



US005388228A

United States Patent [19]

Heath et al.

[11] Patent Number: 5,388,228
[45] Date of Patent: Feb. 7, 1995

- [54] **COMPUTER SYSTEM HAVING DYNAMICALLY PROGRAMMABLE LINEAR/FAIRNESS PRIORITY ARBITRATION SCHEME**
- [75] Inventors: **Chester A. Heath**, Boca Raton, Fla.;
James O. Nicholson, Austin, Tex.;
James D. Reid, Boynton Beach, Fla.;
Frederick E. Strietelmeier, Austin, Tex.
- [73] Assignee: **International Business Machines Corp.**, Armonk, N.Y.
- [21] Appl. No.: **65,057**
- [22] Filed: **May 20, 1993**

Related U.S. Application Data

- [63] Continuation of Ser. No. 725,223, Jun. 26, 1991, abandoned, which is a continuation of Ser. No. 102,690, Sep. 30, 1987, abandoned.
- [51] Int. Cl.⁶ **G06F 13/00**
- [52] U.S. Cl. **395/325; 395/725; 364/241.9; 364/242.6; 364/242.7; 364/242.8**
- [58] Field of Search **395/325, 725**

References Cited

U.S. PATENT DOCUMENTS

3,947,824	3/1976	Doehle et al.	340/172.5
4,334,288	6/1982	Booher	364/900
4,384,323	5/1983	Ahuja	
4,473,880	9/1984	Budde et al.	
4,493,022	1/1985	Nicholas et al.	395/325
4,499,538	2/1985	Finger et al.	364/200
4,559,595	12/1985	Bourdreau et al.	364/200
4,597,054	6/1986	Lockwoot et al.	364/900
4,611,297	9/1986	Dudley et al.	364/900
4,639,859	1/1987	Ott	364/200

OTHER PUBLICATIONS

Beaston, John "Multiprocessor Bus Is Ready To Meet 32-Bit Applications of Future", Electronics, Mar. 1984, pp. 126-131.

Article: Intel (Microprocessor & Peripheral Handbook) 1987 addition, pp. 2-196 & 2-197.

Article: Intel (Microprocessor & Peripheral Handbook) 1987 addition, pp. 2-289.
 Article: Intel (Microprocessor & Peripheral Handbook) 1987 addition, pp. 2-292 & 2-293.
 Article: Intel (Microprocessor & Peripheral Handbook) 1987 addition, pp. 3-163-3-168.
 Article: "The Mutual Exclusion Problem: Part II Statement and Solutions", Leslie Lamport, Journal of the Association for Computing Machinery, vol. 33, No. 2, Apr. 1986, pp. 327-348.


(List continued on next page.)

Primary Examiner—Parshotam S. Lall
Assistant Examiner—Ayni Mohamed
Attorney, Agent, or Firm—Robert S. Babayi

[57] ABSTRACT

A computer system having a plurality of devices which transmit and receive information over a channel is presented. The system includes, in the preferred embodiment, a central arbitration control circuit and a local arbiter associated with each device contending for channel access. Each local arbiter, corresponding to a device which desires channel access, generates a channel request signal to the central control circuit. At the appropriate time when the channel becomes available, the central control circuit generates an arbitrate signal. All local arbiters, then contending for channel access, compare the priority level on the arbitration bus with the priority value of the device it is arbitrating on behalf of, with the winning device gaining access to the channel. Each of the local arbiters contains a programmable circuit which enables the arbiter to operate either utilizing a linear priority arbitration technique or a fairness priority arbitration technique. Thus, each device is dynamically programmable to operate in a linear mode or a fairness mode depending on user and/or application needs. Additionally, each local arbiter includes a circuit which detects whether the device is a burst transfer device or a single transfer device.


2 Claims, 9 Drawing Sheets



OTHER PUBLICATIONS

- Article: "Current Packet Radio Network Protocols", John Jubin, 1985 IEEE Proceedings, pp. 86-92.
- Article: "Realization of Global Fairness by Local Control in Live Simple Nets", I. Suzuki et al, Midwest Symposium on Circuits & Systems (28) 1985, pp. 245-248.
- Article: "Arbitration and Control Acquisition in the Proposed IEEE 896 Futurebus", D. M. Taub, Aug. 1984 IEEE Micro, pp. 28-41.
- Article: "Futurebus Anticipates Coming Needs", John Theus et al, Jul. 1984 Electronics Magazine, pp. 108-112.
- Article: "Wideband Local Nets Enter the Computer Arena", W. E. Burr et al., May 1984 Electronics Magazine, pp. 145-150.
- Article: "High-Speed Network Uses Fiber Optics", H. Ikeman et al, Oct. 22, 1984 Electronics Week, pp. 95-99.
- Article: "Fair Termination of Communicating Processes", O. Grumberg et al, Proceedings of Distributed Computers (3rd annual ACM symposium) 1984, pp. 254-265.
- Article: "Distributed Algorithms for Ensuring Fair Interprocess Communications", A. P. Sistla, Principals Of Distributed Computers (3rd annual ACM symposium) 1984, pp. 266-277.
- Article: "A Fully Distributed Arbiter for Multiprocessor Systems", G. Cioffi et al, Microprocessing and Microprogramming vol. 11, Nr. 1 (Jan. 1983) pp. 15-22.
- Excerpt: "Standard Specification for S-100 Bus Interface Devices", K. A. Elmquist et al, IEEE Task 696.1/D2, reprinted from Computer vol. 12, No. 7, 1979.
- Article: "Variable Architecture Computer", P. A. Smith, IBM Technical Disclosure Bulletin, vol. 13, No. 9, Feb. 1971, pp. 2777-2778.
- Article: "Priority Preemption in a Token Ring Communication System", D. W. Andrews et al, IBM Technical Disclosure Bulletin, vol. 27, No. 7B, Dec. 1984, pp. 4305-4308.
- Article: "Swap-Out Candidate Selection for Storage Management", A. D. Childs et al, IBM Technical Disclosure Bulletin, vol. 28, No. 9, Feb. 1986, pp. 4171-4173.
- Article: "Maintaining Fairness with Priority Tokens in ECL Environment", K. K. Sy, IBM Technical Disclosure Bulletin, vol. 26, No. 2, Jul. 1983, pp. 798-800.
- Article: "Delay-Guaranteed Service for Store-And-Forward Networks", P. Kermani, IBM Technical Disclosure Bulletin, vol. 26, No. 5, Oct. 1983, pp. 2615-2619.

FIG. 1






FIG. 8





FIG. 3










FIG. 5A

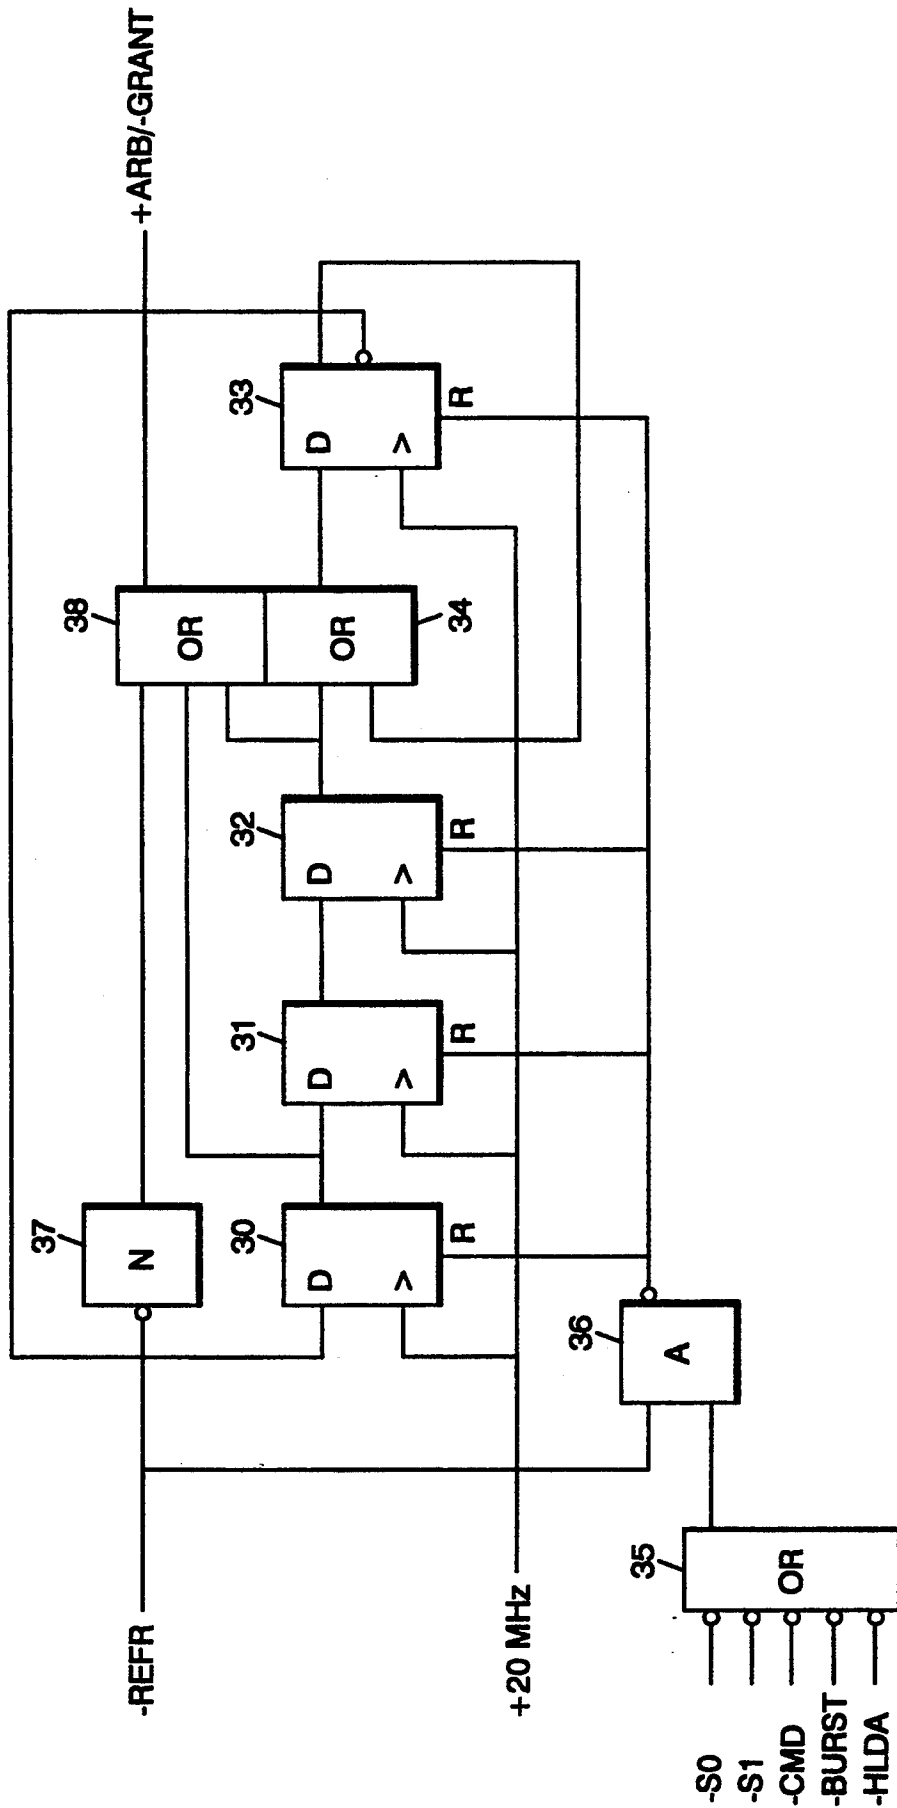


FIG. 5B




FIG. 5C


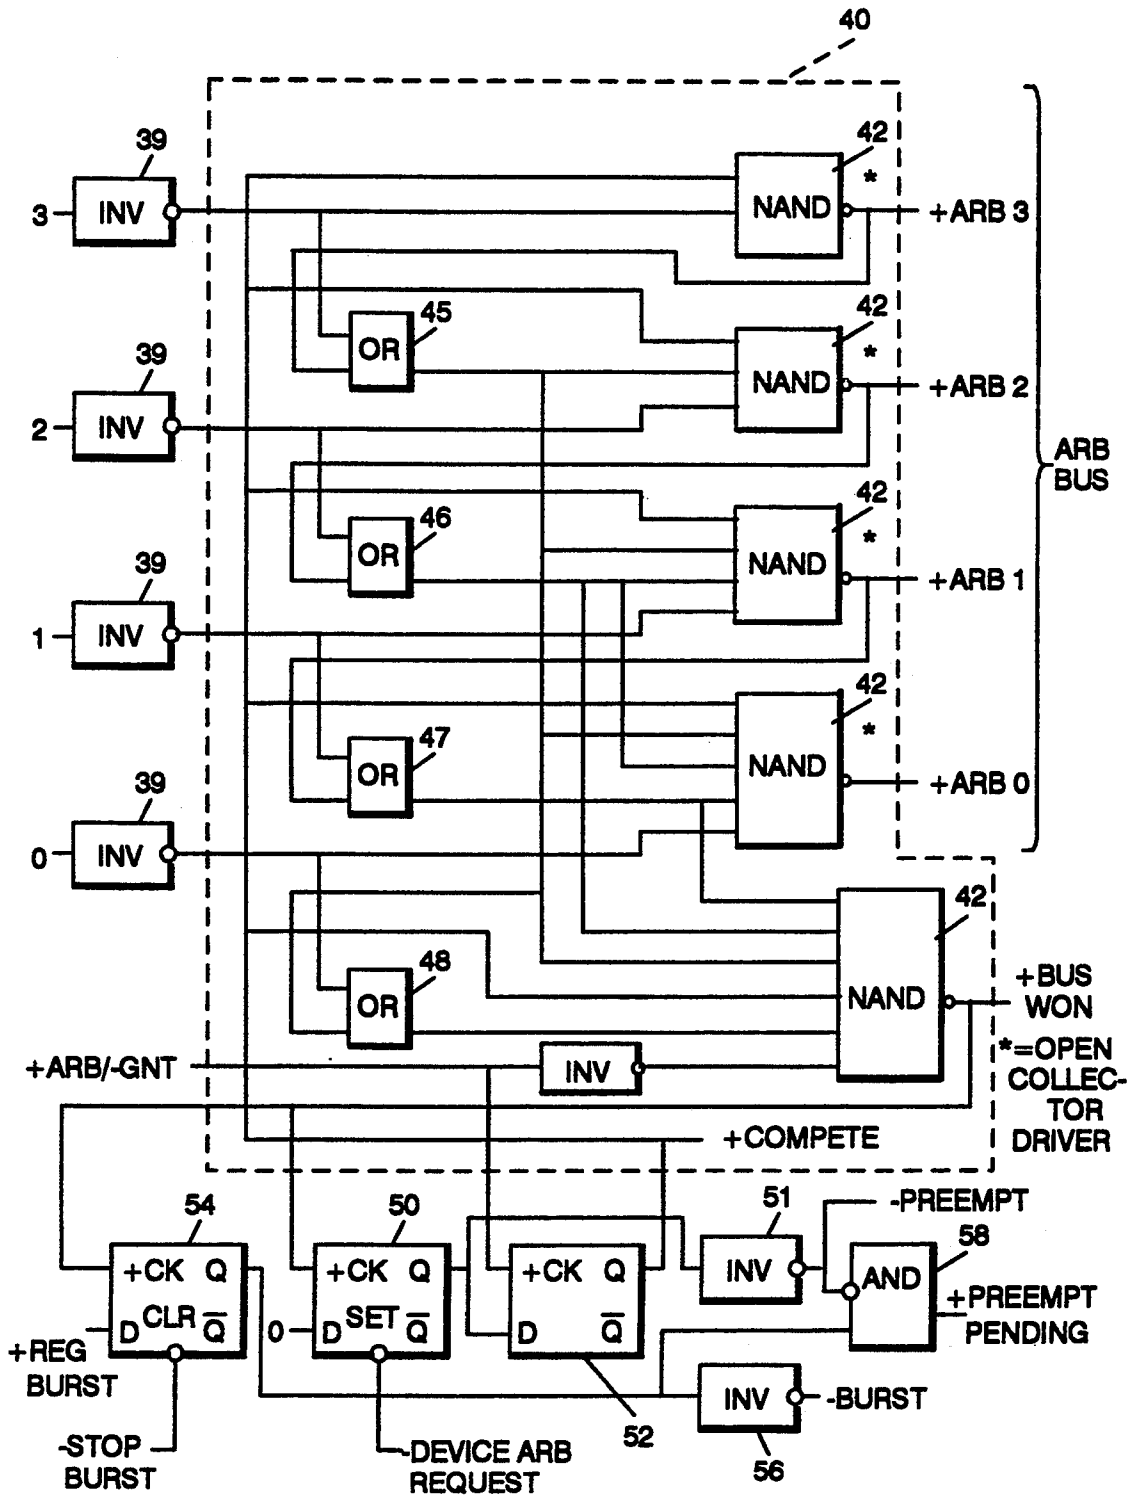


FIG. 6






FIG. 7

**COMPUTER SYSTEM HAVING DYNAMICALLY
PROGRAMMABLE LINEAR/FAIRNESS
PRIORITY ARBITRATION SCHEME**

This is a continuation of patent application Ser. No. 07/725,224, filed Jun. 26, 1991 now abandoned, which itself is a continuation of patent application Ser. No. 07/102,690 filed Sep. 30, 1987 now abandoned.

BACKGROUND OF THE INVENTION

The present invention relates to arbitration schemes for use in computer systems. More particularly, the invention provides for an arbitration scheme having a local arbiter for each device in the system which requires use of the channel, the local arbiter being programmable to provide either a linear priority arbitration technique or a fairness priority arbitration technique.

Presently, there are various arbitration schemes in existence which determine the order in which multiple devices in a computer system may gain access to the channel. The majority of schemes generally fall into one of two categories, namely a centralized arbitration scheme or a distributed arbitration scheme. In both schemes each device has a priority value associated with it which is compared to the priority values of other devices contending for use of the channel, with the device having the highest value winning access to the channel. However, in the centralized scheme the winning device is determined at a central device, such as a CPU, while in a distributed scheme each device makes its own determination if it is the winning device (i.e. highest priority device requesting access to the channel).

The most common arbitration schemes, whether centralized or distributed, utilize a linear priority arbitration technique in which the highest priority device generally maintains control of the bus as long as it may require, and thus may prevent lower priority level devices from gaining access to the bus for inordinate periods of time. This might ultimately result in overrun conditions in the lower priority devices.

Another solution is commonly referred to as a fairness priority arbitration technique in which each device is given an opportunity to gain access to the bus, after the highest priority device has relinquished control of the bus and before the highest priority device may regain control of the bus. One such technique prevents the device presently in control of the bus from contending for the next bus cycle, and generally works well for single transfer devices. However, this will allow first and second devices to trade off but together both devices may lock out a third device. Another type of fairness priority arbitration technique is shown in the Intel Manual for the 8289 device and is illustrated on page 2-196 of that document, the same technique again being illustrated in the Intel Manual for the 82289 device on page 3-167. In both instances, a centralized arbitration scheme is implemented in which a rotating priority resolving technique is utilized. A special circuit rotates priority between requesting devices, thus allowing each device an equal chance to use the multi-master system bus over a given period of time. Therefore, all the devices contending for the bus are required to participate in the rotating priority technique.

It is readily apparent from the aforesaid discussion that both the centralized arbitration scheme and the distributed arbitration scheme were originally designed

so that all the devices working in the scheme operate either utilizing the linear priority arbitration technique or the fairness priority arbitration technique. However, with the addition of burst transfer devices intermingled with single transfer devices it became desirable to have a hybrid priority arbitration technique in which some devices may operate under the linear priority arbitration technique and other devices operate under the fairness priority arbitration technique.

One such hybrid technique, described in the article entitled "Arbitration and Control Acquisition in the Proposed IEEE 896 Futurebus" by D. M. Taub, found in the August 1984 issue of IEEE Micro, describes an arbitration scheme in which potential masters (PM's) are divided into two classes. The first class contains priority modules, whose need for the bus may be particularly urgent, for instance because of having to cope with realtime operations. The second class contains fairness modules, whose need for the bus is less urgent. Priority modules issue a request for the bus whenever they need it, and compete in the arbitration process the next time it occurs. On the other hand, once a fairness module has had control of the bus, it desists from issuing further bus requests until there are no requests left to service. However, each PM is preselected to be either a fairness module or a priority module and thus is not dynamically programmable to accommodate the varying needs of the user and/or different computer applications.

The present invention in the preferred embodiment contemplates an improved hybrid system by providing a dynamically programmable circuit in each local arbiter which allows either the linear priority arbitration technique or the fairness priority arbitration technique to be initiated depending upon the particular needs of the user and/or the application software being utilized.

Accordingly, it is an object of the present invention to provide a computer system having a dynamically programmable arbitration circuit for initiating either a linear priority arbitration technique or a fairness priority arbitration technique for each device requiring access to the system bus.

It is another object of the present invention to provide a computer system having a distributed priority arbitration scheme in which each device requiring access to the system bus employs a local arbiter associated therewith, the local arbiter having a circuit which is programmable to allow the local arbiter to initiate either a linear priority arbitration technique or a fairness priority arbitration technique.

It is yet another object of the present invention to provide a computer system having a distributed dynamically programmable linear/fairness priority arbitration scheme which accommodates a mixture of single transfer devices and burst transfer devices.

SUMMARY OF THE INVENTION

In accordance with the above and other objects of the present invention there is provided in the preferred embodiment a computer system having a distributed dynamically programmable linear/fairness priority arbitration scheme. The system comprises a local arbiter associated with each device which requires access to the system bus and a central arbitration control circuit. The local arbiters are interconnected via an arbitration bus and each local arbiter is set to a priority value representative of the priority level of its corresponding device. The central arbitration control circuit monitors

the availability of the system bus and generates a +ARB/-GNT signal to each local arbiter which starts the arbitration cycle and allows the winning device to be granted access to the bus at the appropriate time. Each of the local arbiters includes a priority value compare circuit which compares the highest priority value on the arbitration bus to its own value to determine which device requesting access to the bus has the highest priority. The local arbiter further includes a circuit portion which is used to detect whether the associated device is a burst transfer device or not and a programmable circuit so that the associated arbitration circuit may operate either using a fairness arbitration priority technique (fairness mode) or a linear priority arbitration technique (linear mode). Thus, if a burst transfer device has control of the bus and is programmed to operate in the fairness mode, as soon as another device desires access to the bus the burst device will relinquish control of the bus after an appropriate number of transfers, allowing the next highest priority level requesting device to gain control of the bus. Since each local arbiter includes the programmable arbitration mode select circuitry, each device may operate either in a fairness mode or a linear mode, regardless of whether the device is a burst transfer device or single transfer device, depending upon the needs of the user and/or the specific requirements of a particular application software program.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram showing a computer system employing the present invention;

FIG. 2 is a conceptual diagram used to explain the operation of the present invention;

FIG. 3 is a block diagram illustrating the relationship between the central arbitration control circuit and the local arbiters of the present invention;

FIGS. 4A through 4F is a timing chart used to explain the operation of the block diagram of FIG. 3;

FIG. 5A is a circuit diagram illustrating a first embodiment of the central arbitration control circuit of FIG. 3;

FIG. 5B is a circuit diagram illustrating a circuit from generating the processor HOLD REQUEST signal;

FIG. 5C is a timing chart used to explain the operation of the block diagram of FIG. 5A;

FIG. 6 is a schematic drawing of the local arbiter 28 of FIG. 3;

FIG. 7 is a schematic drawing of the programmable linear/fairness mode selection circuitry in the local arbiter 28 of FIG. 3;

FIG. 8 is a circuit diagram illustrating an alternate embodiment of the central arbitration control circuit of FIG. 3.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT OF THE INVENTION

Referring first to FIG. 1 of the accompanying drawings, there is shown a block diagram of a computer system in which the present invention is used to advantage.

A CPU 10 communicates with a main memory 15, bus controller 16 and math co-processor 14 via a system bus 26. Communication between the CPU 10 and its associated peripheral devices is through a bus controller 16, the latter being coupled to the peripheral devices through the channel 25. In the present example, the

peripheral devices include an auxiliary memory 17, two communications devices 18 and 19, a hard file 20, an optical disk 21, and two floppy disks 22 and 23. Other peripherals can of course be used as well as system needs dictate. The peripheral devices are represented generically by DMA slave 24, but may also include devices which are actually bus masters.

A DMA controller 12 is provided to allow at least selected ones of the peripheral devices direct memory access. For this purpose the channel 25, or at least a portion thereof, is branched to the DMA controller 12. Each peripheral allowed DMA access is provided with an arbitration circuit 28, and each peripheral having an arbitration circuit 28 is assigned an arbitration (priority) level. As will be explained in more detail later, a central arbitration control circuit 11 generates appropriate control signals to allow the distributed arbitration circuits 28 to arbitrate on behalf of their respective devices for control of the channel 25.

Referring now to FIG. 2, the channel 25 is further shown as being broken down into an address bus, a data bus, a control bus and an arbitration bus. The arbitration bus is comprised of four priority lines, namely ARB0 through ARB3, a +ARB/-GNT line, a -PREEMPT line, and a -BURST line.

Referring now to FIG. 3, the relationship between the central arbitration control circuit 11 and various local arbiter units 28A through 28C will be described. In general, when a device wants access to the channel 25 to transfer data, the local arbiter circuits 28A through 28C will receive a request signal from the particular device to which the arbitration circuit is related. The request signal is converted to a -PREEMPT signal which is generated by the local arbiter and transmitted to the central arbitration control circuit and each of the local arbiters over the -PREEMPT line of the arbitration bus. It should be noted in the specific embodiment of this invention that the -PREEMPT lines are OR'ed together and thus it is irrelevant to the central arbitration control circuit 11 which particular device is requesting an arbitration signal. As described earlier, each local arbiter circuit 28 has the capability of generating a -BURST signal over the -BURST line if the device which that particular local arbiter 28 is associated with happens to be a Burst device, such as the hardfile 20 or floppy disk drive 22 and 23. The central arbitration control circuit 11 generates the +ARB/-GNT signal at an appropriate time as determined by the CPU 10 HOLD signal and +REFRESH memory signal from a refresh controller (not shown) well known to those skilled in the art, in response to a -PREEMPT signal from one or more of the local arbiter circuits 28A through 28C.

Thus, when any one of the devices desires to contend for use of the channel 25, it generates a request signal to its corresponding local arbiter 28, which then generates a -PREEMPT signal over the -PREEMPT line or the arbitration bus. Then at the appropriate time when the bus becomes available, as determined by the HOLD signal from the CPU 10 and the +REFRESH signal from the refresh controller, the central arbitration control circuit 11 generates the +ARB state of the +ARB/-GNT signal over the arbitration bus to each one of the local arbiters 28. In response to the +ARB state, each of the local arbiters 28 which desires access to the channel 25 drives its priority level onto respective lines ARB0-ARB3 of the arbitration bus. Then each of the local arbiters 28A, 28B and 28C, desiring

access to the channel 25 compares its designated priority level with the priority level on the arbitration bus and takes itself out of contention for the bus 25 in the event its priority level is lower than that being driven onto the arbitration bus. Thus, at the end of the arbitration cycle only the local arbiter 28 having the highest priority level during that arbitration cycle remains in contention for the channel 25 and thus gains control of the channel 25 when the -GRANT state is received from the central arbitration control circuit 11 over the +ARB/-GRANT line. If the device corresponding to the particular arbiter circuit 28A, 28B or 28C happens to be a Burst device, the -BURST signal is activated on the -BURST line to the central arbitration control circuit 11 so that burst transfers may be made in the event that the channel 25 is available for burst transfers.

A better understanding of the operation of the system illustrated in FIG. 3 may be appreciated by referring to the example represented by the timing charts of FIGS. 4A-4F. For purposes of illustration, assume that a Burst mode device has gained control of the channel as illustrated by (a) of FIG. 4E, and that another device then requests preemption as shown by (b) of FIG. 4F. Once the second device activates -PREEMPT, the Burst device presently in control of the channel completes the current transfer as illustrated in (c) of FIG. 4D. Upon completion of the current transfer, the Burst device that is relinquishing control of the channel removes its -BURST signal from the -BURST line as shown in (d) of FIG. 4E. It will not participate in the next bus arbitration cycle. The central arbitration control circuit 11 then places the +ARB/-GNT signal in the GRANT state as illustrated in (e) of FIG. 4A. It also represents the beginning of another arbitration cycle and arbitration for the channel 25 begins at (f) of FIG. 4C. After the +ARB/-GNT signal goes low (-GNT), control of the channel is granted to the new device as illustrated in (g) of FIG. 4A. The new device which has gained control of the channel 25 then removes the -PREEMPT signal in response to the -GRANT signal as illustrated in (h) of FIG. 4F.

Referring now to FIG. 5A, a more detailed circuit description of the first embodiment of the central arbitration control circuit 11 is illustrated. The central arbitration control circuit 11 comprises a modified Johnson ring timing chain including counters 30 through 33 and OR gate 34, OR gate 35, NAND gate 36, inverter 37 and OR gate 38. Assuming the bus begins in an idle condition with the CPU 10 "owning" the bus, but not using it, the circuit operation will be described hereinafter in conjunction with the timing chart of FIG. 5C. In the aforesaid condition, +ARB/-GRANT is then active low, and the arbitration priority levels ARB 0 through ARB 3 all have a value of 1. The modified Johnson ring timing chain is held reset by the CPU 10 +HLDA signal through the OR gate 35 and the NAND gate 36. When a device needs access to the bus, the -PREEMPT signal is activated. As shown in FIG. 5B, the -PREEMPT signal going active results in the output of OR gate 39 going positive representing a PROCESSOR HOLD REQUEST (+PROC HRQ) signal. The +ARB 0 through +ARB 3 signal and a +GRANT signal are also input to the OR gate 39 to insure that the CPU 10 will not interfere with bus transfers by other devices. The +PROC HRQ signal causes the CPU 10 to deactivate +HLDA which results in the reset signal (output from OR gate 35) being removed from counters 30 through 33. It should be understood

that inputs -S0, -S1, -CMD and -BURST must be inactive in order for +HLDA to remove the reset signals from the aforesaid counters 30 through 33, as illustrated in FIG. 5C. The -S0 signal represents the WRITE cycle, and the -S1 signal represents the READ cycle. The -CMD signal is generated by the current bus master a specified time period after -S0 or -S1 (in the preferred embodiment a minimum of 55 nanoseconds). During READ cycles the -CMD instructs the slave device to place READ data onto the bus and during WRITE cycles -CMD is activated for validation of WRITE data.

On the next (20 MHZ) clock pulse, after +HLDA is deactivated, the counter 30 output is set causing the output of OR gate 38 to go high (+ARB) indicating an arbitration timing period. The OR gate 38 output remains high until the output of counter 32 goes low sometime after the output of counter 33 has gone high. This establishes a 300 nanosecond timing pulse for the +ARB/-GNT signal. The output from counter 33 remains set until the device begins a bus cycle by either activating -S0 or -S1. The output is then reset, and the counters 30 through 33 are ready to begin timing again at the end of the current bus cycle. If -BURST is active, the timing chain is held reset until the end of the last cycle. If no devices are requesting bus service, the bus returns to the idle condition and control is returned to the processor. HLDA is reactivated and the bus is now available for processor operations.

Referring now to FIG. 8, a circuit description of an alternate embodiment of the central arbitration control circuit 11 is illustrated. The central arbitration control circuit 11 is comprised of a D flip-flop 70, the Q output of which represents the +ARB/-GNT signal, a \bar{Q} output which represents a signal to begin processor HOLD and REFRESH memory on the positive edge, a D input connected to the output of AND gate 72, and a clock input connected to an arbitration clock having a 300 nanoseconds high/7.5 microsecond low timing waveform. The flip-flop 70 also has a clear input connected to an output of the NAND gate 74. The AND gate 72 has one input connected to the -BURST line of the arbitration bus and its other input connected to an inverter 76 which has its input connected to the -PREEMPT line of the arbitration bus. NAND gate 74 has one input connected to an inverter 78 which has its input connected to the arbitration clock, and a second input connected to the +REFRESH signal from the refresh controller and processor HOLD COMPLETE (pHLDA) signal generated by the CPU 10 of FIG. 1. The central arbitration circuit 11 operates as follows: When the -PREEMPT signal becomes active, the output of the inverter 76 becomes high. If the BURST signal is not active then the output of the NAND gate 72 will go high setting the data input of the arbitration latch 70 high. The next positive edge of the arbitration clock steps the high level on the input to the latch 70 thereby moving the Q output (+ARB/-GNT) signal high. If the bus has been owned by the processor 10 the HOLD request to the processor 10 is issued at the \bar{Q} bar output of the latch 70. Additionally, if a refresh cycle to memory is due it can also be issued at this time. After 300 nanoseconds, the arbitration clock will go low and the inverter 78 will generate a high level input to the NAND gate 74. After REFRESH is complete and processor HOLD operations are completed (as indicated by pHLDA signal from the CPU 10 going active), the other input of the

NAND gate 74 goes high and thus the output of the NAND gate 74 goes low. The low output from the NAND gate 34 resets the arbitration latch 70 and sets the +ARB/-GNT line to the GRANT state completing the operation. If the -PREEMPT line remains active, another arbitration cycle will be inserted on the next positive transition of the arbitration clock unless -BURST is asserted by the local arbiter that wins the competition for the GRANT. If -BURST is asserted by the granted arbiter then arbitration will be inhibited indefinitely. The present system therefore imposes a rule that -BURST may not be held longer than approximately 8 microseconds after -PREEMPT is activated. -PREEMPT becomes active periodically approximately an average 16 microseconds for REFRESH. The net result is that the arbitration is not suspended longer than 24 microseconds worst case.

Referring now to FIG. 6, a logic diagram of one of the arbitration circuits 28 is illustrated. The arbitration level assigned to a particular arbitrating device is set in a register not shown, but which stores the values ARB3 through ARB0. This may be done using any of a number of well known techniques. Preferably, the CPU 10 addresses the particular device associated with the local arbiter 28 through a pre-assigned port so that the arbitration level can be set by software. This may be done through the BIOS (basic input-output system), POST (power on self-test) at reset, the operating system, or the applications programs desired. Otherwise, it is possible to input the channel priority assignment value with hardware switches.

It should be understood that the portion of the local arbitration circuit 40, as illustrated in FIG. 6 which actually compares the priority level of its associated device with the priority level being driven on the priority arbitration bus is identical to that which is disclosed and discussed in detail in American National Standards/IEEE Standard Number ANSI/IEE Standard 696-1983, published by the Institute of Electrical and Electronic Engineers, Inc., Jun. 13, 1983. In general, each device incorporated in the arbitration scheme is provided with such a circuit, and the circuits are interconnected to each other via an ARB BUS. The ARB BUS in the example shown has four data lines ARB0-ARB3, allowing for sixteen different arbitration levels. Any desired number though can be used. During an arbitration time period designated by the +ARB/-GNT signal, all devices desiring to gain control over the bus (those having the +COMPETE signal set to the appropriate state) transmit their arbitration (priority) levels onto the ARB BUS. This occurs when the +ARB/-GRANT signal goes high. The value then held in the priority register is gated to the NAND gates 42 onto the respective line of the ARB BUS. At the end of the arbitration period, the value on the ARB BUS lines ARB0 through ARB3 will be the value of the arbitration level of the device having the highest level. The output of the AND gate 44 of the winning device will indicate a BUS WON indication to its corresponding device. It should be understood that the local arbiter 28 of the present invention begins the arbitration cycle in response to the +ARB/-GRANT signal, as opposed to the HOLD and pHLDA signals illustrated in the IEEE article.

It should be further understood that whenever the channel has not been granted to the CPU 10 that it, like other devices, will not be active on the channel. When

the CPU 10 is not active on the channel it may be held inactive entirely or may be active on its local bus.

A better understanding of the operation of the local arbiter 28 may be appreciated by the following example. For purposes of the example, assume that the priority level in the priority register is 6, and therefore the ARB 3 bit would be equal to 0, the ARB 2 bit would be equal to 1, the ARB 1 bit would be equal to 1 and the ARB 0 bit would be equal to 0. This value is inverted to 1001 by the inverters 39. Now assume that the local arbiter is competing against another arbiter which has an assigned arbitration priority level of 7. The device associated with the local arbiter 28 having a priority level of 6 generates a -DEVICE ARB REQUEST for use of the bus, which sets the Preempt latch 50 and -PREEMPT active through the inverter 51, and waits for an arbitration cycle. The central arbitration control circuit 11 responds to the active -PREEMPT and the +ARB/-GNT line goes into the arbitration state as explained earlier. The COMPETE latch 52 is set by the Q output of the Preempt latch 50 on the positive edge of the ARB/-GRANT signal driving the +COMPETE signal active. The other competing arbiter sets ARB3 equal to 0, ARB2 equal to 1, ARB1 equal to 1, ARB0 equal to 1. The first NAND gate 42 associated with the ARB3 line drives the ARB3 line low because COMPETE equals 1 and the inverted ARB3 bit equals 1. The OR gate 45 does not disable lower OR stages because one of its inputs is high, namely the inverted ARB2 bit. If neither of the inputs to this gate were 1, then all lower order NAND gate drivers would be disabled. The NAND gate 42 associated with the ARB2 bus line sets its output high because the inverted ARB2 bit is low. The OR gate 46 therefore does not disable the lower order stages because one of its outputs is high (the ARB2 bit is high on the arbitration bus). If this bit on the arbitration bus were driven low by an opposing arbiter then all lower order stages of this arbiter would be disabled and this arbiter would not win the competition for the channel. The NAND gate 42 associated with the ARB1 bus line sets its output high because the inverted ARB1 bit is low. The OR gate 47 therefore does not disable lower order stages because one of its inputs is high (the ARB1 bit is high on the arbitration bus). The NAND gate 42 associated with the ARB0 bus line drives the ARB0 line low because COMPETE equals 1 and the inverted ARB0 bit equals 1, and it is not disabled by any higher ordered compare. The OR gate 48 therefore does not disable lower ordered stages because one of its inputs is high (the inverted ARB0 bit is high). The opposing arbiter loses the competition at this point because its ARB0 bit is high and the ARB0 bus line signal is low. Both inputs to its OR gate 48 are low and the AND gate 44 is therefore disabled. After a prescribed settling time the +ARB/-GRANT goes low to the GRANT state. Because no stage in the arbiter has disabled an input to AND gate 44 because the +ARB/-GRANT signal is low, COMPETE is high, the AND gate 44 output is high indicating a BUS WON (see ACK of FIG. 3) signal. The Preempt latch 50 is clocked by the BUS WON signal, and the -PREEMPT signal becomes inactive. The operation of the arbitration circuits 28 as explained so far only applies to those devices which generally only transmit one word of data before relinquishing control of the bus. However, as mentioned earlier, some of the devices in the system such as the floppy disk 22 and 23 and the hardfile 20 are able to transmit bursts of data, and thus

the local arbitration circuitry must take the burst transfers into account. Thus, if a burst transfer is requested by a device, then the Burst latch 54 is set by the BUS WON signal and the -BURST signal is asserted from the inverter 56. Normally, as long as the -BURST signal is asserted from the inverter 56, additional arbitration cycles are inhibited. The Burst latch 54 is deactivated when the device gets to its end of transfer or a -STOPBURST signal condition is asserted on the clear input of the Burst latch 54.

Still referring to FIG. 6, it is readily apparent that if a bursting device is assigned a very high priority number it could maintain control of the channel 25 for inordinate amounts of time and possibly cause overrun conditions in other devices. Thus, it becomes necessary to implement a fairness circuit which forces the high priority bursting device off of the channel 25 at an appropriate time and to allow other lower priority devices access to the channel 25 in order to prevent the overrun conditions. In order to implement fairness the AND gate 58 is added to produce an output signal when the local arbiter 28 is driving the -BURST signal and the -PREEMPT signal is also active. For this example, -PREEMPT is not set for a future cycle while making a burst transfer on the channel. The existence of this condition indicates that a device is requesting arbitration but the arbitration signal is being withheld by the central arbitration control circuit 11 as long as -BURST is driven by the local arbiter 28.

In order to fully implement the fairness technique, the circuitry of FIG. 7 must be added. The fairness circuit includes a NAND gate 60 having one input connected to the +DEVICE ARB signal and a second input connected to the Q output of the latch 62. The latch 62 has its set input connected to the output of AND gate 64 and its clear input connected to the output of the AND gate 66. The AND 64 has one input connected to the line which receives a fairness ON or OFF bit which is normally generated by the power on self test (POST) software. Of course other applications software could be used to dynamically change the fairness bit according to the needs of the particular program. The other input to the NAND gate 64 is connected to the output of the NAND gate 58. AND gate 66 has three inputs, namely the status bits S0 and S1 and the -PREEMPT signal. The status bits are those commonly generated by the CPU 10 and are well known to those skilled in the art. They become inactive (high) during idle periods between bus cycles. The fairness circuit of FIG. 7 operates as follows:

When the AND gate 64 detects a positive fairness bit (typically during POST) and the other input to the AND gate 64 detects a "+PREEMPT PENDING" condition from the output of the AND gate 58, the latch 62 is set. Any device requests for arbitration are then interrupted by the NAND gate 60 and are not propagated to the Preempt latch 50 whenever the latch 62 is set. When the -PREEMPT signal becomes inactive (high) indicating that no other devices are requesting an arbitration cycle, and the status bits S0 and S1 are inactive (high) the latch 62 is cleared. It should be understood that the active period in status is long enough that the -PREEMPT line has time to charge to a high level so that all local arbiters 28 monitoring the -PREEMPT signal will be allowed to compete during the next arbitration cycle. After the latch 62 has been cleared the request for arbitration from each device is then initiated through the Preempt latch 50 as previ-

ously explained. It should be noted that for all conditions, that as long as the +ARB/-GRANT signal is low, the arbitration bus is stable and the BUS WON signal is stable. If the +ARB/-GRANT signal goes high, BUS WON goes inactive within 2 gate delays. The Compete latch 52 will not be set unless the -PREEMPT signal is previously driven by the local arbiter 28. The local arbiter 28 will continue to drive the -PREEMPT signal until it eventually wins the competition. The -PREEMPT signal is therefore a global indication of arbitration requests throughout the system.

It is to be understood that while modifications can be made to the structures and teachings of the present invention as described above, such modifications fall within the spirit and scope of the present invention as specified in the claims appended hereto.

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

1. An arbitration subsystem for a computer system in which plural devices, having predetermined priorities for access to a channel and requiring access thereto, arbitrate concurrently, upon receipt of a GRANT signal presented concurrently to all said devices, with the arbitrating device having highest priority assuming control of said channel to transmit or receive information; said subsystem comprising:

means in at least one said device for setting programmably the respective device to operate relative to other said devices in either a fair arbitration mode or a linear arbitration mode;

condition indicating means in each said device for indicating a preempt condition to other said devices when the respective device requires access to said channel;

monitoring means in each said device for monitoring preempt conditions indicated by other said devices; and

arbitration logic in each said at least one device for conditionally arbitrating for access to said channel, relative to other said devices requiring access to said channel, when the respective device requires access to said channel and said GRANT signal is received; wherein, the arbitration logic in each said at least one device programmably settable to said fair arbitration mode is required to meet the following additional condition before arbitrating for control of said channel when the respective device is set to said fair arbitration mode:

said preempt condition monitored by said monitoring means in the respective device must be in a predetermined state; and

wherein each said at least one device settable to said fair arbitration mode is operable to transfer information relative to said channel in single (fixed length) transfer and burst (variable length) transfer modes, and additional condition which must be met before the respective device's arbitration logic can arbitrate while the device is in said fair mode is:

the respective device is operating in said burst mode, and said monitored preempt condition was either inactive when the respective device last accessed said channel or has become inactive since said last access.

2. An arbitration subsystem for a computer system, in which plural devices having predetermined priorities for access to a channel, and requiring access thereto, arbitrate concurrently for said access to said channel

11

upon receipt of a GRANT signal presented concurrently by said system to all said devices, and in which the arbitrating device having highest priority for access assumes control of said channel to transmit or receive information thereon; said subsystem comprising: 5

means in at least one of said devices for programmably setting the respective device to operate in either a fair arbitration mode or a linear arbitration mode;

condition indicating means in each said device for continuously indicating a preempt condition to other said devices while the respective device requires access to said channel;

monitoring means in each said device for continuously monitoring preempt conditions indicated by other said devices: 15

arbitration logic in each said device for arbitrating relative to other said devices for access to said channel when the respective device receives said GRANT signal and simultaneously requires access to said channel; and 20

inhibit logic, in each said at least one device settable to said fair mode, for inhibiting operation of the

25

30

35

40

45

50

55

60

65

12

respective device's arbitration logic and preempt condition indicating means when the respective device is set to said fair arbitration mode and said preempt condition monitored by the monitoring means in said respective device has been continuously active since the respective device last had access to said channel; and

wherein each said at least one device operates in single (fixed length) transfer and burst (variable length) transfer modes, and said

inhibit logic in each said at least one device conditionally inhibits operation of the respective device's arbitration logic only when the following conditions are met:

- 1) the respective device is set to said fair mode; and
- 2) said preempt condition monitored by the respective device's said monitoring means has been continuously active since said respective device last had access to said channel; and
- 3) said respective device is operating in said burst mode.

* * * * *