

**BUS ARCHITECTURE FOR INTEGRATED DATA
AND VIDEO MEMORY**

RELATED APPLICATIONS

This application is related to U.S. Patent Application Serial No. 07/886,045 entitled "Multiple Bus Architecture", filed May 19, 1992.

5

BACKGROUND OF THE INVENTION

1. **Field of the Invention:**

10

The present invention relates to the field of computer systems and associated memory structures. More particularly, the present invention relates to a bus architecture designed specifically for integrating data and video random access memory (RAM) modules.

15

2. **Art Background:**

General purpose computer systems usually contain one or more buses for interconnecting the various active components operating within the computer system. A bus may be generally described as a group of signal lines over which information is transferred from any of several sources to any of several destinations. Physically, a bus is a collection of wires routed on the computer motherboard and terminating at sockets into which may be inserted a variety of active devices. The active devices are usually constructed in modular form and generate and receive the signals transmitted over the

25

collection of wires forming the bus. Well known bus systems include VME bus®, microchannel®, multibus®, MBus, and SBus.

Typically, devices communicating over a bus utilize a set of control signals referred to collectively as a bus protocol to coordinate transfer of signals transmitted over the bus. Many of such control signals are common to each of the above named bus systems and to bus systems generally. However, most bus protocols further include a number of specific control signals permitting manipulation of data transmitted over the bus in a manner unique to the particular bus employing such signal. For example, some buses may be particularly designed to accommodate a number of bus masters, and which bus operation must accommodate signals issued by the masters seeking to control any of several slaves operating on the bus. Alternatively, another bus may be more generally designed to accommodate any of several arbitrary devices operating on the bus, and does not favor any particular operating characteristics of the devices connected to the bus. Further, all buses are constrained to operate within the physical design limits of the bus, in terms of bandwidth, data rate, and clock frequency.

Also in the past, memory devices operating on a bus were typically grouped into modules according to the type of memory function to be served. Typically, one form of memory, for example main memory, would be coupled to the bus to which is connected the processor, or processors in a multiprocessing system processors. Depending on the specific organization of such a computer system, additional memory components forming memory for other purposes, such as non-volatile memory or memory dedicated to video frame buffers, may be grouped into modules coupled to buses other than the bus connecting main

memory to the processors. Further, additional memory may be later added to a computer system and be grouped onto yet a different bus.

In such a computer system having memory distributed over a number of
5 buses, access to the various memory modules can be hampered by the number
of data paths, controllers and interface chips through which signals must
travel from a processor seeking to retrieve or write data from or to a memory
device operating on a remote bus system. Moreover, having distributed
memory modules requires more area within the computer system in order to
10 accommodate the several bus data paths and memory groupings. Another
consequence of having distributed memory groupings is that the memory
modules inserted in such a memory group are specific to that memory group
and cannot be interchanged with memory modules from another grouping.
That is to say, a memory module from main memory generally cannot be
15 interchanged with a memory module for a video frame buffer operating on a
different bus, principally because the control signals to the video frame buffer
memory are different from the control signals directed to the main memory
module. The generic problem in all of these various distributed bus systems is
that the number of signal lines to accommodate the various types of control
20 signals and signal conditions can become large. This adds to the general
complexity, and therefore the cost and reliability of the overall computer
system.

SUMMARY OF THE INVENTION

A bus architecture for integrated data and video memory is disclosed. A high speed dedicated memory bus is coupled to a memory controller. The memory controller is in turn coupled to a multiple processor bus interconnecting one or more processors. Memory modules in the form of low profile single-end-line-memory-modules (SIMM) incorporating dynamic random access memory (DRAM), video RAM (VRAM), and static non-volatile RAM (SRAM) are coupled to the memory bus. Bus control signals and address and data lines from the memory controller are shared by all memory modules operating on the memory bus. Certain control signals invoke specific operations on memory modules, depending on the type of memory module receiving the control signal. Due to the high functionality of the VRAM SIMM memory modules, a plurality of check-bit-write bits provide a combination of write enable and access to I/O registers on the VRAM SIMM. A single busy line can be asserted by any memory module requiring the SIMM controller to back off and reaccess the memory module at a subsequent time.

The memory bus architecture is designed to match the memory bus bandwidth to the bandwidth of the multiprocessor bus and the data rate of the processors coupled thereto.

The memory modules incorporated into the integrated data and video memory are constructed around a consistent protocol of control signals pin out. The VRAM SIMMs further incorporate video buffering and conversion functions, thereby relieving the memory controller of service overhead associated with these functions. The VRAM SIMM further incorporates a

video output block in order to directly drive a video monitor. Multiple frame buffers may be accommodated by the integrated data and video memory bus.

Integrating all forms of memory into a single data and video memory
5 architecture permits simplified, consistent access to memory modules, and
therefore simplified updating of peripheral processes connected to the
computer system. The SIMM modules incorporate a high number density
connector, permitting high functionality to be incorporated into the SIMMs.
The SIMMs are organized to permit memory expansion in full data path-width
10 increments, and to accommodate higher capacity RAM chips when they
become available.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood more fully from the detailed
5 description given below and from the accompanying drawings of the preferred
embodiment of the invention in which:

Figure 1 is a block diagram of the integrated data and video memory bus.

10 Figure 2 is a block diagram of a memory controller for the integrated data
and video memory bus in Figure 1.

Figure 3 is a block diagram of a video random access memory (VRAM)
single in-line memory module (SIMM) coupled to the memory bus.

15

Figure 4 is a block diagram of a dynamic RAM (DRAM) SIMM coupled to
the memory bus.

Figure 5 is a block diagram of a nonvolatile static RAM (SRAM) SIMM
20 coupled to the memory bus.

Figure 6 is a timing diagram illustrating a VRAM page mode read,
according to the protocol of the present invention.

25 Figure 7 is a timing diagram illustrating a VRAM page mode write,
according to the protocol of the present invention.

Figure 8 is a timing diagram illustrating a CAS before RAS refresh cycle, according to the protocol of the present invention.

Figure 9 is a timing diagram illustrating loading of an address into
5 VRAM supporting autoincrement read/write, according to the protocol of the present invention.

Figure 10a is a timing diagram illustrating a direct address mode read cycle for a half word, according to the protocol of the present invention.
10

Figure 10b is a timing diagram illustrating a direct address mode read cycle for a half word, according to the protocol of the present invention.

Figure 11a is a timing diagram illustrating an indirect address mode read
15 cycle for one or more words, according to the protocol of the present invention.

Figure 11b is a timing diagram illustrating an indirect address mode write cycle for one or more words, according to the protocol of the present invention.
20

Figure 12 is a timing diagram illustrating a write cycle to I/O space for special I/O operations, according to the protocol of the present invention.

Figure 13 is a timing diagram illustrating an assertion of an internal
25 control signal within a video buffer to block RAS issued by the memory controller when VRAM is being refreshed/reloaded, according to the protocol of the present invention.

Figure 14 is a timing diagram illustrating an assertion of an internal control signal within the video buffer to terminate a memory access by the memory controller when VRAM requests service, according to the protocol of the present invention.

Figure 15 is a timing diagram illustrating an assertion of internal control signals within the video buffer to terminate an I/O access by the memory controller when VRAM requests service, but where the memory controller is permit to complete its access.

Figure 16 is a timing diagram illustrating an assertion of a data block control signal within the video buffer to block generation of internal control signals.

Figure 17 is a timing diagram illustrating buffering by the video buffer of CAS, RAS, and check-bit-write bits (CBWs) issued by the memory controller.

Figure 18 is a timing diagram illustrating the function of the video buffer during VRAM, refresh, according to the protocol of the present invention.

Figure 19 is a timing diagram illustrating the function of the video buffer during VRAM reload, according to the protocol of the present invention.

Figure 20a and 20b are exemplary timing diagrams illustrating respectively read and write accesses to an SRAM SIMM, according, to the protocol of the integrated data and memory bus.

Figure 21 illustrates the distribution of Column Access Strobe (CAS) and Row Access Strobe (RAS) control signals to all SIMM slots.

5 Figures 22a and 22b illustrate the high density connector.

Figure 23 is a table of the control signals used to implement the protocol for the integrated data and video memory bus.

10

Figure 24 illustrates the pin numbering used in the high density SIMM connector.

DETAILED DESCRIPTION OF THE INVENTION

A bus architecture for integrated data and video memory is disclosed. In the following description, for purposes of explanation, specific numbers, times, signals etc., are set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well known circuits and devices are shown in block diagram form in order not to obscure the present invention unnecessarily.

10

The preferred embodiment of the integrated data and video memory bus is designed to be used with complimentary metal-oxide-semiconductor (CMOS) technology. As such, certain power specifications and timing protocols are optimized to reduce power consumption and to provide sufficient stabilization time of control signals drive by the CMOS devices and to provide sufficient stabilization time of control signals driven by the CMOS devices. It will be apparent, however, to those skilled in the art that the specifications can or may be changed without departing from the scope of the present invention.

15

Further, throughout the following detailed description, certain bus lines will be described as being "active high" or "active low". This terminology will be used to indicate whether a particular bus control or data or address line in question is asserted or not, and therefore relevant or not to the bus operation. It will be apparent however, that the assertion of an active high or active low control signal is a matter of design choice. Although the preferred embodiment has been designed to maximize data transfer rate and accuracy of stability of control signals, it will be appreciated that changing the design of the bus to

20

25

accommodate other alternatives to active high and active low is within the scope of the present invention.

Similarly, the preferred embodiment of the memory bus operates in a completely synchronous fashion, wherein all bus signals change on the rising edge of clock pulses. The synchronous design is desirable to achieve high data transfer rates, and synchronous bus designs enable simplified clock pulse edge detection. However, it will be appreciated that the use of the rising edge of the clock pulses is a matter of design choice, and that the present invention can be made synchronous upon the falling edge of clock pulses.

A. MEMORY SYSTEM - Memory Bus

Reference is now made to Figure 1, wherein a block diagram overview of a multiple bus architecture in a multiple processor high performance computer system incorporating the integrated memory bus of the present invention are shown. The presently preferred embodiment may be found in certain of the computer systems manufactured by Sun Microsystems, Mountain View, California.

In Figure 1, a multiplicity of time-multiplexed bus signal lines 10 form the memory bus. Signal lines 10 are routed from a memory controller 20 and coupled to a multiplicity of single in-line memory modules (SIMMs) 30. Signal lines 10 consist of address lines ADDR[11:0], data lines DATA[127:0], and a numerosity of control lines 15 which together permit memory controller to access SIMMs 30 operating on the bus according to the integrated memory

protocol described herein. The function and effect of each of control signals 15 will be described in detail below. In order to flexibly accommodate different needs of different users in configuring computer systems, SIMMs 30 may consist of any of a variety of memory forms, including dynamic random access
5 memory (DRAM), video RAM (VRAM), and nonvolatile static RAM (SRAM).

Throughout the following discussion, it is contemplated that any combination of the above types of memory will be coupled to the integrated memory bus. As will be described and demonstrated in terms of specific
10 examples, the memory bus of the present invention incorporates uniform control signals and signal pin outs so that, with certain exceptions, SIMMs 30 may be arbitrarily placed in the SIMM slots.

Memory controller 20 is coupled to one or more processors 42 via a
15 multiple processor bus 40. In the presently preferred embodiment, multiple processor bus 40 is an Mbus, as may be found in computer systems manufactured by Sun Microsystems, Mountain View, California, and as described in the Mbus Specification, Revision 1.1. A bus interface 45 couples bus 40 to a second, general purpose bus 50. In the presently preferred
20 embodiment, general purpose bus 50 is an Sbus, as also may be found in computer systems manufactured by Sun Microsystems and as described in the Sbus Specification, Revision 1.1. However, other processor as well as general purpose buses are available, and are within the scope of the present invention. A system clock 47 distributes clock signals to all devices coupled to the memory
25 bus, as well as the other buses 40 and 50. A variety of I/O and peripheral devices 52 may be intercoupled to bus 50, thus forming a complete computer system.

B. MEMORY SYSTEM - Memory Components

5 Reference is now made to Figure 2, which illustrates a block diagram of the memory controller 20 used to control access to memory devices operating on the integrated data and video memory bus shown in Figure 1. In Figure 2, memory controller 20 receives address and data signals transmitted over multiple processor bus 40 through an MBus slave interface 20a. MBus
10 slave interface 20a incorporates facilities to operate as an MBus slave only, and implements the full Level 2 MBus protocol. After being received by MBus slave interface 20a, data and address signals are subsequently transmitted to several internal logic blocks, and subsequently retransmitted to signal lines 10. It is further intended that memory controller 20 should include sufficient
15 functionality to permit the memory bus to operate according the protocol described herein. The memory controller 20 of the presently preferred embodiment includes check-bit generation and error detection block 20b, data bus, check bits, and address block 20c, bus arbiter and write buffer 20c, generation of control signals and refresh block 20e, status and control registers
20 20f, and video I/O access block 20g.

Referring now to Figure 3, a block diagram of a video random access memory (VRAM) SIMM (VSIMM) 32 coupled to the memory bus signal lines 10 is illustrated. In Figure 3, VSIMM 32 receives address signals (ADDR[11:0]),
25 data signals (DATA[127:0]), together with other shared control signals issued by memory controller 20, through an on-board video buffer 35. Video buffer 35 is in turn coupled to a multiplicity of VRAM chips 34, and to a display interface

36 having an output coupled to a display device (not shown). Multiple VSIMMs 32 may be coupled to memory bus 10, subject to address space constraints discussed in more detail below. Of particular interest in Figure 3 are the check-bit-write signal lines CBW[7:0] and CBW[15:8], which are commonly distributed across memory bus 10, but which have particular application to VSIMMs 32. In particular, CBW[8] used as an active low byte write enable for VRAMs 34 within VSIMM 32, whereas CBW[7:0] permit loading of a base address into an autoincrement register of a VRAM 34. Use of CBW[7:0] and CBW[15:8] will be illustrated in more detail in connection with Figures 9 - 12. As shown in Figure 3, the VSIMM 32 contains a video buffer 35 that handles video refresh, and a memory display interface 36 that muxes data from the VRAM into an RGB compatible form. Video buffer 35 can issue a VBSY- signal to command memory controller 20 to release control of the VSIMM 32 whenever VRAMs 34 require refresh or reload cycles. The display interface 36 is contemplated to include a digital-analog-converter on the VSIMM 32, as well as pixel clock generation circuitry.

Similarly, and with reference to Figure 4, a block diagram of a dynamic random access memory (DRAM) SIMM (DSIMM) 72 coupled to the memory bus signal lines 10 is illustrated. In Figure 4, DSIMM 72 also receives address signals (ADDR[11:0]), data signals (DATA[127:0]), together with other shared control signals issued by memory controller 20, through an on-board buffer 75. Buffer 75 is in turn coupled to a multiplicity of DRAM chips 74. As in the case of VSIMMs 32 above, multiple DSIMMs 72 may be coupled to memory bus 10, subject to address space constraints discussed in more detail below. In Figure 3, the check-bit-write signal lines CBW[7:0] and CBW[15:8], which are commonly distributed across memory bus 10 are used as error correction code

bits (ECC) with application to DSIMMs 72 in accordance with generally known standards for ECC. VBSY- signal is not used by DRAMs 74, and is not issued by DSIMMs 72.

5 Referring now to Figure 5, a block diagram of a nonvolatile static RAM (SRAM) SIMM (NVSIMM) 82 consistent with the protocol of the present invention is coupled to the memory bus signal lines 10. In Figure 5, NVSIMM 82 receives address signals (ADDR[11:0]), data signals (DATA[127:0]), together with other shared control signals 15 issued by memory controller 20, through
10 an on-board buffer/latch 86. Buffer/latch 86 is in turn coupled to a multiplicity of SRAM chips 84. As in the case of VSIMMs 32 and DSIMMs 72 above, multiple NVSIMMs 82 may be coupled to memory bus 10, subject to address space constraints discussed in more detail below. In the presently preferred embodiment shown in Figure 5, nonvolatile NVSIMMs 82, are capable of
15 emulating DRAM SIMM operation on the integrated memory bus 10. Therefore, the memory controller 20 communicates with NVSIMM 82 just like another DRAM SIMM 72 installed in as SIMM 30 in Figure 1. To achieve this result, SRAMs 84 have been installed in place of DRAMs 74 and control signal logic 87 is provided to adjust the timing between NVSIMM 82 to the memory
20 controller 20. Further, a separate address decoder 83 may be necessary to provide proper addressing to SRAMs 84 using the shared multiplexed address lines ADDR[11:0]. Also, the NVSIMM 82 may be supplied with a battery and backup circuit 88 to protect the data stored in NVSIMM 82, in case of power loss.

25

Referring briefly to Figures 21-24, all address, data, and control signals used to implement the memory bus protocol are illustrated. In Figure 21,

particular SIMMs 30 are selected on the basis of two Column Access Strobe (CAS) signals and eight Row Access Strobe (RAS) control signals to all SIMM slots as illustrated. However, the present bus architecture can accommodate up to eight CAS- signals to provide additional selectivity of memory components used within SIMMs 30. In Figures 22a and 22b, the high density connector used on each SIMM is illustrated. In contrast to prior art single readout connector designs, the connector used by all SIMMs 30 embraces a double readout design wherein signals are independently routed on each side of the SIMM board. In Figure 23, a table of the control signals used to implement the protocol for the integrated data and video memory bus. Although all control signals are routed to all SIMM slots location, the video control signals pertain only to VSIMMs 32 and are routed only between video buffer 35 and display interface 36. Finally, Figure 24 illustrates the pin numbering used in the high density SIMM connector illustrated in Figures 22a and 22b, and which is shared by all SIMMs 30 operating on the memory bus 10.

Thus, as presently preferred, and as illustrated in Figures 1-5, the memory bus is 128 bits wide for data, and further has 16 error correcting code (ECC) checkbits, yielding a total bus datapath width of 144 bits. In particular, the present embodiment of the memory system uses 80 nanosecond DRAMs 74 and VRAMs 34, thereby providing approximately a 20 MHz data rate for the memory bus. Memory controller 20 is an application specific integrated circuit (ASIC) that in the preferred embodiment interfaces between the 144 bit datapath memory bus (10) and the 64 bit MBus (40). Memory controller 20 services data requests transmitted over bus 40 (MBus) by MBus masters (42) from main memory consisting of one or more DSIMMs 72 (Figure 4). As presently preferred, memory controller 20 operates at the 40 MHz MBus data

rate, and can receive data from any of the SIMMs 30 forming the memory system after RAS- (row access strobe) assertion within 4 clock cycles, or 100 nanoseconds. Actual read access times are 7 clocks, and line read time is at least 10 clocks. Cycle time is at least 11 clocks. Memory controller 20 can generate
5 128 byte burst transfers with sustained 32 Byte transfers at a rate of 87 megabytes per second. The memory controller-to-MBus read latency from MAS- (MBus address strobe) to first data is 7 MBus clocks. Issuance of RAS- in combination with other control signals to form the operational protocol of the present invention is described in more detail below in connection with Figures 6
10 through 20b.

The maximum amount of main memory is 512 MB, and is determined by there being twenty-nine address bits available in the present implementation of memory controller 20 for addressing memory in a
15 multiplexed column/row manner. The other three bits are used to differentiate I/O spaces in memory space, e.g., between VRAM and DRAM address spaces. Each SIMM 30, and accordingly each SIMM position, has 64 MB of address space allocated to it, i.e., position 0 starts at address 0, position 1 starts at address 0x4000000, and position 2 starts at address 0x8000000. Thus, regardless
20 whether there is a 16 MB or a 64 MB SIMM 30 in position 0, the SIMM 30 in position 1 responds to addresses starting at 64 MBytes from 0 or 0x4000000. A boot programmable read-only-memory (PROM) (not shown) determines which SIMM 30 is installed in which position by observing the response from each position while writing and then subsequently reading from appropriate
25 addresses for VSIMMs and DSIMMs.

Because total addressable memory space for the memory bus and associated memory controller is very large (> 2 GB), it is possible to have two memory controllers in the present architecture. As presently configured, the first memory controller 20 responds to addresses 0x00000000 to 0x0F000000, whereas the second would respond to addresses 0x10000000 to 0x1F000000. Two memory controllers would allow for two banks of SIMM sockets, therefore increasing a system's memory capacity from eight SIMMs 30 to sixteen. However, in the present implementation of the memory bus architecture, there is only one memory controller 20 on the motherboard, because there physically is only one bundled datapath, namely signal lines 10, from the memory bus 10 to the connector for memory controller 20. It is possible to have another memory controller coupled on the MBus (40), but in that case the second memory controller would have to have its own local memory on its MBus card.

15

As presently preferred, memory system has eight positions to receive up to eight SIMMs 30, as shown in Figure 1. SIMM positions 0,2,4, and 6 can contain only DSIMMs 72. The DSIMMs 72 are double sided and may contain 16 megabytes (MB) of DRAM memory using thirty-six 1megabit-by-4 (1Mbx4) TSOP DRAMs 74. Alternatively, DSIMMs 72 may comprise 64 MB when using thirty-six 4Mbx4 TSOP DRAMs 74. SIMM positions 1,3,5,7 can contain either VSIMMs 32 or DSIMMs 72. The VSIMMs 32 currently are available in two sizes, 4 MB consisting of thirty-two 256Kx4 VRAMs 34, and 8 MB consisting of thirty-two 256Kx8 VRAMs 34. VSIMMs 32 can also be built with thirty-two 128Kx8 VRAMs, thirty-two 256Kx8 VRAMs, thirty-two 256Kx16 VRAMs, and thirty-two 512Kx8 VRAMs.

25

The presently configured VSIMMs 32 require an extended length video socket in addition to the normal SIMM socket, to route the additional functionality associated with the video output. The present embodiment currently has only two such extended video sockets, in SIMM positions 1 and 3. Thus, although the memory bus architecture supports multiple frame buffers (VSIMMs 32) residing on the memory bus, physical space limitations within the present implementation limits the number of VSIMMs 32 as frame buffers to only two, in the absence of extra flex-cabling to route the video signals to the back panel, from positions 5 and 7. However, it is possible to have more than two frame buffers in the present implementation, where the additional VSIMMs 32 reside on the MBus 40, or on the SBus 50. Thus, the memory bus architecture can presently support a maximum of 512 MB of DRAM when using 64 MB DSIMMs 72. Specific DSIMMs 72 supported by the memory controller 20 are enumerated below. Similarly, specific VSIMMs 32 supported by the memory controller 20 are also enumerated below.

Because the memory controller 20 in the preferred embodiment of the present invention incorporates non-continuous memory addressing, software can use a mirroring scheme to identify the memory configuration at boot time. Unidentified address space accesses will cause an MBus time-out in accordance with the MBus protocol. As stated above, the memory controller 20 supports a variety of SIMMs 30, including DRAM, VRAM, and SRAM modules, and memory controller 20 can directly support up to four VSIMMs 32 in four slots. Given the many SIMM slot/board combinations (forty-eight in the presently preferred embodiment), identification of any SIMM 30 installed in any position is accomplished through software. Configurations for slot-specific SIMMs 30,

for example VSIMM 32, can be set prior to standard SIMM accesses. As presently preferred, the SIMM position probe determined configurations can also be used for system memory mapping.

- 5 The MBus slave interface 20a shown in Figure 2 which implements the full level-2 MBus slave protocol as described in the SPARC MBus Interface Specification, revision 1.1, is subject to the following clarifications. All eight MBus transaction sizes are supported. All MBus defined transaction types are supported. The MBus-acknowledges issued by memory controller 20 are
- 10 presently limited to those without MRTY- (MBus retry), because retries are not generated or serviced by the memory controller 20. However, other implementations of memory controllers may include MRTY-. During coherent transfers, reflective updates are not supported. CI operations can be disabled to allow other devices to own the MBus CI acknowledge. An asserted
- 15 MIH, in the MIHDEL window, aborts the coherent MBus transfer acknowledge and relinquishes bus ownership.

C. MEMORY SYSTEM - Memory Controller Transaction Semantics

20

1. DRAM Control

- The operations cycle timing relationships are shown in the timing diagrams illustrated in Figures 6-20. CAS- is used to select a double word from
- 25 the quad word data bus. Each SIMM 30 is selected by one of eight RAS- signals (one RAS- per 128 bits). The twelve bit address can support up to 16 Mb asymmetric DRAM components in column/row address form. Partial double

word accesses are accomplished with an internal read modify write cycle.
Coherent transfers extend their response window by the programmed MIHDEL.
Wrapped read burst transfers supply the first double word on the transfer's
eight byte aligned address. Stores use the DRAM early write cycle transfer.

5

2. VRAM Control

No error correction code (ECC) check bits or parity are used for VRAM
10 transfers. Therefore the byte write enables (WE-) are multiplexed on the
check-bit write (CBW) signal lines to save pins. The parallel port of
each VRAM 34 is accessed with the same timing constraints as DRAMs 74.
As shown in Figure 8, the CBWs [15:0] are deasserted during refresh cycles.
Each VRAM 34 may be accessed via a serial port over a 128 bit data path
for both two and four MB frame buffers by asserting an SIRQ- signal, or over
15 a 256 bit data path for the eight and sixteen MB frame buffers. The
256 bit data path width would be used to support double buffered true
color for high resolution monitors. The current implementation requires
a 619 ohm pull-up on MRDY- for VRAM read operations.

20

3. VSIMM Preemptive Protocol - Overview

As previously shown in Figure 3, VRAM 34 is located on a VSIMM 32,
and is serviced by the video buffer 35 on VSIMM 32. In efficient modular
25 fashion, the video buffer 35 is responsible for performing VRAM refresh and
reload cycles. Memory controller 20 VSIMM requests are arbitrated with the
video buffer 35 using the preemptive protocol consistent among all SIMMs 30.

The requesting video buffer 35 will synchronously assert VBSY- no sooner than the second rising edge after the memory controller 20 RAS- is asserted. The memory controller 20 then non-destructively preempts any in-progress memory controller 20 transfer cycle as the video buffer 35 transitions out of its
5 idle window. The memory controller 20 relinquishes VSIMM ownership to the video buffer 35 by deasserting RAS- for two consecutive clocks. Once the video buffer 35 acquires VSIMM ownership, it blocks all VRAM controls shared with memory controller 20 and logically overdrives them (adhering to the VRAM access protocol). The video buffer 35 also then deasserts the tri-state
10 enable for the separate memory controller 20 CBW' buffer.

The maximum period from the end of the idle window to a completed service is called the ACTIVE window. The VSIMM 32 synchronously responds with VBSY- to an initial memory controller 20 RAS-slot request while the
15 video buffer 35 is actively servicing. VBSY- remains asserted (active low) until the VSIMM servicing is completed and the video buffer 35 is no longer in an ACTIVE window or blocking shared signals, thus ensuring a memory controller 20 transfer does not get unknowingly blocked. After retaining VSIMM ownership, the memory controller 20 restarts the interrupted transfer.
20 The present protocol permits memory controller 20 to monitor VBSY- as it transitions out of its interrupted transfer state and reasserts RAS- to determine whether, during a memory controller 20 refresh to a different SIMM position, the video buffer 35 assumed VSIMM ownership. Periodic VSIMM refresh and shift reload service, by the video buffer 35, is scheduled to occur after the
25 minimum programmable delay of twelve pending cycles. The active window for video buffer 35 can be extended from a minimum of twelve clocks to accommodate the system configuration (e.g. the active window size should be

programmed larger than the minimum coherent delay). Further details of this protocol should be referenced from the VSIMM protocol spec.

5 4. Protocol/Interface signals

The control signals used by the ECC memory controller 20 and all SIMMs 30 to implement the memory bus protocol are:

- D[127:0] Bi-directional multiplexed data bus signal lines.
- 10 CBW[15:0] Byte write enables. During write accesses to VRAMs 34 these lines specify the bytes to be written and are driven high for VRAM reads. For register I/O accesses CBW[7:0] are used as a bi-directional data bus and CBW[8] is used as the read/write- signal. All register cycles are byte transfers.
- 15 IOSEL- VRAM 34 or register access select. When asserted this signal selects the register I/O space. Otherwise the access is to VRAM 34. This signal is used to qualify RAS- (see Figures 9-16).
- 20 VBSY- Video busy used to implement the preemptive protocol. When asserted in response to a RAS-, the video buffer 35 is either busy with a reload or refresh cycle, or seeks to perform such a cycle. (see Figures 13-20b).
- RAS- Row address strobe for accesses to VRAMs 34, when IOSEL- is negated (high). With IOSEL- asserted low, RAS- is used as a chip select for register I/O transfers (see Figures 13-20b).

25	IOSEL-	RAS-	Function
	0	0	Register I/O chip-select
	1	0	VRAM access, RAS- asserted

	0	1	VSIMM busy, no RAS- or chip-select.
	1	1	VSIMM idle
	CASIN0-	VRAM Column address strobe long word 0.	
	CASIN1-	VRAM Column address strobe long word 1.	
5	A[11:0]	Multiplexed row/column address bus for VRAM accesses.	
	A[8:7]	Used as cycle type specifier for Register I/O cycles as follows:	
	A[8]	A[7]	Cycle
	0	0	AH. CBW[7:0] is used to transfer IOA[15:8]
	0	1	AL. CBW[7:0] is used to transfer IOA[7:0]
10	1	0	Direct mode Data on CBW[7:0].
	1	1	Auto-increment mode DATA on CBW[7:0].
	MCLK	MBus system clock. All transfers are synchronized to this clock.	
15	ROE-	VRAM output enable. This signal is used to enable VRAM read data on the D[127:0] lines. It is not used for register I/O cycles.	

All VRAM transfers from the memory controller 20 are initiated by the assertion of RAS- with IOSEL- negated (high).

20

5. VSIMM Preemptive Protocol for VRAM I/O (IOSEL- negated)
 - Read Transaction Semantics

25

Reference is now made to Figure 6, which illustrates the timing during a VRAM page mode read, according to the protocol of the present invention. In Figure 6, a memory read operation begins with an MBus data request from an MBus address, which is performed by asserting MAS- active low together with

the address from which data is requested from being driven on the MBus address lines MAD [63:0]. If the memory controller 20 finds this data within its memory space, it puts the associated Row Address (RowAd) onto the memory bus address lines ADDR[11:0], and then asserts RAS- to active low. Then, the
5 memory controller 20 puts the associated Column Address (ColAd) on the memory bus address lines ADDR[11:0] while asserting CAS-. After the assertion of CAS-, the first data returns on the memory bus data lines DATA[127:0]. Using presently preferred VRAMs 34, data is returned approximately 100 nanoseconds after the initial assertion of RAS-. The data is
10 buffered through the memory controller 20 in one cycle and sent out onto the multiplexed data/address lines (MAD [63:0]) of MBus (40). Finally, an acknowledgment MRDY- signal (not shown) is asserted while the data is returned on the MBus 40 to acknowledge that the data transfer was successfully completed.

15

5. VSIMM Preemptive Protocol for VRAM I/O (IOSEL- negated)
- Write Transaction Semantics

20 Reference is now made to Figure 7, which illustrates the timing during a VRAM page mode write, according to the protocol of the present invention. In Figure 7, a memory write transaction begins with a MBus write request consisting of a main memory address asserted on the MBus address lines MAD [63:0], along with MAS-. In the next clock cycle, the data to be
25 written is asserted on the MBus address lines MAD [64], the first four bytes being asserted on the MAD lines for two cycles, and then each consecutive four bytes being asserted on the MAD lines for one cycle. During the four cycles that

each four bytes is asserted on the MBus address lines, an MRDY- signal (not shown) is asserted to acknowledge that this is a Valid Data Transfer. Once MRDY- is asserted, memory controller 20 asserts RAS-, along with the specified Row Address (RowAd). On the fifth cycle, the data is asserted on the memory bus 10, along with write enable (WE-), which remains asserted for the next three cycles. On the third cycle, memory controller 20 asserts Column Address (ColAd0), and three cycles later, CAS- is asserted at cycle 6, in which the first 144 bits of data are written. A new Column Address (ColAd1) is asserted in clock cycles 7 and 8, for which the data (that was asserted on cycles 7 and 8) is written when CAS- is asserted in clock cycle 8.

5. VSIMM Preemptive Protocol for VRAM I/O (IOSEL- negated)
- Refresh

15

With reference to Figure 8, the protocol supports the CAS-before RAS- refresh cycle. However the video buffer 35 may ignore this cycle if it is programmed to perform refresh cycles for VRAMs 34.

20

6. VSIMM Preemptive Protocol for Register I/O (IOSEL- asserted)

All register I/O cycles are initiated by the assertion of IOSEL- followed by RAS-. RAS- in this case is used as a chip select for reading and writing the device registers. There are two types of register cycles, 1) direct mode accesses consisting first of an address transfer followed by one or more data byte transfers and 2) indirect mode accesses whereby there is no address transfer.

The responding device automatically increments an address in its preloaded auto increment address register at the end of each data byte transfer. Indirect mode accesses are always done to I/O address (IOA) page 15. The memory controller 20 encodes auto increment mode data on A[8:7]. The addressed device is decoded from the previously loaded auto increment address register discussed below.

Referring now to Figure 9, an illustration of loading an address in an auto increment address register is shown. This cycle type is used to load an auto increment address register in the responding device. It is loaded from the CBW[7:0] lines during two successive data byte transfers of the address to be loaded. Any write to IOA 0xn8nn (bit 11 set) will cause the address on the CBW[7:0] lines during the address transfer to be loaded in the auto increment address register. The data transfer is suppressed in this cycle type.

Referring now to Figures 10a and 10b, an illustration of direct mode read and write operations are shown, respectively. Direct mode register cycles are used primarily to setup device registers. The present protocol supports byte, half-word, word, and double word transfer sizes. An address high (AH) and address low (AL) address transfer cycle is followed by one or more data byte cycles as determined by the transfer size. The responding device increments the address in its address register at the end of each data byte transfer.

Referring now to Figures 11a and 11b, an illustration of indirect mode read and write operations supported by the present protocol are respectively shown. Indirect mode accesses uses the auto increment address register for decoding the addressed device. This addressing mode is used primarily to

initialize look up tables stored in VRAMs 34. This register would have been previously loaded with a device address as described in connection with Figure 9 above. The memory controller 20 then performs page 15 read or write accesses to IOA[15:0]. The auto increment address register is incremented after
5 each data byte transfer.

Referring now to Figure 12, an illustration of a special I/O cycle access is shown. Special I/O cycles are used to access and initialize devices with long chip select access times , or those requiring excessive data setup and hold times.
10 The memory controller 20 and present protocol supports special cycles for accessing, for example, the digital-to-analog converter or equivalent devices contained in VSIMM 32.

15 7. Preemptive Protocol

The VBSY- signal is asserted by the video buffer 35 to preempt a memory controller 20 access cycle in progress to any VRAM 34 if video buffer 35 needs to perform a video reload or refresh operation. VBSY- is also asserted in response
20 to detecting a RAS- while a reload or refresh operation is already in progress. The memory controller 20 in these cases would deassert RAS- and wait for the negation of VBSY- (transition to high) before retrying its cycle.

The video buffer 35 asserts VBSY- based on the assertion of RAS- and
25 whether or not video buffer 35 is currently active or must perform an access cycle to VRAMs 34. The video buffer 35 negates VBSY- when video buffer 35 has completed its access to the VRAMs 34. This must be done such that RAS-

precharge is maintained at the VRAMs 34. The present protocol supports video buffer 35's implementation of REQUEST and ACTIVE windows. The REQUEST window is an service request generated internally by video buffer 35, whereas the ACTIVE window represents the state where the memory

5 controller 20 is blocked from accessing the VRAMs 34. The decision to enter the ACTIVE or blocking window for the REQUEST state is based on whether the current memory controller 35 cycle is to VRAMs 34 memory space or to register I/O space. The ACTIVE state is entered one clock cycle after the latest negation of RAS- or IOSEL-. Thus the active state is controlled by RAS- for

10 access cycles to memory space and IOSEL- for VRAM register I/O cycles.

7.1 VBSY Assertion and negation rules

15 Referring now to Figure 13, an illustration of assertion of VBSY- is shown, wherein video buffer 35 is in ACTIVE state and RAS- is detected as issued by memory controller 20. Generally, the video buffer 35 asserts the VBSY- signal one clock cycle after the RAS- input is asserted if the video buffer 35 is performing or must soon perform a memory cycle. In Figure 13, the video

20 buffer 35 asserts VBSY- one clock cycle after detecting a RAS- if within its ACTIVE window. VBSY- is negated (transitions high) two cycles after it negates RAS- to the VRAMs 34. The memory controller 20 may reassert RAS- two clock cycles after it negates RAS to the VRAMs 34. The memory controller 20 may reassert RAS- two clock cycles after detecting VBSY- negated. RAS-

25 precharge timing is thereby maintained at the VRAMs 34.

Referring now to Figure 14, an illustration of assertion of VBSY- is shown, wherein a VRAM access cycle by memory controller 20 is already in progress, and where video buffer 35 must perform a refresh or reload cycle (i.e., video buffer 35 enters REQUEST state). Generally, if a memory controller 20
5 memory cycle is already in progress, and the video buffer 35 needs to perform a VRAM operation, VBSY- interrupts the memory controller 20. The memory controller 20 will end its access within a predefined number of clock cycles. One clock cycle after RAS- is deasserted, the video buffer 35 will perform its memory cycle. VBSY- in this case is used to preempt the memory controller
10 20's access cycle to VRAM 34 already in progress. The memory controller 20 will terminate its cycle within a predefined number of clock cycles. According to the implementation of the present memory controller 20, this limit is currently set at thirteen (13) cycles. One clock cycle after RAS- is deasserted the video buffer 35 may enter its ACTIVE state. The video buffer 35 is responsible
15 for meeting the RAS- precharge specification for VRAMs 34. Thus the assertion of its RAS- will not occur before four cycles after the negation of RAS-.

Referring now to Figure 15, an illustration of assertion of VBSY- is
20 shown, wherein there is already in progress a register I/O cycle by memory controller 20, and wherein video buffer 35 enters its REQUEST state. If an memory controller 20 register cycle is in progress when the video buffer 35 needs to perform a VRAM operation, VBSY- interrupts the memory controller
20 . This case is similar to the previous case except that the ACTIVE- assertion
25 is delayed until the deassertion of IOSEL-. In this case, VBSY- is also used to preempt the memory controller 20's register I/O cycle in progress. The memory controller 20 will terminate its cycle within a predefined number of

clock cycles, and is currently set at thirteen (13) cycles. Once clock cycle after IOSEL- is deasserted, the video buffer 35 may enter its ACTIVE state.

Referring now to Figure 16, an illustration of special protocols available
5 on the integrated memory bus to guarantee reliable glitch-free operation of
RAS- and chip select during the transition of ownership between the memory
controller 20 and video buffer 35. In each case, VBSY- asserted for one clock at
the end of its ACTIVE window due to maintain consistency with the protocol
illustrated in Figure 13 above. In Figure 16, the VBSY- decision window must
10 extend one cycle beyond the ACTIVE window because the memory controller
20 was blocked from the video buffer 35. If VBSY- were not asserted, the
memory controller 20 would complete its access when in fact the data or
address cycle was blocked for register I/O access or VRAM access, respectively.

15 Still referring to Figure 16, the present protocol support a data_block
signal generated by the video buffer 35 to block the address and data lines. A
rasin_block signal is internally generated to block issuance of other internal
control and blocking signals: RAS-, CAS- and chip_select. Rasin_block would
be negated at the end of MCLK 4 if RAS- is not asserted. Otherwise rasin_block
20 remains asserted until the latter of RAS- or IOSEL- negated.

7.2 Special Protocols to support Video Buffer Functions

The video buffer 35 allows a single memory system to have
25 multiple VSIMMs 32 with independent video timing. Each VSIMM 32
contains the components needed to allow a video monitor to be connected to
the system. All of the VSIMMs 32 in the system are controlled by the memory

controller 20 on the motherboard. With reference to Figure 17, the video buffer 35 generally acts as a simple buffer for the multiplexed address lines ADDR[11:0], the byte write enables CBW[15:0], and the RAS- and CAS- lines which are generated by the memory controller 20. The video buffer 35 asserts DTOE- when the memory controller 20 asserts ROE-, enabling the outputs of the VRAMs 34 for read cycles. Occasionally, the video buffer 35 must take control of the outputs to the VRAMs in order to perform refresh and reload cycles. As shown in Figure 17, when the refresh and reload cycles are in progress or are imminently required, and RAS- is asserted by the memory controller 20, the video buffer 35 asserts the VBSY- signal to memory controller 20 to signal that the VRAMs 34 cannot be accessed. The video buffer 35 is also used to control register operations for its own internal registers and the registers of the other components on the VSIMM 32. The register operations are described in another portion of this document.

15

The protocol of the present invention further supports resolution of control of VRAMs 34 when refresh and reload operations must be performed. The video buffer 35 refreshes VRAMs 34 at a rate determined by the bits in an internal control register. The video buffer 35 counts clock cycles and asserts an internal refresh request at the appropriate time. If refresh enable is set, refresh attempts by the memory controller 20 are blocked. If the refresh enable is cleared and the memory controller 20 is programmed to perform refresh to the VSIMM 32, proper operation will result.

25

Referring now to Figure 18, an illustration of the operative protocol when the video buffer 35 is performing a refresh cycle and the memory controller 20 asserts RAS-, the video buffer 35 asserts the VBSY-

signal. The memory controller 20 will then deassert the RAS- and will wait until VBSY- is deasserted before trying again. If the video buffer 35 needs to perform a refresh cycle when the memory controller 20 is already performing a memory or register I/O cycle, VBSY- will be asserted and held. The memory
5 controller 20 will preempt its memory cycle with-in nine clock cycles. When the memory controller 20 deasserts RAS-, the video buffer 35 will perform the refresh cycle. VBSY- is deasserted when the refresh operation is completed. Note that if RAS- is not asserted, VBSY- will remain tristated throughout the refresh cycle.

10

Referring now to Figure 19, the protocol supports video buffer 35 reload cycles to VRAMs 34 as needed to keep the VRAM shift registers loaded properly. When the video buffer 35 is performing a reload cycle and the memory controller 20 asserts RAS-, the video buffer 35 asserts the VBSY- pin.

15 The memory controller 20 will then deassert the RAS- and will wait until the VBSY- is deasserted before trying again. If the video buffer 35 needs to perform a reload cycle when the memory controller 20 is already performing a memory cycle, VBSY- will be asserted and held. The memory controller 20 will preempt its memory cycle within eleven clock cycles. When the memory controller 20
20 deasserts the RAS-, the video buffer 35 will perform the reload cycle. VBSY- is deasserted when the reload operation is completed.

D. Memory System Architecture and Implementation

25

1. DSIMM

As previously stated, the present memory bus architecture has a 128 bit datapath width, together with 16 bits of ECC (Error Correcting Code).

The memory bus can have two RAS lines per SIMM 30, two CAS lines, one WE*, and 11 bits of address. The DSIMMs 72 are double sided, using TSOP DRAMs, and are spaced 300 mils apart. There are presently two variants of DSIMMs 72: a 16 megabyte version using 1Mbx4 DRAMs 74 and a 64 megabyte version using 4Mbx4 DRAMs. Each DSIMM 72 contains a buffer 75 to distribute the RAS-, CAS0-, CAS1-, WE-, and A<0-11> signals to the DRAMs 74. The present memory system implementation uses 80 nanosecond DRAMs 74 to achieve 100 nanosecond data access time from RAS- generation. The memory controller 20 supports the following densities of DRAMs 74 on DSIMMs 72:

RAM	Density	# of Rams	Capacity	Type
1Mbit	(256Kx4)	36	4 MByte	ECC DRAM
4Mbit	(1Mx4)	36	16 MByte	ECC DRAM
16Mbit	(4Mx4)	36	64 MByte	ECC DRAM

Although the present memory bus architecture supports two RAS lines per SIMM position, in the present implementation the RAS0 and RAS1 signals on the sockets for all SIMMs 30 are tied together on the motherboard, and the OE- is tied to ground on the DSIMM 72. Because the present architecture is arranged in full datapath width arrangements, DSIMMs 72 can be individually added one at a time, each DSIMM 72 having 144 bits of data and is double sided, using 36 DRAMs per DSIMM 72.

2. VSIMM

The present implementation of a VSIMM 32 contains thirty-two VRAMs 34, a video buffer 35, and a display interface 36. The video buffer 35 buffers RAS-, CAS-, WE-, the address bits ADDR[11:0] and the CBWs [15:0].

CBWs are used as ECC bits in DSIMMs 72, and as byte write enables in the VSIMMs 32. The video buffer 35 also provides refresh and reload circuitry for the VRAMs 34 on that VSIMM 32. The display interface 36 provides facility to multiplex the serial shift data from the VRAMs 34 and send the multiplexed data to a digital-to-analog converter (not shown) to form a red, green, and blue (RGB) output signal for the display monitor's color gun.

The memory controller 20 supports the following densities of VRAMs on VSIMMs 32:

10

RAM	Density	# of Rams	Capacity	Type
2Mbit	(128Kx8)	16	2 MByte	128b VRAM
15 1Mbit	(256Kx4)	32	4 MByte	128b VRAM
2Mbit	(256Kx8)	32	8 MByte	256b VRAM
4Mbit	(256Kx16)	16/32	8/16 MByte	256b VRAM
4Mbit	(512Kx8)	32	16 MByte	256b VRAM

20

In the present implementation of the memory bus, there are two versions of the VSIMM 32. There is a four megabyte version using 256Kx4 VRAMs 34, and an eight megabyte version using 256Kx8 VRAMs. The present implementation uses 80 nanosecond VRAMs 34 to provide for 100 nanosecond data access times following generation of RAS- by memory controller 20.

25

3. NVSIMM

30

a. Timing Adjustment Between NVSIMM and Memory Controller

Since low power SRAMs 84 may be used in an NVSIMM 82, and also because the memory controller 20 operates at a faster speed with DRAM SIMMs, there must be some timing adjustment between the memory controller 20 and NVSIMM 82 for proper read and write operations. With
5 reference to Figures 20a and 20b, at the beginning of each operation, whether it is a read or write, the memory controller 20 will be instructed to relinquish the NVSIMM when NVSIMM 82 asserts the VBSY- signal active low. This feature provides more clock cycles to NVSIMM 82 to complete its operation properly. VBSY- signal is synchronously asserted one clock cycle after a first RAS- by
10 NVSIMM 82, and it is deasserted one clock cycle after the removal of the first CAS- signal. During the second cycle which is the retry cycle, VBSY- is asserted two clock cycles after the assertion of RAS- by NVSIMM and deasserted one clock cycle after RAS- is released. The data will be latched by the memory controller 20 at the rising edge of retry CAS- in a read operation, as shown in
15 Figure 20a.

b. NVSIMM Interface Timing

NVSIMM 82 operation, and all outputs from the NVSIMM 82, are
20 synchronous to the rising edge of MCLK except SIRQ- (battery fail). There are four modes of access to the NVSIMM 82: write, read, read-modify-write, and refresh. During a refresh cycle, no operation is done on the NVSIMM 82. For both read and write operations, the NVSIMM 82 asserts VBSY- one cycle after assertion of RAS- signal. If at the next clock cycle no CAS- is asserted, as
25 mentioned above then NVSIMM 82 is in a write operation. VBSY- is deasserted (high) one cycle after deassertion of retry RAS-.

As presently preferred, NVSIMM 82 can support either 8 byte or 16 byte transaction sizes, which is determined by CAS0- and CAS1-. NVSIMM 82 will go to a wait loop mode and ignores any refresh cycles, if a refresh cycle is detected before assertion of retry RAS- signal. NVSIMM 82 proceeds the cycle
5 after the refresh is done. In the present embodiment, it takes 14 clock cycles, including precharge, for NVSIMM 82 to perform the read operation (Figure 20a). Similarly, in the present embodiment it takes 16 clock cycles, including precharge, for NVSIMM 82 to perform the write operation, shown in Figure 20b. There will be extra cycles if refresh occurs between first RAS- and retry
10 RAS-. There will be extra cycles if refresh occurs between first RAS- and retry RAS-.

Finally, the SRAMs 84 do not require refresh cycle, and NVSIMM 82 will ignore all bus traffic associated with refresh during idle mode, read or write
15 operations. NVSIMM 82 uses the Pin 92 (SIRQ-) on the memory bus 10 to report the status of its battery 86. The value of this signal is stored in one of the registers in a controller on the mother board. Software can access this register and verify the status of the battery. Each NVSIMM 82 may be equipped with its own exclusive battery fail signal (Pin 92).

20

The foregoing has described a bus architecture for integrated data and video memory. It is contemplated that changes and modifications may be made by one of ordinary skill in the art, to the device components and arrangements of elements of the present invention without departing from the
25 spirit and scope of the invention.

CLAIMS

1. A computer bus structure, said computer bus structure comprising the elements of:

a processor bus, said processor bus coupled to at least one processor; and a memory bus coupled to said processor bus, said memory bus comprising,

a plurality of address lines;

a plurality of data lines;

a plurality of control lines;

a plurality of memory module connectors, each of said memory module connectors coupled to said address, data, and control lines;

at least one data memory module coupled to one of said memory module connectors, said data memory module storing data;

at least one video data memory module coupled to one of said memory module connectors; said video data memory module storing video data; and

a memory controller, said memory controller coupling said memory bus to said processor bus, said memory controller receiving memory requests over said processor bus and transmitting said memory requests to said data memory module and said video data memory module over said memory bus, said memory controller receiving register I/O requests over said processor bus and transmitting said register I/O requests to said video data memory module over said memory bus.

2. The computer bus structure as claimed in claim 1 wherein said data memory module comprises a single in-line memory module having a plurality of dynamic random access memory devices (DRAMs).

3. The computer bus structure as claimed in claim 1 wherein said video data memory module comprises a single in-line memory module having a plurality of video random



access memory devices (VRAMs).

4. The bus structure as defined by claim 1, wherein said video data memory module is further coupled to a video display device, such that video data stored in said video memory module is directly provided to said video display device.

5. The bus structure as defined by claim 4, wherein said control signals includes a plurality of check bit write signals (CBWs) for transferring write enable signals to a plurality of video memory devices disposed on said video memory module.

6. The bus structure as defined by claim 5, wherein said write enable signals are valid when said CBWs are in a logical low state.

7. The bus structure as defined by claim 5, wherein said CBWs transfer a predetermined error correction code to a plurality of memory devices disposed on said data memory module.

8. The bus structure as defined by claim 7 wherein said memory devices disposed in said data memory module comprise dynamic RAM (DRAM).

9. The bus structure as defined by claim 7 wherein said data memory module and said video memory module have a data path width equal to a data path width of said memory bus.

10. The bus structure as defined by claim 9 wherein said data path width of said data memory module and said video memory module is 144 bits wide.



11. The bus structure as defined by claim 9 wherein said memory bus operates at a frequency substantially equal to 20 MHz.

12. The bus structure as defined by claim 9 wherein a plurality of processors are coupled to said processor bus.

13. The bus structure as defined by claim 12 wherein said memory bus control signals comprise a VBSY signal for indicating refresh cycles for said video memory devices, RAS and CAS signals for selecting said memory devices, Write Enable signals for indicating a write transfer, and IOSEL signals for accessing a register address space of said memory bus means.

14. The bus structure as defined by claim 13 wherein said VBSY signal decouples said memory controller means from said video memory module when said video memory devices require refresh or reloading.

15. A method for integrating data memory and video memory access requests in a computer system, said computer system comprising a processor bus having at least one processor and a memory bus coupled to said processor bus through a memory controller, said method comprising the steps of:

coupling at least one data memory module and at least one video memory module to said memory bus;

sensing a memory access request in said memory controller sent by a processor over said processor bus, said access request sent over said processor bus using a first protocol;

translating said memory access request received over said processor bus using said memory controller into a memory access request for said memory bus and



transmitting said memory access request over said memory bus using a second protocol;

if said memory access request over said processor bus is targeted for said data memory module, said memory controller accessing said data memory module according to said memory access request received over said processor bus;

if said memory access request over said processor bus is targeted for said video memory module, said memory controller accessing said video memory module according to said memory access request received over said processor bus.

sensing a register I/O access request in said memory controller sent by a processor over said processor bus, said register I/O access request sent over said processor bus using said first protocol;

translating said register I/O access request received over said processor bus using said memory controller into a register I/O access request for said memory bus and transmitting said register I/O access request over said memory bus using said second protocol;

16. The method as claimed in claim 15, further comprising the steps of:

coupling a video display device to said video memory module;

transferring video data from said video memory module to said video display device.

17. The method as claimed in claim 15 wherein said second protocol comprises the step of transferring a plurality of check bit write signals (CBWs) to video memory devices disposed in said video memory module, each CBWs indicating a write enable for one of said video memory devices.



18. The method as claimed in claim 17, wherein said second protocol further comprises the step of transferring said CBWs to data memory devices disposed in said data memory module, said CBWs indicating a predetermined error correction code.

19. The method as claimed in claim 18, wherein said second protocol further comprises the step of transferring address signals for accessing said data memory module and said video memory module, a VBSY signal for indicating refresh cycles for said video memory devices, RAS and CAS signals for selecting said memory devices, Write Enable signals for indicating a write transfer over said memory bus, and IOSEL signals for accessing a register address space of said memory bus.

20. The method as claimed in claim 19 wherein said VBSY signal causes said access request over said memory bus to terminate and said access request over said memory bus is subsequently retried.

