



(11) (21) (C) **2,241,841**
(86) 1996/08/12
(87) 1997/07/10
(45) 1999/10/26

(72) CHATTER, Mukesh, US

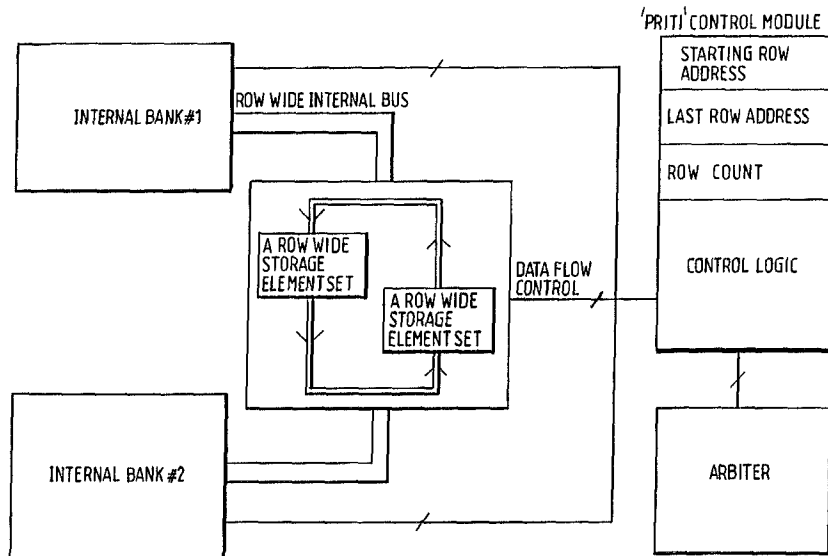
(73) CHATTER, Mukesh, US

(51) Int.Cl.⁶ G06F 13/16

(30) 1995/12/29 (581,467) US

(54) **SYSTEME ET ARCHITECTURE DE MEMOIRE RAM
DYNAMIQUE POLYVALENTE A CAPACITE ELEVEE
POSSEDANT UNE ANTEMEMOIRE INTERNE ET DES ACCES
MULTIPLES ET PROCEDE**

(54) **HIGH PERFORMANCE UNIVERSAL MULTI-PORT
INTERNALLY CACHED DYNAMIC RANDOM ACCESS
MEMORY SYSTEM, ARCHITECTURE AND METHOD**



(57) Nouvelle architecture de mémoire RAM dynamique économique à capacité élevée possédant une antémémoire interne et des accès multiples, désignée "AMPIC DRAM" et, de ce fait, architecture unique de système éliminant les limitations habituelles importantes de la largeur de bande du système. L'invention concerne également un moyen de transfert de blocs de données à l'intérieur de la puce, d'ordres d'intensité plus rapides que ceux du processus classique. La puce effectue également l'interconnexion de nombres considérablement plus élevés de ressources associées à des performances sensiblement améliorées et de façon

(57) A novel low cost/high performance multi-port internally cached dynamic random access memory architecture called "AMPIC DRAM", and consequentially a unique system architecture which eliminates current serious system bandwidth limitations. It also provides a means to transfer blocks of data internal to the chip, orders of magnitude faster than the traditional approach. The chip also interconnects significantly higher numbers of resources with substantially enhanced performance and at notably lower cost. A system configuration based on this novel architecte can work equally efficiently for both main





(11) (21) (C) **2,241,841**
(86) 1996/08/12
(87) 1997/07/10
(45) 1999/10/26

beaucoup plus économique. Une configuration de système basée sur cette nouvelle architecture peut fonctionner avec une efficacité égale pour les deux fonctions de mémoire principale et en tant que mémoire de graphiques, ce qui permet d'obtenir une architecture de mémoire unifiée présentant une réelle économie et une capacité élevée.

memory functions and as graphics memory, thus providing a truly low cost, high performance unified memory architecture.



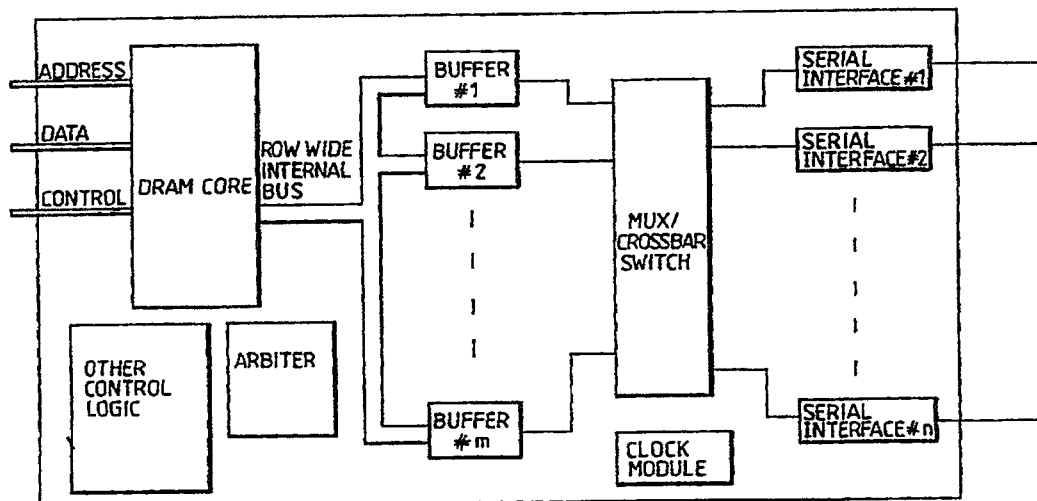
PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

| | | |
|--|------------------|--|
| <p>(51) International Patent Classification ⁶ : G11C 7/00, G06F 13/16</p> | <p>A1</p> | <p>(11) International Publication Number: WO 97/24725 (43) International Publication Date: 10 July 1997 (10.07.97)</p> |
| <p>(21) International Application Number: PCT/IB96/00794 (22) International Filing Date: 12 August 1996 (12.08.96) (30) Priority Data: 08/581,467 29 December 1995 (29.12.95) US (71)(72) Applicant and Inventor: CHATTER, Mukesh [IN/US]; 53 Godfrey Lane, Milford, MA 01757 (US). (74) Agent: RINES, Robert, Harvey; c/o MacLeod & Co, Bledington Grounds, Bledington, Gloucestershire OX7 6XL (GB).</p> | | <p>(81) Designated States: AL, AM, AT, AT (Utility model), AU, AZ, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, 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).</p> <p>Published <i>With international search report.</i></p> |

(54) Title: HIGH PERFORMANCE UNIVERSAL MULTI-PORT INTERNALLY CACHED DYNAMIC RANDOM ACCESS MEMORY SYSTEM, ARCHITECTURE AND METHOD



(57) Abstract

A novel low cost/high performance multi-port internally cached dynamic random access memory architecture called "AMPIC DRAM", and consequentially a unique system architecture which eliminates current serious system bandwidth limitations. It also provides a means to transfer blocks of data internal to the chip, orders of magnitude faster than the traditional approach. The chip also interconnects significantly higher numbers of resources with substantially enhanced performance and at notably lower cost. A system configuration based on this novel architecte can work equally efficiently for both main memory functions and as graphics memory, thus providing a truly low cost, high performance unified memory architecture.

* (Referred to in PCT Gazette No. 43/1997, Section II)

HIGH PERFORMANCE UNIVERSAL MULTI-PORT INTERNALLY CACHED
DYNAMIC RANDOM ACCESS MEMORY SYSTEM, ARCHITECTURE AND
5 METHOD

The present invention relates to dynamic random access memory technology (DRAM),
being more specifically directed to novel DRAM system architectures that eliminate
current system bandwidth limitations and related problems and provide significantly
10 enhanced system performance and at reduced cost, enabling substantially universal
usage for many applications as a result of providing unified memory architecture.

A vast number of such system designs, especially in networking/communications, are
performance-limited due to contentions among various resources to access the system
15 main memory (almost always DRAMs). The large number of contentions is a direct
result of the use of a single bus architecture where a bus interconnects the CPU, the
main memory and the I/O resources. This and similar past and current architectures
also retard the capability of the CPU to manage much greater numbers of I/O
interfaces, due to serious bandwidth limitations.

20

Similar system bandwidth limitations have also caused graphics/multimedia designers
to separate the graphics memory from the main memory, thereby adversely affecting
the system cost. Specialized DRAMs for graphics applications have also been
developed further to enhance video data bandwidth capabilities. Whilst some system
25 architectures using various types of DRAMs have been proposed which will allow a
common memory to be used as both main memory and graphics memory, such have
had limited success because of inefficiency for one operation or another.

Thus the availability of a low cost, high performance unified memory architecture has largely remain elusive until the advent of the present invention that provides an innovative DRAM architecture and consequentially a unique System Architecture later
5 described, which largely eliminates these problems and as a result provides considerably more data bandwidth capability to interconnect significantly higher numbers of resources and with substantially enhanced performance, at notably lower cost. A system configuration based on this novel architecture will work equally efficiently for both main memory function and as graphics memory, thus further
10 achieving a truly low cost, high performance unified memory architecture. This chip solution has been hence named 'AMPIC DRAM', standing for A Multi Port Internally Cached DRAM.

BACKGROUND OF INVENTION

15

As before discussed, most of the high performance systems of this character tend to use bus-based architecture, where a single system bus interconnects the CPU, the main memory and the I/O resources as shown in later-described Fig 1, (the terms 'main memory' and 'system memory' as herein used, being so used interchangeably). This is
20 relatively straight forward design and provides room for expansion; but it has serious limitations. Whenever the CPU or the peripherals need to access the main memory (generally implemented with DRAM), an arbitration takes place for access to the system bus. Thus the amount of concurrent activity in the system is limited by the overall capacity of the external bus.

25

As the speed of the CPU increases, the system bus bandwidth must correspondingly increase to realize the full potential of the system. Increasing the bus bandwidth,

however, is much more difficult and very costly to the point of becoming technically impossible or prohibitively expensive. In addition, the number of I/O resources which can be on the bus is also limited by the bandwidth. It should be noted indeed that while theoretically the single bus allows a high order of expandability, the real operation due to contentions, drastically limits such expansion.

This problem is prevalent among all types of applications. While illustrative networking and graphics applications are presented hereafter as examples for better understanding of these problems, the invention is in no way limited to these exemplary areas only.

A Networking Application Example

A typical networking equipment (also called interconnectivity equipment) such as switches, routers, bridges, hubs, etc. interconnect multiple networks such as ATM, SONET, Token Ring, FDDI, Ethernet, Fiber Channel, etc. as shown in later-described Fig 2. A typical design includes a high performance CPU and a large amount of main memory generally implemented with the use of a traditional DRAM as represented in later-described Figs 3 and 4. Data from various networks is transferred to the main memory in the form of packets (a packet is a collection of bytes), processed by the CPU and then, in general, forwarded to their respective destination networks.

All the networks mentioned above (ATM, SONET, Fiber Channel, Token Ring, FDDI, etc.) provide different means of transporting data from one point to another. They differ in hardware, software and data transfer speeds. Interconnect equipment is required to allow the users on one of these networks to communicate with the users on another network with different protocol, seamlessly.

In a typical interconnectivity equipment, the network interfaces are implemented with a network interface controller (also commonly known as network controller), unique for each type of interface. Thus Ethernet has a different network interface than for
5 Fiber Channel or ATM (Figs 3 and 4).

In the example presented in Fig 4, exemplary only to explain a typical dataflow of this illustrative system configuration, the following typical parameters may be involved:

- 10 a. the system bus is 32 bit wide (4 bytes);
- b. four traditional DRAMs, 2M x 8 have been configured as 2M x 32;
- c. there are four network interfaces Fiber Channel, ATM, Ethernet, and FDDI;
- 15 and
- d. the packet size is 1024 bytes.

Consider the case, when a packet is sent by a user on, for example, the Ethernet
20 network to a user, for example, on the FDDI network. This packet is received by the interconnect equipment Ethernet interface controller and is analyzed by the controller chip, with only the relevant information content stored in its conventional local FIFO (First In First Out) memory, for subsequent transfer to the main memory. As there are multiple devices on the system bus including the CPU and various network
25 controllers, arbitration takes place among all the active resources to acquire the system memory bus. After the Ethernet controller acquires the bus via arbitration, data is then forwarded to the system memory using the 32 bit wide system bus interface. As there

are 1024 bytes in the packet and 4 bytes are transferred to the main memory per transfer, 256 such transfers are required to move the packet. If the network controller is permitted to make only one 4 byte transfer for every acquisition of the bus, minimum 256 arbitration cycles are also be required. (The number of arbitrations may
5 be smaller if burst transfer capability is provided by the network controller. As an example, a 16 byte burst transfers capability for every acquisition, need minimum 64 arbitration cycles).

After this packet is stored in the main memory, it is processed by the CPU (primarily
10 the header information) and redirected towards the FDDI port, in this example. Now the reverse process takes place. The data is picked up by the FDDI interface controller from the main memory and transferred to the chip internal FIFO memory. This also requires 256 transfers and a corresponding number of arbitrations. Data is then concurrently transferred from the FDDI controller chip to its network.

15

FDDI operates at 100 mega bits per second, Ethernet at 10/100 mega bits per second, ATM at nearly 600 mega bits per second, Token ring at 16 mega bits per second, and Fiber Channel at 800 mega bits per second.

20 The large number of transfers and the time spent in arbitration consume a significant part of the available data bandwidth and also reduce the CPU access frequency. As the number of network interfaces increases or higher speed interfaces are added, the amount of time available to each resource, including CPU, decreases, thus clipping the peak performance of the system. It also forces the designers to go for even higher
25 performance CPUs and associated expensive components, thus driving up the cost. The number of networks which can be connected via this prior art type of system bus also remains low due to these severe limitations, and this problem becomes worse as

more and more networks at higher and higher speeds are added to accommodate, for example, the Internet related expansion activity.

A Graphics/Multimedia Application Example

5

Using, again for illustrative background purposes, a graphics/multimedia application, there are two major memory related functions in such graphics systems;

10

- a. updating the screen memory for the picture to be displayed, and
- b. retrieving the screen memory at very high rates to update the Cathode Ray Tube (CRT) or other screen display.

15

The first task requires a large number of frequent data transfers, called 'BitBlt', from one place in the memory to another; but this requirement tends to be bursty in nature. This consumes a considerable portion of the system bandwidth and thus has necessitated the use of separate memory to store graphics data, as shown in later-described Fig 5, thereby adversely affecting the system cost. Consider an example where 16 rows of the screen memory need be updated and conventional 2M x 8

20 DRAM components are used. To transfer 16 rows worth of data in the new location, the number of data transfers required are:

25

Number of Rows to be transferred (16) x Number of Columns in each row (1024) = 16,384.

Obviously a corresponding number of arbitrations also needs to be made for the system bus. This large number of transfers, moreover, is required in a short time, thus

consuming most of the available data bandwidth in a small time slot, starving the CPU and other I/O resources. Before the present invention, however, no practical breakthrough has been offered in this area by the existing DRAM manufacturers to alleviate this problem.

5

Repeated screen memory retrieval, of course, is necessary to load and refresh the CRT display, and depending on the type of display (VGA, super VGA, etc.), the bandwidth requirements for such update vary, but tend to be in the order of hundred megabytes or above per second. Unlike 'BitBlt', the CRT update requirements are continuous in nature and, similar to 'BitBlt', also utilize the system bandwidth considerably.

10

As an example, consider the following case:

a. display size 1024 x 768 pixels;

15

b. non interlaced - updated 72 times a second; and

c. 8 bits per pixel for each of the three Red, Green and Blue colors.

20 The bandwidth required in terms of bytes per second is:

$$1024 \times 768 \times 72 \times 8 \times 3/8 = 170 \text{ MBytes.}$$

It is a significant requirement which is difficult to meet with the conventional DRAM, if also used as main memory. This, therefore, led to the development of more expensive specialized DRAMs and one of the popular and extensively used such special DRAM is 'Video DRAM' also called 'VRAM'. Most of the VRAMs are dual

25

port with some exceptions where a third port has been added. A typical VRAM has a system interface similar to the traditional DRAMs, but is also has a row wide buffer (called SAM - Serial Access Memory) inside the chip, which interacts with the outside world via separate but equal number of data pins as available for the system interface as shown in later-described Fig 6. As an example, a 256K x 8 VRAM also has an additional port which is 8 bits wide to stream refresh data to the CRT continuously. The 'SAM' buffer has a fix connection with the external display interface. In operation, the CPU (or the system bus master) accesses the VRAM via the system data interface and a screen image is stored or updated in the VRAM. Then the screen data of one entire row is moved to the 'SAM' buffer in one access. This data is then subsequently transferred to the display via the SAM I/O interface, identical in width to the system interface.

Such VRAMs provide an acceptable solution for the case where the design has to interact with only one graphics source/destination. They are, however, more expensive than traditional DRAMs due to the larger number of additional pins and the bigger silicon die, and the architecture provides a very rigid structure. Expandability to more interfaces with more devices is severely limited as the number of pins grows significantly. The 'SAM' connection to the external I/O interface is fixed and the data port size is also predetermined. This approach also does not solve the problem of speeding the huge data movement requirements. Thus the VRAMs provide an acceptable solution but only because of lack of any better alternative, until the present invention.

VRAMs (specifically the 3 port version) were also occasionally touted towards networking applications but have rarely been so used due to their before-stated rigid I/O structure, very limited capability to interconnect number of resources (to be

precise - only two), bigger board space, more expensive construction and higher power consumption.

State of The Art in System Configurations & Associated Problems

5

In summary, therefore, the typical prior art and current system configurations based on traditional DRAMs (such as before-mentioned Fig 3) are burdened with the limitations described above. The problems associated with the networking applications, where a large number of high speed networks need be interconnected at low cost, have
10 remained largely unsolved, thus resulting in higher price per network interfaces.

While the system bandwidth limitations and the constant CRT display update requirements led to the development of video DRAMs, primarily intended for graphics applications and generally used in configurations such as shown in Fig 6, as
15 previously discussed, and while this configuration performs better than compared to the traditional DRAMs, as performance requirements on both the main memory and the display data bandwidth increased, the deleterious side effect of this configuration added cost to the system due to separation of main memory from graphics memory and also due to the higher cost of the VRAMs themselves.

20

Another type of a prior art graphics DRAM was also developed by RAMBUS Corporation dubbed 'RDRAM', which operates at 250 Mhz, and works fine for the graphics applications and may even be more cost effective for high end graphics/multimedia environment than VRAMs; but this approach still requires the
25 retention of two separate buses.

In the PC market (60% of all the chips are consumed by this segment), cost is of paramount importance. A search then was initiated to find a configuration dubbed 'Unified Memory Architecture', which will share the common bus for both graphics and main memory functions and consume only one type of memory device.

5

One possibility is to use VRAMs as both the graphics and the main memory components; but the added cost of the components neutralizes the gains achieved by having a common memory. Another potential solution is to use the before-described RDRAM, which has a smaller number of pins per chip than VRAM, thus resulting in lower power consumption, smaller real estate and relatively lower cost. But unfortunately because of their block-oriented protocol and interface limitations, it is highly inefficient with non-localized main memory accesses, and thus does not render itself well to the concept of 'Unified Memory Architecture'. Use of RDRAM also poses a host of significant electrical engineering design challenges related to the emissions, noise interference and PCB layout, making the design task very difficult.

10

15

Thus the search has still been on for a more viable low cost/high performance unified memory architecture capable to meet the diversified requirements of both the main and the graphics memory accesses, equally efficiently.

20

It is to the effective solution of this pressing problem, indeed, that the present invention is directed, the invention being believed to be a breakthrough in the evolution of a new type of DRAM architecture and method that;

25

a. provides much higher system data bandwidth with architectural innovations rather than sheer device speeds;

SUBSTITUTE SHEET (RULE 26)

SUBSTITUTE SHEET (RULE 26)

- b. moves the large amount of data to/from multiple I/O resources with minimum impact on the system bandwidth;
- c. interconnects considerably larger numbers of I/O resources compared to the
5 traditional approaches;
- d. moves large blocks of data internal to the chip in at least an order of magnitude faster time frame and with negligible impact on the system bandwidth;
- 10 e. is configurable to accommodate different data transfer rates of the I/O resources;
- f. reduces the latency time between the reception of incoming packet and its
subsequent transmission;
- 15 g. provides low pin count;
- h. provides reasonable cost;
- 20 i. provides for low power consumption;
- j. provides a simple system interface to minimize the design effort, and
- k. works equally efficiently for both the main memory and the graphics
25 requirements, thus providing a truly 'Unified Memory Architecture' and substantially universal approach.

OBJECTS OF INVENTION

An objective of the invention accordingly, is to provide a new and improved dynamic random access (DRAM) system, architecture and method utilizing a novel multi-port
5 internally cached DRAM structure that obviates current system bandwidth limitations and related problems, while providing significantly enhanced system performance at reduced cost, and which thereby enable substantially universal usage for myriads of applications.

10 A further object is to provide such a novel system in which the transfer of blocks of data internal to the chip is an order of magnitude faster than traditional approaches, and with the facility to interconnect significantly higher numbers of resources with substantially enhanced performance and at notably lower cost.

15 Still a further object is to provide a system configuration based on this novel architecture that works equally efficiently in both main memory functions and as graphics memory - a truly high performance unified memory architecture.

Other and further objectives will be explained hereinafter and are more particularly
20 delineated in the appended claims.

SUMMARY

In summary, however, from one of its viewpoints, the invention embraces for use in a
25 system having a master controller such as a central processing unit (CPU) having parallel data ports and a dynamic random access memory (DRAM) each connected to and competing for access to a common system bus interface, an improved DRAM

architecture comprised of a multi-port internally cached DRAM(AMPIC DRAM) comprising a plurality of independent serial data interfaces each connected between a separate external I/O resource and internal DRAM memory through corresponding buffers; a switching module interposed between the serial interfaces and the buffers; and a switching module logic control for the connecting of the serial interfaces to the buffers under a dynamic configuration by the bus master controller, such as said CPU, for switching allocation as appropriate for the desired data routability.

Preferred and best mode designs and techniques are hereafter explained in detail.

10

DRAWINGS

The invention will now be described in convention with the appended drawings, in which Figs 1-6 illustrate prior art techniques as follows:

15

Fig 1 is a block diagram of a typical prior art single bus parallel architecture;

Fig 2 shows a typical prior art network configuration;

20 Figs 3 and 4 illustrate a typical prior art networking equipment employing DRAMs for use in a configuration such as that of Fig 2;

Fig 5 is a block diagram of a prior art configuration of a graphics application with separate memories and using traditional DRAMs; and

25

Fig 6 is a similar diagram of a typical architecture of a graphics applications using VRAMs;

Fig 7 is a block diagram of a system architecture constructed in accordance with the present invention and embodying a multiple-port internally cached '(AMPIC) DRAM' of the invention;

5

Fig 8 is a similar view of a partial top level architecture of the 'AMPIC DRAM' of Fig 7, showing multiplex/crossbar switching between supplemental serial interfaces, buffers and the DRAM core;

10 Fig 9 illustrates details of an illustrative serial data-multiplexer implementation in Fig 8;

Fig 10 shows an example of the multiple serial interfaces configured as ports;

15 Fig 11 is a diagram of an exemplary 2-bit port of the AMPIC DRAM and associated control line;

Figs 12 and 13 are diagrams of examples of a serial data transfer format, with Fig 13 applied to the illustration of the 2 bits per port;

20

Fig 14 is a block diagram of an example of a partial top level of a two-bank 'AMPIC DRAM' control module architecture with a later-described parallel row internal transaction intervention (called PRITI), without storage elements;

25 Fig 15 shows the sequence of operations of the PRITI transfers with internal data transfer from the two banks;

Fig 16 is similar to Fig 14 but with two row wide sets of storage elements provided for the 'PRITI' capability;

Fig 17 illustrates the 'PRITI' transfer, with two row wide sets of storage elements of
5 Fig 16, showing the sequence of operations and internal data exchange between the two banks;

Fig 18 is similar to Fig 17, but employs only one row wide set of storage elements;

10 Fig 19 presents an example of a useful pinout for an 'AMPIC DRAM' of the invention with an exemplary 9 serial interfaces;

Fig 20 illustrates an exemplary networking equipment architected with AMPIC
DRAMS of the invention with a 32 bit wide system bus operating with the CPU;
15

Fig 21 is a similar diagram for a graphics application;

Fig 22 is a similar diagram illustrating a four bank system configuration with each
bank is connected to different network interfaces;
20

Fig 23 is similar to Fig 22 but uses two banks of 'AMPIC DRAMs' and two banks with
traditional DRAMs;

Fig 24 is also similar to Fig 22 but with two banks for graphics and two for other
25 applications;

Fig 25 is a diagram of still a further modification of an 'AMPIC DRAM' architecture in which there are two internal banks and the before-mentioned 'PRITI' capability, wherein one bank is for main memory usage and the other bank for graphics or other applications.

5

Fig 26 is a modification of the AMPIC DRAM system of before mentioned Fig 19, adapted for use with a so-called 'PARAS' interface and access, described in copending US patent application serial no. 08/320,058, filed October 7, 1994, and with a low pin count for the integrated memory architecture. [This application discloses a method of and apparatus for improving the accessing capability of asynchronous and synchronous dynamic random access memory devices by a novel interfacing and accessing procedure in which the same pins are used for each of row, column and data accessing and in both the write and read cycles; such enabling effective increasing of the data bandwidth and addressing range in substantially the same size packages with fewer pins.]

10

15

Fig 27 is a block diagram of an example of a partial top level of a multi-bank 'AMPIC DRAM' control module architecture with the before-described parallel row internal transaction intervention (PRITI), with a one row wide set of storage elements.

20

PREFERRED EMBODIMENT(S) OF THE INVENTION

It is now in order to describe the invention for eliminating the bandwidth and other previously described bottlenecks in its new 'AMPIC DRAM'-centered solution by dramatically reducing the number of transfers and corresponding arbitrations on the system bus, thus substantially improving the over all system performance and with a much faster internal data transfer capability. Other benefits include the system

25

expandability with much lower impact on the data bandwidth, resulting in reduced system cost as before stated.

5 Referring to Fig 7, a CPU unit, so labelled, using conventional primary parallel port data, is shown connected to a system bus, also connected with the main memory unit containing the later-described 'AMPIC DRAM' of the invention, and into which supplemental serial interface inputs (#1 thru #n) are connected from respective input/output (I/O) resources #1 through #n.

10 In this illustration, multiple one bit wide independent serial interfaces are thus provided on the 'AMPIC DRAM' to transport data between the I/O resources and the main memory. These serial interfaces are in addition to the primary parallel ports for the system bus interface for use by the central processing unit CPU or similar master controller device. The number of such serial interfaces is limited only by the device
15 technology, pin count, power consumption and cost, etc. The serial data received or data to be transmitted via these interfaces #1 through #n, is stored inside the 'AMPIC DRAM' in small respective buffers #1 through #n, as more fully shown in Fig 8. For practical considerations, this could range from 64 bytes to 512 bytes, but in theory is limited by the layout of the sense amplifiers. In a traditional layout, it is limited to the
20 number of data bits available in one row of the DRAM core. Thus, if 1024 bytes are available per row access, then the maximum buffer size 'AMPIC DRAM' could be designed for 1024 bytes per internal bank. If 'm' is the number of buffers and 'n' is the number of serial interfaces, then the number of packet buffers 'm' (where the term "packet buffers" is used interchangeably with the term "buffers") is greater than or
25 equal to the number of serial interfaces 'n'. The upper limit on 'm' is restricted by the technology limitations rather than by the architecture.

In accordance with a preferred embodiment, a multiplexer and/or crossbar switch logic or combination (mux/crossbar switch in Fig 8) connects the 'n' serial interfaces to the 'm' buffers. The connection between each serial interface and a buffer is dynamically configured by the CPU (or the current system bus master) and is altered as appropriate
5 for data routability.

A functional block diagram of one possible serial data interface implementation is shown in Fig 9, for the configuration of 4 serial interfaces and 4 buffers. There are, however, multiple ways to implement the desired architecture, though the basic idea
10 remains the same. When data need be moved in Fig 8 between the packet buffers and the DRAM core, arbitration among the various active packet buffers and the CPU is performed. No arbitration, however, is required to receive or transmit data from/to the packet buffers via the serial interfaces.

15 The incoming packet buffer can be redefined as output packet buffer and data rerouted to its destination without even performing the intermediate step of transferring the data between the buffers and the core DRAM. This reduces the latency involved in reception of an incoming packet and its subsequent transmission to its destination. This is possible only because of the 'AMPIC DRAM' capability to assign any buffer to
20 any serial interface via the Mux/Crossbar switch module.

In 'AMPIC DRAM' architectural organization, multiple one bit wide serial interfaces can also be configured as a narrow width bus (termed "port") of sizes such as 1, 2, 4 or 8, etc, but in theory there is no such limitation. It can be any number from 1 to 'n'
25 permissible by the device technology and is also subject to implementation. Once multiple serial interfaces have been grouped and defined as a port, they all connect to a common packet buffer, as more particularly shown in Fig 10, wherein a 1 bit wide

port is the same as a serial interface. This allows even faster data transfers and maintains flexibility at the same time, and is very useful in interfacing with resources operating with different bandwidth and data transfer requirements. It also follows that each packet buffer, Fig 8, has the capability to interface with a maximum of 'n' serial
5 interfaces simultaneously, if defined as a port. The buffers get configured for the same port size identical to the port to which it is connected (also sometimes termed docked).

The serial data flow on each port is controlled by its respective control line. Each port
10 interface consists of one control and a set of serial data interfaces. As an example, if each serial port is only 1 bit wide, then one control line is used per one bit of data line. If two serial interfaces are configured as one port, then one control line is used for the two bit port and so on, as shown in Fig 11. To minimize the pin count, moreover, a serial interface control line is also configurable as a serial data interface line so long as
15 the criterion of one control line per port is met. The association of each control line to its port is configurable. The purpose of the control lines is to control the data flow between the I/O resources and the serial ports. If pin count is not a concern for the manufacturer, one certainly could provide separate control pins. For some applications, furthermore, control pins are not even required, and the I/O resource
20 parallel interface to the system bus is sufficient to exchange the control information.

The data transfer format between an I/O resource and the 'AMPIC DRAM' serial port is such that each memory chip (part of the same external bank) receives and transmits data bits on its port simultaneously, as illustrated in Fig 12. This can be better
25 explained with an example. Assume the four 'AMPIC DRAM's of 2M x 8 form with a 32 bit wide system interface as in Fig 12 (with a port being defined as a one bit serial interface). Each of the four chips receives data simultaneously. The chip0 receives

bit0, chip1 receives bit8, chip2 receives bit16; and chip3 receives bit24. In the next cycle, all the bit numbers will be incremented by 1. This will continue until all 32 bits have been transferred, so that each chip received its 8 bits. Once completed, this process will be repeated for, as illustrated, in Fig 12, the next 32 bits, and so on.

5

Consider, now, another example when a port is defined as consisting of two serial interfaces, as in Fig 13. A total 8 bit interface is then provided to the I/O resource, which must provide two bits each, to each 'AMPIC DRAM' simultaneously. The ordering of the bits is such that chip0 receives bit0 and bit1, chip1 receives bit8 and bit9, chip2 receives bit16 and bit17; and chip3 receives bit24 and bit25 concurrently. In the next cycle, all the bit numbers will be incremented by 2. This will continue until all 32 bits have been transferred, so that each chip received its 8 bits. Once completed, this process will be repeated for the next 32 bits, and so on.

15 It should be noted that this architecture of the invention does not prevent the I/O resources, such as network controller chips, to share the parallel system bus, if so desired. It could be useful for tasks such as controller configuration and status management.

20 Preferably, the 'AMPIC DRAM' is provided with one master clock pin as shown in Fig 19, and every serial interface is configured to operate at a multiple or submultiple of this clock rate, allowing flexibility to accommodate a variety of resources. It is also possible to provide more than one independent clock instead of one master clock, the limitation being only the device technology, pin count, and cost constraints. It should
25 be noted that the clock frequency assignment is a characteristic of the serial interface and not of the buffers. Thus, any of the 'm' buffers can be docked to any of the serial ports and operate at that port speed.

The 'AMPIC DRAM' configurability of the invention, in addition, allows switching of the serial interface from one buffer to another buffer without interrupting the transfers. This has a number of important applications in networking and graphics. One buffer
5 can be loaded with row wide data in one access, while the other one is being used to transmit the information. It should be noted, furthermore, that while a similar two buffer approach has been implemented in VRAMs by other vendors, called split buffer transfers, this is quite different from the present invention in that the external I/O
10 interface with the buffer is always fixed and of the same width as the VRAM system data width. The 'AMPIC' 's multiplexer/crossbar switch module of this invention, totally eliminates all such limitations.

There can be more than one internal bank in the AMPIC DRAM connected via a row wide bus such that all the buffers reside on this bus; or in another embodiment,
15 separate groups of buffers may be provided for each internal bank.

The impact on the system bandwidth due to the frequent 'BitBlt' operations above-described can be reduced considerably by having more than one internal DRAM core bank, as in Fig 14, connected via a row wide interface such that when a transfer need
20 be made from one internal bank of the memory to the other internal bank, the appropriate row address of each bank is strobed with corresponding 'RAS' signals simultaneously. After the data from the row being read is available at the sense amplifiers, it is written to the other bank. The direction control is provided by the internal logic, with Fig 15 showing the sequencing of such an operation. After a row
25 transfer is completed, another transfer can be initiated, and this process can go on until it is complete. This type of transfer is named 'PRITI' (Parallel Row Internal Transaction Intervention). Obviously when such an internal transfer is going on, other

accesses to the accessed banks are not permitted. It should be noted that the transfer on the serial interfaces can also go on in parallel to this internal transfer. A similar concept, later contrasted from the features of the present invention, is disclosed in US Patent No. 5473566, issued on December 5, 1995.

5

Through this novel technique, in one row access time, a massive amount of data can be transferred. As an example, consider an 'AMPIC DRAM' with two internal banks, each being 1M x 8. The internal structure for each bank is 1K rows, each being 8 bit wide. With 'PRITI' capability, in one 'RAS' cycle, 8K bits can be transferred internally. This is a great advance over current-day existing approaches, where 8 bits of data are transferred via the system bus interface and in the best case will take 1K cycles and corresponding arbitrations. If there are 'r' rows and 'c' columns, then the total number of transfers required with the 'PRITI' capability of the invention is 'r'; whereas total transfers required in the traditional approach are 'r x c'.

15

The process is identical for any number of banks. If there are 'm' internal banks connected via a row wide interface, the 'PRITI' module is capable of transferring data from a bank to more than one remaining bank, simultaneously. This is very useful when a broadcast packet is moved from one bank to all other internal banks. With this invention, no row wide registers or latches (also termed a row wide set of storage elements) are required to perform this operation, thus resulting in a very cost effective implementation.

20

The top level internal structure of the 'PRITI' module is shown in before-described Fig 14. 'PRITI' is loaded with the starting row addresses of each bank and the transfer count. After configured, it arbitrates to acquire the internal buses of both the banks. There can be many variations on this basic concept, such as the 'PRITI' module being

25

configured for a predetermined number of burst transfers once it acquires the right to access the rows, or release the bus after every transfer to allow other resources to share the DRAM core.

5 An alternative embodiment of this invention involves using a row wide set of storage elements as labelled in Fig 27 (or any implementation capable of performing a logically equivalent task) to perform a data exchange operation. As an example, for a 1M x 1 DRAM having 1024 bit wide row, the said set of storage elements would contain 1024 storage elements. Here, a row of an internal bank is accessed with read
10 operation, and the retrieved data at the sense amplifiers for the bank (call it bank2) is stored in the row wide set of storage elements. Data is then retrieved from the other bank (bank1) and is then written to bank2. Subsequent to this operation, data from the storage elements is written to the bank1. The diagram of Fig 18 shows a suitable sequencing of such an operation. This implementation requires less circuit than the
15 later-described approach with two sets of storage elements and still permits data exchange, though at the cost of somewhat slower execution. This approach is a universal approach for banks 1 to 'm'. This capability permits a massive information exchange in a very short time, an exceptionally useful tool for multimedia/graphics applications. This implementation obviously requires more circuit than the original
20 approach due to the addition of a set of storage elements, but here it is not necessary to save the original data before new data is moved at its place.

Another modification of this invention uses two sets of storage elements as labelled in Fig 16 (or any circuit capable of performing a logically equivalent task) to perform a
25 data exchange operation. Here, a row in each of two internal banks is accessed simultaneously with read operation, and the retrieved data at the sense amplifiers is stored in the row wide set of storage elements, as indicated in Fig 16. Data so

retrieved and then stored is subsequently written back to both the rows simultaneously. The diagram of Fig 17 shows exemplary sequencing of such an operation.

- 5 The 'PRITI' approach of the invention, therefore, is not restricted to two internal banks only, and is equally applicable for any multibank organization inside a DRAM chip. It is also possible to add 'PRITI' capability even in a traditional type DRAM without the rest of the 'AMPIC DRAM' architecture. A more sophisticated 'PRITI', furthermore can also have the transfer boundaries defined in terms of columns in addition to rows
10 requiring additional registers to load the column addresses.

Unlike the system of said US Patent No 5473566, in accordance with the preferred embodiment of the invention, only one row wide set of storage elements is required, and not one for each internal bank. This is what renders the approach of this invention
15 suitable for universal application and keeps the DRAM construction relatively inexpensive.

Further modifications to this invention:

- 20 a. more than one row wide bus is implemented along with their own set of storage elements connecting multiple banks, thus allowing more than one parallel 'PRITI' transfers. In general, if there are 'm' banks, then the maximum number of possible row wide buses without redundancy is 'm/2'. If a set of such above-described storage elements is used per bus, it follows that only 'm/2' sets of storage elements are
25 required to make 'm' separate simultaneous 'PRITI' transfers, one for each bank;

b. If the number of banks is large, then the banks can be subgrouped on separate buses. As an example, in an 8 bank configuration, 4 banks can reside on one bus while the other four are on the second bus with their own above-described 'PRITI' transfer implementation, and these two sub groups are then connected via another bus
5 with any of the above-described 'PRITI' transfer capabilities.

While a row wide bus is used in this description, a fraction of the row wide bus is also valid, if required to reduce the cost. It should be noted, moreover, that the memory does not have to be a DRAM to take advantage of these particular capabilities.

10

The 'AMPIC DRAM' Impact on Interface Design

The chip has somewhat different pin out to reflect its unique architecture. One possible pinout for a 2M x 8 chip with 9 serial interfaces is shown exemplarily in Fig
15 19, with the added pins requiring changes in the interface design of the 'AMPIC DRAM'-based main memory.

A 'WAIT' signal is provided for the system bus interface, Fig 19, whenever an internal transfer between the buffers and the DRAM core is taking place. The CPU (or other
20 master controller) can either use it to delay the start of access or in an alternative implementation, the access cycle can be extended to allow for the internal transfer to complete, before proceeding with this access. A master clock pin ("master clock"), as explained earlier, is also provided.

25 As the 'AMPIC DRAM' is highly configurable, a mechanism is required to differentiate between a regular DRAM core access versus either a configuration

command or buffer transfer interaction. The approach in Fig 19 is to provide an additional control signal to signify a command or data access.

During command cycle, the command instruction can be carried over the data lines as
5 they are not used during 'RAS' cycle. This is specifically useful for internal transfer
commands, where a DRAM core address need be provided along with the buffer ID.
This scheme allows use of traditional signals 'RAS' and 'CAS' to provide the core
DRAM address, where the data lines will have the buffer number or any other
10 additional information/instruction. In fact it is possible to issue two commands; one,
when 'RAS' goes active, and then when 'CAS' is asserted. There are multiple known
ways to implement this access mechanism, such also being a function of the device
technology and cost considerations.

While the 'AMPIC DRAM' of the invention has more pins than the traditional DRAMs
15 due to serial ports, if the earlier proposed 'PARAS' type DRAM model of said pending
application is used, one could have this DRAM with only a marginal increase in
number of pins.

Networking Application Example with 'AMPIC DRAM'

20

As before explained, in accordance with the invention, serial interfaces/ports are
provided between each network controller and the main memory. The data movement
between the controllers and the main memory is primarily serial. The serial data
received from a network controller, or data to be transmitted to a network controller, is
25 stored in a packet buffer, assigned to it by the system bus master. This discussion, of
course, assumes that the network controllers are capable of sourcing or receiving serial
data streams in the format required by the new system architecture.

Considering the same earlier example of a 32 bit wide bus, four 2M x 8 AMPIC DRAMs instead of traditional DRAMs, row wide packet buffer and four network interfaces, a user on, for example, the Ethernet network is sending a packet of 1024
5 bytes to another user on, for example, the FDDI network. In this novel system architecture with 'AMPIC DRAM' based main memory, as shown in Fig 20, the data is to be received by the serial port on the 'AMPIC' connected to the Ethernet controller. No arbitration is required and no main memory bandwidth is consumed on the transfers. After the data transfer has been completed (each of the four 'AMPIC
10 DRAMs' would receive 256 bytes), it can be transferred entirely to the DRAM core in only one access after this packet buffer acquires the internal bus via arbitration. When a row address is provided to the DRAM core, its sense amplifiers have all the data bits of this row available. Thus, the entire packet buffer can be stored in one access. If the size of the packet buffer is less than one row wide, then multiple, though still few,
15 accesses are required.

This is a most significant advantage over the current-day existing parallel bus solutions before-described, where 256 accesses and corresponding arbitrations were required.

20

After this packet is transferred to the 'AMPIC DRAM' core, it is processed by the CPU and redirected towards the FDDI port in this example. Now the reverse process takes place. The packet is transferred from the core to the appropriate packet buffer in a single access requiring arbitration. This data is subsequently transferred from the
25 packet buffer to the FDDI controller via the serial ports and then concurrently moved from the FDDI controller chip to its network. Once again, this reverse process would

arbitrate only once for its transfer, whereas 256 transfers and corresponding arbitrations are required in existing designs.

Additional advantages available with this new DRAM of the invention also reside in
5 the considerable gain achievable by the fact that a broadcast packet can be loaded in
all the appropriate buffers in one access only, and subsequently transferred to the
whole network, and that multiple rows can be accessed sequentially and loaded into
different buffers and subsequently transferred via their ports, providing significantly
higher performance.

10

Graphics/Multimedia Application Example with 'AMPIC DRAM'

As before stated, a significant bandwidth of any graphics system is spent in 'BitBlt'
operation, where a large chunk of data from one area of the memory needs to be
15 moved to another. This consumes a considerable portion of the system bandwidth and
because of it, in general, DRAMs used for graphics purposes are kept separate from
the main system memory. This necessity, however, adversely impacts the system cost.
The present invention also provides an approach to eliminate the necessity of two
separate buses, as presented later.

20

Using the earlier example of prior art Fig 5 where 16 rows of a display need be
updated, but this time the DRAM component is 'AMPIC DRAM' of the same 2M x 8
size and with two internal banks, each of size 1M x 8, and equipped with 'PRITI'
capability, 16 rows worth of data can be transferred in the new location, with the
25 number of data transfers for the new DRAM being fortuitously precisely the same as
the number of rows:

Number of Transfers = Number of rows to be transferred (16).

This again is a vast improvement over the 16384 transfers and associated arbitrations currently required for the conventional DRAM-based designs, and represents a 3 order
5 of magnitude improvement in the performance. It also reduces the impact on the system bandwidth proportionately. In but one row access time, massive amounts of data can be transferred in accordance with the present invention. The novelty of the architecture, both at system and chip level, allow unique DRAM configurations to provide enhanced system performance.

10

The 'AMPIC DRAM' can also be configured as previously stated to provide the graphics screen data to the display at high speeds.

Consider, for example, the following illustrative parameters applied to the
15 configuration of Fig 21:

- a. 5 serial interfaces are available per chip;
- b. 32 bit wide system bus implemented with 4 such chips;
- 20 c. 4 serial interfaces per chip have been defined as a port and used to transfer the display data; and
- d. data retrieval rate per port is at 100 Mhz (can be faster than this rate).

25

In this example, since 4 serial interfaces per chip are used for graphics, a 16 bit wide graphics interface is implemented, which can provide data at the rate of 2 bytes every

SUBSTITUTE SHEET (RULE 26)

SUBSTITUTE SHEET (RULE 26)

clock, thus providing 200 Mbytes per second bandwidth, sufficient for most graphics applications. If an 'AMPIC DRAM' with 9 serial interfaces is used, the chip could be configured for an 8 bit wide bus to provide more display data bandwidth.

5 Further Examples of 'AMPIC DRAM'-Based System Configurations

In the example used above, one external bank with a 32 bit wide system bus was used. Some applications, however, may use more than one external bank, such as 4 banks, 32 bit wide each of 'AMPIC DRAM', as shown in Fig 22. This architecture permits
10 that different network interfaces can be connected to each bank, if so desired. This can increase the network interconnectivity to a much larger level within reasonable cost as compared to the prevailing solutions. As an illustration, if 9 serial interfaces were provided on each 'AMPIC DRAM' and each bank was connected to 4 network
15 interfaces, then a total of 16 networks can be connected. This represents a major gain compared to the existing technology which generally peaks out at 4 to 5 interfaces.

It also follows from this architecture, that when more than one external bank is used, an 'AMPIC DRAM' serial port from one bank can be connected to a serial port of
20 another bank. This provides an added path between banks such that packet buffers can be used to transfer data rapidly between external banks.

It is not necessary, moreover, to use all 'AMPIC DRAMs' in a system configuration. Some applications may mix 'AMPIC DRAMs' with existing type of DRAMs, as
25 proposed in the modifications of Fig 23.

In still another system configuration, the 'AMPIC DRAM' can be used to provide both the graphics or display interface, as depicted in Fig 24, and to connect to other types of I/O resources such as, for example, video cameras, or a satellite interface or like.

5 Unified Memory Architecture

In an ideal world, it would be most advantageous, as before explained, to have a common memory chip for both the graphics and the main memory functions and still provide the necessary performance. This approach dubbed as 'Unified Memory
10 Architecture' (UMA) is currently under intense debate and some proposed solutions have been offered including the before mentioned RDRAM chip. While this proposal uses a smaller number of pins, thus resulting in lower power consumption, smaller real estate and relatively lower cost, the packet based protocol and the interface limitations earlier discussed, do not work efficiently as main memory, where accesses tend to be
15 non-localised.

Another potential solution is the use of the previously described VRAMs for both the main memory and graphics memory; but the added cost does not justify the change.

20 Before the present invention, therefore, unfortunately no reasonable solution has come close to meeting the requirements of the broad cross-section of the PC market.

A system level solution based on the 'AMPIC DRAM' of the invention, previously discussed in connection with the embodiment of Fig 19, provides the configurable
25 serial interfaces and 'PRITI' capability that certainly fills this void. It has more signal pins (but the number of power and ground pins may be smaller) than the one proposed by the RAMBUS, but less than VRAMs, and is equally efficient at both the

operations. With this implementation of the invention, indeed, both the graphics and main memory functions can reside in the same memory with negligible bandwidth reduction and thus meet the elusive goal of the 'Unified Memory Architecture'.

5 Another alternate embodiment for such solution is to have two internal banks in the chip. One may follow the 'AMPIC' model for graphics or similar applications, while a second internal bank, possibly larger, may resemble the traditional DRAM based memory, and with both banks sharing the 'PRITI' capability of the invention, as in Fig 25. This potent integration provides the best of both worlds; one bank appears like
10 main memory, while the other bank appears as an optimized graphics memory. This chip architecture allows massive amounts of data transfers between the two internal banks with very little impact on the system bandwidth due to the 'PRITI' capability, and thus provides all the necessary capabilities needed to allow a common chip and one single bus, universally for all types of applications.

15

Further improvement to this invention may be made to reduce the pin count and thus cost, if the interface access mechanism of said copending application called 'PARAS' DRAM is used along with the 'AMPIC' organization, as in Fig 26, to realize a memory
20 chip, optimized for both the requirements: main memory and graphics at the lowest possible cost. Consider the example of a 2M x 8 DRAM with 'PARAS' interface, the number of pins saved is 8, which then can be used to provide the serial interfaces. If only 5 serial interfaces are implemented, then the number of pins in this high performance/low cost chip is comparable to the traditional DRAMs, of course with the tremendous additional benefits previously explained. The advantages of such an
25 integrated DRAM are:

- a. It provides significantly enhanced system data bandwidth with architectural innovations, rather than sheer device speeds.
- b. It is capable of moving the large amount of data to/from multiple I/O resources
5 with minimum impact on the system bandwidth.
- c. It is configurable to accommodate different data transfer rates of the I/O resources.
- 10 d. It is capable of moving large blocks of data internal to the chip in several orders of magnitude faster time frames and with negligible impact on the system bandwidth.
- e. The device has low pin count for the functionality provided.
- 15 f. It is relatively low cost due to reduction in the number of pins.
- g. Relatively low power consumption is required.
- h. The architecture reduces the latency time between the reception of an incoming
20 packet and its subsequent transmission.
- i. It interconnects a considerably larger number of I/O resources as compared to the traditional approach.
- 25 j. The system design interface is nearly identical to the existing DRAMs, thus minimizing the design cycle.

k. It works equally efficiently for both the main memory and the graphics requirements, thus providing a Unified Memory Architecture.

Further modifications will occur to those skilled in this art including, among others,
5 providing logic to enable offloading the same packet buffer to similarly defined other ports, and the capability to string buffers, or applying the serial interfacing and buffer switching to other memory devices than DRAMS, and such are considered to fall within the spirit and scope of the invention as defined in the appended claims.

CLAIMS

1. For use in a system having a master controller such as a central processing unit (CPU) having parallel data ports and a dynamic random access memory (DRAM) each
5 connected to and competing for access to a common system bus interface, an improved DRAM architecture comprising a multi-port internally cached DRAM(AMPIC DRAM) comprising a plurality of independent serial data interfaces each connected between a separate external I/O resource and internal DRAM memory through corresponding buffers; a switching module interposed between the serial
10 interfaces and the buffers; and a switching module logic control for the connecting of the serial interfaces to the buffers under a dynamic configuration by the bus master controller, such as said CPU, for switching allocation as appropriate for the desired data routability.
- 15 2. A system as claimed in claim 1 and in which the switching module comprises one or more multiplexers or cross bar switches, or a combination thereof.
3. A system as claimed in claim 1 and in which the said memory is the DRAM
20 core main memory of the system.
4. A system as claimed in claim 3 and in which the buffers are data packet buffers and means is provided for performing arbitration for bus interface access amongst the various active packet buffers and the CPU, and with no arbitration required to receive or transmit data from or to the packet buffers via the serial interfaces.

25

5. A system as claimed in claim 1 and in which the AMPIC DRAM switching module assigns any buffer to any serial interface and without any intermediate step of transferring data between the buffer and the core DRAM.
- 5 6. A system as claimed in claim 1 and in which each independent serial interface is one bit wide.
7. A system as claimed in claim 6 and in which multiple one bit wide serial interfaces are configured with a common I/O source as a narrow width bus or port,
10 connected to a common buffer.
8. A system as claimed in claim 7 and in which, each buffer has the capability to interface with up to all the serial interfaces simultaneously when defined as a port, and with each buffer configured for the same port size as the port to which it is connected
15 or docked.
9. A system as claimed in claim 1 and which a control line is provided for each port interface to control the serial data flow between the corresponding I/O resources and the serial port.
20
10. A system as claimed in claim 1, and in which a plurality of the AMPIC DRAM chips is provided, each connected between the bus interface and an I/O resource serial interface and with one or multiple serial interfaces serving as a port therewith.
- 25 11. A system as claimed in claim 1 and in which at least two internal DRAM core banks are provided connected via a row wide interface having row address strobing

with simultaneous corresponding RAS signal cycle facility such that after the data from a row on one bank is read, it is written to at least one other bank.

12. A system as claimed in claim 11, wherein a direction control is provided by the
5 internal logic, such that after a row transfer is completed, another transfer can be initiated, with the resulting parallel row internal transaction intervention (PRITI) continuing to completion.

13. A system as claimed in claim 12 and in which there is provided means for
10 denying the access to the DRAM core during such internal transfer, while permitting transfer on the serial interfaces during the internal transfer.

14. A system as claimed in claim 13 and in which two row wide sets of two storage
15 elements are interfaced between said banks and provided with means for accessing a row in each bank simultaneously with said operation, storing in said storage elements, and then writing back to both sources simultaneously.

15. A system as claimed in claim 13 and in which a row wide set of a storage
20 elements is disposed on the row wide bus interface and is provided with the row data of one bank stored, and written to the at least one other bank after it has written data to the said one bank.

16. A system as claimed in claim 1 and in which the AMPIC DRAM chip in
25 addition to separate serial interface pins is provided on the bus interface side with address, data, RAS, CAS, WRITE, WAIT, CMD/DATA and MASTER CLOCK pins; the CPU using the WAIT signal whenever an internal transfer takes place between the buffers and the DRAM core, either to delay the start of access (WAIT) or to extend

the access cycle to allow for the completion of the internal transfers before processing with the access; a command control signal for access being applied over the data lines not used during the RAS cycle: the RAS and CAS lines providing the core DRAM address while the data line provide the buffer number or additional instruction
5 information; and the master clock controlling the serial interfacing.

17. A system as claimed in claim 1 adapted for a networking application involving a plurality of network controllers each serially interfaced with one or more banks of corresponding main memory AMPIC DRAMs, in turn connected to one side of the bus
10 interface, and with a CPU connected to the other side of the bus interface, wherein the data movement between the controllers and the main memory is primarily serial, and the serial data received from the network controller as data to be transmitted to a network controller is stored in a packet buffer assigned to it by the system bus master
CPU.

15

18. A system as claimed in claim 17 and in which, as one network controller is sending data to a second network controller, the data received by the serial port of the said one network controller is transferred to the corresponding DRAM-based main memory without arbitration or consumption of main memory bandwidth, and with the
20 data transfer applied to the corresponding DRAM core in only one access following the acquisition of the bus by the packet buffer via arbitration, and the row address data is provided to the DRAM core; and wherein the packet transferred to the AMPIC DRAM is thereupon processed by the CPU and redirected towards the said second network controller port and with data transferred to the corresponding packet buffer in
25 a single access after arbitration, and then transferred via the corresponding serial port to the said second network controller and its network.

19. A system as claimed in claim 1 adapted for a graphics/multimedia application involving the transfer of a minimum of a number of rows of display and from a plurality of a display serial interface ports to a plurality of memory AMPIC DRAMs in turn connected to one side of said bus interface, and with a CPU connected to the other side of the bus interface, wherein the data movement between the display ports and the memory is primarily serial, and the serial data received from the display interface is stored in packet buffers assigned to the system bus master CPU.

20. A system as claimed in claim 19 and in which means is provided for transferring said number of rows of display by the same number of data transfer operations.

21. A system as claimed in claim 17 and in which one or more further external banks of conventional traditional DRAMS, absent serial interfacing, are also connected to the bus interface.

22. A system as claimed in claim 17 and in which one or further external banks of AMPIC DRAMs are provided also connected to the bus interface and serially interfaced with graphics display data ports.

23. In a system having a parallel data port CPU controlling a system bus interface to which one or more memory-equipped DRAM units or bank(s) of DRAM units are connected, a method of obviating DRAM system bandwidth limitations, materially increasing data transfer speed, materially reducing bus arbitration requirements, enabling increased I/O resource interfacing and at reduced cost and lower power consumption, that comprises, equipping each DRAM unit with at least one supplemental serial data port for interfacing through a corresponding serial interface

with an external I/O data resource; providing internally of each DRAM a buffer for each of the serial interfaces and with a switching module interposed between the buffers and the serial interfaces; and causing the CPU to control the connecting of the serial interfaces to the buffers by dynamically configuring the switching of the
5 switching module as appropriate for the desired data routability.

24. A method as claimed in claim 23 and in which the switching is effected by multiplexing or cross bar switching or both.

10 25. A method as claimed in claim 23 and in which said switching assigns buffers to any serial interface and without any intermediate step of transferring data between the buffers and the DRAM memory, and arbitrating for bus interface access amongst the buffers and the CPU, but with no arbitration required to receive or transmit data from or to the buffers via the serial interfaces.

15


26. An improved DRAM architecture as claimed in claim 1 for a chip internally containing at least two banks of DRAMs and the switching module and buffers; parallel row internal transaction intervention means for internal data row transfers and connected such that at least one bank serially connected to an I/O graphics display
20 uses another bank primarily as the main memory core; such that the CPU accessing one or both banks, moves data in the said another bank, and the data is moved between the banks under the control of parallel row internal transaction intervention means, thereby providing a chip suitable for unified memory architecture.

25 27. For use in a system having a master controller such as central processing unit (CPU) having parallel data ports and a random access memory unit each connected to and competing for access to a common system bus interface, an improved memory

unit architecture comprising a multi-port internally cached memory unit comprising a plurality of independent serial data interfaces each connected between a separate external I/O resource and internal memory in the unit through corresponding buffers; a switching module interposed between the serial interfaces and the buffers; and a
5 switching module logic control for the connecting of the serial interfaces to the buffers under a dynamic configuration by the bus master controller, such as said CPU, for switching allocation as appropriate for the desired data routability.

- 28 For use in a system for controlling data-accessing of a common bus in an improved DRAM architecture comprising a multi-port internally cached DRAM (AMPIC DRAM) comprising a plurality of independent serial data interfaces each connected between a separate internal I/O resource and internal DRAM memory through corresponding buffers competing for access to a common internal bus; a switching module interposed between the serial interfaces and the buffers; and a switching module logic control for the connecting of the serial interfaces to the buffers under a dynamic configuration of switching allocation as appropriate for the desired data routability among the interfaces.
- 29 A system as claimed in claim 28 wherein one or both of the number of I/O resources and the number of AMPIC DRAM switches is increased to expand the amount of data to be switched and the internal data switching capacity due to the effective corresponding increase in internal bus size.

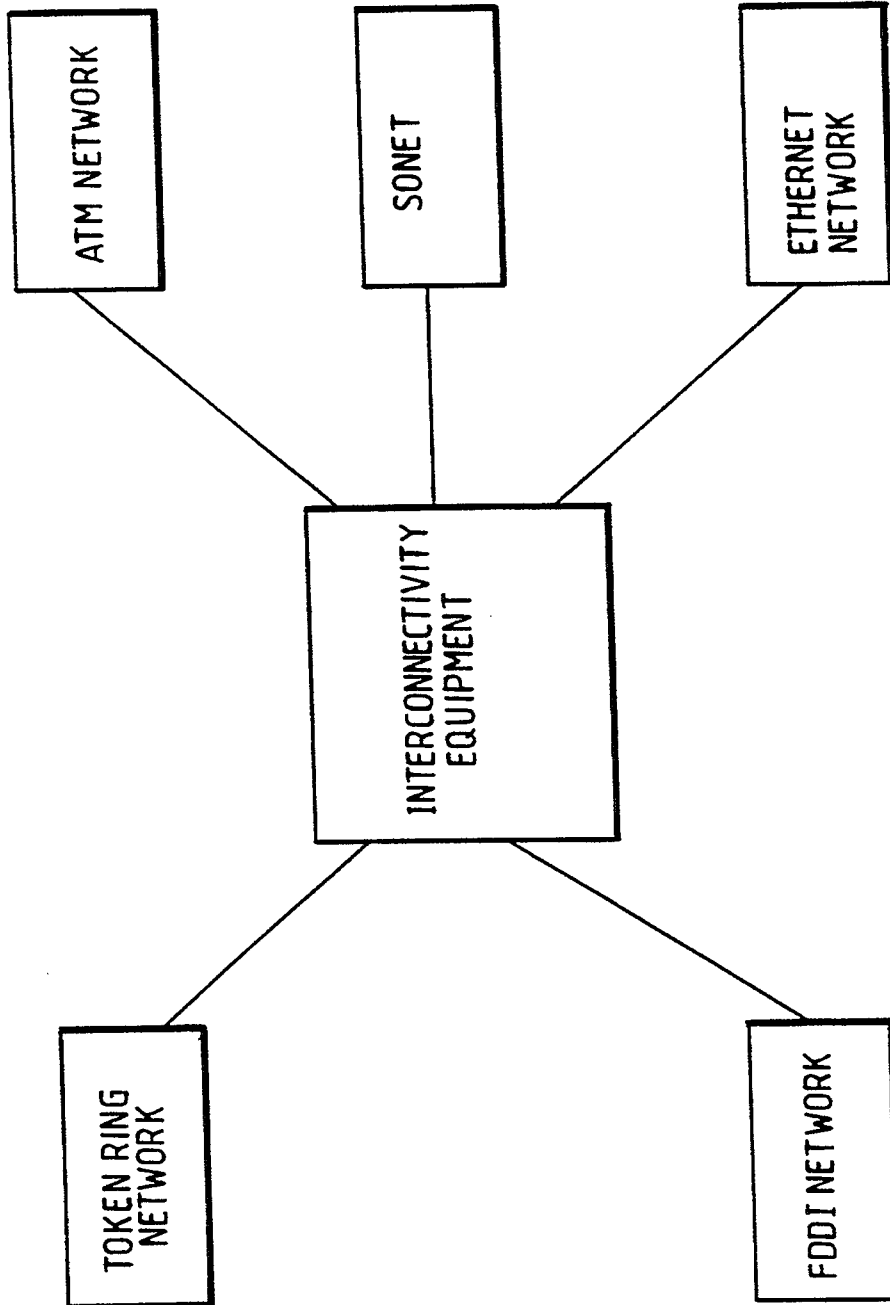
1/27



PRIOR ART

Fig.1.


2/27



PRIOR ART

Fig.2.

3/27



PRIOR ART
Fig.3

4/27





Fig.4.


5/27



PRIOR ART

Fig.5.

6/27



PRIOR ART

Fig.6.

7/27




Fig.7.

8/27

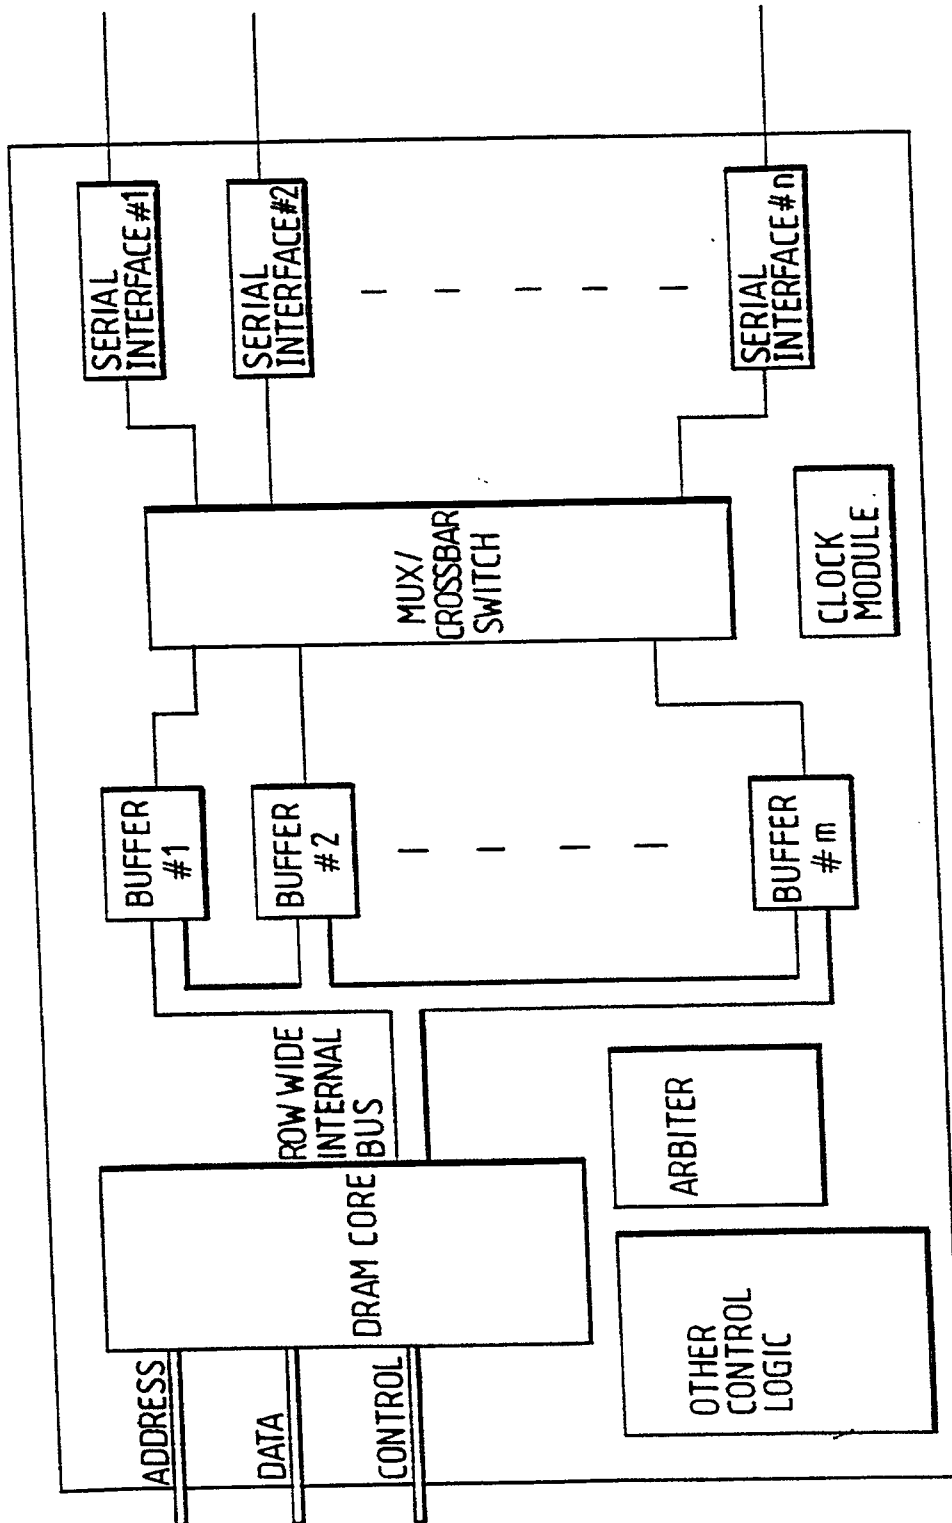


Fig.8.

9/27




Fig.9.

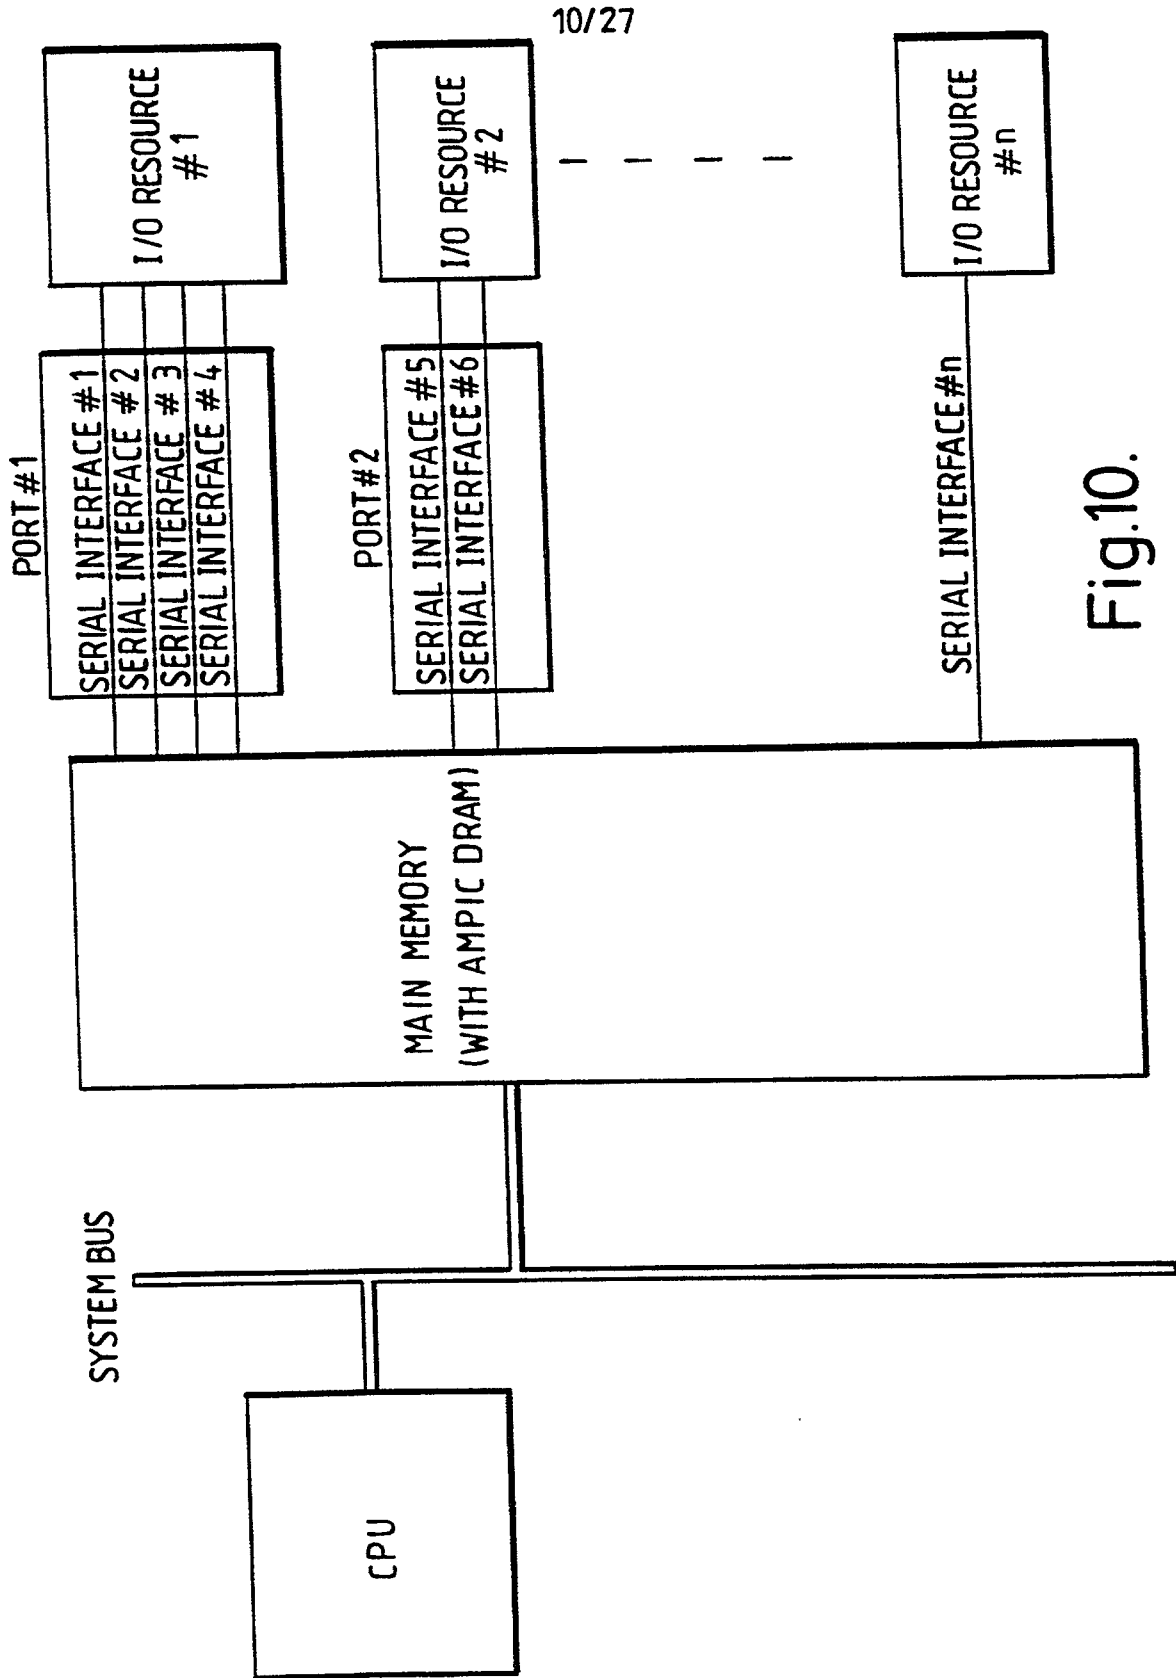


Fig.10.

11/27

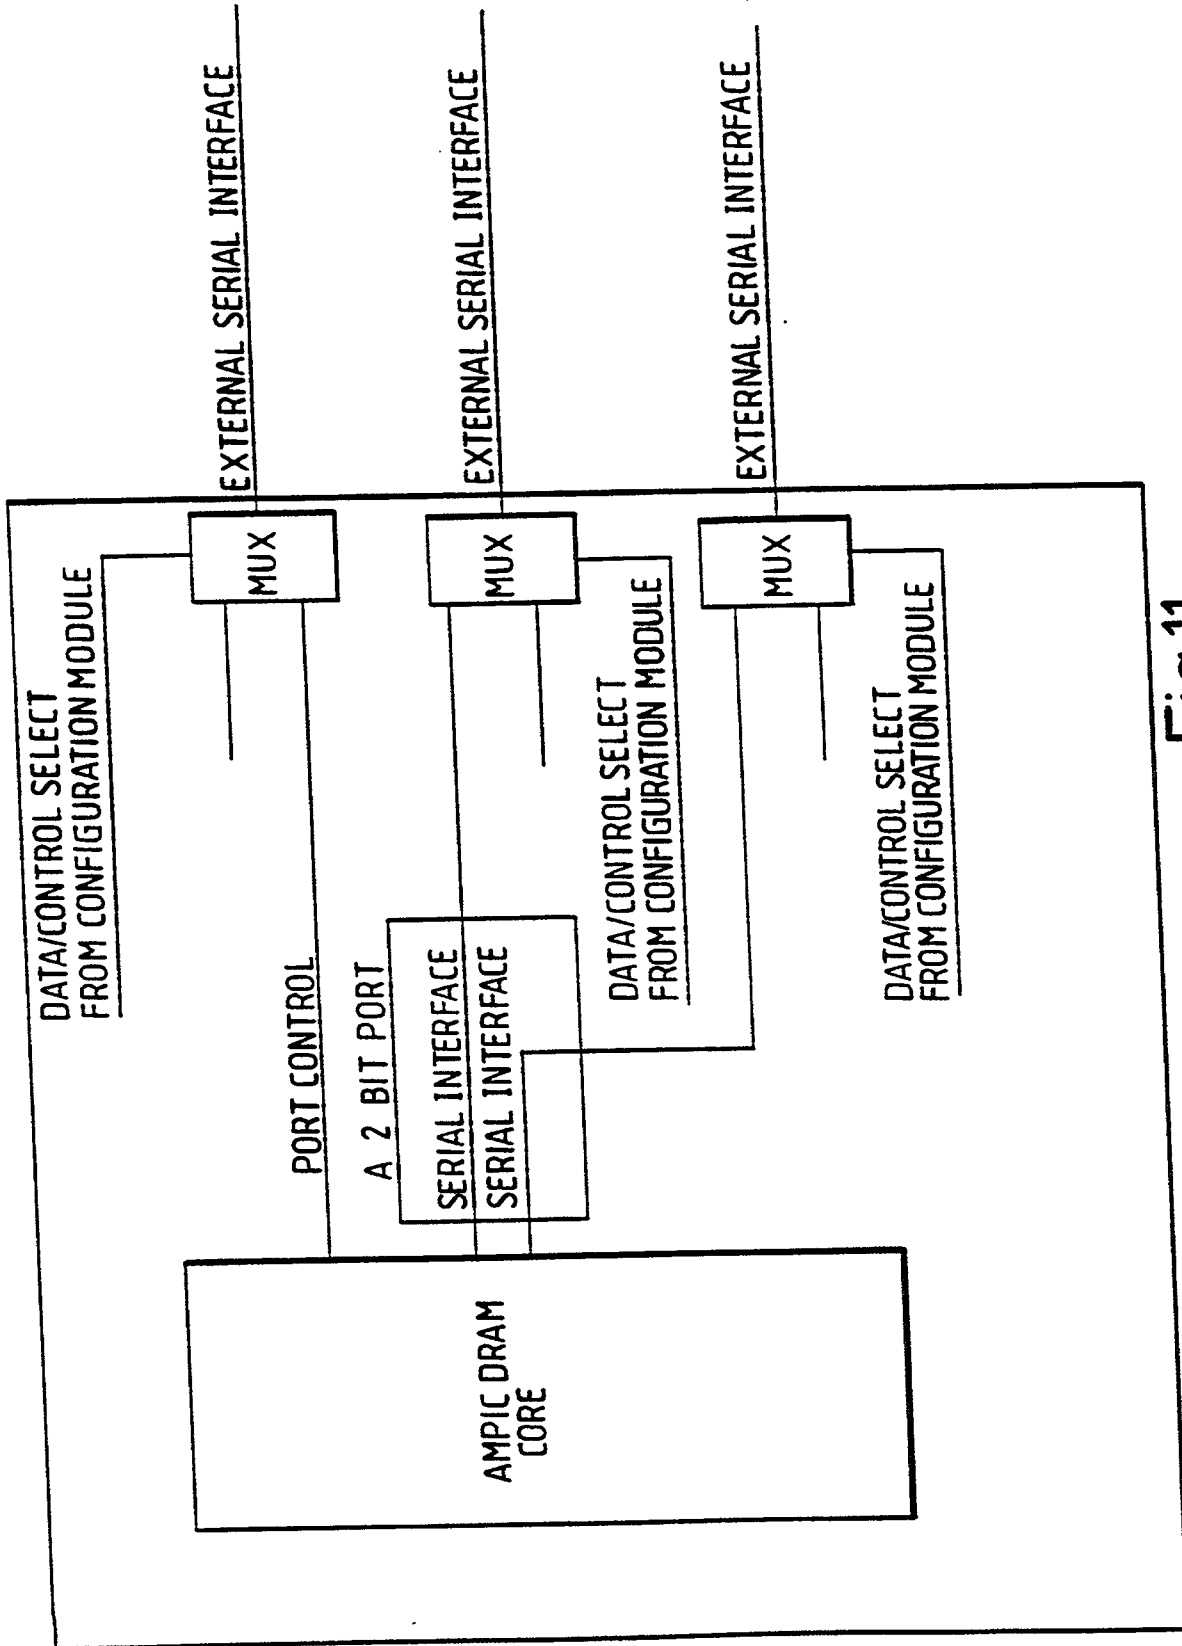


Fig.11.

12/27




Fig.12.

13/27




Fig.13.

14/27




Fig.14.

15/27




Fig.15.




Fig.16.

17/27

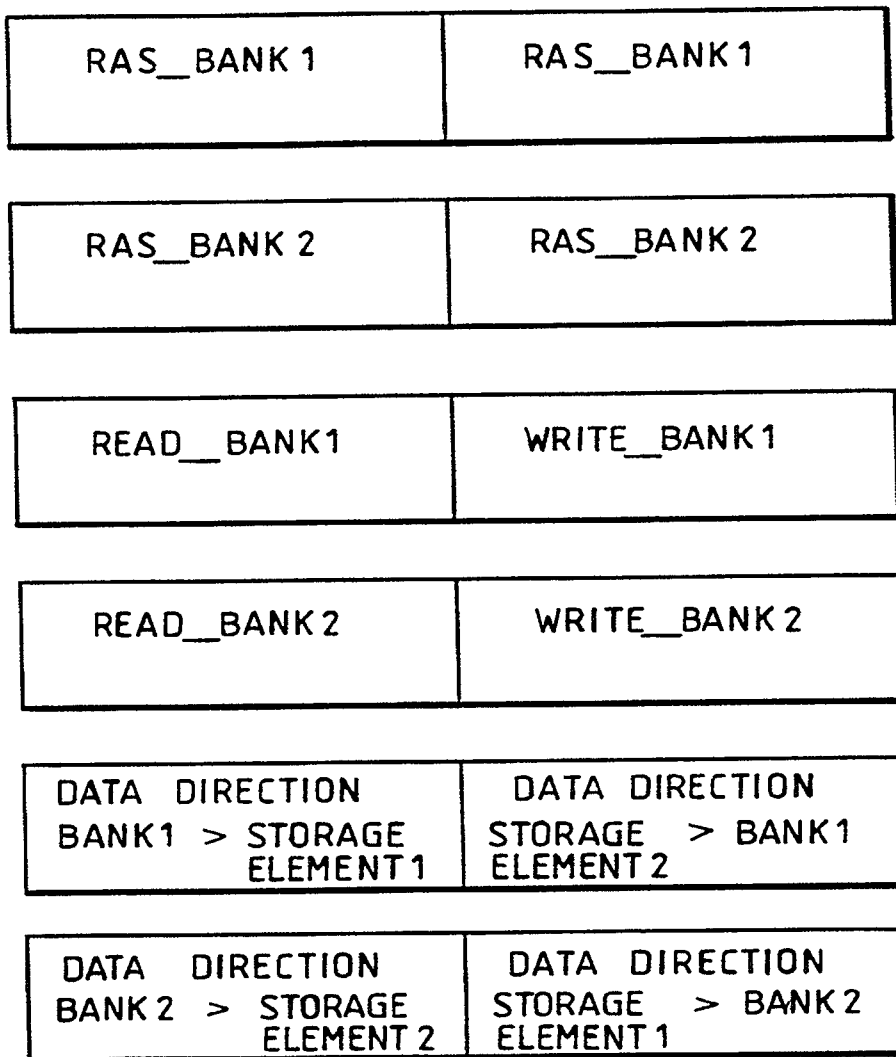


Fig.17.

18/27

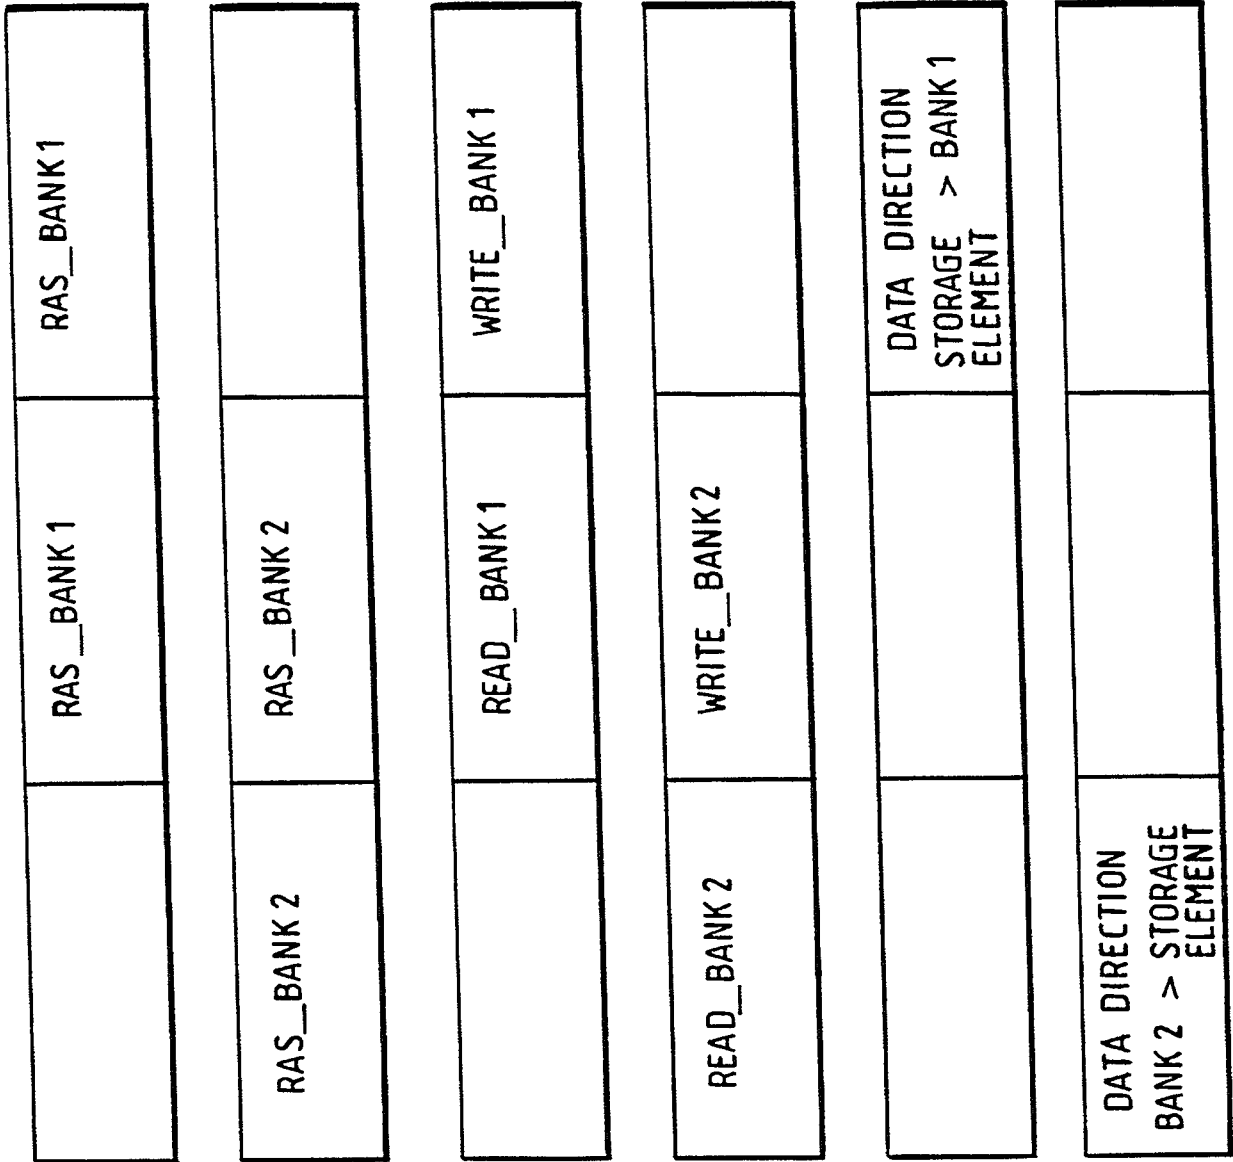


Fig.18.

19/27




Fig.19.

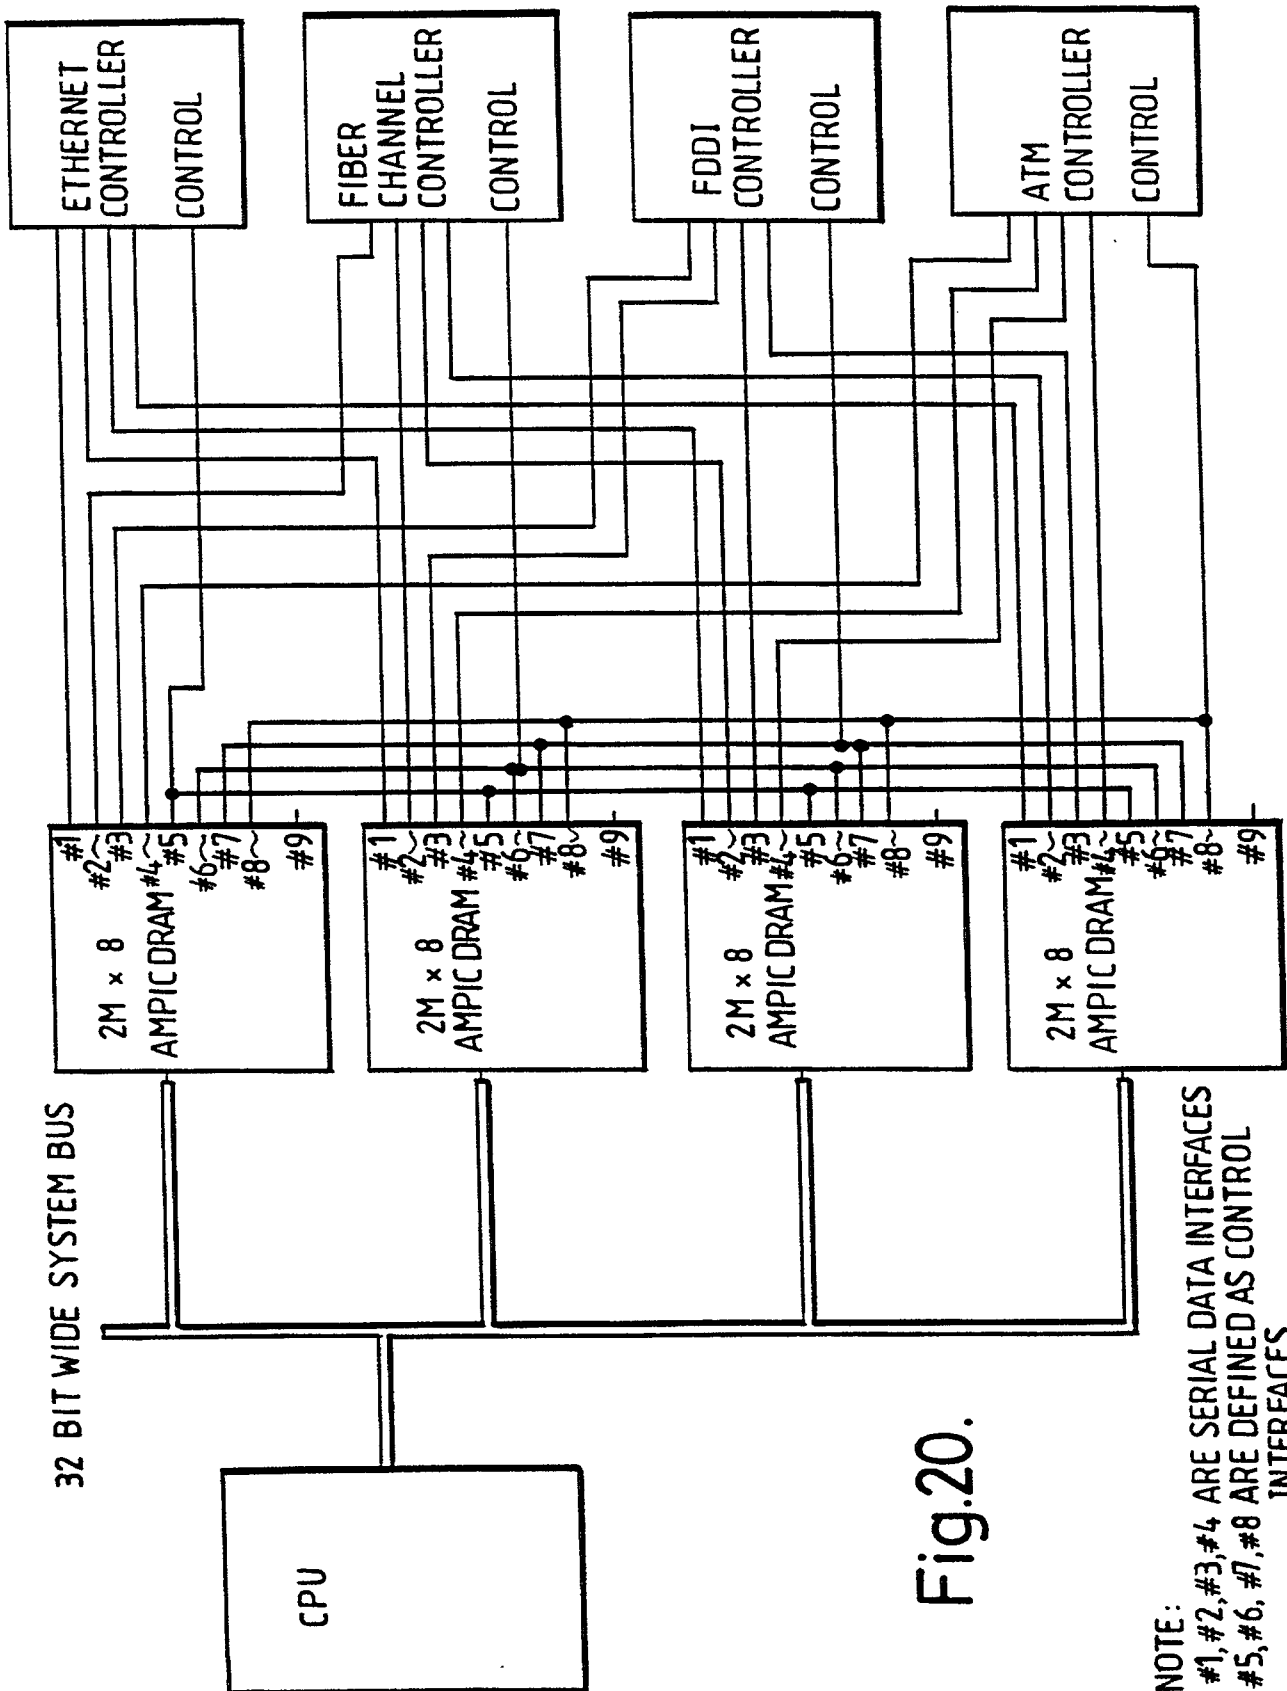


Fig.20.

NOTE:
 #1, #2, #3, #4 ARE SERIAL DATA INTERFACES
 #5, #6, #7, #8 ARE DEFINED AS CONTROL INTERFACES

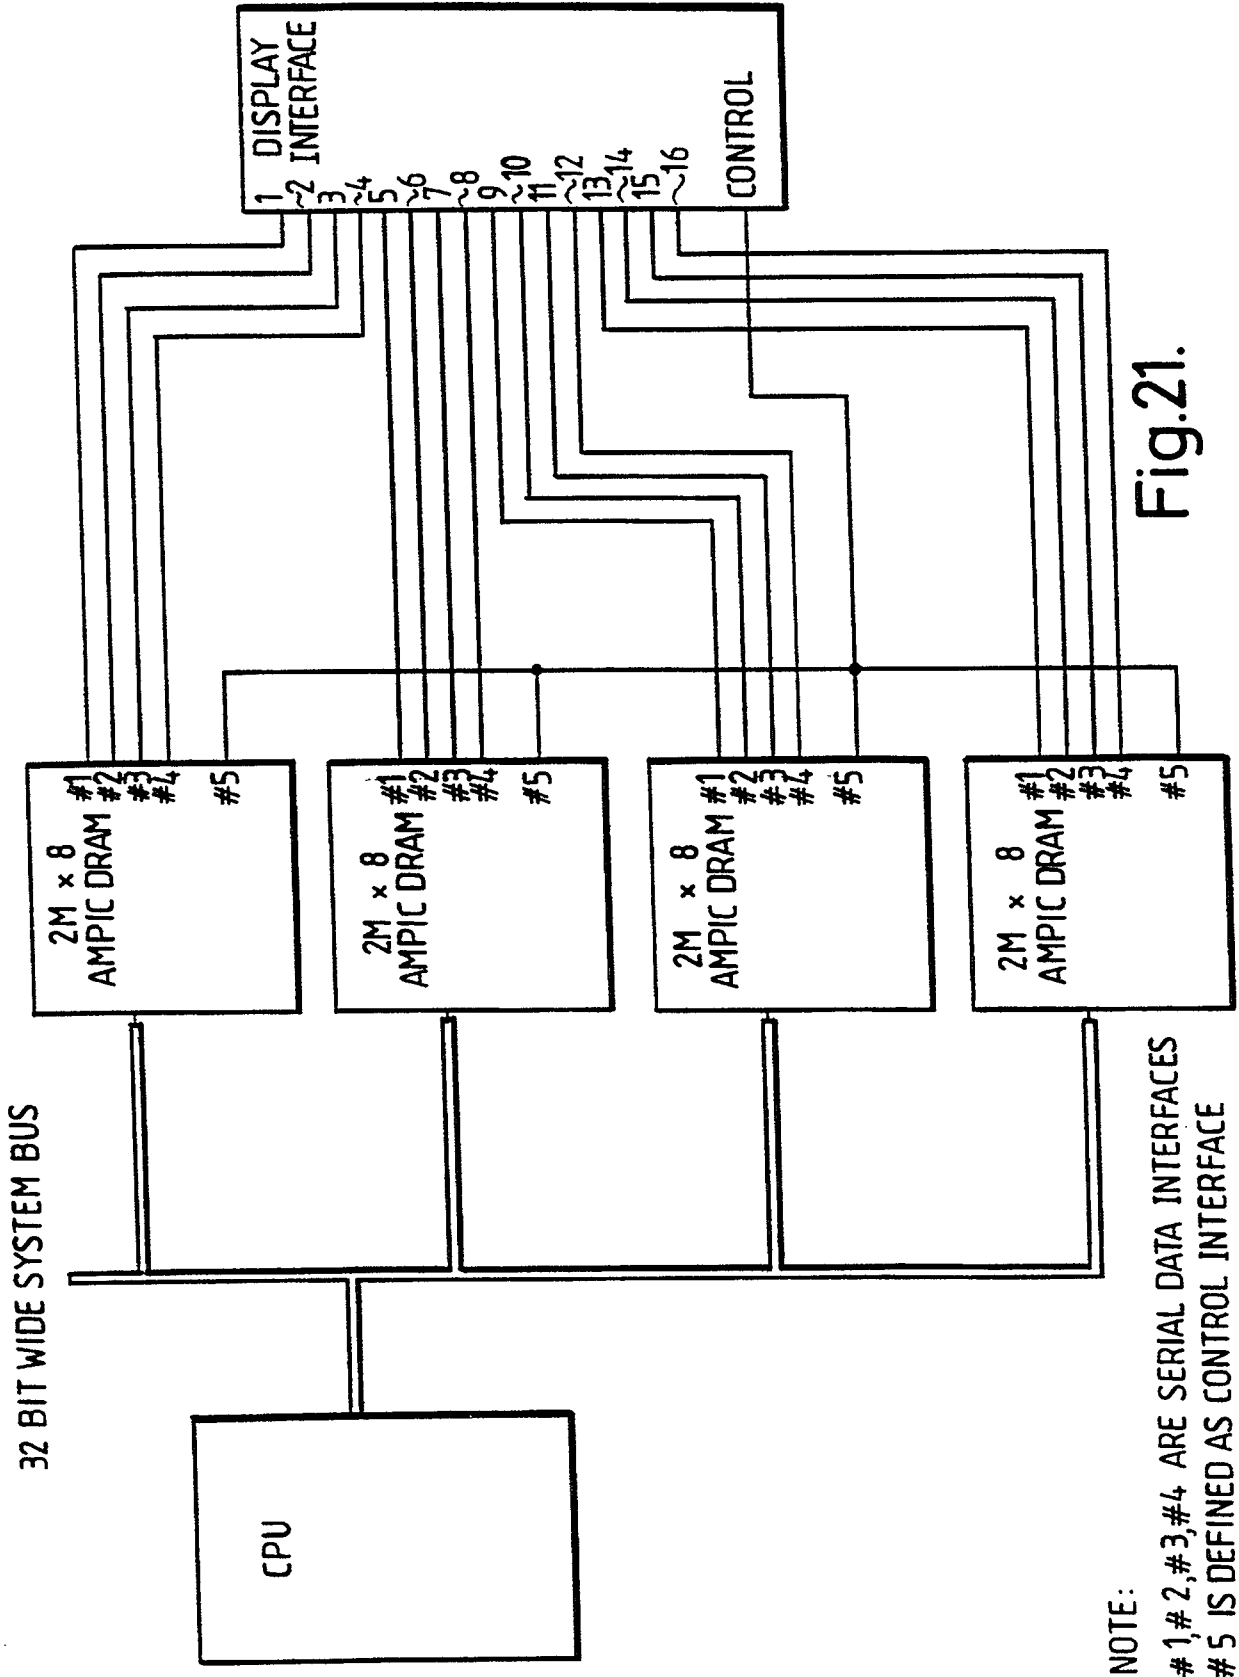


Fig.21.

NOTE:
 # 1,# 2,# 3,# 4 ARE SERIAL DATA INTERFACES
 # 5 IS DEFINED AS CONTROL INTERFACE




Fig.22.




Fig.23.

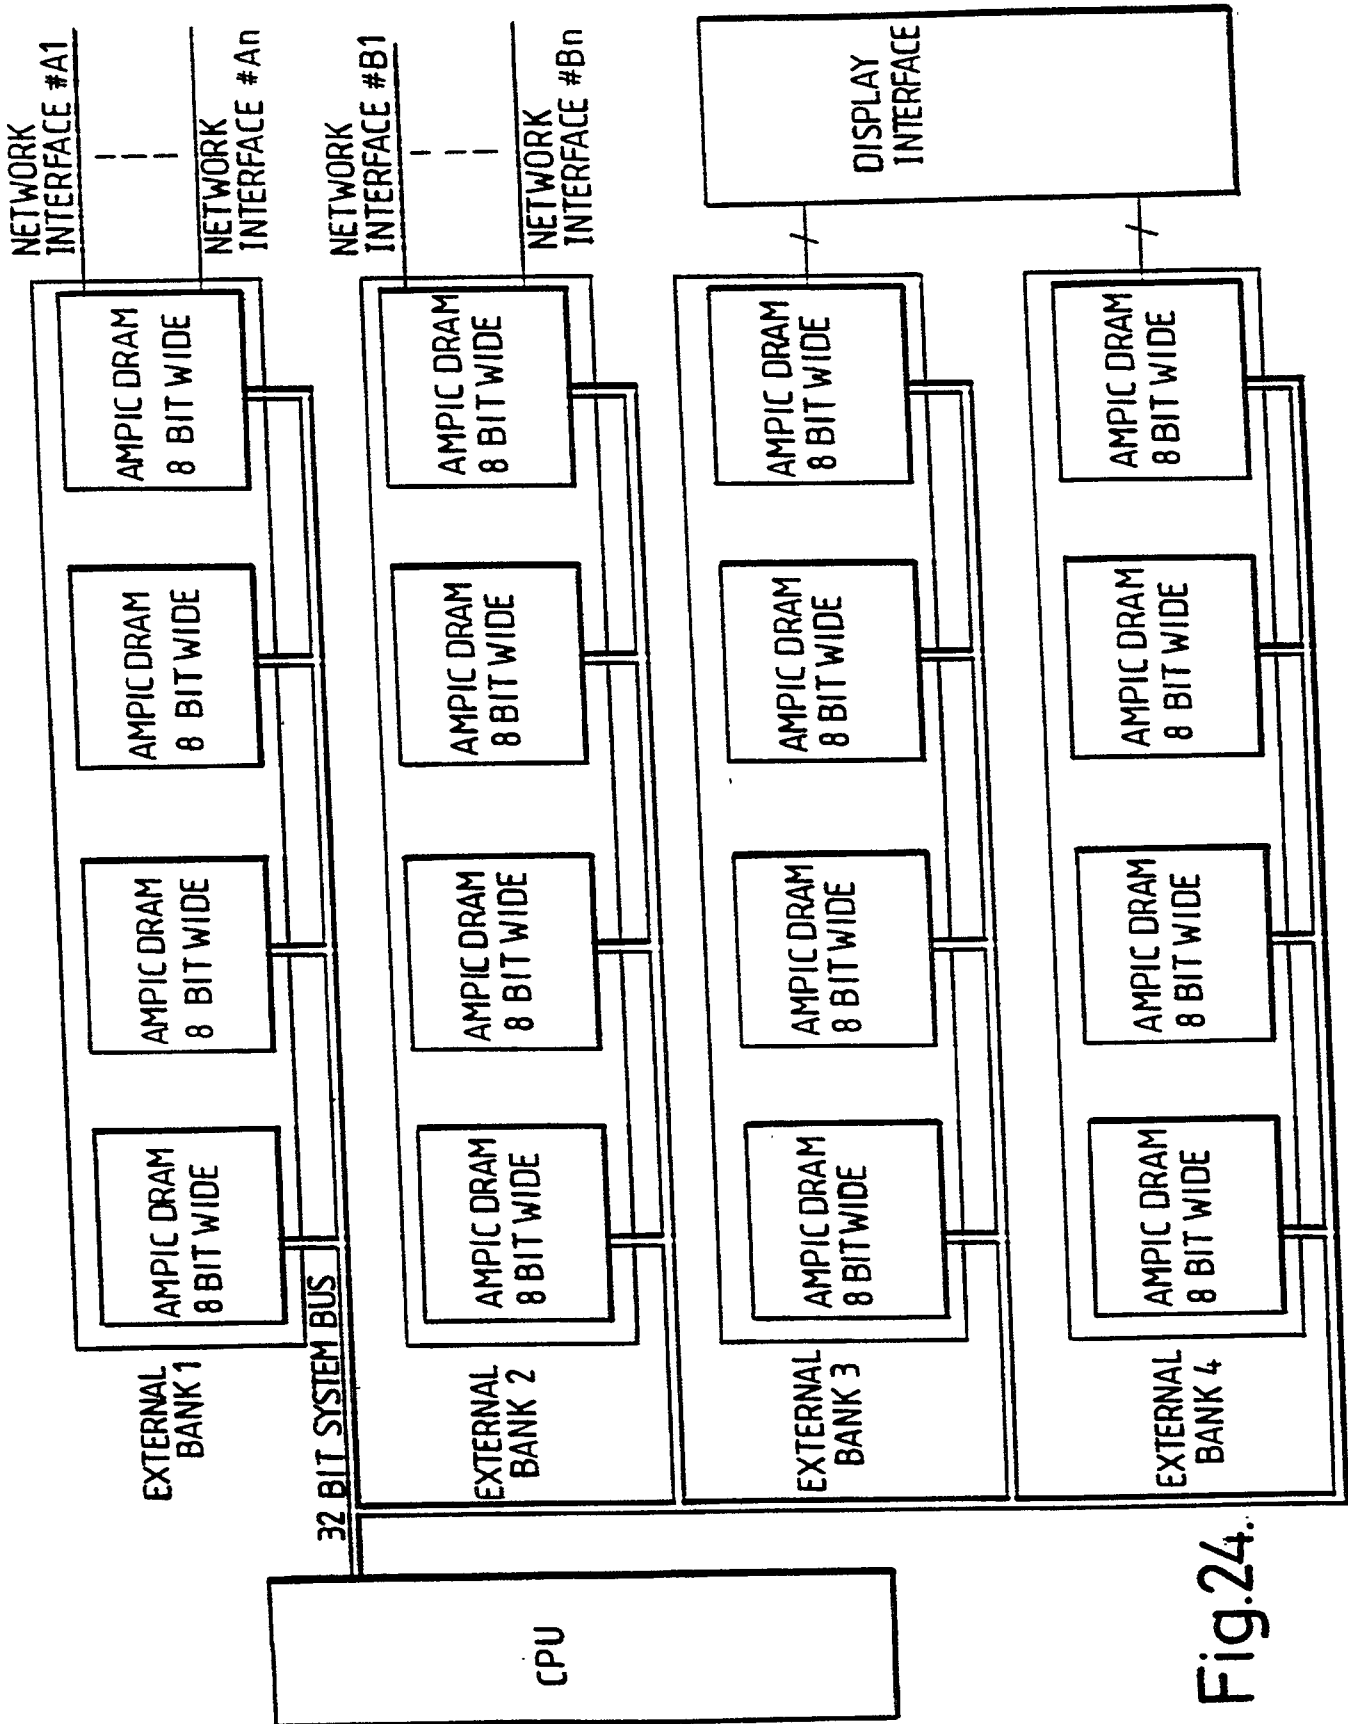


Fig.24.

25/27

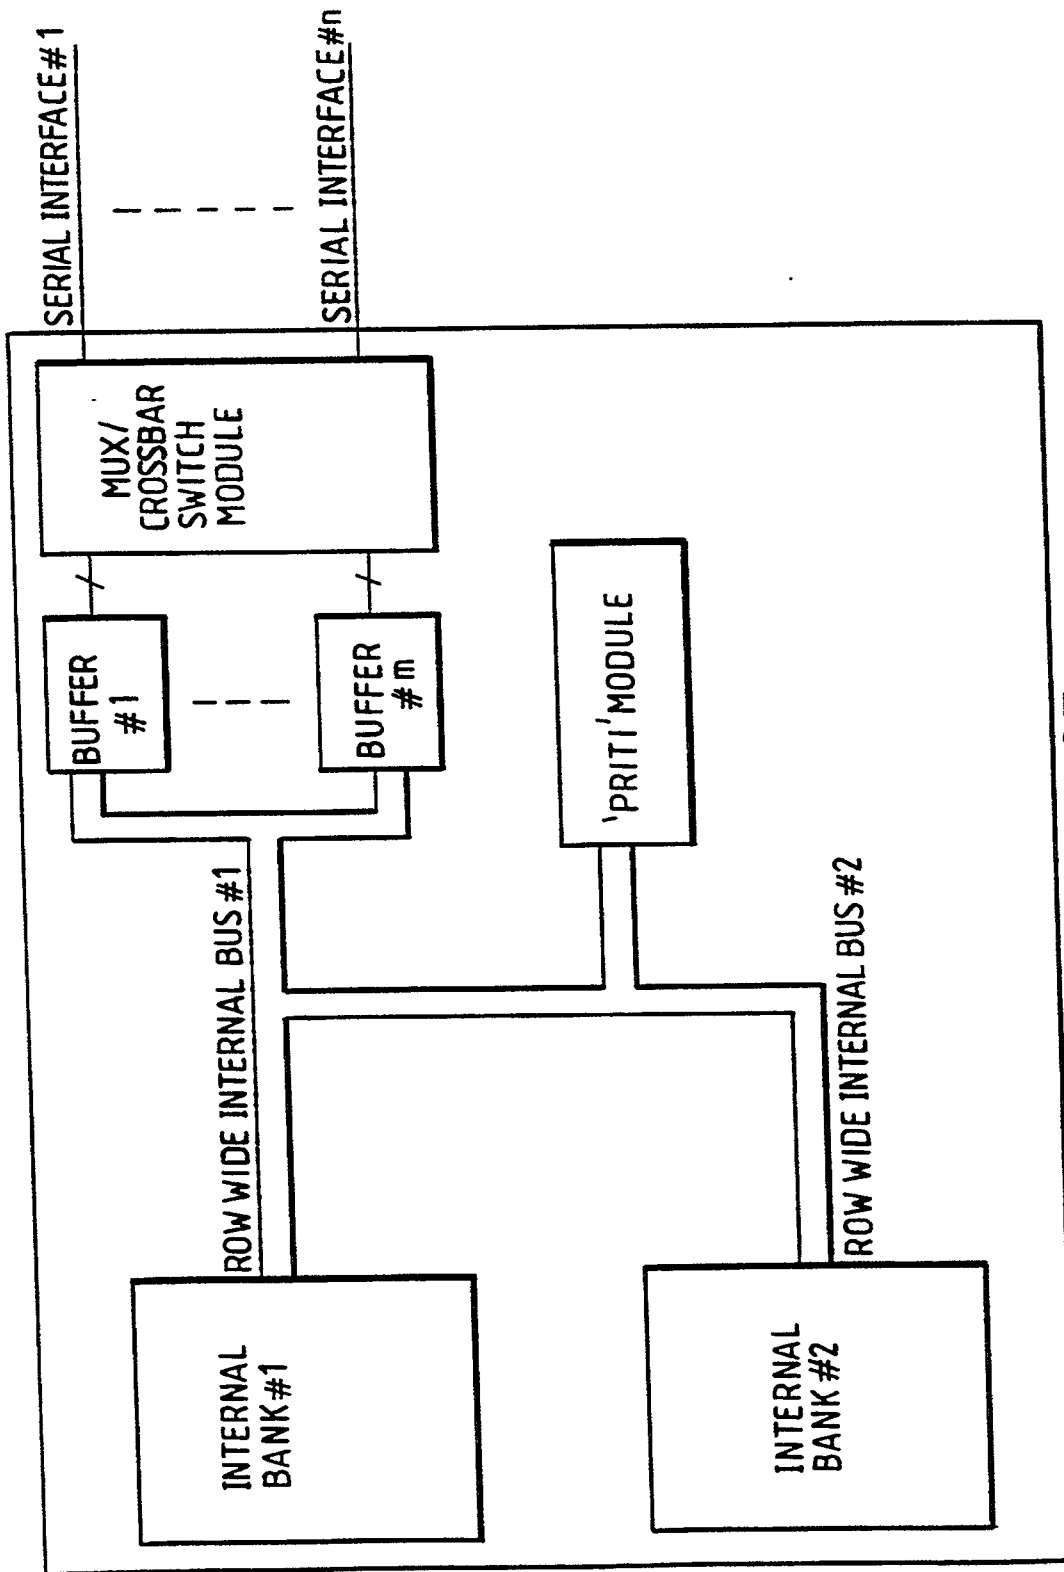


Fig.25.

26/27




Fig.26.




Fig.27.

