



US005712664A

United States Patent [19]

[11] Patent Number: **5,712,664**

Reddy

[45] Date of Patent: **Jan. 27, 1998**

- [54] **SHARED MEMORY GRAPHICS ACCELERATOR SYSTEM**
- [75] Inventor: **Chitranjan N. Reddy**, Milpitas, Calif.
- [73] Assignee: **Alliance Semiconductor Corporation**, San Jose, Calif.
- [21] Appl. No.: **136,553**
- [22] Filed: **Oct. 14, 1993**
- [51] Int. Cl.⁶ **G09G 5/00**
- [52] U.S. Cl. **345/200; 345/201; 395/508; 395/519**
- [58] **Field of Search** 345/133, 132, 345/189-191, 204, 203, 200, 201; 395/162, 164, 166, 185, 375, 129, 725, 141, 508, 519; 348/584, 441, 586; 365/200

- 5,392,393 2/1995 Deering 395/162
- 5,396,586 3/1995 Van Aken 395/141
- 5,537,128 7/1996 Keene et al. 345/200

FOREIGN PATENT DOCUMENTS

- 4-211293 8/1992 Japan .

OTHER PUBLICATIONS

- TMS34010 User's Guide, Texas Instruments, pp. 1-5 through 1-7, 1986.
- IBM Technical Disclosure Bulletin, 35(1A):45-46, Jun. 1992.
- Tandy. "TRS-80 Color Computer Technical Reference Manual" pp. 17-21, 1981.

Primary Examiner—Steven Saras
Attorney, Agent, or Firm—Limbach & Limbach L.L.P.

[57] ABSTRACT

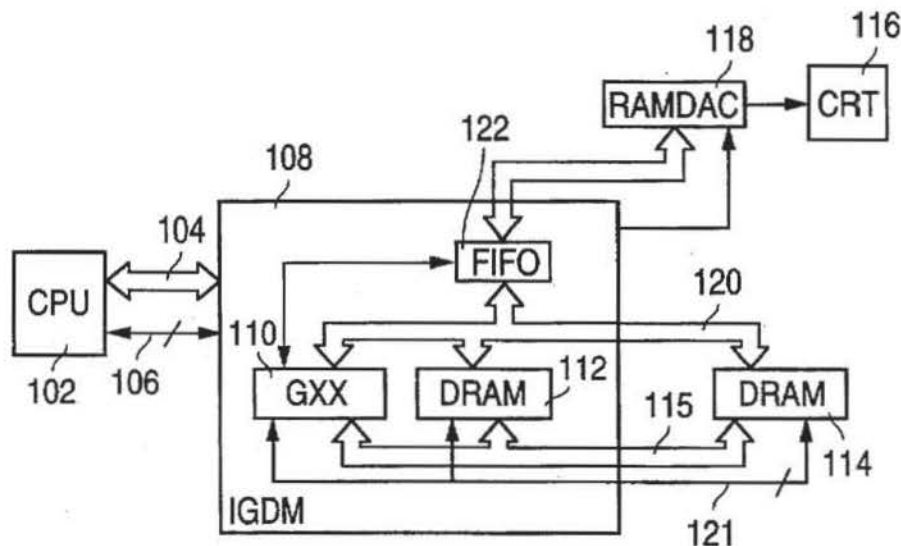
A shared memory graphics accelerator system that provides graphics display data to a display includes a central processing unit for generating graphics display data and graphics commands for processing the display data. An integrated graphics display memory element includes both a graphics accelerator connected to receive display data and graphics commands from the central processing unit and an on-chip frame buffer memory element. The on-chip frame buffer memory element is connected to receive display data from the graphics accelerator via a display data distribution bus. An off-chip frame buffer memory element is also connected to the display data distribution bus to receive display data from the graphics accelerator. The graphics accelerator selectively distributes display data to the on-chip frame buffer memory element and to the off-chip frame buffer memory element based on predetermined display data distribution criteria.

3 Claims, 2 Drawing Sheets

[56] References Cited

U.S. PATENT DOCUMENTS

4,191,956	3/1980	Groothuis	345/201
4,228,528	10/1980	Cenker	365/200
4,812,836	3/1989	Kurakake et al.	345/201
4,816,815	3/1989	Yoshiba	345/201
4,951,232	8/1990	Hannah	345/189
4,956,708	9/1990	Itagaki	348/441
5,031,092	7/1991	Edwards et al.	395/375
5,083,294	1/1992	Okajima	365/200
5,202,962	4/1993	Matsuo et al.	345/133
5,258,843	11/1993	Truong	348/586
5,293,540	3/1994	Trani et al.	348/584
5,297,148	3/1994	Harari et al.	365/200
5,303,334	4/1994	Synder et al.	395/129
5,319,388	6/1994	Mattison et al.	345/200
5,321,806	6/1994	Meinerth et al.	395/162
5,363,500	11/1994	Takeda	345/201
5,386,573	1/1995	Okamoto	395/725



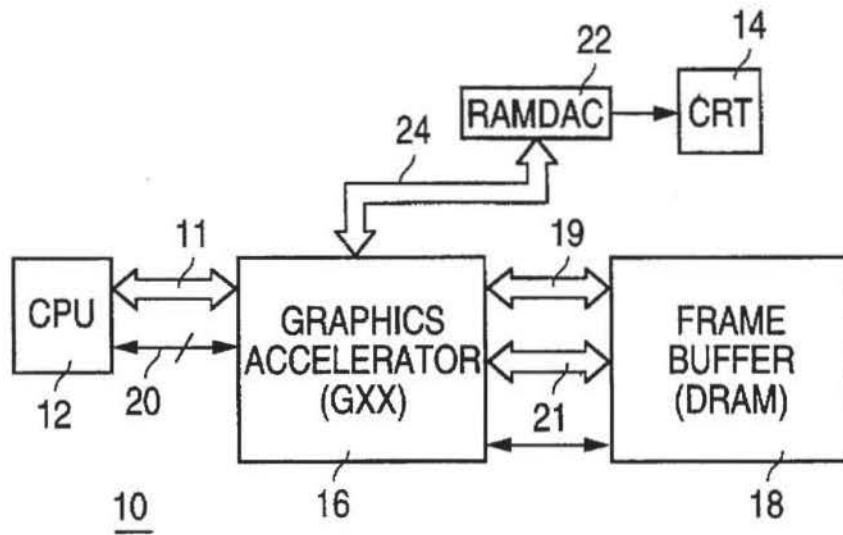


FIG. 1
(PRIOR ART)

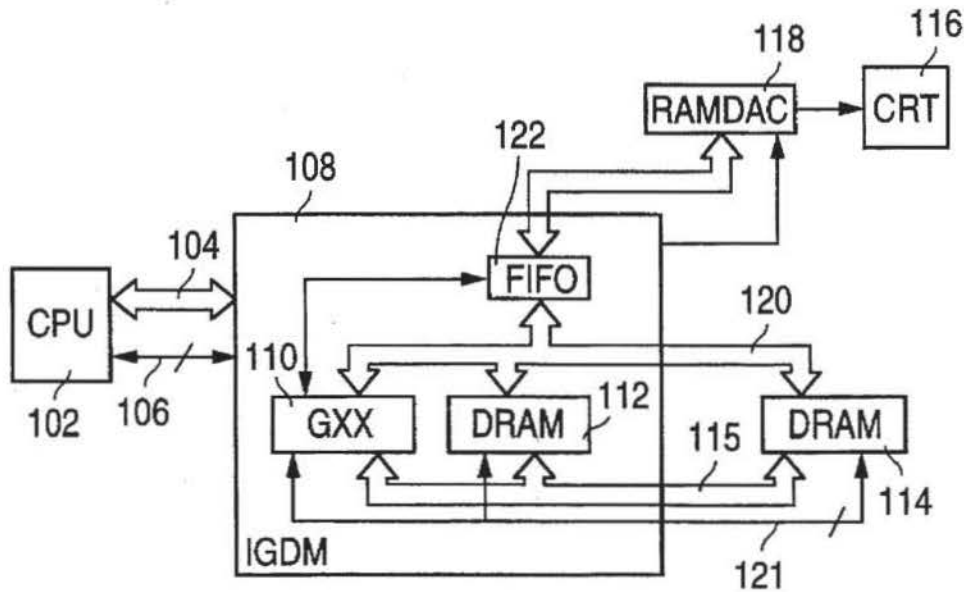


FIG. 2

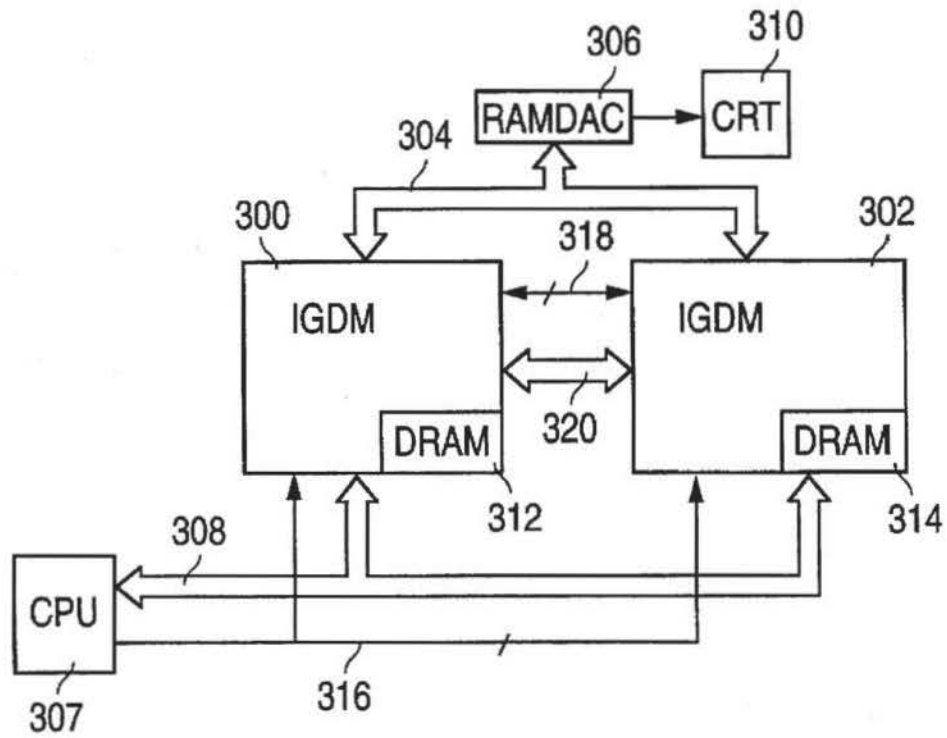


FIG. 3

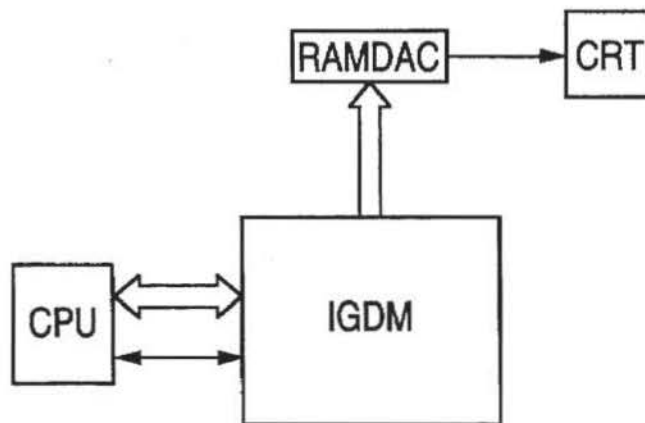


FIG. 4

SHARED MEMORY GRAPHICS ACCELERATOR SYSTEM

BACKGROUND OF THE INVENTION

1. Field of the Invention

The present invention relates to the visual display of a computer graphics image and, in particular, to a graphics display system that integrates both a graphics accelerator engine and a portion of the graphics frame buffer memory on the same monolithic chip.

2. Discussion of the Prior Art

A video graphics system typically uses either VRAM or DRAM frame buffers to store the pixel display data utilized in displaying a graphics or video image on a display element such as a CRT.

A VRAM frame buffer includes two ports that are available for the pixel data to flow from the memory to the display. One port is known as the serial port and is totally dedicated to refreshing the display screen image. The other port is a random access port that is used for receiving pixel updates generated by a CPU or a graphics accelerator engine. A typical VRAM arrangement allocates 99% of the available bandwidth to the random port thereby allowing the system to display fast moving objects and to support large display CRTs.

However, in a DRAM-based video system, the pixel data updates and the screen refresh data contend for a single frame buffer memory port. This contention reduces the amount of bandwidth available for pixel data updates by the CPU and the graphics engine, resulting in a lower performance graphics display system.

However, in most applications, the DRAM solution is preferable to the VRAM solution at the expense of lower performance, because DRAMs are cheaper than VRAMs.

FIG. 1 shows a conventional graphics display system wherein a CPU 12 writes pixel display data on data bus 11 to be displayed on the CRT screen 14 through a graphics accelerator (GXX) 16 onto a DRAM frame buffer 18 via data bus 19. The CPU 12 also provides certain higher level graphics command signals 20 to the graphics accelerator 16 to manipulate the display data stored in the DRAM frame buffer 18.

The graphics accelerator 16 retrieves display data from the frame buffer 18 via data bus 19 utilizing reference address bus 21, processes the retrieved display data based on the CPU command signals 20 and writes the new pixel data back to the frame buffer 18.

The pixel data is displayed on the CRT 14 through a random access memory digital-to-analog converter (RAMDAC) 22 that receives the data via a data display bus 24.

The graphics accelerator 16 also reads display data from the frame buffer 18 via data bus 19 and sends it to the RAMDAC 22 via the data display bus 24 to meet the periodic refresh requirements of the CRT display 14.

Thus, as illustrated in FIG. 1, the bandwidth of the data bus 19 is shared by three functions: display refresh, CPU display data update, and graphics accelerator display manipulation. As the display size (i.e., the number of pixels to be displayed on the CRT screen 14) increases, the display updates and display manipulation functions are reduced because of the bandwidth limitations of the data bus 19 caused by the fixed refresh requirements of the CRT 14.

While these limitations can be addressed by increasing the data bus width or by increasing its speed, both of these

solutions have either physical or practical limitations. Increasing the bus width increases the silicon area and the package pin count. Increasing the speed of the bus requires utilization of more complex silicon process technology.

SUMMARY OF THE INVENTION

The present invention provides a graphics display system that enhances performance by integrating a portion of the frame buffer storage space and the graphics accelerator engine on the same chip while at the same time maintaining the flexibility to expand the frame buffer size as needed.

Generally, the present invention provides a shared memory graphics accelerator system that provides display data to a display element. The shared memory graphics accelerator system includes a central processing unit that generates both display data and graphics commands for processing the display data. An integrated graphics display memory element includes both a graphics accelerator that receives display data and graphics commands from the central processing unit and an on-chip frame buffer memory element that is connected to receive display data from the graphics accelerator via a display data distribution bus. An off-chip frame buffer memory element is also connected to the data distribution bus to receive display data from the graphics accelerator. The graphics accelerator selectively distributes the display data to the on-chip memory element and to the off-chip memory element based on predefined display data distribution criteria.

The above-described integrated solution increases the performance of the graphics display system because display data retrieval from the on-chip frame buffer is much faster than from an external frame buffer and the DRAM timing constraints are reduced, thus achieving improved system performance. This integrated solution also allows the display memory size to be expanded by adding external memory so that large displays can be accommodated on an as-needed basis. Also, the frame buffer space can be distributed among several integrated solutions, thereby increasing both the display bandwidth and the parallel processing capability between the CRT display and the CPU.

A better understanding of the features and advantages of the present invention will be obtained by reference to the following detailed description and accompanying drawings which set forth an illustrative embodiment in which the principals of the invention are utilized.

DESCRIPTION OF THE DRAWINGS

FIG. 1 is a schematic diagram illustrating a conventional graphics subsystem.

FIG. 2 is a schematic diagram illustrating a shared memory graphics accelerator system in accordance with the present invention.

FIG. 3 is a schematic diagram illustrating a shared memory graphics accelerator system in accordance with the present invention in a distributed display arrangement.

FIG. 4 is a schematic diagram illustrating a shared memory graphics accelerator system in accordance with the present invention but with no expansion memory.

DETAILED DESCRIPTION OF THE INVENTION

The present invention addresses the data bus bandwidth problem common to conventional DRAM-based graphics display systems by integrating a portion of the display data frame buffer memory space on the graphics accelerator chip

and thereby allowing simultaneous access to both on-chip DRAM frame buffer data and off-chip DRAM frame buffer data while maintaining the flexibility to increase the display data memory size externally to meet a variety of CRT display size requirements.

FIG. 2 shows a shared memory graphics accelerator system 100 that includes a central processing unit (CPU) 102 that sends pixel display data via address/data bus 104 and graphics command signals via a control bus 106 to a single integrated graphics display memory (IGDM) 108. Those skilled in the art will appreciate that the bus widths are CPU-dependent.

The integrated graphics display memory element 108 includes a graphics accelerator (GXX) 110 that receives the pixel display data and distributes it between an on-chip DRAM frame buffer 112 and an off-chip DRAM frame buffer 114 via a display data distribution bus 120, using a common address bus 115. The data distribution between on-chip memory 112 and off-chip memory 114 is based upon user defined criteria loaded onto the integrated graphics display memory element 108 during power-up. This information can be stored either in the CPU hard disk or in a boot-up EPROM. This distribution of the pixel display data is optimized for maximum CPU updates onto the on-chip display buffer DRAM 112 and the off-chip DRAM 114 and, at the same time, for supporting a maximum display size refresh on the CRT display 116.

By splitting the display frame buffer into an on-chip DRAM portion 112 and an off-chip DRAM portion 114, the graphics accelerator engine 110 can double the pixel read data to a RAMDAC 118 by simultaneously accessing on-chip and off-chip frame buffer display data and multiplexing it onto the distributed data bus 120 using control signals 121. A FIFO memory 122 provides a buffer between the RAMDAC 118 which requires continuous display data input and the distributed data bus 120, which is shared for display update, display manipulate and display refresh operations.

It is also possible for the graphics accelerator engine 110 to read on-chip DRAM 112 at a much faster rate that it can read off-chip DRAM 114, thereby making more CPU 102 update time available for on-chip DRAM 112. This increase in CPU update bandwidth can, for example, be translated into a faster moving image portion which can be stored onto the on-chip DRAM 112 and a slower moving portion which can be stored onto the off-chip DRAM 114. Those skilled in the art will appreciate that this distribution of the load can be implemented many different ways between the on-chip DRAM 112 and the off-chip DRAM 114 to meet the performance requirements of the total graphics display system.

Those skilled in the art will also appreciate that successful implementation of the integrated graphics display memory element 108 described above requires that the on-chip DRAM frame buffer 112 have substantially different characteristics than a monolithic DRAM used for data storage.

A typical monolithic DRAM requires a 200 nsec. refresh cycle every 15.6 μ sec., which is equivalent to a 1.28% refresh overhead. During this refresh time, no data may be read from the DRAM; the time is used primarily for refreshing the DRAM cell data. This refresh overhead time needs to be constant (or as small as possible) with increasing chip density. Unfortunately, chip power dissipation must be increased with increasing chip density in order to maintain constant overhead.

For the integrated graphics display memory element 108, the on-chip DRAM frame buffer memory 112 is imple-

mented with substantially increased refresh frequency (much less than 15.6 μ sec.) to reduce the on-chip power dissipation. For example, a 16 Mbit on-chip DRAM frame buffer memory 112 could have one 200 nsec. refresh cycle every 2 μ sec., which translates to a 10% refresh overhead. While this refresh overhead is a significant portion of the total available bandwidth, with improved on-chip DRAM access time resulting from integration of the DRAM 112 with the graphics accelerator 110, overall system performance is improved significantly. Those skilled in the art will appreciate that, as more of the system sub-blocks, such as the RAMDAC 118, are integrated with the graphics accelerator 110 and the on-chip DRAM frame buffer memory 112, the refresh overhead is optimized with respect to improved on-chip DRAM access time and increased on-chip power dissipation to provide improved total system performance. Furthermore, increased refresh frequency permits smaller memory storage cell capacitance which reduces total chip size.

Thus, the on-chip DRAM 112 has a substantially higher refresh frequency than the monolithic off-chip DRAM 114. The integrated graphics display memory element 108 includes means for supporting the multiple refresh frequency requirements of the on-chip DRAM 112 and the off-chip DRAM 114.

In some low power applications, average power dissipation can be reduced by increasing both the memory cell size and the refresh interval. Another way to reduce power is to increase the number of DRAM sense amplifiers, but this solution increases chip size.

Those skilled in the art will appreciate that the FIG. 2 configuration of system 100 can be implemented utilizing available integrated circuit technology.

FIG. 3 shows two integrated graphics display memory elements (IGDM) 300 and 302 connected in parallel between a display data output bus 304 and RAMDAC 306 and to CPU 307 via an address and data bus 308, without any external memory, to display a contiguous image on the CRT screen 310 using a frame buffer DRAM 312 on-chip to integrated graphics display element 300 and a frame buffer DRAM 314 on-chip to integrated graphics display element 302. Thus, the two integrated graphics display memory elements 300 and 302, provide the total frame buffer storage space for pixel display data to be displayed on the CRT screen 310. Each of integrated graphics display memory elements 300 and 302 can receive CPU instructions via the CPU control bus 316 and can display portions of the required image on the CRT screen 310. Also the two integrated graphics display memory elements 300 and 302 can communicate with each other via the control signal bus 318 and address/data path 320 to split the image or redistribute the load among themselves without CPU intervention, thereby increasing the total system performance.

One possible example of load sharing in the environment of the FIG. 3 system could arise when one integrated graphics display memory element works on even lines of the CRT display while the other integrated graphics display memory element is drawing odd lines on the CRT screen 310. Those skilled in the art will recognize that it is also possible to subdivide the CRT screen 310 even further into multiple small sections with each section being serviced by a corresponding integrated graphics display memory element; these integrated graphics display memory elements can be cascaded to display a contiguous image on the CRT screen 310.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.