出力される。一方、入出力ゲートＩＯＧ０から並列的に出力される４ビットのデータは出力レジスタＲＥＧ１２，ＲＥＧ１４，ＲＥＧ１６，ＲＥＧ１８に並列的にラッチされ、ラッチされたデータは出力ゲートＢＣ１２，ＢＣ１４，ＢＣ１６，ＢＣ１８を介して順次直列的に出力バッファＢＣ２０から入出力端子ＩＯＰ０に与えられる。その他のレジスタバッファＲＢＵＦ１～ＲＢＵＦ１５も同様に構成される。このように構成されたレジスタバッファＲＢＵＦ０～ＲＢＵＦ１５は、入出力端子ＩＯＰ０～ＩＯＰ１５から入力される１６ビットのデータを４組の直列データ毎に６４ビットの並列データに変換して、コマンドバス２９や入出力ゲートＩＯＧ０～ＩＯＧ１５に与える。また、レジスタバッファＲＢＵＦ０～ＲＢＵＦ１５は、入出力ゲートＩＯＧ０～ＩＯＧ１５から与えられる６４ビットの並列データを１６ビット単位の４組のデータに直列変換して、入出力端子ＩＯＰ０～ＩＯＰ１５から外部に出力させる。この例において、レジスタバッファＲＢＵＦ０～ＲＢＵＦ１５は、モジュールデータバス２７及びコマンドバス２９と入出力端子ＩＯＰ０～ＩＯＰ１５との間の速度変換手段を実現する。

メモリモジュールＭＯＤｃに対する外部からのアクセス制御形態は、Rambus（ランバス）に代表されるようなプロトコル制御形態とされ、ＳＤＲＡＭデバイスに対するアクセス制御形態とは相違される。前記プロトコル制御回路２６は入出力端子ＩＯＰ０～ＩＯＰ１５から供給されるコマンドパケットに含まれている読出しリクエストや書き込みリクエスト及びアドレス情報を解読して監視し、自らの動作が選択されたこ

Micron Technology Inc. et al.
Ex. 1010, p. 585

とをその解読結果から認識すると、当該解読結果を用いて、前記ＳＤＲＡＭデバイスＭ０～Ｍ１２５に対するアクセス制御情報９を出力する。このアクセス制御情報９は、第１図と同様のＳＤＲＡＭデバイスにおけるシングルデータ転送レートに対応されるアクセス制御信号及びアドレス信号である。前記コマンドパケットが書き込みリクエストの場合、当該コマンドパケットの後に続く書き込みデータは入出力端子ＩＯＰ０～ＩＯＰ１５から入力される。

前記クロックドライバ２５は外部から供給されるクロック信号ＣＬＫ１を入力して内部クロック信号ＣＬＫ２を生成する。この例では内部クロック信号ＣＬＫ２は外部クロック信号ＣＬＫ１に対して２分周されている。前記内部クロック信号ＣＬＫ２は、前記ＳＤＲＡＭデバイスＭ０～Ｍ１５及びプロトコル制御回路２６などに供給され、それらの動作基準クロック信号とされる。前記ＳＤＲＡＭデバイスＭ０～Ｍ１５の機能は第１図で説明した内容と同じである。

前記タイミングジェネレータ２８は、クロックドライバ２５から前記クロック信号ＣＬＫ１，ＣＬＫ２を入力すると共に、ＳＤＲＡＭデバイスＭ０～Ｍ１５への外部データ入出力動作の指示をプロトコル制御回路２６から入力して、前記レジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５及び入出力ゲートＩＯＧ０～ＩＯＧ１５の制御信号を生成する。第７図において前記制御信号は、入力制御信号φ１０Ｗ、ラッチ制御信号φ１０－１Ｗ，φ１０－２Ｗ，φ１０－３Ｗ，φ１０－４Ｗ、ゲート出力信号φ１１Ｗ、ゲート出力信号φ１３Ｒ，ラッチ制御信号φ１２Ｒ、ゲート出力信号φ１２－１Ｒ，φ１２－２Ｒ，φ１２－３Ｒ，φ１２－４Ｒとされる。特に、プロトコル制御回路２６は、入出力端子ＩＯＰ０～ＩＯＰ１５から供給されるコマンドパケットを常に監視しなければならないため、スタンバイ状態でもレジスタバッファＲＢＵＦｃ０～Ｒ

Micron Technology Inc. et al.
Ex. 1010, p. 586

ＢＵＦｃ１５は直列・並列変換動作を怠ることはない。

前記メモリモジュールＭＯＤｃも例えば第４図に示されるようなマザーボードＰＣＢに適用することができる。特にこの場合、前記メモリバス１２はメモリモジュールＭＯＤｃの入出力端子ＩＯＰ０～ＩＯＰ１５をメモリコントローラＢＣＯＮＴにインタフェースされると共にクロック信号ＣＬＫ１をメモリモジュールＭＯＤｃに供給するための信号線を有すれば充分とされる。また、メモリコントローラＢＣＯＮＴは、Rambus（ランバス）に代表されるようなプロトコル制御方式でメモリモジュールＭＯＤｃをアクセス制御する機能を有する。換言すれば、メモリバス１２にＲＤＲＡＭを実装したメモリモジュールを接続することも可能である。

第８図にはメモリモジュールＭＯＤｃのアクセス動作タイミングの一例が示される。第８図において、メモリコントローラＢＣＯＮＴは、時刻ｔ０～ｔ１の期間にモジュールライトコマンドのようなコマンドパケットをクロック信号ＣＬＫ１に同期してメモリバス１２に出力する。その後、規定のタイミングに従って時刻ｔ２～ｔ３の期間にクロック信号ＣＬＫ１に同期してモジュール書き込みデータをメモリバス１２に出力する。

プロトコル制御方式によるアクセスをサポートするメモリモジュールＭＯＤｃは、メモリバス１２上のコマンドパケットを監視する。すなわち、入出力端子ＩＯＰ０～ＩＯＰ１５から供給される情報をレジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５で並列信号に変換し、コマンドバス２９を介してプロトコル制御回路２６に供給している。プロトコル制御回路２６は、それによって与えられる情報を解読し、自らのアクセスが指定されたかを判定する。コマンドパケットには、メモリアクセスモードを指定する情報と共にロウアドレス及びカラムアドレス情報な

Micron Technology Inc. et al.
Ex. 1010, p. 587

どメモリサイクルに必要なアクセス制御情報が含まれている。

例えば、プロトコル制御回路２６が、前記時刻ｔ０～ｔ１の期間に供給されたコマンドパケットによりメモリモジュールＭＯＤｃの動作が指示されていることを認識すると、当該プロトコル制御回路２６は、コマンドパケットの解読結果に従って、時刻ｔ２に同期させ、／ＣＳ２＝“Ｌ”、／ＲＡＳ２＝“Ｌ”、／ＣＡＳ２＝“Ｈ”、／ＷＥ＝“ＩＩ”によってメモリチップアクティブコマンドとロウアドレス信号ＲＡｉをアクセス制御情報９としてＳＤＲＡＭデバイスＭ０～Ｍ７に並列的に出力する。ＳＤＲＡＭデバイスＭ０～Ｍ７は、クロック信号ＣＬＫ２の立ち上がりエッジに同期して時刻ｔ２に、メモリチップアクティブコマンドを認識し、ロウアドレス系の動作を開始する。

その後、プロトコル制御回路２６は、前記コマンドパケットの解読結果に従って、時刻ｔ４に同期させて、／ＣＳ２＝“Ｌ”、／ＲＡＳ２＝“Ｈ”、／ＣＡＳ２＝“Ｌ”、／ＷＥ＝“Ｌ”と、カラムアドレス信号ＣＡｉをアクセス制御情報９としてＳＤＲＡＭデバイスＭ０～Ｍ１５に並列的に出力する。このとき、レジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５は、メモリコントローラＢＣＯＮＴが時刻ｔ２～ｔ３の期間に直列的に出力した書き込みデータＤ(i)，Ｄ(i+1)，Ｄ(i+2)，Ｄ(i+3)を６４ビットの並列データに変換し、モジュールデータバス２７に供給している。これにより、ＳＤＲＡＭデバイスＭ０～Ｍ１５には、データ入出力端子ＭＤＱ０～ＭＤＱ６３を介して６４ビットのデータＤ(i)，Ｄ(i+1)，Ｄ(i+2)，Ｄ(i+3)が書き込まれる。

書き込みの後、第８図の例では同一ロウアドレスに対するリード動作が行われる。即ち、メモリコントローラＢＣＯＮＴは、時刻ｔ５～ｔ６に期間にクロック信号ＣＬＫ１に同期して供給されるモジュールリードコマンドとしてのパケットコマンドによってその動作が指示されて

Micron Technology Inc. et al.
Ex. 1010, p. 588

いることを認識すると、当該コマンドパケットの解読結果に従い、時刻ｔ７に同期させて、／ＣＳ２＝“Ｌ”、／ＲＡＳ１＝“Ｈ”、／ＣＡＳ１＝“Ｌ”、／ＷＥ＝“Ｈ”によりメモリチップリードコマンドを出力すると共に、前記解読結果から得られたカラムアドレス信号ＣＡｊをＳＤＲＡＭデバイスＭ０～Ｍ１５に並列的に出力する。ＳＤＲＡＭデバイスＭ０～Ｍ１５は、それに応答し、クロック信号ＣＬＫ２の立ち上がりエッジに同期する時刻ｔ７に、メモリチップリードコマンドを認識して、カラムアドレス系の動作を開始する。カラムアドレス系の動作の結果、時刻ｔ８にＳＤＲＡＭデバイスＭ０～Ｍ１５のデータ入出力端子ＭＤＱ０～ＭＤＱ６３からモジュールデータバス２７にリードデータD(j), D(j+1), D(j+2), D(j+3)が並列的に出力される。この出力動作に同期して、プロトコル制御回路２６は、前記レジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５にリードデータD(j), D(j+1), D(j+2), D(j+3)の並列・直列変換動作を指示する。これにより、モジュールデータバス２７上の並列データD(j), D(j+1), D(j+2), D(j+3)は１６ビットづつ入出力端子ＩＯＰ０～ＩＯＰ１５から直列的に出力される。この出力動作は、前記ダブルデータレート形式と同じように、クロック信号ＣＬＫ１の立ち上がりエッジ及び立ち下がりエッジの双方に同期して行われる。

このように、上記メモリモジュールＭＯＤｃは、データの直列・並列／並列・直列変換機構を構成するレジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５とプロトコル制御回路２６等を有するプロトコルコンバータ２３を含むから、ＲＤＲＡＭの代わりにＳＤＲＡＭデバイスを用いてRambus（ランバス）に準拠するようなアクセスプロトコルをサポートすることができる。上記以外のメモリ仕様とモジュール仕様の組み合わせについても同様に実現することが可能である。したがって、メモリモジ

Micron Technology Inc. et al.
Ex. 1010, p. 589

ュールＭＯＤｃに実装するＳＤＲＡＭデバイスのような半導体記憶装置毎に、プロトコルデコーダや内外の位相同期の為のＰＬＬ、ＤＬＬなどの複雑な制御手段を搭載する必要はない。よって、ＥＤＯＤＲＡＭ，ＳＤＲＡＭ，ＲＤＲＡＭ，Syclink ＤＲＡＭなど多種の半導体記憶装置を品種展開しなくとも、外部とのインタフェース仕様の異なる種々のメモリモジュールを容易に提供することができる。これは、各種仕様のメモリモジュールの為の半導体メモリを開発するコストを削減可能にする。

また、データの直列・並列／並列・直列変換機構を構成するレジスタバッファＲＢＵＦｃ０～ＲＢＵＦｃ１５とプロトコル制御回路２６等を有するプロトコルコンバータ２３は、ＳＤＲＡＭデバイスＭ０～Ｍ１５とは別の半導体集積回路で形成することができる。プロトコルコンバータ２３は、ＳＤＲＡＭデバイスＭ０～Ｍ１５とは異なるプロセスで形成できるため、より高速のＣＭＯＳ論理プロセス、バイポーラプロセス、化合物半導体プロセスなど、マザーボードＰＣＢのメモリバス上でのデータ伝送に最適な半導体プロセスを使用することができる。

《速度変換アダプタ》

前記速度変換技術はメモリバス上にも適用することができる。例えば、第９図に例示されるように、メモリコントローラＢＣＯＮＴとメモリバス１２との接続を速度変換アダプタ３０を介して行う。速度変換アダプタ３０は、前記コンバータ３，４と同じようにレジスタバッファを用いた速度変換手段を備える。メモリコントローラＢＣＯＮＴと速度変換アダプタ３０との間は１６ビットのデータバス３１に接続される。速度変換アダプタ３０はデータバス３１に直列的に伝達される１６ビットのデータを３２ビットのデータサイズに並列変換して３２ビットのデータバス３２に供給し、データバス３２の３２ビットデータを１６ビット

Micron Technology Inc. et al.
Ex. 1010, p. 590

に直列変換してデータバス３１に直列的に伝達する。データバス３２には例えば前記メモリモジュールＭＯＤａが複数個結合されている。

メモリコントローラＢＣＯＮＴの動作周波数（４×ｆＭＨｚ）はメモリモジュールＭＯＤａ(1)～ＭＯＤ a(n)の動作周波数（２×ｆＭＨｚ）の２倍とされる。前述の説明から明らかなように、メモリモジュールＭＯＤａに搭載されているＳＤＲＡＭデバイスＭ０～Ｍ７の２倍の周波数でメモリモジュールＭＯＤａの入出力動作が行われる。従って、図９の例では、メモリコントローラＢＣＯＮＴの入出力の動作周波数はＳＤＲＡＭデバイスＭ０～Ｍ７の動作周波数（ｆ）の４倍になる。このようなシステム構成により、メモリコントローラＢＣＯＮＴは高速に動作するが、そのピン数を減らすことができる。

第１０図の例はメモリバス１２の間に速度変換アダプタ３０を介在させた例であり、前段のメモリバス１２Ｂは１６ビットのデータバス３１を備え、後段のメモリバス１２Ａは３２ビットのデータバス３２を備える。速度変換アダプタ３０の機能は第９図と同じである。前段のメモリバス１２Ｂには前記メモリモジュールＭＯＤａが複数個結合されている。前記メモリバス１２ＢにはメモリモジュールＭＯＤａａとメモリデバイスＭＥＭが結合されている。メモリモジュールＭＯＤａａは前記メモリモジュールＭＯＤａに対して入出力データが１６ビットにされた点が相違される。メモリデバイスＭＥＭはデータ入出力ビット数が１６ビットのＳＤＲＡＭ等の半導体メモリである。

このシステム構成によれば、小容量だが高速に動作するメモリチップやメモリモジュールをメモリバス１２Ｂに接続し、大容量だが速度の遅いメモリモジュールをメモリバス１２Ａに接続して、システムに最適な性能（アクセス速度、メモリ容量）を達成することができる。

前記速度変換アダプタ３０を採用したシステムは、ＭＯＤｂ，ＭＯＤ

Micron Technology Inc. et al.
Ex. 1010, p. 591

ｃなどその他のメモリモジュールにも適用できることは言うまでもない。

《メモリモジュールの外部インタフェース方式》

第１１図には外部インタフェースを光で行うようにしたメモリモジュールの一例が示される。第１１図に示されるメモリモジュールＭＯＤｄは第１図のメモリモジュールＭＯＤａに対して外部インタフェース信号を光とした点が相違される。すなわち、前記コンバータ３，４、クロックドライバ５及び制御信号バッファ６の外部インタフェース部分に光電変換回路３Ｌ，４Ｌ，５Ｌ，６Ｌが配置されている。前記光電変換回路３Ｌ，４Ｌ，５Ｌ，６Ｌは、ハーフミラー４０を介して光バス４１に接続される。８Ｌはアクセス制御情報８を伝達する光ファイバ、ＬＤＱ０～ＬＤＱ７はデータを伝達する光ファイバ、ＬＣＬＫはクロック信号ＣＬＫ１を伝達する光ファイバである。

光電変換回路３ＡＬ，４Ｌはピンフォトダイオード有する光レシーバと、フォトダイオードを有する光トランスミッタとを有する。光電変換回路５Ｌ，６Ｌはピンフォトダイオード有する光レシーバを備える。尚、第１図と同一回路ブロックには同一符号を付してその詳細な説明は省略する。

第１２図にはメモリモジュールＭＯＤｄを用いたシステム構成例が示される。前記メモリコントローラＢＣＯＮＴと光バス４１との接続は入出力アダプタ４２を介して行われる。入出力アダプタ４２は、前述と同様の速度変換機能を有するコンバータと光電変換回路の機能を有する。メモリモジュールＭＯＤｄへのデータ書き込みに際しては、入出力アダプタ４２から送られる光信号の一部を所定のメモリモジュールＭＯＤｄに取り込む。メモリモジュールＭＯＤｄからのデータ読み出しに際しては、メモリモジュールＭＯＤｄからの光をハーフミラー４０で反

Micron Technology Inc. et al.
Ex. 1010, p. 592

射させて入出力アダプタ４２の送信する。

ハーフミラー４０は、この例のようにメモリモジュールＭＯＤｄ上に搭載しなくとも、マザーボード上に設置して、反射した光がメモリモジュール内の光電変換回路の受光／発光部分に照射されるように配置してもよい。また、プロセッサボード上の光の伝送媒体としては、光ファイバ４１でも良いが、コンピュータの主記憶のように伝送距離が短く、光の強度の劣化が問題ない場合には、空中を伝送させてもよい。

このように、プロセッサボード上を光で伝送した場合、光は電気に比べて信号の減衰や劣化が少ないため従来の電気によるバス接続に比べて飛躍的に多数のメモリモジュールＭＯＤｄを実装することが可能になる。また、光同士は相互干渉がなく、また、他の電気信号との間の干渉もないため、メモリバスから他の信号にノイズが載ったり、あるいは逆に、他の信号からノイズを受けて誤動作するような危険性を少なくすることができる。また、従来のバスではマザーボード上の配線のインダクタンスがアンテナのように働き、電磁波を外部に放出する問題が起きやすかったが、光伝送とすることにより、そうした問題も回避することができる。

第１３図には外部インタフェースを多値変調信号で行うようにしたメモリモジュールの一例が示される。第１３図に示されるメモリモジュールＭＯＤｅは第１図のメモリモジュールＭＯＤａに対して外部インタフェース信号を多値変調信号とした点が相違される。すなわち、前記コンバータ３，４の外部インタフェース部分に多値変復調回路３Ｓ，４Ｓが配置されている。多値変復調方式には例えばＱＰＳＫ等を採用できる。メモリバス５１に含まれるデータバスＳＤＱ０～ＳＤＱ１５には多値変調信号が伝送される。これにより、各半導体記憶装置毎に、変復調の為の複雑な信号処理回路を設ける必要はない。尚、第１図と同一回路

Micron Technology Inc. et al.
Ex. 1010, p. 593

ブロックには同一符号を付してその詳細な説明は省略する。

第１４図にはメモリモジュールＭＯＤｅを用いたシステム構成例が示される。前記メモリコントローラＢＣＯＮＴとメモリバス５１との接続は入出力アダプタ５２を介して行われる。入出力アダプタ５２は、前述と同様の速度変換機能を有するコンバータと多値変復調回路の機能を有する。メモリモジュールＭＯＤｅへのデータ書き込みに際しては、入出力アダプタ５２から送られる変調信号は多値変復調回路３Ｓ，４Ｓで復調されてコンバータに取り込まれる。メモリモジュールＭＯＤｅからのデータ読み出しに際しては、コンバータ３，４の出力が多値変復調回路３Ｓ，４Ｓで変調されてメモリバス５１に供給される。

マザーボード上を多値変調信号でデータ伝送した場合、通常のデジタル信号に比べて伝送信号の周波数スペクトラムを狭めることが可能となる。例えば正弦波に近い伝送波形を用いることも可能になる。したがって、同じメモリバスを用いても、波形制御がやり易くなり高い周波数まで伝送することが可能となる。

《コンバータと冗長メモリデバイス》

第１５図には冗長メモリデバイスを有するメモリモジュールＭＯＤｆの一例が示される。基本的な構成は第１図と同じであり、１個のＳＤＲＡＭデバイスＭ８を余計に搭載している。そして、９個のＳＤＲＡＭデバイスＭ０～Ｍ８の内のどの８個を前記モジュールデータバス７Ｌ，７Ｈに接続するかを選択するセレクタ６１が設けられ、前記セレクタ６１による選択動作を決定する為の救済アドレス情報が設定される冗長プログラム回路６０及び救済アドレス判定回路６２が設けられている。冗長プログラム回路６０は、救済すべきアドレス（不良ビットのアドレス）情報がプログラマブルに設定されると共に、その救済アドレスにおいて不良を有するＳＤＲＡＭデバイスの番号情報（No.0～No.7）がプロ

Micron Technology Inc. et al.
Ex. 1010, p. 594

グラマブルに設定される。一つの不良アドレスに対して不良ＳＤＲＡＭデバイスの数は１個に制限される。冗長プログラム回路６０はヒューズプログラム回路又は電気的に書き込み可能な揮発性半導体メモリによって構成することができる。

前記救済アドレス判定回路６２は、制御信号バッファ６に供給されるアドレス情報Ａｄｄｒｅｓｓ－１と前記冗長プログラム回路６０に設定された救済すべきアドレスとを比較する。比較結果が一致の場合には、当該救済すべきアドレスと対の不良ＳＤＲＡＭデバイスの番号に相当するＳＤＲＡＭデバイスをセレクタ６１によってコンバータ３，４から切り離し、これに代えてＳＤＲＡＭデバイスＭ８がセレクタ６１によってコンバータ３，４に接続される。前記セレクタ６１、冗長プログラム回路６０、救済アドレス判定回路６２、コンバータ３，４、及び第１５図には図示を省略したクロックドライバ５は、１個の半導体集積回路化されている。その他の構成は第１図と同じであるからその詳細な説明は省略する。

前記冗長プログラム回路６０の初期状態においてセレクタ６１による選択状態は第１６図のようにされる。ＳＤＲＡＭデバイスＭ０～Ｍ７の不良アドレスと不良デバイス番号に応じて冗長プログラム回路６０がプログラムされる。例えば、第１７図に例示されるように、あるアクセスアドレスに関する不良ＳＤＲＡＭデバイスがＭ４であるとき、セレクタ６１は当該ＳＤＲＡＭデバイスＭ４をモジュールデータバス７Ｌから切り離し、これに代えてＳＤＲＡＭデバイスＭ８をモジュールデータバス７Ｌに接続する。

これにより、不良のメモリチップを用いてメモリモジュールの良品を提供できる。これらはメモリチップの入出力がそのままモジュールの入出力端子に接続されている従来のメモリモジュールでは実現できない。

Micron Technology Inc. et al.
Ex. 1010, p. 595

《コンバータとＥＣＣ》

第１８図にはＥＣＣ回路を有するメモリモジュールの一例が示される。基本的な構成は第１図と同じであり、ＥＣＣ回路７０とＳＤＲＡＭデバイスＭ８を余計に搭載している。ＥＣＣ回路はコンバータ３，４とモジュールデータバスとの間に配置される。ＥＣＣ回路は、コンバータ３，４を通って並列化された６４ビットのデータに対する誤り訂正コードを生成する。生成された誤り訂正コードはＳＤＲＡＭデバイスＭ８に格納される。読出し動作時において、ＥＣＣ回路７０は、ＳＤＲＡＭデバイスＭ８から読出された前記誤り訂正コードを用いて、前記ＳＤＲＡＭデバイスＭ０～Ｍ７から読出されたデータの誤り検出及び訂正を行い、その結果をコンバータ３，４に与える。

これにより、信頼性の高いメモリモジュールを提供することができる。これらはメモリチップの入出力がそのままモジュールの入出力端子に接続されている従来のメモリモジュールでは実現できない。

以上本発明者によってなされた発明を実施例に基づいて具体的に説明したが、本発明はそれに限定されるものではなく、その要旨を逸脱しない範囲において種々変更可能であることは言うまでもない。例えばメモリモジュールに利用する半導体記憶装置それ自体の種類は上記の例に限定されず適宜変更可能である。メモリモジュールの入出力端子のビット数、メモリモジュールに実装する半導体記憶装置の数も適宜変更可能である。また、マザーボード上の入出力インターフェースはＣＭＯＳ、ＧＴＬ（Gunning Transceiver Logic）、ＴＴＬ等の何れの形式であっても本発明は適用可能である。ＳＳＴＬ（Stub Series Terminated Transceiver Logic）インターフェースの場合にはモジュールの入出力端子とコンバータ３，４との間にスタブ抵抗を配すればよい。

Micron Technology Inc. et al.
Ex. 1010, p. 596

産業上の利用可能性

本発明は、パーソナルコンピュータやワークステーションのメモリモジュールなどに広く適用することができる。

Micron Technology Inc. et al.
Ex. 1010, p. 597

請 求 の 範 囲

1．配線基板と、前記配線基板に設けられた複数個の半導体記憶装置と、並列動作される前記複数個の半導体記憶装置のデータ入出力端子が個別に接続されるデータバスと、前記配線基板に設けられた複数個の外部データ入出力端子と、前記データバスから前記データ入出力端子に与えるデータを保持し、また、前記データ入出力端子から前記データバスに与えるデータを保持するレジスタバッファと、を含んで成るものであることを特徴とするメモリモジュール。

2．前記データバスのビット数mは前記外部データ入出力端子のビット数nの整数倍であることを特徴とする請求の範囲第1項に記載のメモリモジュール。

3．前記複数個の半導体記憶装置は外部から供給されるクロック信号に同期動作され、前記外部データ入出力端子からのデータ入出力動作周波数ｆ1は前記半導体集積回路のデータ入出力動作周波数ｆ2よりも高く、前記並列動作される複数個の半導体集積回路と前記データバスとの間での最高データ転送レートｎ×ｆ2と前記外部データ入出力端子を介する入出力動作の最高データ転送レートｍ×ｆ1とが実質的に等しいことを特徴とする請求の範囲第2項に記載のメモリモジュール。

4．前記データバスに接続されて並列動作される複数個の半導体記憶装置を複数組有し、前記複数組の中の一組の半導体記憶装置が選択的に動作可能にされて成るものであることを特徴とする請求の範囲第3項に記載のメモリモジュール。

5．予備の半導体記憶装置と、欠陥のある半導体記憶装置を前記予備の半導体記憶装置に選択的に置き換える救済手段とを更に含んで成る

Micron Technology Inc. et al.
Ex. 1010, p. 598

ものであることを特徴とする請求の範囲第４項に記載のメモリモジュール。

６．マイクロプロセッサと、マイクロプロセッサに結合されメモリバスを制御するメモリコントローラと、前記メモリバスに結合されたメモリモジュールと、を含んで成り、

上記メモリモジュールは、配線基板と、前記配線基板に設けられた複数個の半導体記憶装置と、並列動作される前記複数個の半導体記憶装置のデータ入出力端子が個別に接続されるデータバスと、前記配線基板に設けられた複数個の外部データ入出力端子と、前記データバスから前記データ入出力端子に与えるデータを一時的に保持し、また、前記データ入出力端子から前記データバスに与えるデータを一時的に保持するレジスタバッファと、を含んで成るものであることを特徴とするデータ処理システム。

７．配線基板と、前記配線基板に設けられ外部からのクロック信号に同期動作される複数個の半導体記憶装置と、並列動作される前記複数個の半導体記憶装置のデータ入出力端子が個別に接続されるデータバスと、前記配線基板に設けられた複数個の外部端子と、前記データバスから前記外部端子に与えるデータを所定ビット数単位で並列から直列に変換し、前記外部端子から前記データバスに与えるデータを所定ビット数単位で直列から並列に変換する速度変換手段と、を含んで成るものであることを特徴とするメモリモジュール。

８．前記速度変換手段は半導体集積回路化されて成るものであることを特徴とする請求の範囲第７項に記載のメモリモジュール。

９．メモリモジュールに対するアクセス制御情報がデータ入出力用とは異なる専用の外部端子を介して供給されるタイミング制御回路を有し、このタイミング制御回路は、前記アクセス制御情報に基づいて前

Micron Technology Inc. et al.
Ex. 1010, p. 599

記複数個の半導体記憶装置に内部アドレス信号及び内部アクセス制御信号を並列的に供給するものであることを特徴とする請求の範囲第７項又は第８項に記載のメモリモジュール。

１０．メモリモジュールに対するアクセス制御情報がデータ入出力用と兼用された外部端子から前記速度変換手段を介して供給されるプロトコル制御回路を有し、このプロトコル制御回路は、前記アクセス制御情報を解読して内部アドレス信号及び内部アクセス制御信号を生成し前記複数個の半導体記憶装置に並列的に供給するものであることを特徴とする請求の範囲第７項又は第８項に記載のメモリモジュール。

１１．マイクロプロセッサと、前記マイクロプロセッサに結合されメモリバスを制御するメモリコントローラと、前記メモリバスに結合されたメモリモジュールとを含み、

前記メモリコントローラは、アクセス制御情報とアクセスデータとを前記メモリバスの異なる信号線に伝達して前記メモリモジュールをアクセス制御するものであり、

前記メモリモジュールは、配線基板と、前記配線基板に設けられ外部からのクロック信号に同期動作される複数個の半導体記憶装置と、並列動作される前記複数個の半導体記憶装置のデータ入出力端子が個別に接続されるデータバスと、前記配線基板に設けられた複数個の外部端子と、前記データバスから前記外部端子に与えるデータを所定ビット数単位で並列から直列に変換し、前記外部端子から前記データバスに与えるデータを所定ビット数単位で直列から並列に変換する速度変換手段と、を含んで成るものであることを特徴とするデータ処理システム。

１２．マイクロプロセッサと、前記マイクロプロセッサに結合されメモ

Micron Technology Inc. et al.
Ex. 1010, p. 600

リバスを制御するメモリコントローラと、前記メモリバスに結合されたメモリモジュールとを含み、

前記メモリモジュールは、配線基板と、前記配線基板に設けられ外部からのクロック信号に同期動作される複数個の半導体記憶装置と、並列動作される前記複数個の半導体記憶装置のデータ入出力端子が個別に接続されるデータバスと、前記配線基板に設けられた複数個の外部端子と、前記データバスから前記外部端子に与えるデータを所定ビット数単位で並列から直列に変換し、前記外部端子から前記データバスに与えるデータを所定ビット数単位で直列から並列に変換する速度変換手段と、該メモリモジュールに対するアクセス制御情報がデータ入出力用と兼用された外部端子から前記速度変換手段を介して供給されるプロトコル制御回路とを有し、

前記プロトコル制御回路は、前記アクセス制御情報を解読して内部アドレス信号及び内部アクセス制御信号を生成し前記複数個の半導体記憶装置に並列的に供給するものであり、

前記メモリコントローラは、アクセス制御情報とアクセスデータとを前記メモリバスの共通の信号線に伝達して前記メモリモジュールをアクセス制御するものであることを特徴とするデータ処理システム。

１３．前記メモリバスは、その途中に、前段から後段に与える情報を所定ビット数単位で直列から並列に変換し、後段から前段に与える情報を所定ビット数単位で並列から直列に変換する速度変換アダプタを有し、前記速度変換アダプタの後段のメモリバスに前記メモリモジュールが結合されて成るものであることを特徴とする請求の範囲第１２項に記載のデータ処理システム。

１４．前記速度変換アダプタの前段のメモリバスにも前記メモリモージ

Micron Technology Inc. et al.
Ex. 1010, p. 601

ュールが結合されて成るものであることを特徴とする請求の範囲第１３項に記載のデータ処理システム。

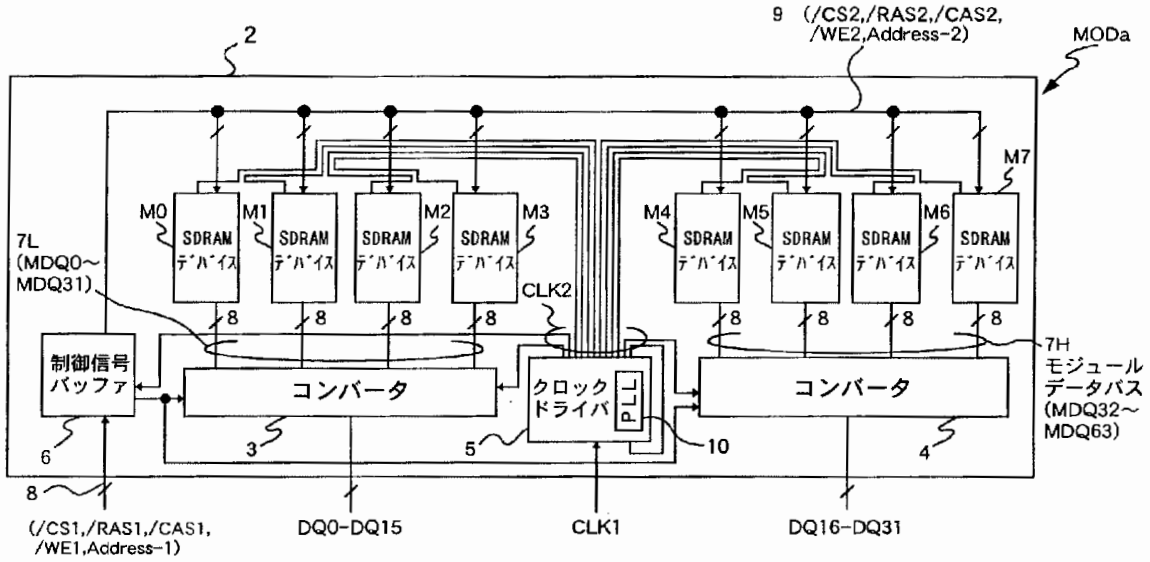
１５．前記外部端子から供給される光信号を電気信号に変換し、前記外部端子から出力する信号を電気信号から光信号に変換する光電変換手段を更に備えて成るものであることを特徴とする請求の範囲第７項に記載のメモリモジュール。

１６．前記外部端子の一部又は全部から供給される変調信号を復調し、前記一部又は全部の外部端子から出力する信号を変調する変復調手段を更に備えて成るものであることを特徴とする請求の範囲第７項に記載のメモリモジュール。

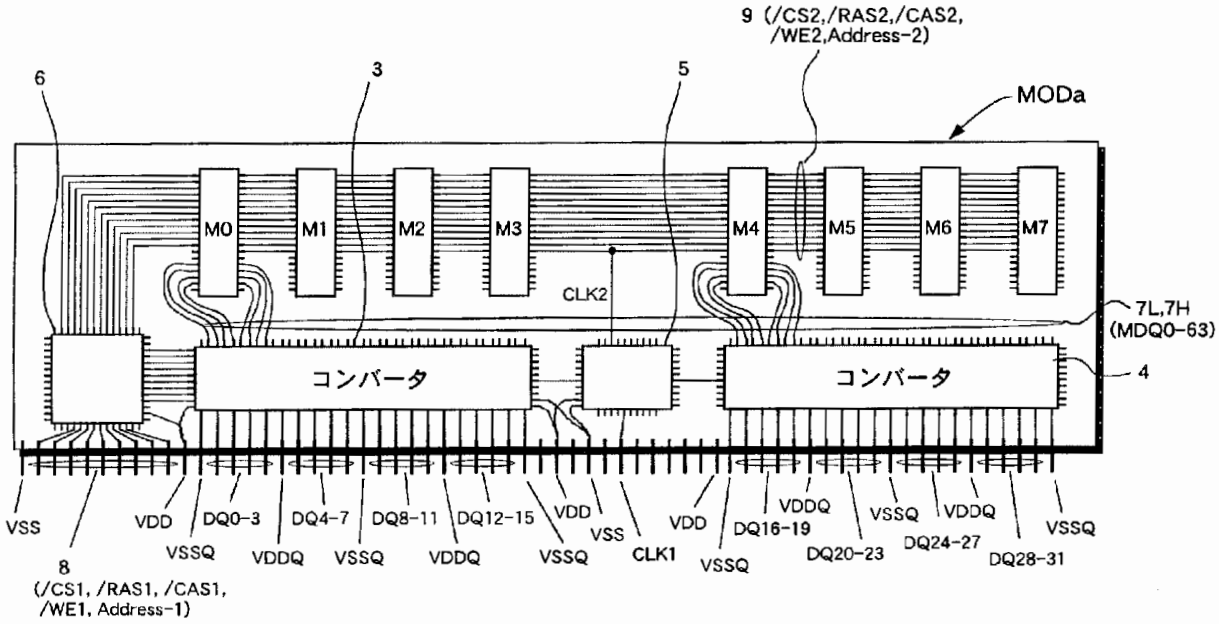
１７．前記速度変換手段の前段又は後段に誤り訂正手段を設けると共に、前記誤り訂正手段で用いる誤り訂正コードを格納する半導体記憶装置を設けて成るものであることを特徴とする請求の範囲第７項に記載のメモリモジュール。

Micron Technology Inc. et al.
Ex. 1010, p. 602

第１図

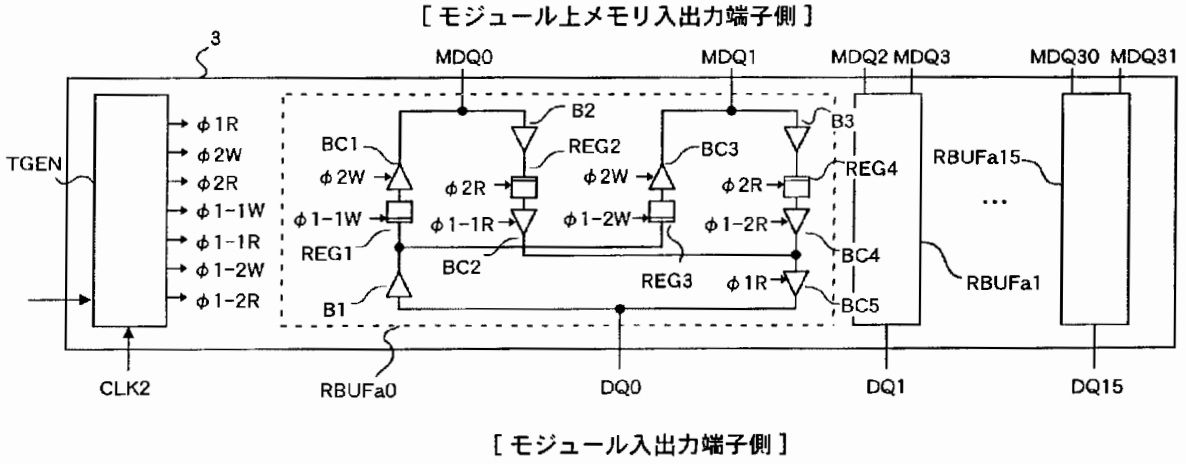

2
9 (/CS2,/RAS2,/CAS2, /WE2,Address-2)
MODa
M0
M1
M2
M3
M4
M5
M6
M7
SDRAM デバイス
7L (MDQ0～ MDQ31)
CLK2
8
制御信号 バッファ
コンバータ
クロック ドライバ
PLL
コンバータ
7H モジュール データバス (MDQ32～ MDQ63)
6
3
5
10
4
8
(/CS1,/RAS1,/CAS1, /WE1,Address-1)
DQ0-DQ15
CLK1
DQ16-DQ31

第２図



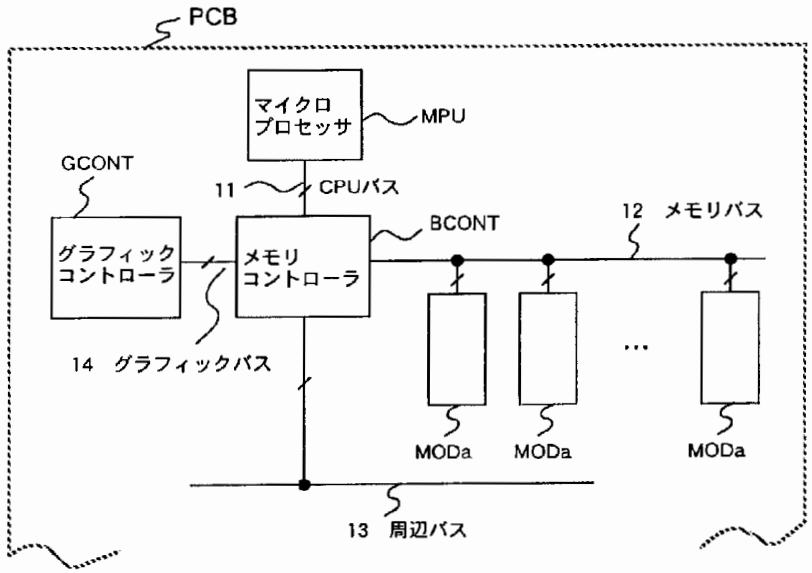
Micron Technology Inc. et al.
Ex. 1010, p. 604

第３図

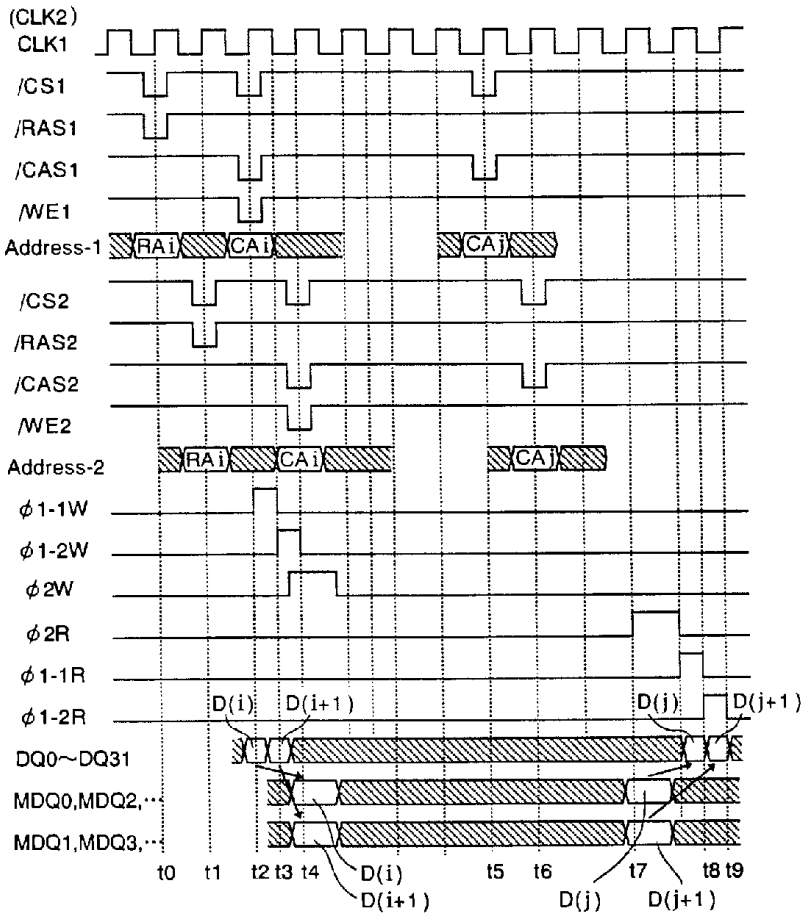


Micron Technology Inc. et al.
Ex. 1010, p. 605

第4図

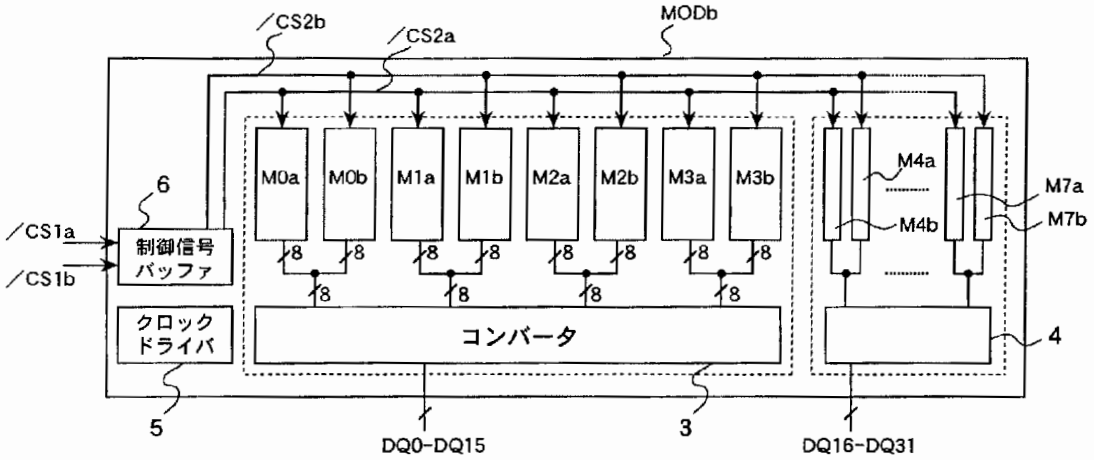

PCB
マイクロ
プロセッサ
MPU
GCONT
11
CPUバス
BCONT
12 メモリバス
グラフィック
コントローラ
メモリ
コントローラ
14 グラフィックバス
MODa
MODa
MODa
13 周辺バス

第5図



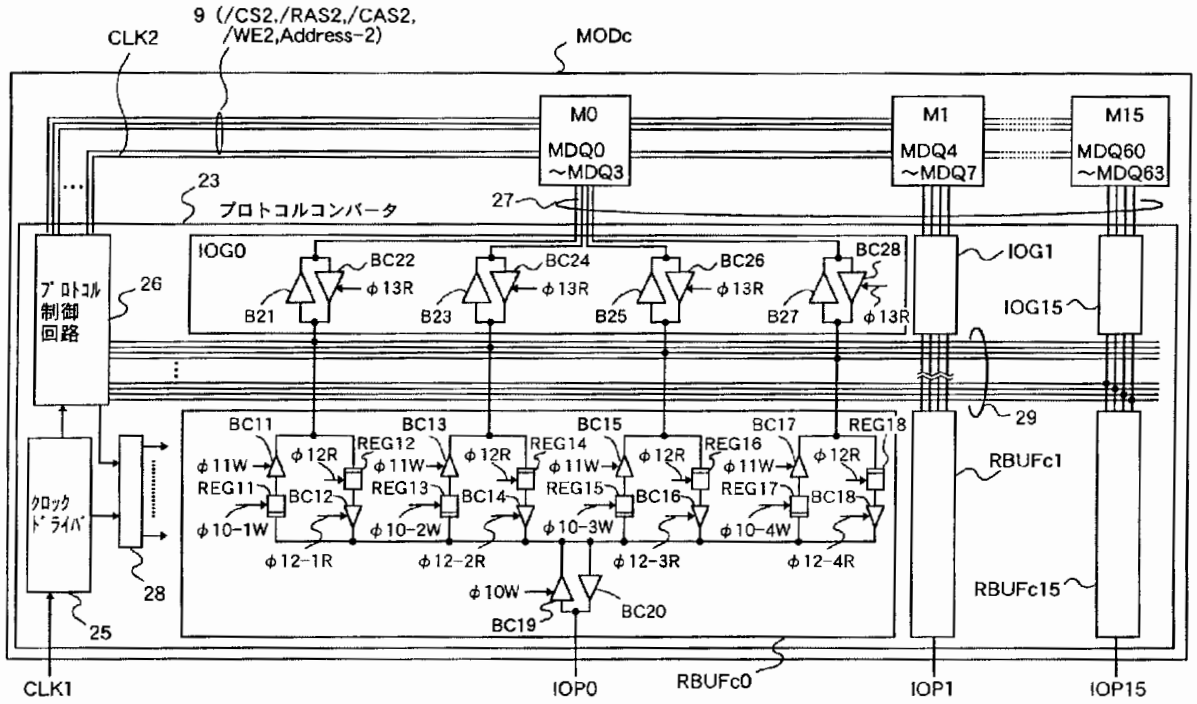
t0: モジュールアクティブコマンド（from バスコントローラ）
t1: メモリチップアクティブコマンド（from 制御信号バッファ）
t2-t3: モジュールライトコマンド＆ライトデータ（from バスコントローラ）
t4: メモリチップライトコマンド＆ライトデータ（from 制御信号バッファ）
t5: モジュールリードコマンド（from バスコントローラ）
t6: メモリチップリードコマンド（from 制御信号バッファ）
t7: メモリチップリードデータ（from SDRAM）
t8-t9: モジュールリードデータ（fromメモリモジュール）

第６図



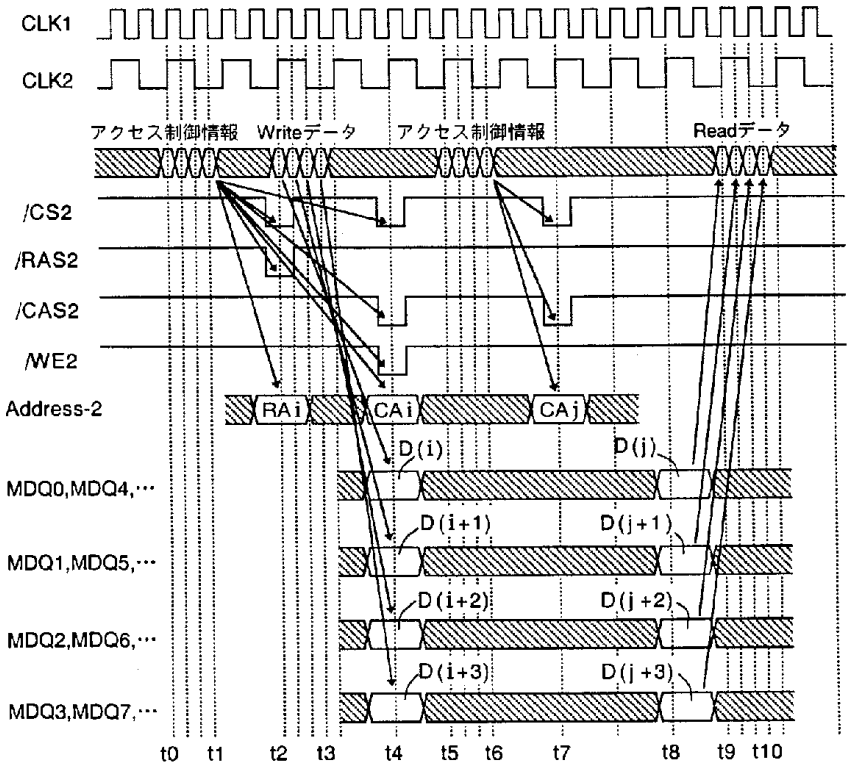
Micron Technology Inc. et al.
Ex. 1010, p. 608

第７図



Micron Technology Inc. et al.
Ex. 1010, p. 609

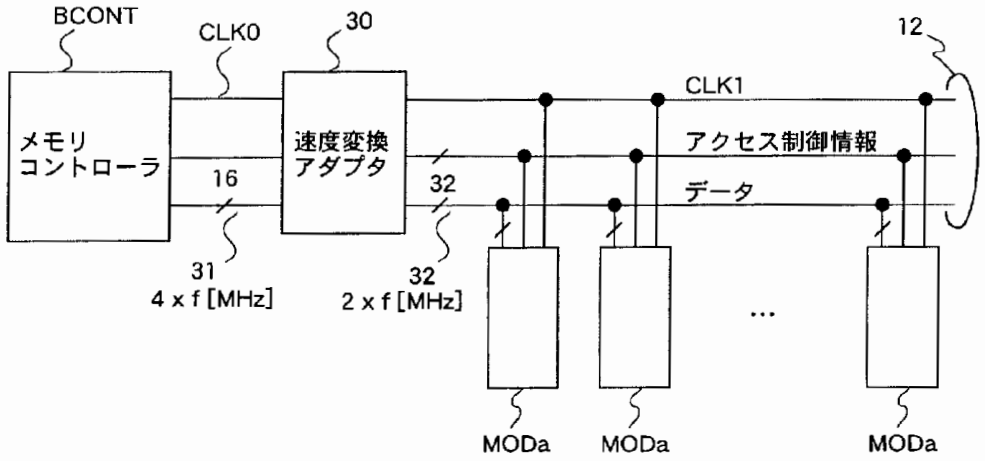
第8図



t0-t1: モジュールライトコマンド（from メモリコントローラ）
t2: メモリチップアクティブコマンド（from プロトコル制御回路）
t2-t3: モジュールライトデータ（from メモリコントローラ）
t4: メモリチップライトコマンド&ライトデータ（from プロトコル制御回路）

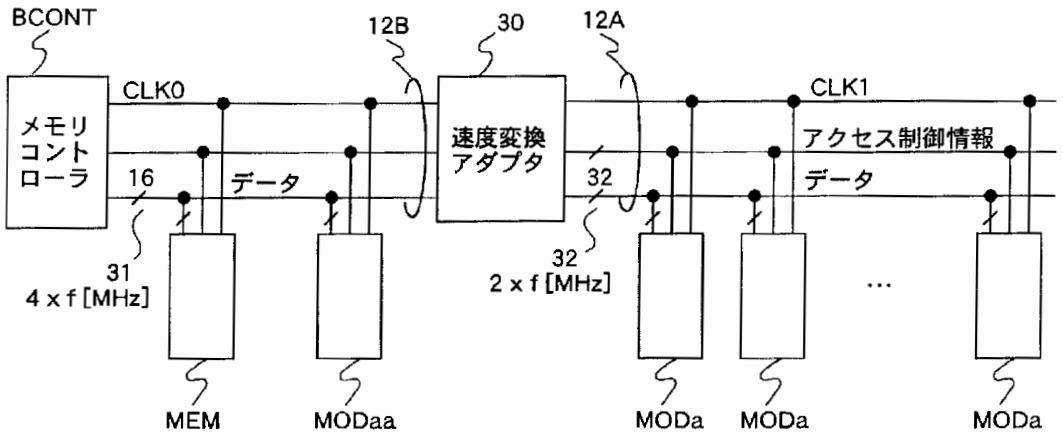
t5-t6: モジュールリードコマンド（from メモリコントローラ）
t7: メモリチップリードコマンド（from プロトコル制御回路）
t8: メモリチップリードデータ（from SDRAMデバイス）
t9-t10: モジュールリードデータ（from メモリモジュール）

第９図



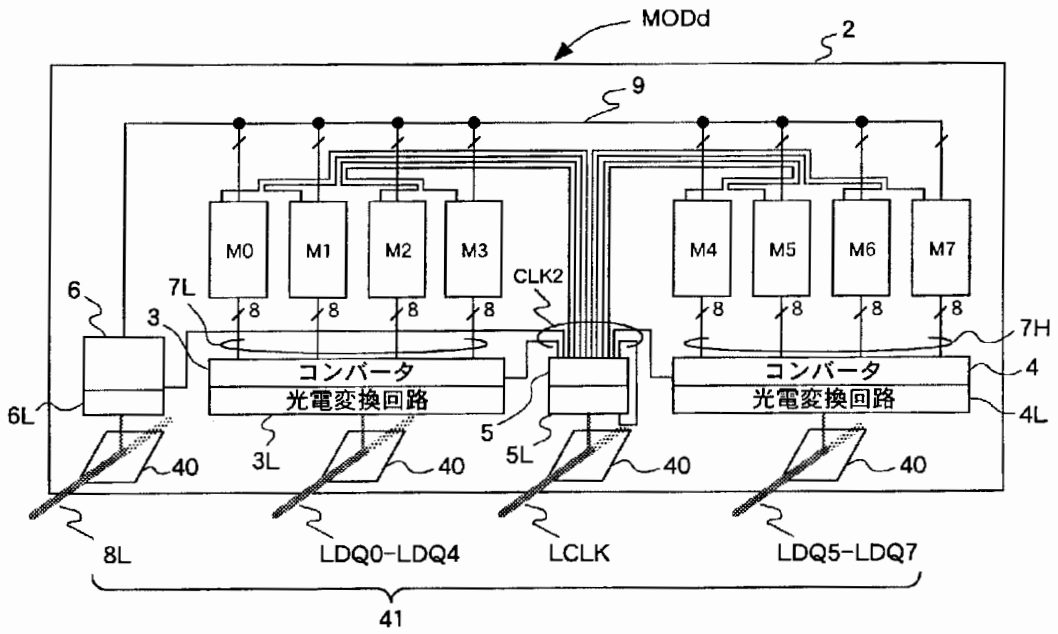
Micron Technology Inc. et al.
Ex. 1010, p. 611

第１０図

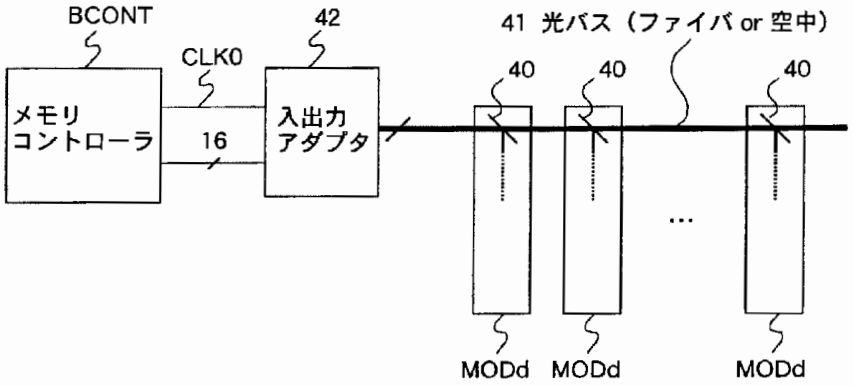


Micron Technology Inc. et al.
Ex. 1010, p. 612

第１１図

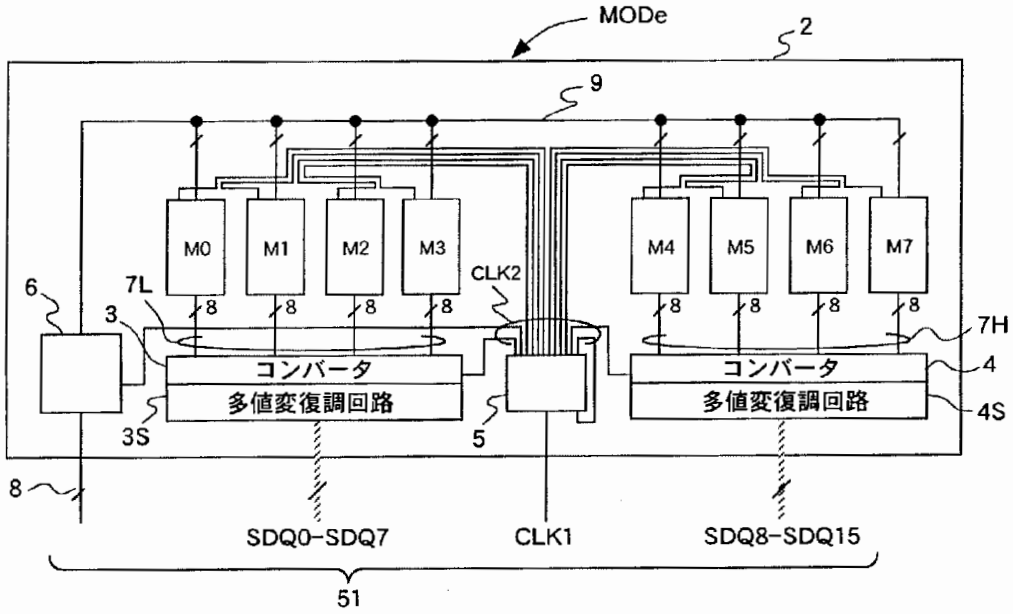

MODd
2
9
M0
M1
M2
M3
M4
M5
M6
M7
CLK2
6
7L
3
8
7H
コンバータ
光電変換回路
4
6L
4L
5
3L
5L
40
8L
LDQ0-LDQ4
LCLK
LDQ5-LDQ7
41

第１２図



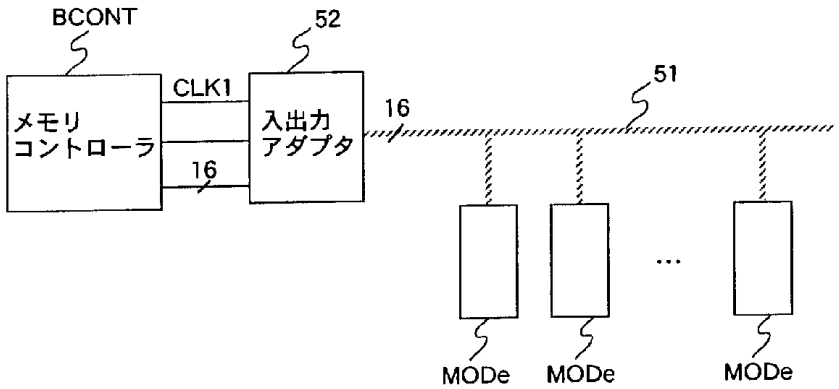
Micron Technology Inc. et al.
Ex. 1010, p. 614

第１３図



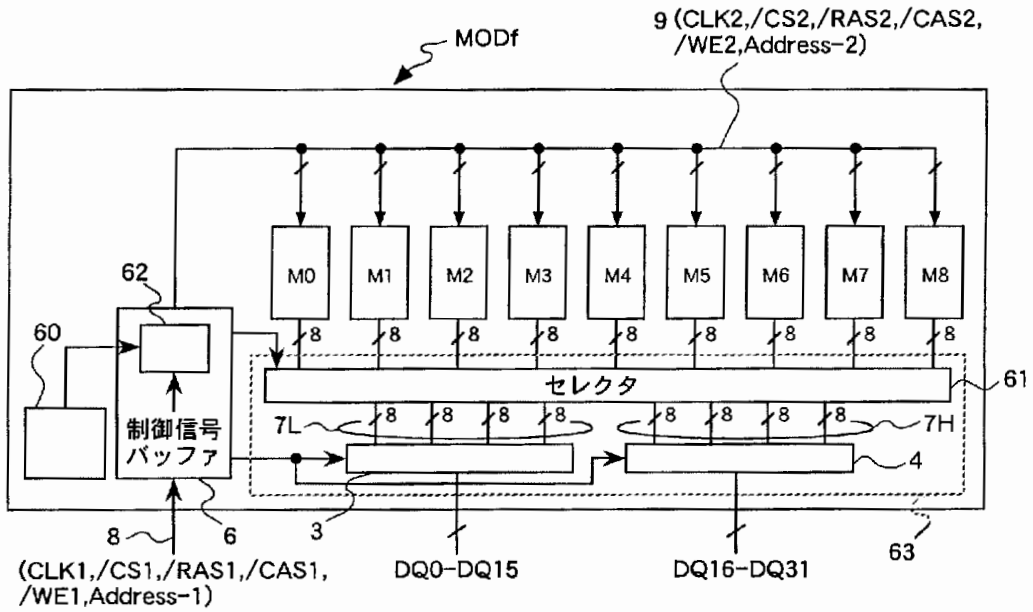
Micron Technology Inc. et al.
Ex. 1010, p. 615

第１４図

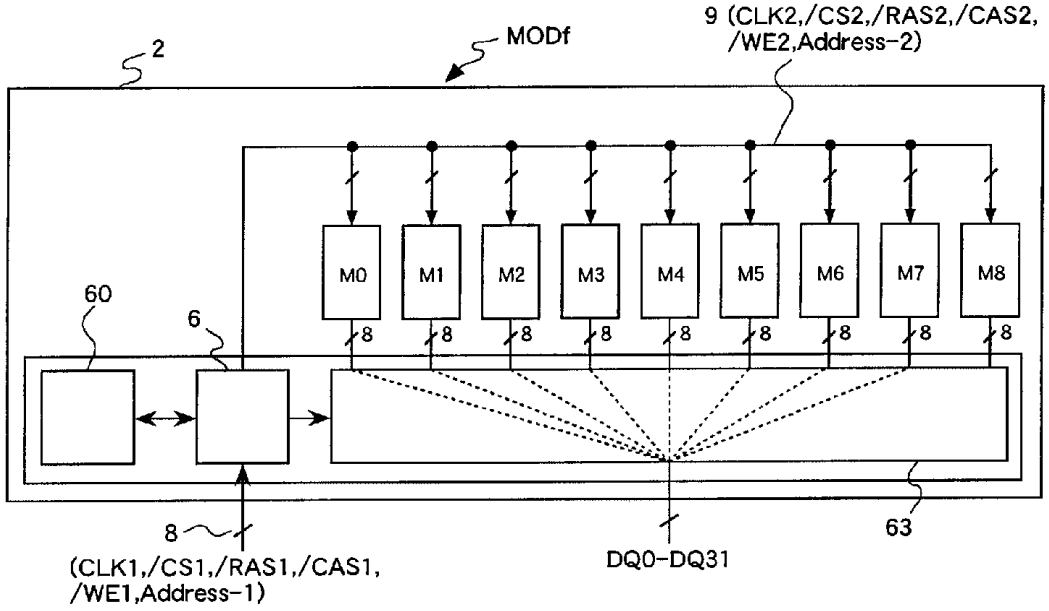


Micron Technology Inc. et al.
Ex. 1010, p. 616

第１５図

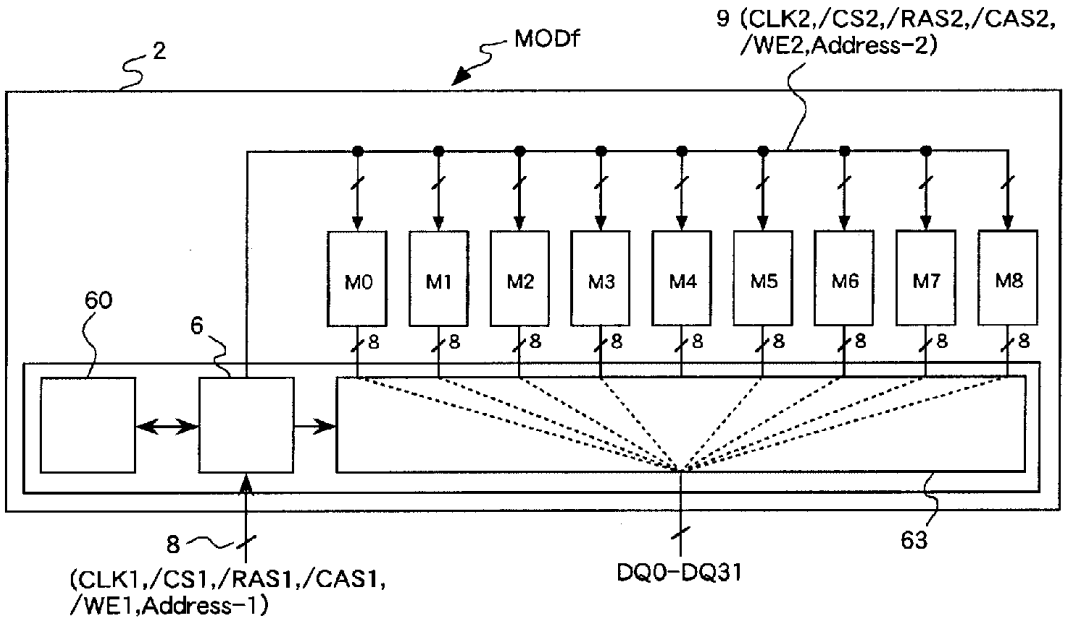

MODf
9 (CLK2,/CS2,/RAS2,/CAS2,
/WE2,Address-2)
M0
M1
M2
M3
M4
M5
M6
M7
M8
62
60
8
セレクタ
61
7L
7H
制御信号
バッファ
4
3
8
6
63
(CLK1,/CS1,/RAS1,/CAS1,
/WE1,Address-1)
DQ0-DQ15
DQ16-DQ31

第１６図

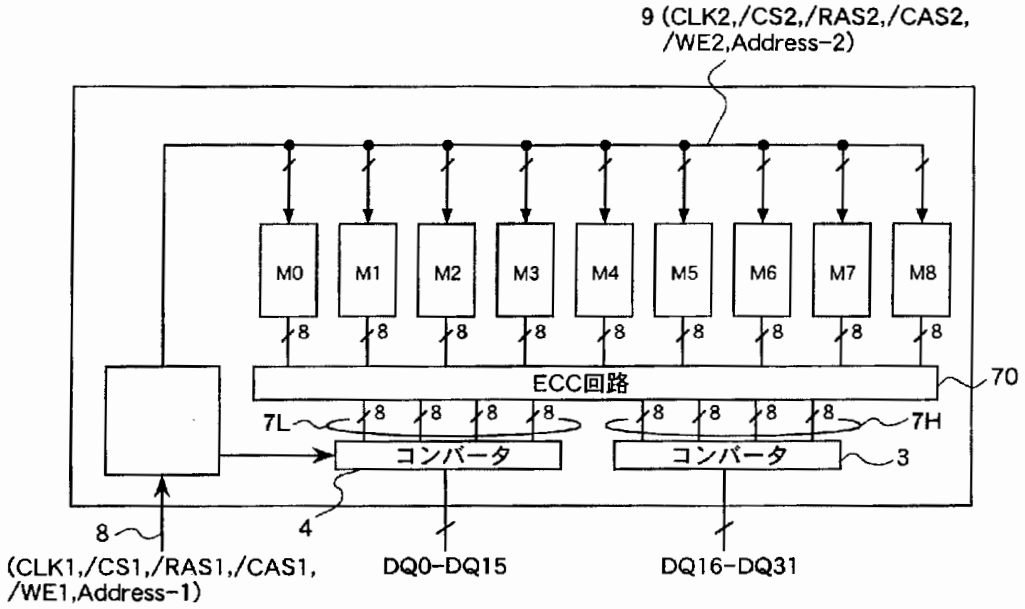


Micron Technology Inc. et al.
Ex. 1010, p. 618

第１７図

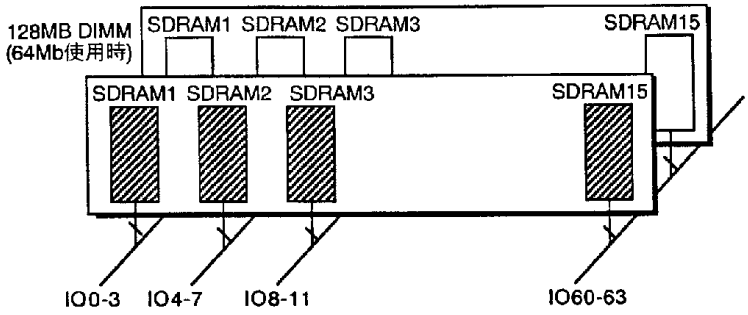

2
MODf
9 (CLK2,/CS2,/RAS2,/CAS2,
/WE2,Address-2)
M0
M1
M2
M3
M4
M5
M6
M7
M8
8
60
6
8
(CLK1,/CS1,/RAS1,/CAS1,
/WE1,Address-1)
DQ0-DQ31
63

第１８図

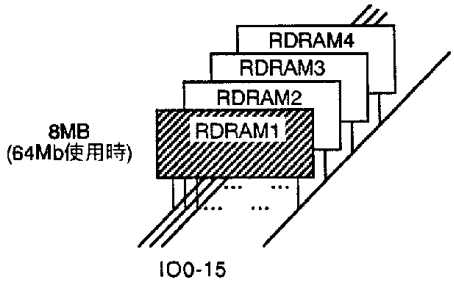


Micron Technology Inc. et al.
Ex. 1010, p. 620

第１９図


128MB DIMM
(64Mb使用時)
SDRAM1 SDRAM2 SDRAM3
SDRAM15
SDRAM1 SDRAM2 SDRAM3
SDRAM15
IO0-3
IO4-7
IO8-11
IO60-63

第２０図


RDRAM4
RDRAM3
RDRAM2
RDRAM1
8MB
(64Mb使用時)
IO0-15

INTERNATIONAL SEARCH REPORT

International application No.
PCT/JP98/02866

A. CLASSIFICATION OF SUBJECT MATTER
Int.Cl[6] G06F12/06

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
Int.Cl[6] G06F12/06

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Jitsuyo Shinan Koho 1922-1996 Toroku Jitsuyo Shinan Koho 1994-1998
Kokai Jitsuyo Shinan Koho 1971-1998 Jitsuyo Shinan Toroku Koho 1996-1998

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X Y A	JP, 62-121978, A (Hitachi ULSI Engineering Corp.), 3 June, 1987 (03. 06. 87) (Family: none)	1-2, 6 3-9, 11, 17 10, 12-16
Y A	JP, 8-77097, A (Ricoh Co., Ltd.), 22 March, 1996 (22. 03. 96) (Family: none)	3-5, 7-9, 17 10-12
Y	JP, 3-286234, A (Matsushita Electric Industrial Co., Ltd.), 17 December, 1991 (17. 12. 91) (Family: none)	4, 5
Y	JP, 1-171047, A (Fujitsu Ltd.), 6 July, 1989 (06. 07. 89) (Family: none)	5, 17
Y A	JP, 63-29357, B2 (Sony/Tektronix Corp.), 13 June, 1988 (13. 06. 88) (Family: none)	7-9 10-16

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:
"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier document but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed
"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
22 September, 1998 (22. 09. 98)	6 October, 1998 (06. 10. 98)
Name and mailing address of the ISA/ Japanese Patent Office Facsimile No.	Authorized officer Telephone No.

Form PCT/ISA/210 (second sheet) (July 1992)

国際調査報告

国際出願番号　PCT／JP98／02866

A.　発明の属する分野の分類（国際特許分類（IPC））
Int Cl⁶ G06F12／06

B.　調査を行った分野

調査を行った最小限資料（国際特許分類（IPC））
Int Cl⁶ G06F12／06

最小限資料以外の資料で調査を行った分野に含まれるもの
日本国実用新案公報　1922－1996年
日本国公開実用新案公報　1971－1998年
日本国登録実用新案公報　1994－1998年
日本国実用新案登録公報　1996－1998年

国際調査で使用した電子データベース（データベースの名称、調査に使用した用語）

C.　関連すると認められる文献

引用文献のカテゴリー＊	引用文献名　及び一部の箇所が関連するときは、その関連する箇所の表示	関連する請求の範囲の番号
X Y A	JP，62－121978，A（日立超エル・エス・アイエンジニアリング）　3．6月．1987（03．06．87）　（ファミリーなし）	1-2, 6 3-9, 11, 17 10, 12-16
Y A	JP，8－77097，A（株式会社リコー） 22．3月．1996（22．03．96）　（ファミリーなし）	3-5, 7-9, 17 10-12
Y	JP，3－286234，A（松下電器産業株式会社） 17．12月．1991（17．12．91）（ファミリーなし）	4, 5
Y	JP，1－171047，A（富士通株式会社） 6．7月．1989（06．07．89）　（ファミリーなし）J	5, 17

☒ C欄の続きにも文献が列挙されている。　☐ パテントファミリーに関する別紙を参照。

＊　引用文献のカテゴリー
「A」特に関連のある文献ではなく、一般的技術水準を示すもの
「E」先行文献ではあるが、国際出願日以後に公表されたもの
「L」優先権主張に疑義を提起する文献又は他の文献の発行日若しくは他の特別な理由を確立するために引用する文献（理由を付す）
「O」口頭による開示、使用、展示等に言及する文献
「P」国際出願日前で、かつ優先権の主張の基礎となる出願の日の後に公表された文献
「T」国際出願日又は優先日後に公表された文献であって出願と矛盾するものではなく、発明の原理又は理論の理解のために引用するもの
「X」特に関連のある文献であって、当該文献のみで発明の新規性又は進歩性がないと考えられるもの
「Y」特に関連のある文献であって、当該文献と他の1以上の文献との、当業者にとって自明である組合せによって進歩性がないと考えられるもの
「&」同一パテントファミリー文献

国際調査を完了した日	国際調査報告の発送日
22．09．98	06.10.98

国際調査機関の名称及びあて先	特許庁審査官（権限のある職員）	
日本国特許庁（ISA／JP） 郵便番号100－8915 東京都千代田区霞が関三丁目4番3号	漆原孝治　印	5B　9366
	電話番号　03－3581－1101　内線　3546	

様式PCT／ISA／210（第2ページ）（1992年7月）

Micron Technology Inc. et al.
Ex. 1010, p. 623

国際調査報告	国際出願番号　ＰＣＴ／ＪＰ９８／０２８６６

Ｃ（続き）．　関連すると認められる文献

引用文献のカテゴリー＊	引用文献名　及び一部の箇所が関連するときは、その関連する箇所の表示	関連する請求の範囲の番号
Ｙ Ａ	Ｐ，６３－２９３５７，Ｂ２（ソニー・テクトロニクス株式会社） １３．６月．１９８８（１３．０６．８８）（ファミリーなし）	7-9 10-16

様式ＰＣＴ／ＩＳＡ／２１０（第２ページの続き）（１９９２年７月）

Micron Technology Inc. et al.
Ex. 1010, p. 624

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
25 July 2002 (25.07.2002)

PCT

(10) International Publication Number
WO 02/058069 A2

(51) International Patent Classification[7]: G11C 5/00

(21) International Application Number: PCT/US02/01575

(22) International Filing Date: 17 January 2002 (17.01.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/262,472 17 January 2001 (17.01.2001) US

(71) Applicant: HONEYWELL INTERNATIONAL INC. [US/US]; Law Department, 101 Columbia Road, P.O. Box 2245, Morristown, NJ 07962-2245 (US).

(72) Inventor: ANDERSON, David; 304 S. Market, Rock Port, MO 64482 (US).

(74) Agents: HALSNE, Eric et al.; Honeywell International Inc., 101 Columbia Road, P.O. Box 2245, Morristown, NJ 07962-2245 (US).

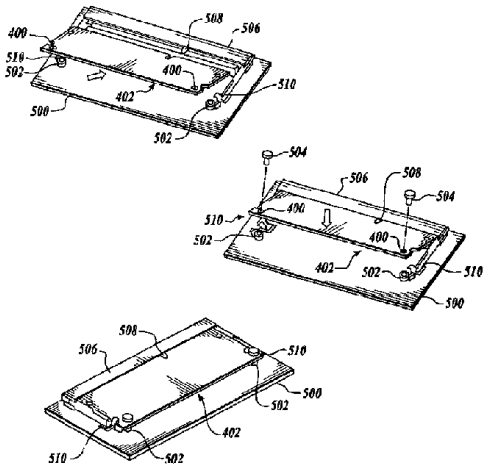
(81) Designated State *(national)*: JP.

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

Published:
— *without international search report and to be republished upon receipt of that report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: ENHANCED MEMORY MODULE ARCHITECTURE



(57) Abstract: A memory module architecture that supports Flash and static memory devices in addition to dynamic memory devices. The module architecture of the present invention preferably redefines standard application of chip select signals on existing module architectures to provide requisite signaling to support Flash and static RAM devices. Use of serial presence detect signaling features of standard memory modules is also modified to provide desired identity and parameters of such an enhanced module. A further aspect of the present invention provides for a support structure to immobilize the loose edge of the memory module opposite the electrical edge connector of the module to further enhance the module's resistance to vibration and mechanical shock by immobilizing the module with respect to rotation about the key in the socket connector.

WO 02/058069 A2

ENHANCED MEMORY MODULE ARCHITECTURE

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0001] The invention relates to memory module architectures and in particular relates to an enhanced memory module configuration to provide additional flexibility in memory module applications in conjunction with PowerPC and other system applications.

2. Discussion of Related Art

[0002] Present-day computing systems often utilize standardized memory modules to add memory capacity to a system. A memory module is generally comprised of a number of memory chips mounted on a printed circuit card adapted for insertion into a standard carrier or socket connector. The module is easily added or removed from a system to permit flexible configuration of memory capacity. The module and socket adhere to a standardized pin assignment for exchange of signals between a system bus and the memory chips. Signals are routed from the carrier connector to the memory chips in accordance with the particular standard used for the module and the particular memory chips selected by the memory module manufacturer.

[0003] A number of standards are defined for such modules. Exemplary of one such standard is a so-called 144-pin Small Outline DIMM module (also referred to as "SO-DIMM") as defined by the JEDEC JESD21-C standard specification (available online at www.jedec.org). The 144-pin SO-DIMM provides a convenient, compact design for use in space-constrained applications such as portable computing devices or embedded systems. The module is a small printed circuit card with memory devices mounted thereon and one edge that provides electrical connectivity of up to 144 signals when the module is inserted into a mated connector on the main system board (i.e., the system motherboard). The printed circuit card signal traces connect standard signals assigned to the 144 pins of the module's edge connector to appropriate pins of the memory chips mounted thereon. A memory controller on the system board exchanges these standard memory control signals with the module through the mated connector.

[0004] In particular, the 144-pin standard module provides for memory chips to be mounted on one or both physical sides of the printed circuit board. A select signal is typically applied to the module by a CPU or memory controller to select the memory chips on one or the other side of the memory module for exchange of data and control signals. This select signal is often referred to as a "row select" signal in regards to the 144-pin standard configuration. When the "front" side is selected by the memory controller, only those chips respond to the exchange of signals with the system board. Similarly, when the "back" side is selected, only those memory chips respond to the memory controller signals.

[0005] Most memory modules and carriers define signal connections for standard DRAM memory devices (dynamic random access memory) mounted on the module. Such signal definitions include signals common to most DRAM devices including, in particular, SDRAMs (synchronous DRAM devices). Other forms of memory, in particular asynchronous memories, are not generally accommodated by the standard signal paths provided for on the standard DIMM memory module. For example, present standard DIMM modules such as the 144-pin SO-DIMM do not support so-called Flash memory devices. Neither are static RAMs (SRAM devices) generally supported by such standard DIMM modules.

[0006] This presents a particular problem to embedded applications where a wide variety of memory devices are often applied to resolve special problems in particular applications. For example, the non-volatile nature of Flash memory devices makes them particularly useful in many embedded applications for storage of computer program instructions while a standard SDRAM may be used for storage and retrieval of program data. Since no single standard memory module supports such a variety of memory devices, it is often necessary to design embedded applications for multiple memory module architectures and associated busses and/or to redesign the system as the memory requirements change or new memory components become available. For example, it is common in many such embedded applications to design a first memory bus for exchange of signals with synchronous memory components (i.e., SDRAMs) and a separate slower memory bus for asynchronous devices such as Flash or static RAMs. This added complexity adds cost and size to the design and complicates portability of a system design to multiple applications that may have different needs for memory configuration.

Micron Technology Inc. et al.
Ex. 1010, p. 627

[0007] Furthermore, it is a problem for many such memory module designs to be used in embedded applications where there is significant vibration or mechanical shock. Standard DIMM modules insert into a mated socket at one edge. Some DIMM modules are designed to latch or lock on two other sides of the module. However, the edge opposing the socket is often unsupported and may lead to failures in applications of high vibration and mechanical shock. In particular it has been found that under certain common vibrational modes, the SO-DIMM module may twist about a pivot point at the key in the DIMM socket. Sufficient twisting due to such vibration can lead to unacceptable, premature DIMM module failure.

[0008] It is evident from the above discussion that a need exists for an improved memory module design that permits a wider variety of memory device types to be used in a standard DIMM socket on a common bus. It is further evident that a need exists for an improved mechanical design for a memory module that better withstands high vibration environments or environments that may entail significant mechanical shock.

SUMMARY OF THE INVENTION

[0009] The present invention solves the above and other problems, thereby advancing the state of the useful arts, by providing a modified memory module design that accommodates a wider variety of memory devices. In particular, the modified memory module architecture of the present invention redefines a number of standard signals to permit utilization of Flash and static memory components in a DIMM module while maintaining backward compatibility to standard DRAM and SDRAM DIMM modules. Still more specifically in one exemplary preferred embodiment, the present invention preferably utilizes row select signals ("CS(0:1)") of a standard 144-pin SO-DIMM module to provide requisite select and control signals for Flash and static memory devices.

[0010] Still more specifically, where a memory controller is capable of multiplexing both synchronous and asynchronous memory device accesses over the same signal paths, the present invention permits use of Flash and other special memory modules that are physically compatible with SO-DIMM configurations in conjunction with typical SDRAM (or other DRAM) SO-DIMM modules over the same multiplexed memory controller bus signals. For example, the Motorola family of embedded PowerPC microcomputers (MPC824x and others) includes such a memory controller -- specifically the MPC107 memory controller. The MPC106 or MPC107 memory controller device

Micron Technology Inc. et al.
Ex. 1010, p. 628

provides such a memory controller feature and is integrated into the MPC824x chip die and available as a separate integrated circuit component. Systems using the MPC106 or MPC107 memory controller as a separate component may also benefit from the present invention. Further, other similarly featured memory controllers are commercially available and systems using such similarly featured memory controller may beneficially apply the features of the present invention.

[0011] In one particularly useful embodiment, the signals of a JEDEC JESD21-C 144-pin SO-DIMM standard memory module are modified to enable support of Flash and SRAM memory devices on the modified standard DIMM module. In particular, the chip select (row select) signals are modified to provide the requisite selection of Flash and SRAM memory components in addition to standard SDRAM (or DRAM) memory components. Setting of jumpers on the system board preferably indicate the presence of such additional memory device types in the SO-DIMM sockets of the system and adapt the signal exchange with a memory controller to conform to the signaling requirements of the special module. A single bus structure supports both synchronous and asynchronous memory device signals and couples the memory controller to multiple SO-DIMM sockets. Each DIMM socket may be populated with a Flash or SRAM special memory module or with a standard SDRAM or DRAM memory module. The design may then be easily ported to other applications requiring a different compliment of memory device types.

[0012] In this first exemplary preferred embodiment, a variety of memory configurations are achievable. In this first exemplary preferred embodiment, the system utilizes a DIMM module similar to that defines as the 144-pin SO-DIMM standard referenced above. In such memory modules, the module may be populated on one or both sides of the module printed circuit card with memory devices. All memory devices on one side typically respond to the same select signal. The select signal is usually referred to as a "row select" as noted above. As used herein, the terms "row" and "side" are used synonymously with respect to such SO-DIMM modules. Those skilled in the art will note however that it is not required that all memory devices on one side of a module respond to a single select signal. Rather, in accordance with the module standard, the physical place of the memory devices on the module is irrelevant with respect to the select signals. Use of the terms "row" and "side" as essentially synonymous is therefore not intended to limit the invention to a particular physical layout of memory modules. Rather, the terms

Micron Technology Inc. et al.
Ex. 1010, p. 629

are both used to broadly refer to a first subset of memory devices on a module that respond to a first row select signal and a second subset of memory devices that respond to a second row select signal.

[0013] With two DIMM sockets in an exemplary system configuration, each socket may be populated with a standard SDRAM SO-DIMM where the DIMM is limited to single row configurations (i.e., SDRAM memory devices responding to only one row select of the module). Alternatively, only one DIMM socket may be populated with a double-row SO-DIMM module. The double-row DIMM may provide SDRAM devices on both rows, SDRAM on one row and Flash or other asynchronous devices on the other. In yet another alternative, a first single row DIMM may provide SDRAM memory devices and a second single row DIMM may provide Flash or other asynchronous memory devices. Jumper configurations on the system printed circuit board can then be used to route the row select signals as appropriate for the particular application.

[0014] In an alternative embodiment of the present invention, use of the serial presence detect ("SPD") signaling standard of DIMMs provides the required identity of the special modules supporting the extended range of memory device types. The SPD feature of memory modules provides for a simple serial interface (usually I^2C protocol compliant) to return to the system certain identification and configuration information regarding the memory module. The memory controller of the system is then programmed to adapt its signaling as required for the sensed memory configuration. An application specific integrated circuit (ASIC) could be used in addition to a memory controller or in place of a memory controller to provide desired address processing. In general, a system CPU would preferably sense the configuration of memory modules in the sockets by reading and decoding the SPD information and program the memory controller and/or ASIC device to provide appropriate address decoding and timing according to the specific compliment of memory modules found in the system.

[0015] Another aspect of the present invention provides a support structure for the DIMM edge opposite the connector edge to enhance the module's resistance to vibration and mechanical shock. The present invention provides a support structure associated with the edge of a SIMM module opposite the connector edge of the DIMM. The support structure prevents a module inserted in the SO-DIMM socket from rotating or twisting within the socket connector. In particular, the present invention preferably prevents the module from twisting about an axis centered at the key of the module socket.

Micron Technology Inc. et al.
Ex. 1010, p. 630

[0016] In one exemplary preferred embodiment representing the best presently known mode of practicing the invention, the module is adapted with holes at the corners, each overlaying a swaged extension nut or standoff on the motherboard or system board. The holes and associated nut or standoff are preferably positioned at the corners of the module opposite the connector edge of the module. The module is inserted in the mated connector and preferably secured with screws or other fasteners to the standoff or extension nut to immobilize the module and prevent undesired twisting of the module in the socket.

[0017] A first feature of the invention provides for a memory module comprising: a printed circuit assembly having connector pads at one edge of the assembly; and a plurality of memory devices mounted on the assembly and electrically coupled to the connector pads such that the printed circuit assembly is adapted to support both synchronous and asynchronous types of the memory devices.

[0018] Another aspect further provides that each of the plurality of memory devices is a synchronous dynamic random access memory device.

[0019] Another aspect further provides that each of the plurality of memory devices is a synchronous Flash memory device.

[0020] Another aspect further provides that each of the plurality of memory devices is an asynchronous Flash memory device.

[0021] Another aspect further provides that each of the plurality of memory devices is an asynchronous static random access memory device.

[0022] Another aspect further provides that each of the plurality of memory devices is an asynchronous fast static random access memory device.

[0023] Another aspect further provides that each of the plurality of memory devices is an asynchronous low power static random access memory device.

[0024] Another aspect further provides that the connector pads include: a first select signal connector pad that selects a first subset of memory devices mounted on the module when a signal is applied thereto such that the first subset of memory devices are synchronous memory devices; and a second select signal connector pad that selects a second subset of memory devices mounted on the module when a signal is applied thereto such that the second subset of memory devices are asynchronous memory devices.

[0025] Another aspect further provides that the connector pads include: a first select signal connector pad that selects a first subset of memory devices mounted on the

Micron Technology Inc. et al.
Ex. 1010, p. 631

module when a signal is applied thereto such that the first subset of memory devices are synchronous dynamic random access memory devices; and a second select signal connector pad that selects a second subset of memory devices mounted on the module when a signal is applied thereto such that the second subset of memory devices are synchronous Flash memory devices.

[0026] A second feature of the invention provides for a system comprising: a system board; a memory bus adapted for exchanging signals between a memory controller and both synchronous and asynchronous memory devices; a memory controller on the system board coupled to the memory bus such that the memory controller is capable of generating signals for control of both synchronous and asynchronous memory devices and such that the memory controller is capable of multiplexing the signals on the memory bus; a first socket connector on the system board for receiving a first memory module such that the first socket connector is coupled to the memory controller through the memory bus; and a first memory module inserted in the first socket connector and electrically coupled to the memory controller such that the first memory module includes a plurality of synchronous or asynchronous memory devices.

[0027] Another aspect further provides for a jumper on the system board for configuring signals exchanged between the memory controller and the first memory module in accordance with the type of memory devices on the first memory module.

[0028] Another aspect further provides that the first memory module provides serial presence detect information used in conjunction with the memory controller to identify the type of memory devices included on the first memory module.

[0029] Another aspect provides for a second socket connector on the system board for receiving a second memory module such that the second socket connector is coupled to the memory controller through the memory bus; and a second memory module inserted in the second socket connector and electrically coupled to the memory controller such that the second memory module includes a plurality of synchronous or asynchronous memory devices.

[0030] Another aspect further provides for a jumper on the system board for configuring signals exchanged between the memory controller and the second memory module in accordance with the type of memory devices on the first memory module and on the second memory module.

Micron Technology Inc. et al.
Ex. 1010, p. 632

[0031] Another aspect further provides that the first memory module provides serial presence detect information used in conjunction with the memory controller to identify the type of memory devices included on the first memory module, and that the second memory module provides serial presence detect information used in conjunction with the memory controller to identify the type of memory devices included on the second memory module.

[0032] A third feature of the invention provides for a system comprising: a system board; a socket connector on the system board for receiving a memory module such that the socket connector includes a key; a memory module having a connector edge inserted in the socket connector and having an opposing edge opposite the connector edge such that the memory module has a notch mated to the key when the memory module is inserted in the socket connector; and a memory module retainer adapted to substantially immobilize the opposing edge with respect to rotation about the key.

[0033] Another aspect further provides that the memory module retainer comprises: a nut affixed to the system board; a hole in the memory module along the opposing edge and aligned with the nut; and a screw inserted through the hole into the nut to substantially immobilize the opposing edge with respect to rotation about the key.

[0034] Another aspect provides that the nut is a swaged extension nut.

[0035] Another aspect provides that the memory module retainer comprises: a half card-cage affixed to the system board such that the half card-cage includes a channel for receiving the opposing edge of the memory module to substantially immobilize the opposing edge with respect to rotation about the key.

[0036] Another aspect provides that the memory module retainer comprises: a standoff pin affixed to the system board; a hole in the memory module along the opposing edge and aligned with the pin such that the standoff pin extends through the hole and locks when the memory module is completed inserted in the socket connector to substantially immobilize the opposing edge with respect to rotation about the key.

[0037] Another aspect provides that the memory module retainer comprises: a standoff pin affixed to the socket connector; a hole in the memory module aligned with the standoff pin such that the standoff pin extends through the hole and when the memory module is completed inserted in the socket connector to substantially immobilize the opposing edge with respect to rotation about the key.

BRIEF DESCRIPTION OF THE DRAWINGS

Micron Technology Inc. et al.
Ex. 1010, p. 633

[0038] Figure 1 is a block diagram of an exemplary preferred embodiment of a system using enhanced memory modules in accordance with the present invention.

[0039] Figure 2 is a schematic diagram showing relevant portions of an exemplary preferred embodiment of a DIMM socket connector with modified row select signals in accordance with the present invention.

[0040] Figure 3 is a block diagram depicting a typical system as presently practiced in the art with separate bus structures for support of different memory types.

[0041] Figure 4 is a dimensioned diagram of a memory module with holes for securing the module the system board in accordance with one exemplary preferred embodiment of he present invention.

[0042] Figure 5 is a sequence of diagrams depicting insertion of a memory module into a connector socket and securing the module to the system board in accordance with one exemplary preferred embodiment of he present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0043] While the invention is susceptible to various modifications and alternative forms, a specific embodiment thereof has been shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that it is not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

[0044] Figure 1 is a block diagram of a system 100 in which an embedded microcomputer interfaces with a variety of memory modules 116 and 118 via memory bus 150 and standard memory module socket connectors 114. Embedded microcomputer 102 preferably includes at least a processor 104 and a memory controller 106 to control and interact with memory modules 116 and 118. In one exemplary preferred embodiment of such a system, embedded microcomputer 102 may be a Motorola PowerPC MPC8245 or other similar members of that embedded microcomputer family. The MPC8245 includes a memory controller for interaction with, and control of, a variety of memory devices.

[0045] Figure 3 shows a system 300 using the same embedded microcomputer 102 adapted as known in the art to interact with standard SO-DIMM SDRAM (or DRAM) memory modules through a synchronous bus structure 350 while interacting with other memory devices (Flash 318 or SRAM 320) through a slower asynchronous memory

Micron Technology Inc. et al.
Ex. 1010, p. 634

bus 352. As noted above, such a structure requires each system application to uniquely adapt to a particular memory configuration. The system cannot be easily re-used for different applications requiring different memory configurations.

[0046] Returning now to figure 1, by contrast and in accordance with the present invention, the memory control signals are all directed over the same multiplexed memory bus 150 to memory modules 116 and 118. The memory modules 116 and 118 are preferably adapted to SO-DIMM standard mechanical dimensions. As noted above, memory controller 106 is capable of multiplexing both synchronous memory control signals and asynchronous memory control signals over the same bus, Such a memory controller is embedded within an MPC8245 and other embedded microcomputers. Further, as noted above, those skilled in the art will readily recognize equivalent system architectures in which the memory controller element 106 is a separate and distinct integrated circuit component not embedded within the same chip die as the processor. For example, the MPC106 or MPC107 memory controller as well as other commercially available memory controller devices provide the desired multiplexing of synchronous and asynchronous memory signals over a common memory interface bus.

[0047] The memory bus signals are preferably directed through standard DIMM memory module connectors 114 to any of a variety of DIMM memory modules. For example, standard 144-pin SO-DIMM memory module 116 may provide standard DRAM or SDRAM memory devices. Flash/SRAM memory module 118 is adapted in accordance with the present invention to provide a 144-pin interface similar to that of the standard 144-pin DIMM 116. DIMM memory module 118 is adapted in accordance with the present invention to apply control signals from memory controller 106 to Flash or static memory devices (i.e., SRAM devices including low power SRAM, high-speed SRAM, etc.). Further, as noted above, a custom DIMM module may be utilized that provides SDRAM (or other DRAM) devices responding to a first row select signal and Flash or other asynchronous devices that respond to a second row select signal.

[0048] Other components may be embedded within microcomputer 102 including, for example, a DMA controller 108, integrated I/O peripheral device controllers 110, and external interface bus controllers 112 (i.e., PCI, PCMCIA, etc.). Those skilled in the art will recognize that the devices shown in figure 1 are intended as one exemplary embodiment of a system that may advantageously apply the enhanced memory architecture of the present invention.

Micron Technology Inc. et al.
Ex. 1010, p. 635

[0049] As noted above, one standard DIMM configuration is defined by JEDEC specification JESD21-C as a 144-pin SO-DIMM memory module. This particular memory module is useful in a number of applications due to its small footprint. The small module is effective in portable systems including avionics systems and other embedded systems to conserve space required for memory expansion.

[0050] In one exemplary preferred embodiment of the invention, the signals of such a standard 144-pin DIMM are adapted to permit use of a wider variety of memory devices in the memory module configuration. Those skilled in the art will recognize that the features of the present invention may be easily adapted to other memory module configurations and standards. The exemplary preferred embodiment of the invention as a 144-pin SO-DIMM structure is not intended to limit the applicability of the invention to other memory modules. Rather, it is intended to express one particularly useful application for embedded and portable system applications.

[0051] The flexible memory architecture of the present invention allows a single system design to be applied to a number of different applications potentially requiring different memory configurations. By way of example, a first system may provide a large system such as an avionics Primary Flight Display. Such systems are often referred to as "Level A" or "mission critical" in the avionics and aerospace industries. Such a large system may use of a full 72-bit (64 + 8 check bits) ECC SDRAM module in a first DIMM module socket of the system to enhance system reliability through error correction. This memory module would preferably be a single row SO-DIMM module selected by a "front-side" select pin of a standard DIMM socket. As noted herein above, use of the term "side", "front-side" or "back-side" is to be understood as synonymous with "row" as defined for the standard 144-pin SO-DIMM module of an exemplary preferred embodiment. This SDRAM module would serve as storage for dynamic data and scratchpad memory for the system. Those skilled in the art will note that the number and size of memory devices presently commercially available and required to implement a 72-bit ECC module would consume both sides of the module. However, expected advancements and market demands will easily allow for such a memory module in a single row configuration. The application program instructions in such a system would preferably be stored in a Flash memory module inserted in the second module socket. The Flash module would also preferably be a single row module selected by a "back-side" select pin on the SO-DIMM standard socket. The width, depth, and timing characteristics

Micron Technology Inc. et al.
Ex. 1010, p. 636

for Flash module memory accesses are preferably programmed into the memory controller by boot code operable at system startup and reset.

[0052] A second system application of the same system board could provide a very small system, such as a Frequency Control Head for a remote radio. The criticality of such a system application may be such that ECC reliability is not required on the data memory. Also the data memory and application size and performance requirements may be such that a 32-bit bus width is sufficient. In this case the data memory would preferably be 32-bit wide SDRAM totaling perhaps 4 megabytes in an SO-DIMM module populated on one side with SDRAM devices. The SDRAM devices are preferably selected by a "front-side" select pin of the DIMM socket. The application program instructions could be stored in 32-bits of SyncFlash™ (a trademarked technology of Micron Technology, Inc.) on the other side of the same module and responding to the "back-side" select signal. SyncFlash™ memory devices provide Flash memory features but present an addressing and control interface substantially similar to that of SDRAM. Both rows of this custom module may therefore be addressed as though containing standard SDRAM synchronous memory devices. Thus a single custom DIMM module may provide all the necessary memory for this small system application and still uses the same SO-DIMM standard socket.

[0053] A third system application of the same system board could provide the same Control Head using standard Flash and SRAM on different rows of a single custom SO-DIMM. Such a module might use 2 2Mx16-bit asynchronous Flash memory devices on one side for a total of 8MB selected by the second row select pin logically ANDed with the un-inverted highest used address bit. This ANDing can be accomplished by single-gate logic available from several logic vendors. The same custom SO-DIMM module might also use 4 1Mx8-bit asynchronous SRAM memory devices for a total of 4MB. This SRAM would preferably also be selected by the second row select signal ANDed with the inverted highest used address bit. In this preferred embodiment, both the synchronous Flash and the asynchronous SRAM would be selected by the second row ("back-side") select signal with the highest used address bit used to differentiate between the two types of memory.

[0054] The flexible memory architecture of the present invention allows all three systems to use the exact same system board design -- a desirable benefit to ease certification of systems in avionics and aerospace industries as well as other high-

Micron Technology Inc. et al.
Ex. 1010, p. 637

reliability applications. All three systems can use the exact same software if the boot code configures the memory controller from SPD data retrieved from the module(s) (as discussed further herein below) -- another benefit to reliability certification requirements. Further, as memory device packaging evolves, a new SO-DIMM module can be easily adapted to the modified signaling of the present invention. The system board, CPU and memory controller need not be re-designed to utilize improved memory devices. Further, use of standard SO-DIMMs for all memory requirements of such systems helps minimize footprint requirements in space constrained applications.

[0055] The table below describes one exemplary preferred embodiment for the connection of standard memory controller signals generated by a typical memory controller to an enhanced DIMM in accordance with the present invention. The signals of the memory controller in the first three columns are expressed in the table as the signals generated by the memory controller element within a Motorola MPC8245 embedded microcomputer (i.e., the MPC107 memory controller). This memory controller is exemplary of a large number of memory controllers that may be useful in controlling memory modules of the present invention. Those skilled in the art will recognize a variety of other memory controller devices, both as separate devices and integrated within a microcomputer, that are capable of controlling a variety of synchronous and asynchronous memory devices over a common bus.

[0056] The first column indicates the signals used in connecting the memory controller to a standard SDRAM module. The second column indicates the signals used in connecting the memory controller to an asynchronous Flash or SRAM memory module. The third column indicates the signals used in connecting the memory controller to a burst/page Fast Flash or SRAM module. When a signal is not used for a particular module type, "NU" so designates an unused connection.

Micron Technology Inc. et al.
Ex. 1010, p. 638

Synch. SDRAM Module	Asynch. FLASH/ SRAM Module	Burst/Page Fast Flash/ SRAM Module	Connector/Module Pin Numbers
D(0:31)	D(0:31)	D(0:31)	138,136,134,132,128,126,124,122,100, 98,96,94,90,88,86,84,54,52,50,48,44, 42,40,38,20,18,16,14,10,8,6,4
D(32:63)	D(32:63)	D(32:63)	137,135,133,131,127,125,123,121,99, 97,95,93,89,87,85,83,53,52,49,47,43,41,39,37,19,17,15,13,9,7,5,3
DQM_n(0:7)	NU	NU	23,25,115,117,24,26,116,118
MA(14)	XA(24)	XA(24)	Not currently supported, would probably be pin 73 on SDRAM module.
MA(13)	XA(23)	XA(23)	72*
MA(12)	XA(22)	XA(22)	70**
SDBA0	XA(11)	XA(11)	106
SDBA1	XA(20)	XA(20)	110
MA(0:11)	XA(0:10,21)	XA(0:10,21)	29,31,33,30,32,34,103,104,105,109,111,112
CB(0:7)	XA(19:12)	XA(19:12)	57,59,77,79,58,60,78,80
RAS_n	NU	NU	65
CAS_n	NU	NU	66
WE_n	WE_n	WE_n	67
CKE(0:1)	NU	NU	62,68***
CLK(0:1)	NU	CLK	61,74****
SDA	SDA	SDA	141
SCL	SCL	SCL	142
DNU	OE_n	OE_n	73 (this is defined for EDO and FP DRAM modules)
S_n(0:1)	CE_n	CE_n and ADV_n	69,71***** (use pin 69 to select single-row SDRAM modules, pin 71 to select FLASH)
VDD	VDD	VDD	11,12,27,28,45,46,63,64,81,82,101,102, 113,114,129,130,143,144
VSS	VSS	VSS	1,2,21,22,35,36,55,56,75,76,91,92, 107,108,119,120,139,140
	RESET_n	RESET_n	72*

Notes regarding above table entries:

* Seldom used pins, MA(13) not supported on 54-pin devices up to 256Mb. There is a reserved pin for higher density parts in this package.

** MA(12) is seldom used, supported ONLY for 256Mb parts.

*** Sometimes used on single row modules (required by PC100/PC133). Loading on MPC824x not affected. Included here to allow possibility of use of COTS SDRAM modules.

**** Second CLK is sometimes used in 1-Row modules to reduce loading and balance clock trees.

***** Sharing FLASH and SRAM select requires address decoding on the module.

Micron Technology Inc. et al.
Ex. 1010, p. 639

[0057] Those skilled in the art will recognize other mappings of memory controller signals appropriate to control the above and other types of memory modules. The above mappings of signals are therefore intended as representative of one exemplary preferred embodiment of the invention to provide a flexible memory architecture for use of SO-DIMM modules. Further, as noted above, those skilled in the art will recognize pin assignments for signals that may be used for other memory module configurations and standards.

[0058] Figure 2 is a schematic drawing showing relevant portions of signals used in a first exemplary preferred embodiment of the present invention. Those skilled in the art will recognize from the tables above that numerous other signals required for accessing each type of memory device are removed from figure 2 to simplify the drawing.

[0059] Figure 2 shows a portion of a system having two sockets (connectors) for insertion of SO-DIMM memory modules. The first of the two sockets 202 may preferably be used for a standard SDRAM (or DRAM) SO-DIMM module to provide standard dynamic RAM features for program instruction and data storage and retrieval. The second socket 200 may preferably be reserved for a special module such as a Flash memory module or static RAM memory module adapted in accordance with the present invention for use in a modified 144-pin SO-DIMM socket.

[0060] In a first exemplary preferred embodiment as shown in figure 2, the row select signals (CS(0:1)) 204 are handled differently from the standards defined for the 144-pin SO-DIMM. In particular, the first signal of the row select signals (CS(0)) is preferably used in accordance with the standards defined for the module to select the front side (the first row) of the addressed module. CS(0) is therefore preferably connected to pin 69 of the standard SO-DIMM socket. In this exemplary preferred embodiment, standard SDRAM (or DRAM) SO-DIMMs are preferably limited to single-row configurations so that the second row select signal is not needed for SDRAM (or DRAM) addressing. Rather the second row select signal is provided by one of a plurality of configurable sources to select the Flash or asynchronous memory device in second socket 200 using pin 71 of the SO-DIMM standard socket. A first source of the second row select signal is the row select signal as normally generated by the memory controller (CS(1) to thereby permit use of any standard SO-DIMM module including double-row modules). Other sources for the second row select signal generation are programmable chip select signals generated by the memory controller of the MPC824x microcomputer

Micron Technology Inc. et al.
Ex. 1010, p. 640

memory controller (or other similarly featured memory controller). Specifically, the F64CS control signal (generated by an MPC8240 with a separate address decoder ASIC) and the RCS(0:3) control signals (generated by an MPC8245) may be used as a second row select signal when an address is decoded as referencing the asynchronous memory device on a second module in the second SO-DIMM socket 200.

[0061] As noted above, in a first exemplary preferred embodiment, the various sources of signals used for the second row select signal are selected by configuring mutually exclusive jumpers (0 Ohm resistors) 273, 274 and 275. In particular, jumper 273 is inserted to enable selection of the second row by use of the CS(1) standard select signal generated by the memory controller. Jumper 274 enables use of the RCS(3) signal generated in response to properly decoded memory access addresses. Jumper 275 enables use of the F64CS signal generated in response to particular memory access addresses. Those skilled in the art will recognize a variety of other programmable chip select signals that may be used in a programmable memory controller to permit flexible generation of the second row select signal applied to the memory modules. The particular signals described in the tables above and shown in figure 2 for providing second row select signals to SO-DIMM modules are therefore merely exemplary of one preferred embodiment of the invention. Numerous other configurations will be apparent to those skilled in the art. The selected source for the second row select signal is then applied via path 210 to pin 71 of the SO-DIMM socket.

[0062] As noted above, several configurations of memory modules are permitted using the signaling features of the present invention as described above. In particular, the features of the present invention fully support use of standard SDRAM (or other DRAM) SO-DIMM modules. In other words, the features of the present invention retain "backward" compatibility to permit use of existing SDRAM (and other DRAM) SO-DIMM modules. In addition, the features of the present invention allow both SO-DIMM sockets of the system to be populated with single-row modules using memory devices other than DRAMs. For example, both modules may provide synchronous or asynchronous Flash or static RAM devices or combinations of the same. Further, special modules may be used providing SDRAM (or other DRAM) on the front side and Flash or SRAM on the other. Those skilled in the art will recognize a variety of memory configurations enabled by the signaling features of the present invention. Further, those skilled in the art will recognize that a similar jumper structure may be employed for the

Micron Technology Inc. et al.
Ex. 1010, p. 641

first row select signal (CS(0)) preferably coupled to pin 69 of the SO-DIMM sockets 200 and 204. Such a structure would permit the added flexibility to allow both rows to utilize asynchronous devices if desired for a particular application.

[0063] As noted above, figure 2 suggests a first exemplary preferred embodiment wherein the particular configuration of memory devices used on the system is identified by jumpers. Those skilled in the art will recognize a jumper as a switch and therefore recognize a wide variety of equivalent switching means including for example DIP switches of various types and other mechanical and electrical switch devices and means.

[0064] In an alternate embodiment representing the best presently known mode of practicing the invention, an enhanced memory module in accordance with the present invention would identify its configuration in returned SPD data. Serial presence detect (SPD) information is usually stored in a read-only memory (i.e., PROM) on standard memory modules. The SPD information includes, among other things, timing, size and geometry parameters of the memory module. These parameters are used by a memory controller to adjust it memory signal timings and other parameters to match the requirements of the module. This SPD information is usually retrieved at boot time or system reset of the system board and is used by low level boot code to configure the memory controller.

[0065] In accordance with the best known mode of practicing the invention, an enhanced module of the present invention returns information sufficient to identify the type of memory elements on the module and other parameters related to control of memory access to the module by a memory controller. Those skilled in the art will recognize that official designations of SPD information bit and byte assignments and values are controlled by standards agencies (i.e., JEDEC in the exemplary preferred embodiment utilizing 144-pin SO-DIMM standards). In particular, JEDEC JESD21-C section 4.1.2 defines a standard content for the first 128 bytes of SPD information returned by a memory module. Byte 2 under this standard defines a memory type. At present 7 out of 255 possible memory type values are defined for existing memory modules. A new memory type would preferably be used in the exemplary preferred embodiment of the present invention to identify a memory module populated with Flash or SRAM memory devices (typically addressed as a flat memory space as distinct from a two dimensional array of bits organized according to column addresses and row addresses). With such a new memory type value defined, other standard fields of the SPD

Micron Technology Inc. et al.
Ex. 1010, p. 642

information could be coded as appropriate for such "flat" memory devices to define any required memory addressing and timing parameters.

[0066] Another aspect of the present invention is the improved immobilization of the DIMM of the present invention to aid in resistance to vibration and shock related failures of the module. As noted above, testing results show that standard SO-DIMM sockets permit some rotation of the module about the mated keys used to help position the module within the socket. Over time, this rotation can wear the edge contacts (signal pads) to the point of premature failure. The invention provides for immobilizing the edge of the module opposite the connector edge that inserts into the socket. The increased immobilization of this edge helps prevent such wear to reduce the potential for premature failure of the module.

[0067] The invention broadly provides for any means for immobilizing the module to reduce the undesired rotation of the module within the socket. In particular, the present invention provides for securing the edge opposite the connector edge to the underlying system board on which the socket is also mounted. This serves to match the immobilization of the module with respect to the system board and that of the socket in which the module is inserted for use. The edge so immobilized is fastened to the system board by any of several fastening structures. Those skilled in the art will recognize a wide variety of fastening techniques and structures that may be equivalently employed to secure the edge of the memory module with respect to the socket and with respect to the system board.

[0068] Figures 4 and 5 depict one exemplary preferred embodiment of the invention to secure the memory module to the system board and thereby reduce movement of the module within the mated socket. Specifically, figure 4 is a dimensioned mechanical drawing showing mounting holes 400 positioned at the corners of a standard SO-DIMM module 402. The mounting holes are preferably positioned at the corners of the edge of the module opposite the connector edge. Figure 5 provides further detail of the mounting holes 400 in the DIMM module 402 mated with a connector 506 mounted on a system board 500. Specifically, figure 5 provides a sequence of three positions of the module 402 as it is mated with connector 506. Module 402 is first tilted at an angle and its edge connector is inserted into the mated slot of connector 506. A key notch in the module is positioned over the mated key 508 in the socket to ensure accurate positioning of the module within the socket. With the connector edge of the module so inserted in the

Micron Technology Inc. et al.
Ex. 1010, p. 643

socket, the module is tilted downward deflecting by cam action retaining spring clips 510 of the socket on each side of the socket 506. When tilted down substantially parallel with the system board. spring clips 510 bias back to normal position to retain module 402 within socket 506. In accordance with the present invention, mounting screws 504 are inserted through mounting holes 400 of the module 402 and mated with standoff nuts 502 affixed to the system board.

[0069] In one exemplary preferred embodiment, the improvement of the present invention secures module 402 within socket 506 by inserting screws 504 through mounting holes 400 into nuts 502 on the system board. Those skilled in the art will recognize that the number and position of mounting holes 400 may be varied as a matter of design choice to achieve the desired effect of reducing the motion of the module within the mated socket.

[0070] Further, those skilled in the art will recognize that a screw and mated standoff nut type fastener structure is but one of a number of equivalent fastening techniques and structures useful to immobilize the module within the socket. Numerous equivalent examples are readily apparent to those skilled in the art as a matter of design choice. For example, standoff nuts 502 may be soldered components secured to the system board to simplify the manufacture of the system board. Further, other types of fastening systems may be employed. For example, a half card-cage with a channel for gripping the edge of the module opposite its connector edge may be affixed to the system board. The module is inserted in the socket and tilted down to "snap" into the side spring clips 510 as shown in figure 5. When in place, the module also snaps into the channel of the half card-cage structure on the system board. Another example uses an elastomer material between the module and the system board to reduce vibrations of the module within its connector socket. Another example provides for a mounting pin affixed to the system board such that the mounting pin enters through the mounting hole and snaps into a locked position to secure the module onto the mounting pin (and hence to the system board). Still another example provides for a pin associated with the socket connector such that the pin inserts through a hole in the module when the module is inserted in the socket. These and other alternative embodiments would be viewed as equivalent means for securing the module to the underlying system board to thereby immobilize the module with respect to the mated connector socket.

Micron Technology Inc. et al.
Ex. 1010, p. 644

[0071] While the invention has been illustrated and described in the drawings and foregoing description, such illustration and description is to be considered as exemplary and not restrictive in character, it being understood that only the preferred embodiment and minor variants thereof have been shown and described and that all changes and modifications that come within the spirit of the invention are desired to be protected.

Micron Technology Inc. et al.
Ex. 1010, p. 645

CLAIMS

What is claimed is:

1. A memory module comprising:

a printed circuit assembly having connector pads at one edge of said assembly; and

a plurality of memory devices mounted on said assembly and electrically coupled to said connector pads wherein said printed circuit assembly is adapted to support both synchronous and asynchronous types of said memory devices.

2. The memory module of claim 1 wherein each of said plurality of memory devices is a synchronous dynamic random access memory device.

3. The memory module of claim 1 wherein each of said plurality of memory devices is a synchronous Flash memory device.

4. The memory module of claim 1 wherein each of said plurality of memory devices is an asynchronous Flash memory device.

5. The memory module of claim 1 wherein each of said plurality of memory devices is an asynchronous static random access memory device.

6. The memory module of claim 1 wherein each of said plurality of memory devices is an asynchronous fast static random access memory device.

7. The memory module of claim 1 wherein each of said plurality of memory devices is an asynchronous low power static random access memory device.

8. The memory module of claim 1 wherein said connector pads include:

a first select signal connector pad that selects a first subset of memory devices mounted on said module when a signal is applied thereto wherein said first subset of memory devices are synchronous memory devices; and

a second select signal connector pad that selects a second subset of memory devices mounted on said module when a signal is applied thereto wherein said second subset of memory devices are asynchronous memory devices.

9. The memory module of claim 1 wherein said connector pads include:

a first select signal connector pad that selects a first subset of memory devices mounted on said module when a signal is applied thereto wherein said first subset of memory devices are synchronous dynamic random access memory devices; and

Micron Technology Inc. et al.
Ex. 1010, p. 646

a second select signal connector pad that selects a second subset of memory devices mounted on said module when a signal is applied thereto wherein said second subset of memory devices are synchronous Flash memory devices.

10. A system comprising:

a system board;

a memory bus adapted for exchanging signals between a memory controller and both synchronous and asynchronous memory devices;

a memory controller on said system board coupled to said memory bus wherein said memory controller is capable of generating signals for control of both synchronous and asynchronous memory devices and wherein said memory controller is capable of multiplexing said signals on said memory bus;

a first socket connector on said system board for receiving a first memory module wherein said first socket connector is coupled to said memory controller through said memory bus; and

a first memory module inserted in said first socket connector and electrically coupled to said memory controller wherein said first memory module includes a plurality of synchronous or asynchronous memory devices.

11. The system of claim 10 further comprising:

a jumper on said system board for configuring signals exchanged between said memory controller and said first memory module in accordance with the type of memory devices on said first memory module.

12. The system of claim 10 wherein said first memory module provides serial presence detect information used in conjunction with said memory controller to identify the type of memory devices included on said first memory module.

13. The system of claim 10 wherein said first memory module includes a plurality of synchronous dynamic random access memory devices.

14. The system of claim 10 wherein said memory module includes a plurality of asynchronous Flash memory devices.

15. The system of claim 10 wherein said memory module includes both synchronous Flash memory devices and synchronous dynamic random access memory devices.

Micron Technology Inc. et al.
Ex. 1010, p. 647

16. The system of claim 10 wherein said memory module includes both synchronous memory devices and asynchronous memory devices.

17. The system of claim 10 further comprising:

a second socket connector on said system board for receiving a second memory module wherein said second socket connector is coupled to said memory controller through said memory bus; and

a second memory module inserted in said second socket connector and electrically coupled to said memory controller wherein said second memory module includes a plurality of synchronous or asynchronous memory devices.

18. The system of claim 17 further comprising:

a jumper on said system board for configuring signals exchanged between said memory controller and said second memory module in accordance with the type of memory devices on said first memory module and on said second memory module

19. The system of claim 17

wherein said first memory module provides serial presence detect information used in conjunction with said memory controller to identify the type of memory devices included on said first memory module, and

wherein said second memory module provides serial presence detect information used in conjunction with said memory controller to identify the type of memory devices included on said second memory module.

20. The system of claim 17

wherein said first memory module includes a plurality of synchronous dynamic random access memory devices, and

wherein said second memory module includes a plurality of asynchronous Flash memory devices.

21. The system of claim 17

wherein said first memory module includes a plurality of synchronous dynamic random access memory devices, and

wherein said second memory module includes a plurality of asynchronous static random access memory devices.

22. The system of claim 17

Micron Technology Inc. et al.
Ex. 1010, p. 648

wherein said first memory module includes a plurality of asynchronous Flash memory devices, and

wherein said second memory module includes a plurality of asynchronous static random access memory devices.

22. A system comprising:

a system board;

a socket connector on said system board for receiving a memory module wherein said socket connector includes a key;

a memory module having a connector edge inserted in said socket connector and having an opposing edge opposite said connector edge wherein said memory module has a notch mated to said key when said memory module is inserted in said socket connector; and

a memory module retainer adapted to substantially immobilize said opposing edge with respect to rotation about said key.

23. The system of claim 22 wherein said memory module retainer comprises:

a nut affixed to said system board;

a hole in said memory module along said opposing edge and aligned with said nut; and

a screw inserted through said hole into said nut to substantially immobilize said opposing edge with respect to rotation about said key.

24. The system of claim 23 wherein said nut is a swaged extension nut.

25. The system of claim 22 wherein said memory module retainer comprises:

a half card-cage affixed to said system board wherein said half card-cage includes a channel for receiving said opposing edge of said memory module to substantially immobilize said opposing edge with respect to rotation about said key.

26. The system of claim 22 wherein said memory module retainer comprises:

a standoff pin affixed to said system board;

a hole in said memory module along said opposing edge and aligned with said pin wherein said standoff pin extends through said hole and locks when said memory module is completed inserted in said socket connector to substantially immobilize said opposing edge with respect to rotation about said key.

Micron Technology Inc. et al.
Ex. 1010, p. 649

27. The system of claim 22 wherein said memory module retainer comprises:

a standoff pin affixed to said socket connector;

a hole in said memory module aligned with said standoff pin wherein said standoff pin extends through said hole and when said memory module is completed inserted in said sockct connector to substantially immobilize said opposing edge with respect to rotation about said key.

28. A system comprising:

a system board;

a socket connector on said system board for receiving a memory module wherein said socket connector includes a key;

a memory module having a connector edge inserted in said socket connector and having an opposing edge opposite said connector edge wherein said memory module has a notch mated to said key when said memory module is inserted in said socket connector; and

memory module retainer means adapted to substantially immobilize said opposing edge with respect to rotation about said key.

29. The system of claim 28 wherein said memory module retainer means comprises:

nut means affixed to said system board;

a hole in said memory module along said opposing edge and aligned with said nut means; and

screw means inserted through said hole into said nut means to substantially immobilize said opposing edge with respect to rotation about said key.

30. The system of claim 29 wherein said nut means is a swaged extension nut.

31. The system of claim 28 wherein said memory module retainer means comprises:

card-cage means affixed to said system board wherein said card-cage means includes a channel for receiving said opposing edge of said memory module to substantially immobilize said opposing edge with respect to rotation about said key.

32. The system of claim 28 wherein said memory module retainer means comprises:

pin means affixed to said system board;

Micron Technology Inc. et al.
Ex. 1010, p. 650

a hole in said memory module along said opposing edge and aligned with said pin means wherein said pin means extends through said hole and locks when said memory module is completed inserted in said socket connector to substantially immobilize said opposing edge with respect to rotation about said key.

33. The system of claim 28 wherein said memory module retainer means comprises:

pin means affixed to said socket connector;

a hole in said memory module aligned with said pin means wherein said pin means extends through said hole and when said memory module is completed inserted in said socket connector to substantially immobilize said opposing edge with respect to rotation about said key.

Micron Technology Inc. et al.
Ex. 1010, p. 651

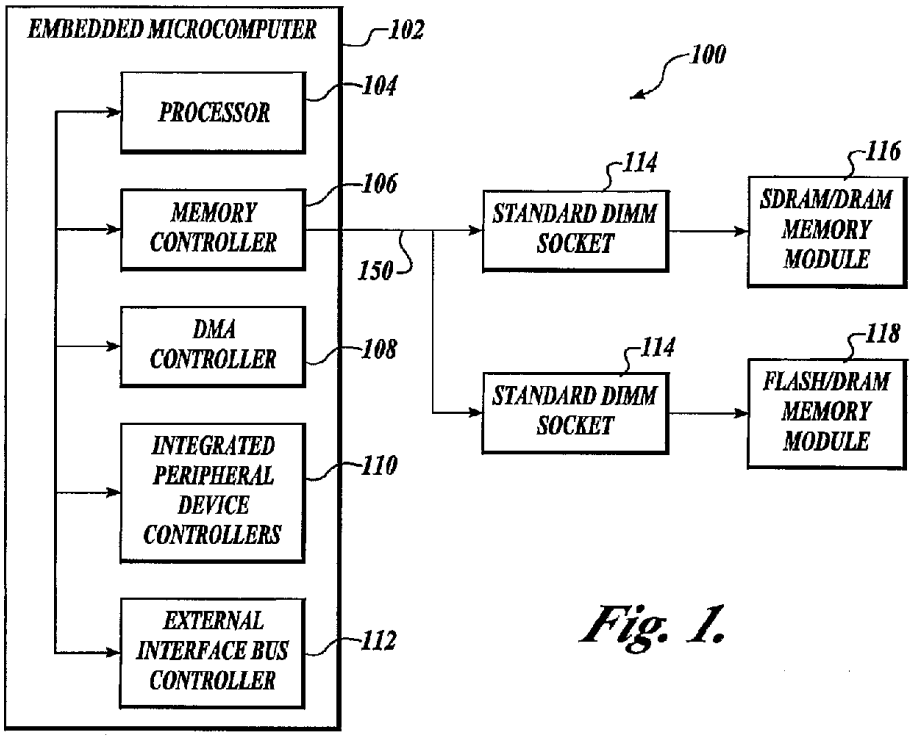

EMBEDDED MICROCOMPUTER
102
100
PROCESSOR
104
MEMORY CONTROLLER
106
DMA CONTROLLER
108
INTEGRATED PERIPHERAL DEVICE CONTROLLERS
110
EXTERNAL INTERFACE BUS CONTROLLER
112
150
STANDARD DIMM SOCKET
114
SDRAM/DRAM MEMORY MODULE
116
STANDARD DIMM SOCKET
114
FLASH/DRAM MEMORY MODULE
118

Fig. 1.

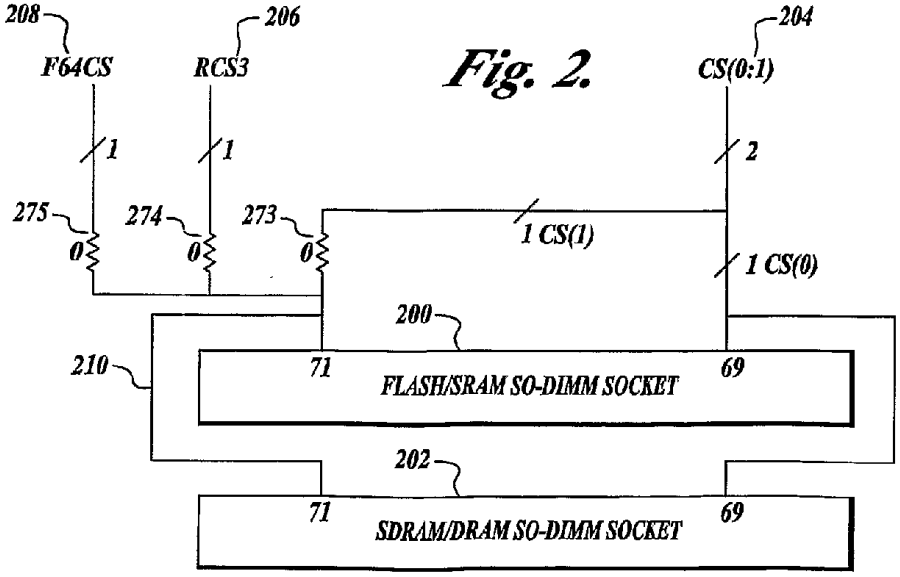

208
F64CS
206
RCS3
204
CS(0:1)
1
1
2
275
274
273
0
0
0
1 CS(1)
1 CS(0)
200
210
71
FLASH/SRAM SO-DIMM SOCKET
69
202
71
SDRAM/DRAM SO-DIMM SOCKET
69

Fig. 2.

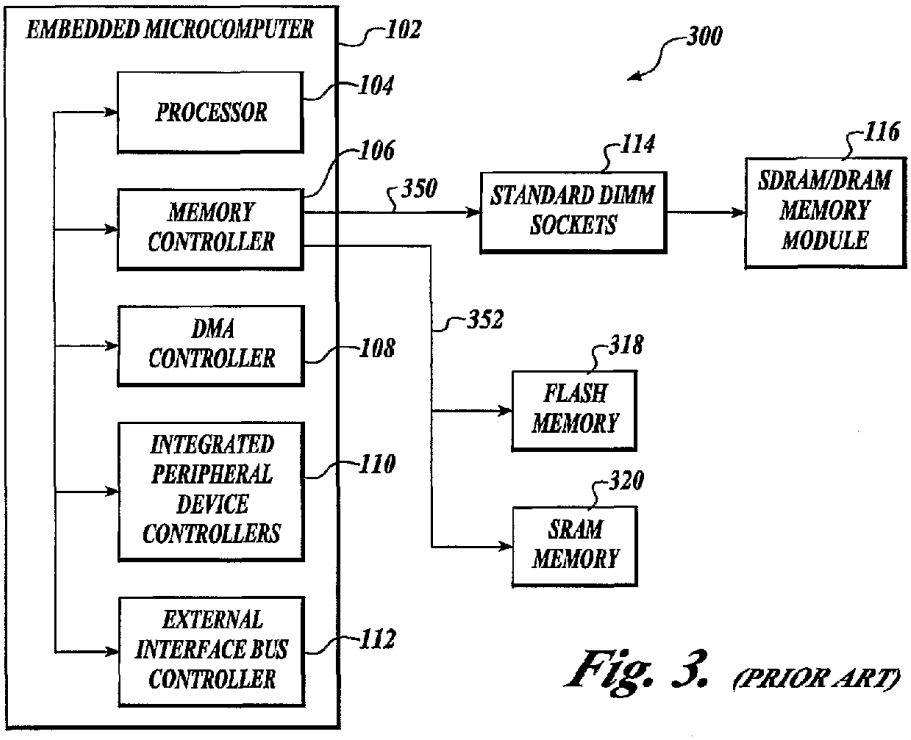


Fig. 3. *(PRIOR ART)*

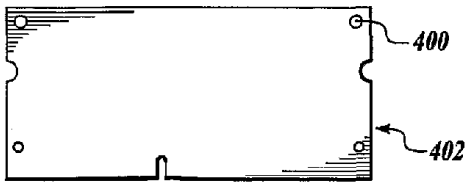


Fig. 4.

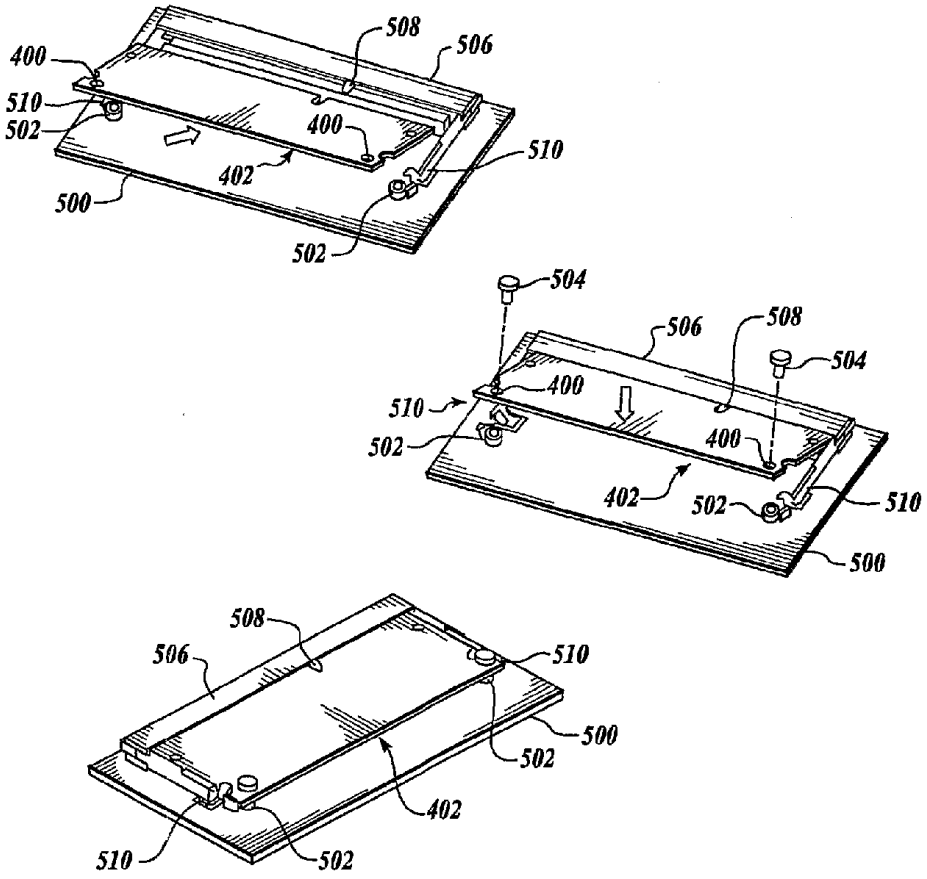

508
506
400
510
502
400
510
402
500
502
504
506
508
504
400
510
502
400
402
502
510
500
508
506
510
500
502
402
510
502

Fig. 5.