



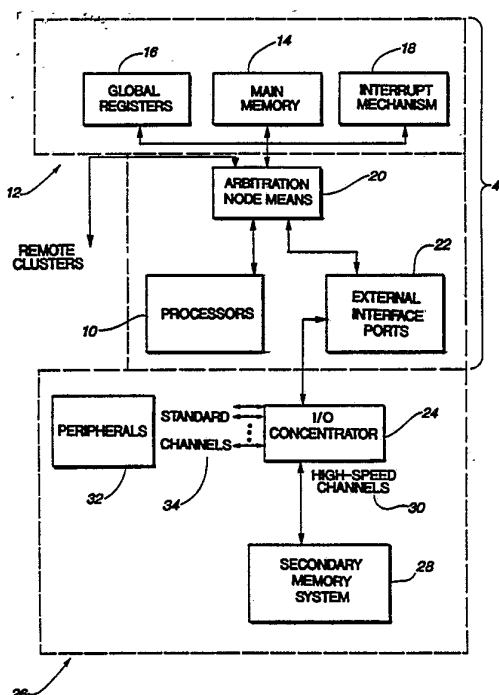
## INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification <sup>5</sup> : <b>G06F 13/00, 12/00, 15/16</b>		<b>A1</b>	(11) International Publication Number: <b>WO 91/20038</b>
			(43) International Publication Date: 26 December 1991 (26.12.91)
(21) International Application Number: PCT/US91/04054 (22) International Filing Date: 10 June 1991 (10.06.91) (30) Priority data: 535,786                      11 June 1990 (11.06.90)                      US (71) Applicant: SUPERCOMPUTER SYSTEMS LIMITED PARTNERSHIP [US/US]; 1414 W. Hamilton Avenue, Eau Claire, WI 54701 (US). (72) Inventors: WILSON, Jimmie, R. ; 3716 Partridge Run, Eau Claire, WI 54701 (US). BEARD, Douglas, R. ; S10505 Lowes Creek Road, Eleva, WI 54738 (US). CHEN, Steve, S. ; Route 5, Box 4300, Chippewa Falls, WI 54729 (US). ECKERT, Roger, E. ; 3130 Eldorado Blvd., #610, Eau Claire, WI 54701 (US). HESSEL, Richard, E. ; 3628 Altoona Avenue, Altoona, WI 54720 (US). PHELPS, And- rew, E. ; 6551 Hillview Road, Eau Claire, WI 54701 (US). SILBEY, Alexander, A. ; 2518 West Princeton Avenue, Eau Claire, WI 54703 (US). VANDERWARN, Brian, D. ; East 2535 Kirk Drive, Eau Claire, WI 54701 (US).		(74) Agents: PEDERSEN, Brad, D. et al.; Patterson & Keough, 615 Peavey Building, 730 Second Avenue South, Min- neapolis, MN 55402 (US). (81) Designated States: AT (European patent), AU, BE (Euro- pean patent), CA, CH (European patent), DE (Euro- pean patent), DK (European patent), ES (European pa- tent), FR (European patent), GB (European patent), GR (European patent), IT (European patent), JP, KR, LU (European patent), NL (European patent), SE (Euro- pean patent). Published With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.	

(54) Title: METHOD AND APPARATUS FOR NON-SEQUENTIAL RESOURCE ACCESS

## (57) Abstract

A method and apparatus for non-sequential access to shared resources (12) in a multiple requestor system uses a variety of tags to effectively re-order the data at its destination. In simplest form, the tag directs switching logic to where in a buffer to locate another tag for direction information or where in a buffer or processor (register) to put the response associated with the tag. For example, loading data from memory (14) requires that the requestor provides a request signal, an address, and a request tag. The request signal validates the address and request tag. The address specifies the location of the requested data in memory (14). The request tag specifies where to put the data when it is returned to the processor. The switching logic (44) for the requestor includes a tag queue for storing the request tags associated with the resource requests, logic means for associating the respective request tag from the tag queue with a resource response, and means for returning the resource response and respective request tag to the requestor. The switching logic (400) associated with the memory (14) includes switching means to route the request into and out of the shared resource, control logic to correctly route the requests, logic to handle multiple decision requests, and logic to store or retrieve the ultimate data entity being requested.



***FOR THE PURPOSES OF INFORMATION ONLY***

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	ES	Spain	MG	Madagascar
AU	Australia	FI	Finland	ML	Mali
BB	Barbados	FR	France	MN	Mongolia
BE	Belgium	GA	Gabon	MR	Mauritania
BF	Burkina Faso	GB	United Kingdom	MW	Malawi
BG	Bulgaria	GN	Guinea	NL	Netherlands
BJ	Benin	GR	Greece	NO	Norway
BR	Brazil	HU	Hungary	PL	Poland
CA	Canada	IT	Italy	RO	Romania
CF	Central African Republic	JP	Japan	SD	Sudan
CG	Congo	KP	Democratic People's Republic of Korea	SE	Sweden
CH	Switzerland	KR	Republic of Korea	SN	Senegal
CI	Côte d'Ivoire	LI	Liechtenstein	SU	Soviet Union
CM	Cameroon	LK	Sri Lanka	TD	Chad
CS	Czechoslovakia	LU	Luxembourg	TG	Togo
DE	Germany	MC	Monaco	US	United States of America
DK	Denmark				

5

10

## METHOD AND APPARATUS FOR NON-SEQUENTIAL RESOURCE ACCESS

### TECHNICAL FIELD

This invention relates generally to the field of memory systems and  
15 memory management for computer and electronic logic systems. More  
particularly, the present invention relates to a method and apparatus for  
non-sequential memory access capable of issuing requests to shared  
resources, particularly main memory, for which the responses may be  
returned out-of-order relative to the sequence in which the requests were  
20 issued.

### BACKGROUND ART

Traditional methods and apparatus in multiple requestor systems  
for accessing shared hardware resources, and in particular main memory,  
25 have required that requests to the shared resource maintain a  
time-ordered or sequenced relationship with one another to prevent  
overwriting or incorrect access to data or instructions by the system. In  
fact, the possibility of an out-of-order access to a shared resource is  
normally thought of as an incorrect access to data and/or instructions and  
30 is referred to in the prior art as a memory access hazard.

In the prior art, one technique used to control access to shared  
resources in a multiprocessor system is to maintain a centralized control  
mechanism that handles requests to the shared resources with a global  
decision algorithm. When the multiprocessor system has few requestors  
35 and few shared resources in a physically small system, this approach is  
viable. In larger multiprocessor systems with many requestors and many  
shared resources, the centralized approach using a global decision

algorithm becomes unwieldy and the decision time begins to impact the overall processing performance of the multiprocessor system.

Another prior art technique is to interlock the various decisions with time tags so that a given decision will be made only when previous  
5 decisions have cleared. The problem with this approach is that the transient time for signals to travel around the system updating these time tags becomes prohibitive and, again, the overall system performance is adversely affected.

In essence, the problem of resource lockout and shared resource  
10 access has been managed in prior art supercomputers by employing centralized control mechanisms that schedule each resource sequentially throughout the multiprocessor system. This approach effectively avoids the problem of memory access hazards at the expense of system performance.

15 In designing new methods and systems for accessing shared resources in a multiprocessor system, there are four important issues to be considered. The first issue is how to maximize the throughput of the shared resources. The second issue is how to maximize the bandwidth between the processors and the shared resources. The third issue is how to  
20 minimize the access time between the processors and shared resources. The last issue is how to avoid access hazards so that predictable results are obtained for resource requests. When a computer processing system does not have optimized solutions to all of these issues, the performance of the system is effectively limited.

25 For example, consider the problem of a processor making three requests to main memory, each request destined for a different section of main memory that, as a result, will be processed by three separate logic mechanisms. In the prior art, each of the requests is required to be consecutive and subsequent requests can not begin until the previous  
30 request is completed. In the three request example, the requirement for sequential access effectively results in two-thirds of the logic associated with the main memory being idle. This limitation is particularly damaging in high performance systems where the objective is to keep the processing elements and shared resources continually busy by providing  
35 them with a steady stream of work to be done in the form of pipelined operations.

In an effort to increase the processing speed and flexibility of such high performance computer processing systems, often referred to as supercomputers, the previously filed parent application to the present invention entitled CLUSTER ARCHITECTURE FOR A HIGHLY  
5 PARALLEL SCALAR/VECTOR MULTIPROCESSOR SYSTEM, PCT Serial No.: PCT/US90/07655, provides an architecture for supercomputers wherein a multiple number of processors and external interfaces can make multiple and simultaneous requests to a common set of shared hardware resources, such as main memory, global registers or interrupt  
10 mechanisms. The prior art techniques that utilize centralized control mechanisms to schedule each resource sequentially throughout the multiprocessor system are unacceptable for this type of cluster architecture for highly parallel multiprocessors. Consequently, a new method and apparatus for memory access is needed to increase performance in a  
15 multiprocessor system by insuring equal and democratic access for all requestors across all shared resources and allowing each shared resource to process data at independent rates, and at the same time avoiding access hazards.

## 20 SUMMARY OF THE INVENTION

The present invention is a method and apparatus for non-sequential access to shared resources in a multiple requestor system. To accomplish this, the present invention uses a variety of tags to effectively re-order the data at its destination. In simplest form, the tag  
25 directs switching logic to where in a buffer to locate another tag for direction information or where in a buffer or processor (register) to put the response associated with the tag. For example, loading data from memory requires that the requestor provide a request signal, an address, and a request tag. The request signal validates the address and request tag. The  
30 address specifies the location of the requested data in memory. The request tag specifies where to put the data when it is returned to the processor.

The apparatus for non-sequential shared resource access in accordance with the present invention is used in a multiprocessor system  
35 having a plurality of processors and a plurality of input/output interfaces that may also access the shared resources of the multiprocessor system. The shared resources in the preferred embodiment include main memory,

global registers and interrupt mechanisms. For non-sequential access within a cluster of tightly-coupled processors, the present invention comprises a request generation means for generating a plurality of resource requests from the processors, switching means operably  
5 connected to the request generation means for receiving the resource requests in a time order in which the resource requests were generated and routing the resource requests to the shared resources, and shared resource means servicing the resource requests as the requested resources become available. Each of the resource requests includes an address for a requested  
10 shared resource, and a request tag designating a location within the requestor where the resource request is to be returned.

The switching logic associated with the switching means includes a tag queue for storing the request tags associated with the resource requests, logic means for associating the respective request tag from the tag queue  
15 with a resource response, and means for returning the resource response and respective request tag to the processor. The switching logic associated with the shared resource includes switching means to route the request into and out of the shared resource, control logic to correctly route the requests, logic to handle multiple decision requests, and logic to store or  
20 retrieve the ultimate data entity being requested.

In an alternative embodiment, the present invention also permits non-sequential access to shared resources external to a tightly-coupled cluster of processors. In this embodiment, a remote cluster adapter is provided with logic means to attach additional routing information to the  
25 request tag to create a new tag which is called the cluster tag. The cluster tag is passed to the target cluster's remote cluster adapter where the cluster tag is stripped from the request and stored in a tag buffer. A new request tag is generated for use internal to the target cluster. When the response is returned to the target cluster's remote cluster adapter the returned request  
30 tag is used to locate the corresponding cluster tag associated with it in the tag buffer. The response and the cluster tag are then returned to the requesting cluster. At the requesting cluster's remote adapter the cluster tag is broken apart into its parts, additional return routing information and request tag. The additional return routing part is used to steer the  
35 response and the requesting tag part back to the requesting switch means.

It is a particular consequence of this system that control is handled locally and decisions concerning access to shared resources are made

quickly and only at the time necessary to insure full utilization of the resource. By maintaining the highest possible system throughput coupled with a very high system bandwidth, the present invention insures that requestors are constantly supplied with data for processing in a minimal  
5 access time, thereby increasing the overall performance of a multiprocessor system for a given set of system bandwidth and throughput parameters.

An objective of the present invention is to provide a method and apparatus for non-sequential memory access capable of issuing requests to  
10 shared resources for which the responses may be returned out-of-order as compared with the time order in which the requests were issued.

Another objective of the present invention is to provide a method and apparatus for memory access that can increase performance in an interleaved shared resource system by insuring that each component of  
15 the shared resource system may operate concurrently in parallel and therefore potentially at different rates.

A further objective of the present invention is to provide a method and apparatus for memory access system for a multiprocessor system that delivers high bandwidth, high throughput, and low latency when the  
20 multiprocessor system is configured with many requestors and many shared resources.

These and other objectives of the present invention will become apparent with reference to the drawings, the detailed description of the preferred embodiment and the appended claims.

25

### DESCRIPTION OF THE DRAWINGS

Figs. 1a, 1b, 1c and 1d are representations of the pipelining request/response memory access techniques of the prior art and the present invention.

30 Fig. 2 is a block diagram of a single multiprocessor cluster of the preferred embodiment of the present invention.

Figs. 3a and 3b are a block diagram of a four cluster implementation of the preferred embodiment of the present invention.

35 Fig. 4 is a block diagram of a single multiprocessor cluster showing the arbitration node means of the preferred embodiment.

Figs. 5a and 5b are a detailed block diagram of the input/output interface in the preferred embodiment of the present invention.

Figs. 6a and 6b are a detailed block diagram of the ports associated with the input/output interface in the processor.

Figs. 7a, 7b, 7c, 7d and 7e are representations of the various request tags of the present invention.

5 Fig. 8 is a block diagram of the external interface ports.

Fig. 9 shows the correspondence between the request tag and the command block word.

Fig. 10 is a detailed block diagram of the NRCA means in the preferred embodiment of the present invention.

10 Figs. 11a and 11b are a detailed block diagram of the MRCA means in the preferred embodiment of the present invention.

#### **DESCRIPTION OF THE PREFERRED EMBODIMENT**

Referring now to Figs. 1a-1d, a representation of the pipelined,  
15 out-of-order access mechanism of the present invention as compared to the prior art is set forth. These figures are applicable to a request/response operation at each level of a memory/shared resource architecture. It will be understood that the shared resource generally accessed in a multiple requestor system is memory and the preferred embodiment of the present  
20 invention may be described in terms of accesses to memory; however, the present invention contemplates that the accesses are made to any type of shared hardware resource. Shared hardware resources in this sense will include memory, global registers, interrupt mechanism, as well as ports, paths, functional units, registers, queues, banks, etc.

25 Fig. 1a shows how a stream of requests and responses would be handled in a prior art system. Because there is no capability of out-of-order access or streaming, each consecutive request must wait for its associated response to be completed before the next request can be initiated. Referring now to Fig. 1b, some prior art vector processors support the ability to make  
30 consecutive requests to load or write a vector register without the need to wait for each response to return. The limited pipeline technique shown in Fig. 1b has been applied to vector processors accessing main memory, but has not been applied to other system resources.

In contrast, Fig. 1c shows an interleaved series of requests and  
35 responses in which all of the requests and their associated responses are time ordered, however, response 1 may be returned before request n is issued. In Fig. 1d the full capabilities of out-of-order access mechanism of



the present invention are illustrated with requests and their associated responses occurring with no particular respect to time ordering. In the memory access system shown here, response 2 is allowed to be returned prior to response 1. The pipeline techniques shown in Fig. 1d have not  
5 been applied in the prior art.

In describing the preferred embodiment, the presentation will begin with a description of the preferred embodiment of the multiprocessor system and then describe the method and apparatus for non-sequential access starting with a description of the various requestors and their  
10 associated ports and proceeding out into the shared resources of the multiprocessor system.

#### Multiprocessor System

Referring now to Fig. 2, the architecture of a single multiprocessor cluster of the preferred embodiment of the multiprocessor system for use  
15 with the present invention will be described. The preferred cluster architecture for a highly parallel scalar/vector multiprocessor system is capable of supporting a plurality of high-speed processors 10 sharing a large set of shared resources 12 (e.g., main memory 14, global registers 16,  
20 and interrupt mechanisms 18). The processors 10 are capable of both vector and scalar parallel processing and are connected to the shared resources 12 through an arbitration node means 20. Also connected through the arbitration node means 20 are a plurality of external interface  
ports 22 and input/output concentrators (IOC) 24 which are further  
25 connected to a variety of external data sources 26. The external data sources 26 may include a secondary memory system (SMS) 28 linked to the input/output concentrator 24 via a high speed channel 30. The external data sources 26 may also include a variety of other peripheral devices and  
interfaces 32 linked to the input/output concentrator 24 via one or more  
30 standard channels 34. The peripheral devices and interfaces 32 may include disk storage systems, tape storage system, printers, external processors, and communication networks. Together, the processors 10, shared resources 12, arbitration node 20 and external interface ports 22  
comprise a single multiprocessor cluster 40 for a highly parallel  
35 multiprocessor system in accordance with the preferred embodiment of the present invention.

The preferred embodiment of the multiprocessor clusters 40 overcomes the direct-connection interface problems of present shared-memory supercomputers by physically organizing the processors 10, shared resources 12, arbitration node 20 and external interface ports 22 into one or more clusters 40. In the preferred embodiment shown in Figs. 3a and 3b, there are four clusters: 40a, 40b, 40c and 40d. Each of the clusters 40a, 40b, 40c and 40d physically has its own set of processors 10a, 10b, 10c and 10d, shared resources 12a, 12b, 12c and 12d, and external interface ports 22a, 22b, 22c and 22d that are associated with that cluster. The clusters 40a, 40b, 40c and 40d are interconnected through a remote cluster adapter 42 that is a logical part of each arbitration nodes means 20a, 20b, 20c and 20d. Although the clusters 40a, 40b, 40c and 40d are physically separated, the logical organization of the clusters and the physical interconnection through the remote cluster adapter 42 enables the desired symmetrical access to all of the shared resources 12a, 12b, 12c and 12d across all of the clusters 40a, 40b, 40c and 40d.

Referring now to Fig. 4, the preferred embodiment of the arbitration node means 20 for a single cluster 40 will be described. At a conceptual level, the arbitration node means 20 comprises a plurality of crossbar switch mechanisms that symmetrically interconnect the processors 10 and external interface ports 22 to the shared resources 12 in the same cluster 40, and to the shared resources 12 in other clusters 40 through the remote cluster adapter means 42. Typically, a full crossbar switch would allow each requestor to connect to each resource. In the present invention, the arbitration node means 20 allows a result similar to a full crossbar switch to be achieved in the situation where there are more requestors than resources. In the preferred embodiment, the arbitration node means 20 is comprised of sixteen arbitration nodes 44 and the remote cluster adapter means 42. The remote cluster adapter means 42 is divided into a node remote cluster adapter (NRCA) means 46 and a memory remote cluster adapter (MRCA) means 48. The NRCA means 46 allows the arbitration node 44 to access the remote cluster adapter means 42 of all other multiprocessor clusters 40. Similarly, the MRCA means 48 controls access to the shared resources 12 of this cluster 40 from the remote cluster adapter means 42 of all other multiprocessor clusters 40.

In this embodiment, the sixteen arbitration nodes 44 interconnect thirty-two processors 10 and thirty-two external interface ports 22 with the

main memory 14, global register 16 and interrupt mechanism 18, and the NRCA means 46. Each arbitration node 44 is connected with the main memory 14 by eight bidirectional parallel paths 50. A single parallel bidirectional path 52 connects each arbitration node 44 with the NRCA means 46. In the preferred embodiment, the same path 52 from each arbitration node 44 is also used to connect the arbitration node 44 with the global registers 16 and the interrupt mechanism 18, although it will be recognized that separate paths could be used to accomplish this interconnection.

Like each of the arbitration nodes 44, the MRCA means 48 is connected with the main memory 14 by eight parallel bidirectional paths 54. Similarly, a single parallel bidirectional path 56 connects the MRCA means 48 with the global registers 16 and interrupt mechanism 18. In the preferred embodiment, a total of six parallel bidirectional paths 58 are used to interconnect the clusters 40. For example, cluster 40a has two paths 58 that connect with each cluster 40b, 40c and 40d. In this manner, the MRCA means 48 allows other clusters 40 to have direct access to the shared resources 12 of this cluster 40.

As shown in Figs. 5a and 5b, arbitration networks 303 and 306 for each of the memory ports 310, the NRCA port 312, and each of the processor ports 314, 315, 316, and 317, comprise the arbitration node 44. Also included in the arbitration node 44 are input port queue 301, crossbars 302 and 307, tag queue 304 and data queue 305. As explained in detail in the parent application, the arbitration networks 303 and 306 use a first-come-first-served, multiple-requestor-toggling system to insure that the oldest reference is processed first. In the case of multiple old references of the same age, a fairness algorithm ensures equal access to the ports 310 and 312 and 315, 316, and 317 that are controlled by that arbitration network 303 or 306, respectively.

In other supercomputers, memory returns come back in the same order the requests were sent out. Thus, there is no ambiguity in where to place the data when it returns, and the memory return logic for the processor is simple. However, restricting memory returns to be in-order also means sacrificing performance because sequencing constraints would cause the shared resources to wait until they were clear of any ordering violations, therefore reducing the amount of concurrent activity. Early returns are those that come back with shorter latency than previously

requested returns because of the nonhomogeneous latency of the memory system. If returns are not restricted to coming back in the same order as the requests were issued, the memory subsystem has to provide a sorting mechanism to guarantee this. This represents a considerable burden  
5 when there are multiple ports requesting data and multiple memory sections returning the data.

In the present invention, memory data returns may come back in a random order relative to the order of the requests. It is a characteristic of the queuing and arbitration networks throughout the multiprocessor  
10 system that whenever requests accumulate in queues, the responses may be returned out of order with respect to the relative time in which they were first entered into the queue. This means that data for early requests may come back later than data for late requests. However, it may not be possible to use the early-arriving data because of sequencing restrictions in  
15 the processor. For example, data involved in arithmetic must be used in the same sequence specified in the original program if reproducible results are desired. Therefore, in the present invention, memory return data is placed in its final destination (a vector register, scalar register, L register, instruction cache buffer, or input/output buffer) without regard to  
20 whether it can actually be used yet. Whether it can be used is based on the return status of previously issued requests.

The present invention provides a method and apparatus for shared resource access wherein all system resources may be accessed using all of the pipeline techniques as shown in Figs. 1b-1d. To accomplish this, the  
25 present invention uses tags and queues to keep track of requests and responses, to effectively re-order the data at its destination. In simplest form, a tag tells the logic where in the buffer to locate another tag for direction information or where in the buffer or processor (register) to put the response associated with the tag. For example, requesting data from  
30 memory requires that the requestor provide a request signal, an address, and a request tag. The request signal validates the address and request tag. The address specifies the location of the response data in main memory. The request tag specifies where to put the data when it is returned to the processor. Although the description of the preferred embodiment is  
35 presented in the context of a multiprocessor system with each processor in the system having multiple ports for accessing shared resources, it will be evident that the present invention is equally applicable to a

multiprocessor system wherein each processor has only a single port to access memory, or a single processor system having multiple ports to access memory.

#### Processor

5       The following paragraphs describe in greater detail how the processor 10 of the preferred embodiment manages scalar, vector, and instruction requests. There are four vector load ports, one scalar load port, and one instruction load port. Each port has a different mechanism for issuing out-of-order returns.

10       Referring now to Figs. 6a and 6b, vector load ports 724, 726, 728, and 730 provide a request signal, an address, and a request tag for memory accesses whose data is destined for the vector registers 718. There are four vector load ports, and each load port supports two outstanding vector loads at any given time. Vector load ports make variable-size requests for  
15       memory returns for a particular vector register 718. Each request group can consist of between 1 and 64 individual requests. Of the two possible outstanding vector loads for each port, only one can be issuing requests at a time. However, because of the non-sequential nature of the memory returns, it is possible that all of the returns for the second load would  
20       arrive before any of the returns for the first load. The vector load port control mechanism 732 must ensure that the data is placed in the proper register, and that the data is not used before all previously issued data for that register is used.

      As shown in Fig. 7a, the request tag accompanying each memory  
25       request and address from a vector port is large enough to indicate the destination of the memory return, i.e. which of the two possible vector registers to put the data in and which location in that vector register to use. When the request is issued, the vector load port control ensures that the proper tag accompanies the request and address. The first component  
30       of the vector tag is the destination register identifier bit. In the preferred embodiment, instead of sending the destination register number as part of the tag, the processor simply sends a single bit indicating one of two registers. The specific register number is established upon instruction issue, and may be different each time the load instruction is issued to the  
35       port. Because only a single destination register identifier is included in the tag, only two loads may be outstanding at any given time. Tag collisions (reuse of the same destination register bit) within a port are prevented by

waiting for all returns for each register to return before reusing its tag bit. In addition, successive loads from each port always use different destination register tag bits. Another type of tag collision, where multiple loads are outstanding for each vector register (but in different vector load ports), is avoided by waiting for the first vector register load to complete before allowing another load for that register to begin.

The second part of the vector tag is the element number of the vector register. This simply indicates which word of a given register should be updated with the memory data.

10 The third, and final, part of the vector tag is the "cancel" indicator, which effectively nullifies a load request with respect to the processor 10. To the memory system, however, cancelled requests are treated just like non-cancelled requests, with a minor exception. The exception is that cancelled out-of-cluster requests are redirected in-cluster for performance reasons. The ability to cancel a request after it is issued results in a reduction of effective memory latency because it eliminates some time-critical handshaking between the portion of the processor 10 and arbitration node 44 that handles requests and the portion of the processor 10 and arbitration node 44 that handles addresses. In this embodiment, 15 the cancel bit allows for address validation and request arbitration to overlap and be accomplished in parallel within the processor 10 and the arbitration node 44. In other words, the arbitration node 44 proceeds to arbitrate for a memory request, for example, at the same time that address control logic within the processor 10 determines whether the address for the memory request is a valid request, i.e. is within the proper address range. If the processor 10 determines that the address is out-of-range, then 20 the cancel bit is set and the request is directed to an in-cluster address.

In one type of operation, the cancel feature is useful for allowing bottom-loading of data in loops. In this software technique, the effect of 30 memory latency is sidestepped by pre-fetching data before you know whether the addresses are valid. In the preferred multiprocessor system, a program may turn off mapping exceptions and access memory at random addresses (even addresses outside its data space) with impunity, and the cancel mechanism makes sure that the program is not allowed to access forbidden data. Thus, the program may access memory before validating 35 the addresses it uses, and the load instructions may be placed at the bottom of the loop ("bottom-loading") rather than the top.

Referring again to Figs. 6a and 6b, the scalar request tag will be described. A single scalar port 722 to memory provides a path for memory data to the S registers 714 and L registers 716. To ensure proper data sequencing, the scalar port, like the vector ports, accompanies each  
5 memory request and address with a request tag.

As shown in Fig. 7b, the scalar request tag indicates the destination register type and register number. The first part of the scalar tag is a register type. This distinguishes between L and S register returns. The second part of the scalar tag is a register number. This indicates which  
10 register number (L or S) to place the return data. Tag collisions are prevented by allowing only one outstanding reference per register, or in the case of L registers, one outstanding reference per register group. The third, and final, part of the scalar tag is the "cancel" indicator, which prevents user programs from accessing data which is outside their allowed  
15 address space. This mechanism is described elsewhere in this document.

As shown in Figs. 6a and 6b, the scalar load port 722 accepts returns destined for either the L registers 716 or the S registers 714. The L/S bit of the return tag determines which destination register type to write when the data returns from memory. When an S register is written, the request  
20 tag performs two functions. First, it causes the register to be unreserved, and this forms the write address for the S register file. When an L register is written, the tag performs the same two functions. However, the L registers are reserved and unreserved on a block basis rather than individual registers.

The request tag scheme outlined for scalar port returns allows L register loads to be interrupted by higher priority loads, such as S register loads. Since no ordering is imposed on the outgoing addresses, and none is imposed on the incoming data stream, it is possible to treat one set of register loads (e.g., L registers) as background activity, and treat another set  
30 (S registers) as foreground, high priority activity. With proper software assistance, then, the L registers can be treated as a software-managed cache accessed by block loads and stores.

Referring again to Figs. 6a and 6b, the instruction request tag will be described. The instruction and input/output port 720 provides a path  
35 from memory to the instruction cache 710 and input/output buffers 712. To ensure proper instruction sequencing, the instruction port, like the other ports, supplies a request tag along with the request and address. The

instruction request tag indicates the buffer number and element number in the buffer for the instruction return as shown in Fig. 7c.

Like the vector port register number, the instruction port buffer number indicator is encoded in a single tag bit. This restricts the number of outstanding buffer requests to two, but simplifies the control and data path associated with instruction loads. Instruction tag conflicts are prevented by waiting for the oldest buffer fill to complete before allowing a new one to start. Multiple outstanding loads for the same buffer are prohibited by the cache replacement policy. An encoding of the tag indicates whether the return destination is the instruction cache buffers 712 or the input/output return buffers 710.

Each vector load port 724, 726, 728, 730 maintains two sets of "back" bits, one for each possible outstanding vector register for that port. As a return comes back from memory, the back bit for that register and element number is set equal to 1 and the data is written into the vector register. However, the control mechanism 732 prevents the data from being used until all elements up to and including the previous element have been returned from memory. This ensures that all elements of a vector register are used in order even though the returns may come back out of order.

When all elements of a vector register have been returned from memory, the back bits for that side are marked "not back," and that set of back bits is made available for another load instruction. Vector and scalar load data whose tag indicates that the request has been cancelled, is discarded and replaced by a non-signalling NaN (Not a Number). This is how the preferred embodiment of the present invention protects memory data from access by programs that do not enable operand mapping exceptions. When instruction data returns from memory 14, the processor 10 uses the return tag to decide where to put it. Instruction returns, unlike data returns, do not use the cancel feature because instruction mapping exceptions are always enabled when in mapped mode. To avoid hazards associated with a non-coherent memory system, the processors 10 and external interface ports 22 require that the arbitration node 44 provide information beyond the tags. This information relates to the sequencing of requests and when those requests are committed to be processed by a particular shared resource. In the preferred embodiment, the technique



used to insure coherency is referred to as, the data mark mechanism. The data mark mechanism is disclosed in greater detail in the parent application.

5           External Interface Ports

In addition to requests issued by the processor 10, the present invention is also capable of servicing resource requests issued from peripheral devices issued through the external interface ports 22. Data fetched from main memory 14 or the global registers 16 by the external  
10 interface ports 22 can return in an order that is different than the order in which it was requested. To accomplish the non-sequential access of the present invention, both the IOCs 24 and the SMSs 28 associated with the external interface ports 22 append a request tag to resource requests made through the external interface ports 22.

15           Referring now to Fig. 8, a detailed description of the external interface ports 22 of the preferred embodiment is provided. The external interface ports 22 accepts packets of command and data words from the main memory 14 via a memory port cable (not shown) that physically connects the cluster channel interface (CCI) 120. Command words are  
20 placed in the command buffer 350 and data is routed into the data FIFO 360. The presence of a command in the command buffer 350 causes control logic 370 in the external interface ports 22 to request access to memory 14 through the arbitration node 44. Data from the word count, command, address, and mtag fields of the command word are loaded into the  
25 respective registers 382, 384, 386 and 388 in preparation for delivery to the arbitration node 44 when the request is acknowledged. A new request tag and address must be computed for every word request made.

For fetch requests, no data is sent but an address and request tag are sent for a number of requests equal to the contents of the command word  
30 count field. The request tags are computed starting with the lower six bits set to zero and incrementing the contents of that field until the proper number of tags have been sent. Similarly, addresses for requests are computed starting with the address presented in the command word and incrementing it as each request is acknowledged.

35           For store requests, the next word in the data FIFO 360 is presented along with the address, tag, and command information. The word count value is decremented after each request. No further requests are made

when the word count value reaches zero. FIFOs 350 and 360 are used to hold commands and data to ensure that, wherever possible, a command and data are always available at the external interface ports 22 to keep the arbitration node 44 constantly busy.

5        Fetched data returns from the shared resources through the transmit register 390. The request tag issued when the request was made is returned with the data. The output of the transmit register 390 is connected to the main memory 14. Control lines associated with the data lines in the cable linking the external interface ports 22 are inserted to  
10        indicate that a valid data word is on the bus for the CCI 120.

Input/output request tags are sent with the command word that precedes the data packet when it travels through the IOC 24. The tags contain a four-bit field indicating which buffer in the IOC 24 will receive the data, and a 6-bit field indicating the location in that buffer where the  
15        data will be stored. The four-bit buffer select field is decoded as shown in Fig. 9. The code 1011 directs data to one of the two SMTc command buffers. The other 6 tag bits indicate which buffer and in what location the returning data word is to be stored. The code 1111 is not used by the IOC 24. It is reserved for the processor instruction cache fetch operations. The IOC  
20        24 shares a memory port with the instruction cache.

The 6 bit field is generated at the external interface ports for each individual data word request when it is made. Requests are made in order starting with the lowest address. Referring to Fig. 7d, the input/output request tags, the 6-bit word identifier and the four bit destination  
25        identifier, are placed in the 10 bit tag field. This tag travels with the request through the memory system and is returned with each word of data to the IOC 24. The request tags are then used in the IOC 24 means to direct the data word to the appropriate buffer and location in that buffer.

Because tags are created in sequential order when requests are made,  
30        using tags to address locations in destination buffers ensures that data is always loaded into the buffer in proper sequence, even if data returns in any arbitrary order. Reading the data from the buffer in sequential order therefore guarantees that data is returned to the destination in proper order.

35

#### Arbitration Node

Referring again to Figs. 5a and 5b, as viewed from the perspective of the shared resource (in this case the paths 50 and 52), each incoming request is arbitrated for by a request arbitration network 303. A similar response arbitration network 306 arbitrates the returning data back to their respective processor ports 315, 316, or 317. For incoming requests, an input queue 301 holds up to sixteen requests that are waiting to be passed through the request arbitration network 303. For returning responses, a data queue 305 holds up to sixty-four responses waiting for the response arbitration network 306 to pass the returning data back into the destination port 315, 316, or 317. Each of the queues 301 and 305 are strategically sized to cover any control latency as data flows between requestors and shared resources. Also as data is returned from a memory section, its associated tag is retrieved from the tag queue 304, and re-attached, before data and tag are loaded into data queue 305. For returns via the NRCA path 52, data and its associated tag are already paired, the NRCA and MRCA means handle associating data and tags.

When the request arbitration network 303 determines that an incoming request is requesting available resources and has the highest priority and is destined for a memory section, the address and data components of that request are placed on the path 50 and routed to the proper memory section. When the request arbitration network 303 determines that an incoming request is requesting available resources and has the highest priority and is destined for the NRCA 46, the address, data, and tag components of that request are placed on the path 52 and routed to the proper shared resource 12. It should be noted that the arbitration network 303 is essentially controlling access to the interconnect wires 50 and 52. Subsequent arbitration networks further along in the request path control access to other shared resources. Data may be returned to the requesting ports 315, 316 and 317 in a different order than requested. The arbitration node 44 receives a set of tags with each load address and queues them for future reference. When data is returned from main memory, the tags are re-attached to the corresponding data word and both data and tags are passed back to the requesting port. The processor 10 uses these tags to place the data in the proper destination and ensure that it is used in the correct sequence.

### Main Memory

Again referring to Figs. 5a and 5b, for memory references the switching logic 400 for each memory section a subsection catch queue 401 in each memory subsection which collects all incoming requests to that particular memory subsection from a given arbitration node 44. Each arbitration node 44 has its own set of catch queues in each memory section 14. A bank request arbitration network 405 will arbitrate among its group of subsection catch queues 401 that have pending requests for that bank 403 on each cycle. Once the request is selected, the selected request is issued to the destination bank 403. If the request is a store, address and data are issued to the bank. If the request is a load, only the address is issued to the bank. If the request is a load&flag, the address and data are issued to the bank. For the load and the load&flag requests, the returning data from the bank 403, is held in a hold queue 406 before a return arbitration network 407 grants the outgoing responses from the memory section.

15

#### Remote Cluster Adapter

Referring now to Fig. 10, for the NRCA means 46 an input queue 610 or 630 collects incoming requests. The input queue 610 holds references destined for the external interface. Arbitration network 612 will arbitrate among its group of input queues 610 that have pending requests for the external resource on each cycle. Once the request is selected, the selected request is issued to the destination resource, with address, data, a new tag called the cluster tag that is made up the request tag and additional source information (See Fig. 7e) being placed on path 58. The input queue 630 holds references destined for the interrupt mechanism, global register, or SETN register. Arbitration network 634 will arbitrate among its group of input queues 630 that have pending requests for that resource 620, 632, or 633 on each cycle. Once the request is granted, it is issued to the destination resource. If data is to be returned to the arbitration node 44 from the global registers 633 or SETN registers 632, a high priority request is presented to the output arbitration network 615 which results in the output path being cleared back to the arbitration node.

Data and tags returning from the MRCA means 48 via port 58 are received into a queue 614. Data from the global registers 633 or SETN registers 632 are pre-arbitrated and immediately place on the return path 52 along with the associated tags. During each clock cycle the response arbitration network 615 arbitrates for the return data path for port 52 or 56.

The appropriate data is selected from the data queue 614 or the global registers 633 or SETN registers 632 and returned to the appropriate port 52 or 56.

Referring now to Figs. 11a and 11b, for the MRCA means 48 there are six ports through which store and load operations from other clusters are received, 520. Each of these ports consist of a receiving queue 500, a tag buffer 502, a return queue 504, and port control logic 501 and 503. Each port operates independently of all of the other ports. All operations that arrive at a port from another cluster are returned to that cluster. This includes stores as well as loads. The queues and arbitration for the MRCA means 48 (506, 507, 508, 509, 510, 511, and 512) essentially operate in a similar manner as the queues and arbitration for arbitration node 44 (301, 302, 303, 304, 305, 306, and 307, respectively).

When an operation arrives at an MRCA port 520 from an outside cluster, the data, address, and tag information are written into a receive queue 500 which is 64 locations deep. After a valid operation has been written into the receiving queue 500, port control logic 501 performs a resource check to determine if the operation can be passed into the MRCA means 48. There are three resources that are checked. The first resource that is checked concerns the availability of tags. As operations are passed into the MRCA means 48, the original cluster tag that arrived with the request is written into a tag buffer 502 and a new 8 bit request tag is generated by tag generator 501. The address of the location in the tag buffer 502 to which the original tag is written becomes the new request tag. This request tag must be unique, therefore once a new cluster tag has been generated and passed into the MRCA means 48 it can not be reused until that operation has returned back to the port from the MRCA means 48. The implementation of this logic requires that the request tags be generated in order. If the next request tag to be generated is still outstanding in the MRCA means 48, the port can not issue its next operation from the receive queue 500. The tag buffer 502 is 256 locations deep.

The second resource that must be checked before an operation can be issued to the MRCA means 48 concerns the availability of locations in the return queue 504. The MRCA means 48 has no mechanism to force the arbitration node 44 to hold a returning operation, therefore the MRCA means 48 must guarantee that there is always a location in the return

queue 504 to store any operation that returns from the arbitration node 44. This return queue 504 is 128 locations deep. Once all of the locations in the return queue 504 have been allocated, no other operations can be issued to the MRCA means 48 until locations become available.

5       The port queue 506 inside the MRCA means 48 is the third resource that must be checked before an operation can be issued from the receive queue 500. The port control logic 501 keeps a running total of the number of operations currently in the port queue 506. If the port queue 506 becomes full, the port control logic 501 must hold issue until locations  
10       become available.

      When an operation returns from the MRCA means 48, the data, if any, is stored directly into the return queue 504. The request tag that is returned along with the operation is used to access the tag buffer and recover the original cluster tag 503. This original cluster tag is extracted  
15       from the tag buffer 502 and stored into the return queue 504 with the data. The port control logic 501 then does a resource check on the cluster at the far end of the inter-cluster path 520. If the far-end cluster has available locations in its receive queue, the return queue 504 is unloaded. Otherwise data is held until queue locations become available.

20       Although the description of the preferred embodiment has been presented, it is contemplated that various changes could be made without deviating from the spirit of the present invention. Accordingly, it is intended that the scope of the present invention be dictated by the appended claims, rather than by the description of the preferred  
25       embodiment.

      We claim:

CLAIMS

1. An apparatus for non-sequential shared resource access in a  
5 multiprocessor system having a plurality of processors, the shared  
resources including main memory, global registers and interrupt  
mechanisms, the apparatus comprising:

request generation means operably connected to each of the  
processors for generating a plurality of resource requests from the  
10 processors, each of the resource requests comprising:

an address for a requested shared resource; and

a request tag designating a location within the  
processor where the resource request is to be returned;

switching means operably connected to the request  
15 generation means for receiving the resource requests in a time  
order in which the resource requests were generated and routing  
the resource requests to the shared resources, the switching means  
including:

tag queue means for storing the request tags associated  
20 with the resource requests; and

logic means for associating the respective request tag  
from the tag queue means with a resource response; and

means for returning the resource response and  
respective request tag to the processor; and

25 means operably connected to the switching means and the  
shared resources for servicing the resource requests as the requested  
resources become available and returning the resource response to  
the switching means in an order in which the resource requests are  
serviced,

30 whereby the resource responses may be returned out-of-order  
as compared with the time order in which the resource requests  
were issued.

2. The apparatus of claim 1 wherein the switching means further  
includes arbitration node means for arbitrating among the resource  
35 requests that will be routed to the shared resource on a given clock cycle of  
the multiprocessor system.

3. The apparatus of claim 2 wherein the logic means for associating the respective request tag further comprises cancel logic means for receiving a cancel indication from the request generation means and for cancelling the resource request prior to the time that the resource request  
5 is routed to the shared resource in response to the cancel indication.

4. The apparatus of claim 3 wherein the cancel logic means indicates that a resource request has been cancelled by returning a non a number value in response to the resource request.

5. The apparatus of claim 1 wherein the switching means further  
10 includes address verification means for verifying the validity of the address of the resource request.

6. The apparatus of claim 5 wherein the logic means for associating the respective request tag further comprises cancel logic means for receiving a cancel indication from the request generation means and for  
15 cancelling the resource request prior to the time that the resource request is routed to the shared resource in response to the cancel indication.

7. An apparatus for non-sequential shared resource access in a multiprocessor system having a plurality of requestors, the requestors including both processors and external interface ports, the shared  
20 resources including main memory, global registers and interrupt mechanisms, the apparatus comprising:

request generation means operably connected to each requestor for generating a plurality of resource requests from the requestor, each of the resource requests comprising:

25                   an address for a requested shared resource; and  
                  a request tag designating a location within the requestor where the resource request is to be returned;  
                  switching means operably connected to the request generation means for receiving the resource requests in a time  
30 order in which the resource requests were generated and routing the resource requests to the shared resources, the switching means including:

tag queue means for storing the request tags associated with the resource requests; and

35                   logic means for associating the respective request tag from the tag queue means with a resource response; and



means for returning the resource response and respective request tag to the requestor; and

5 means operably connected to the switching means and the shared resources for servicing the resource requests as the requested resources become available and returning the resource response to the switching means in an order in which the resource requests are serviced,

10 whereby the resource responses may be returned out-of-order as compared with the time order in which the resource requests were issued.

8. The apparatus of claim 7 wherein the switching means further includes arbitration node means for arbitrating among the resource requests that will be routed to the shared resource on a given clock cycle of the multiprocessor system and address verification means for verifying  
15 the validity of the address of the resource request.

9. The apparatus of claim 8 wherein the logic means for associating the respective request tag further comprises cancel logic means for receiving a cancel indication from the request generation means and for cancelling the resource request prior to the time that the resource request  
20 is routed to the shared resource in response to the cancel indication.

10. The apparatus of claim 7 wherein the requestors and shared resources are organized into a plurality of clusters and the switching means further comprises remote cluster adapter means associated with each cluster for receiving resource requests from requestors in the cluster  
25 which are directed to shared resources in a remote cluster and transmitting those requests to a remote cluster adapter means in the remote cluster and receiving resource responses from the remote cluster, and for receiving resource requests from remote cluster adapter means in the remote clusters which are directed to shared resources in the cluster  
30 and for returning the resource responses to the remote cluster.

11. A multiprocessor system comprising:  
means for establishing queues and pipelines between one or more requestors;  
means for making a request to a resource; and  
35 one or more resource means for responding to said requests such that said requests may be serviced in an order different from the time sequence in which said requests were made..

12. A method for accessing shared resources in a multiprocessor system having a plurality of processors, the shared resources including main memory, global registers and interrupt mechanisms, the method comprising the steps of:

- 5           generating a plurality of resource requests from one of the processors, each of the resource requests comprising:  
            an address for a requested shared resource; and  
            a request tag designating a location within the processor where the resource request is to be returned;
- 10           presenting the resource requests to a switching mechanism associated with the shared resources in a time order in which the resource requests were issued;  
            storing the request tags associated with the resource requests in a tag queue in the switching means;
- 15           servicing the resource requests as the requested resources become available to produce a resource response;  
            returning the resource response to the switching means in an order in which the resource requests are serviced;  
            associating the respective tag from the tag queue with the resource response; and
- 20           returning the resource response and respective tag to the processor,  
            whereby the resource responses may be returned out-of-order as compared with the time order in which the resource requests were issued.
- 25

13. The method of claim 12 wherein the step of servicing the resource requests comprises the steps of :

- arbitrating among the resource requests that will be routed to the shared resource on a given clock cycle of the multiprocessor system; and
- 30           verifying the validity of the address of the resource request.

14. The method of claim 13 wherein the step of servicing the resource request further includes the step of checking for a cancel indication associated with the resource request and for cancelling the resource request prior to the time that the resource request is routed to the shared resource in response to the cancel indication.

35

15. The method of claim 14 wherein the step of cancelling the resource request indicates that the resource request has been cancelled by returning a non a number value in response to the resource request.

16. A method for accessing shared resources in a multiprocessor system  
5 having a plurality of requestors, the requestors including both processors and external interface ports, the shared resources including main memory, global registers and interrupt mechanisms, the method comprising the steps of:

10 generating a plurality of resource requests from one of the requestors, each of the resource requests comprising:

an address for a requested shared resource; and

a request tag designating a location within the requestor where the resource request is to be returned;

15 presenting the resource requests to a switching mechanism associated with the shared resources in a time order in which the resource requests were issued;

storing the request tags associated with the resource requests in a tag queue in the switching means;

20 servicing the resource requests as the requested resources become available to produce a resource response;

returning the resource response to the switching means in an order in which the resource requests are serviced;

associating the respective tag from the tag queue with the resource response; and

25 returning the resource response and respective tag to the processor,

whereby the resource responses may be returned out-of-order as compared with the time order in which the resource requests were issued.

30 17. The method of claim 16 wherein the step of servicing the resource requests comprises the steps of :

arbitrating among the resource requests that will be routed to the shared resource on a given clock cycle of the multiprocessor system; and

35 verifying the validity of the address of the resource request.

18. The method of claim 17 wherein the step of servicing the resource request further includes the step of checking for a cancel indication

associated with the resource request and for cancelling the resource request prior to the time that the resource request is routed to the shared resource in response to the cancel indication.

19. The method of claim 16 wherein the requestors and shared  
5 resources are organized into a plurality of clusters and the step of servicing the resource requests further comprises:


receiving remote resource requests from requestors in the cluster which are directed to shared resources in a remote cluster; and

- 10 transmitting the remote resource requests to a remote cluster and receiving resource responses from the remote cluster.


20. The method of claim 19 the step of servicing the resource requests further comprises:

- 15 receiving remote resource requests from the remote clusters which are directed to shared resources in the cluster; and  
returning the resource responses to the remote cluster for the remote resource requests.


**Fig. 1a**




**Fig. 1b**



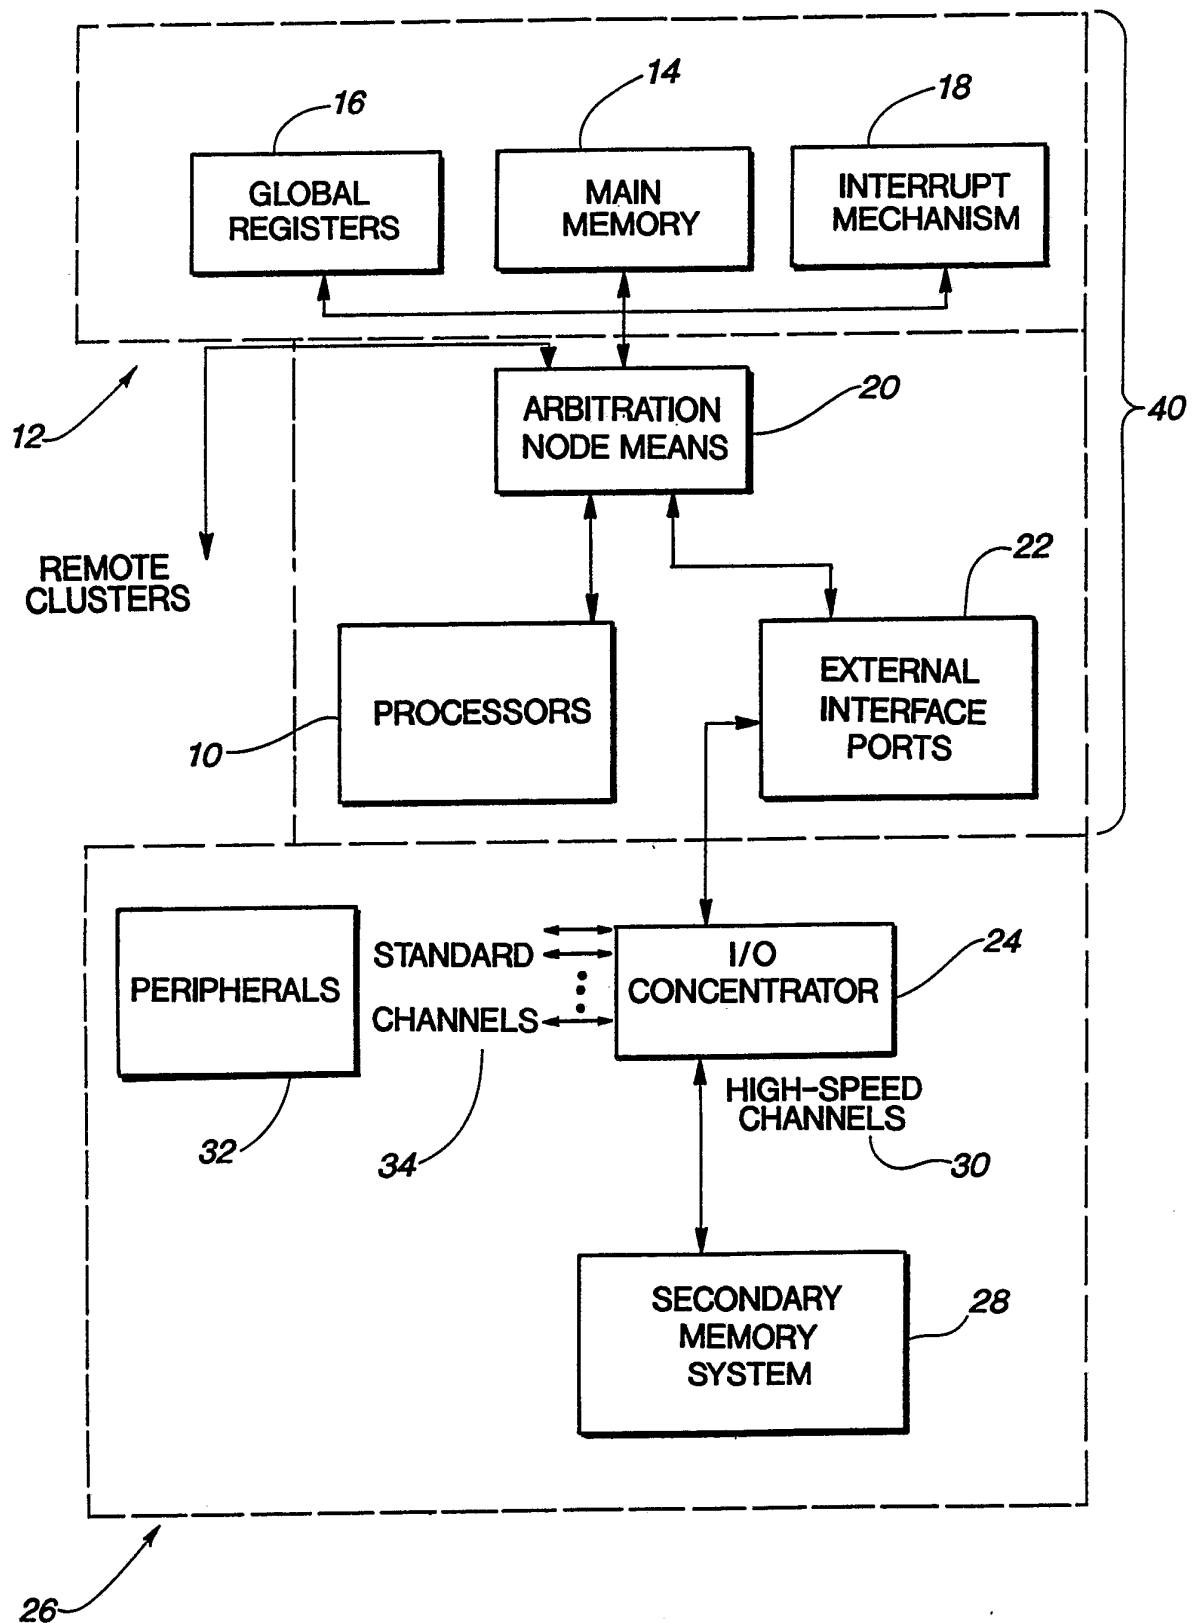
**Fig. 1c**




**Fig. 1d**

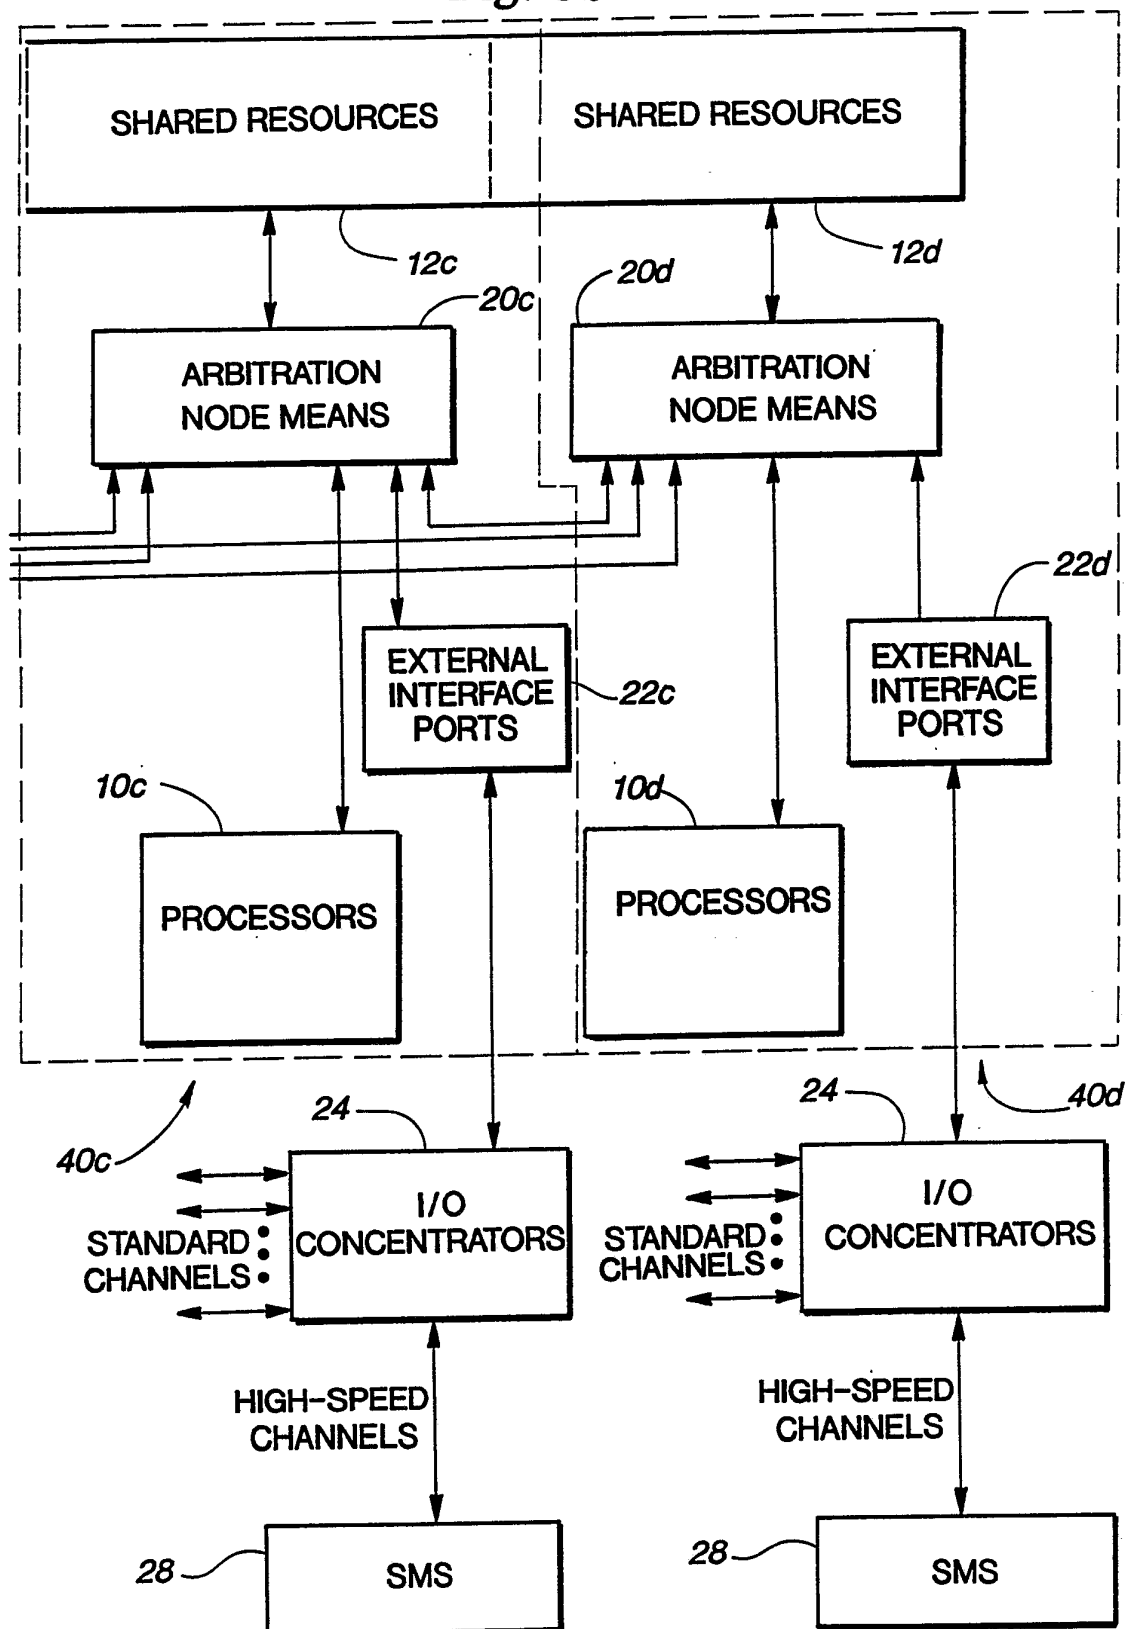


2/17

*Fig. 2*

3/17


*Fig. 3a*

4/17  
*Fig. 3b*




5/17


*Fig. 5*




*Fig. 6*

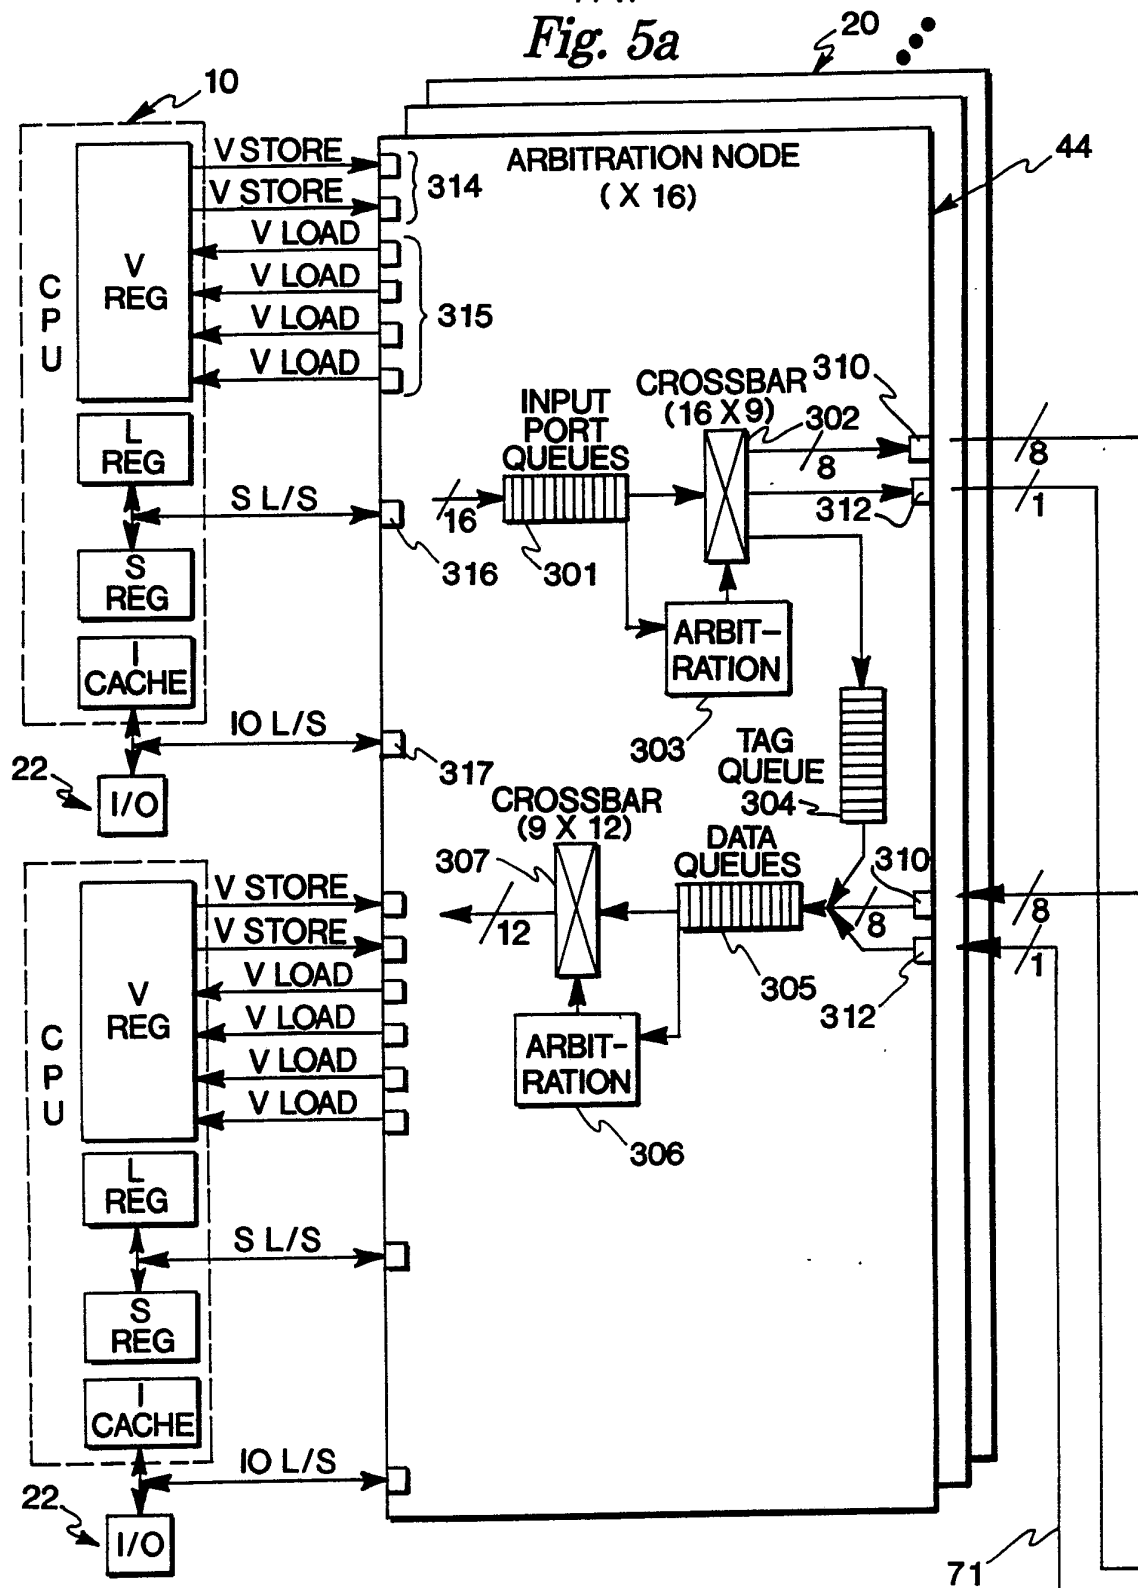


*Fig. 11*



6/17



7/17  
Fig. 5a

8/17

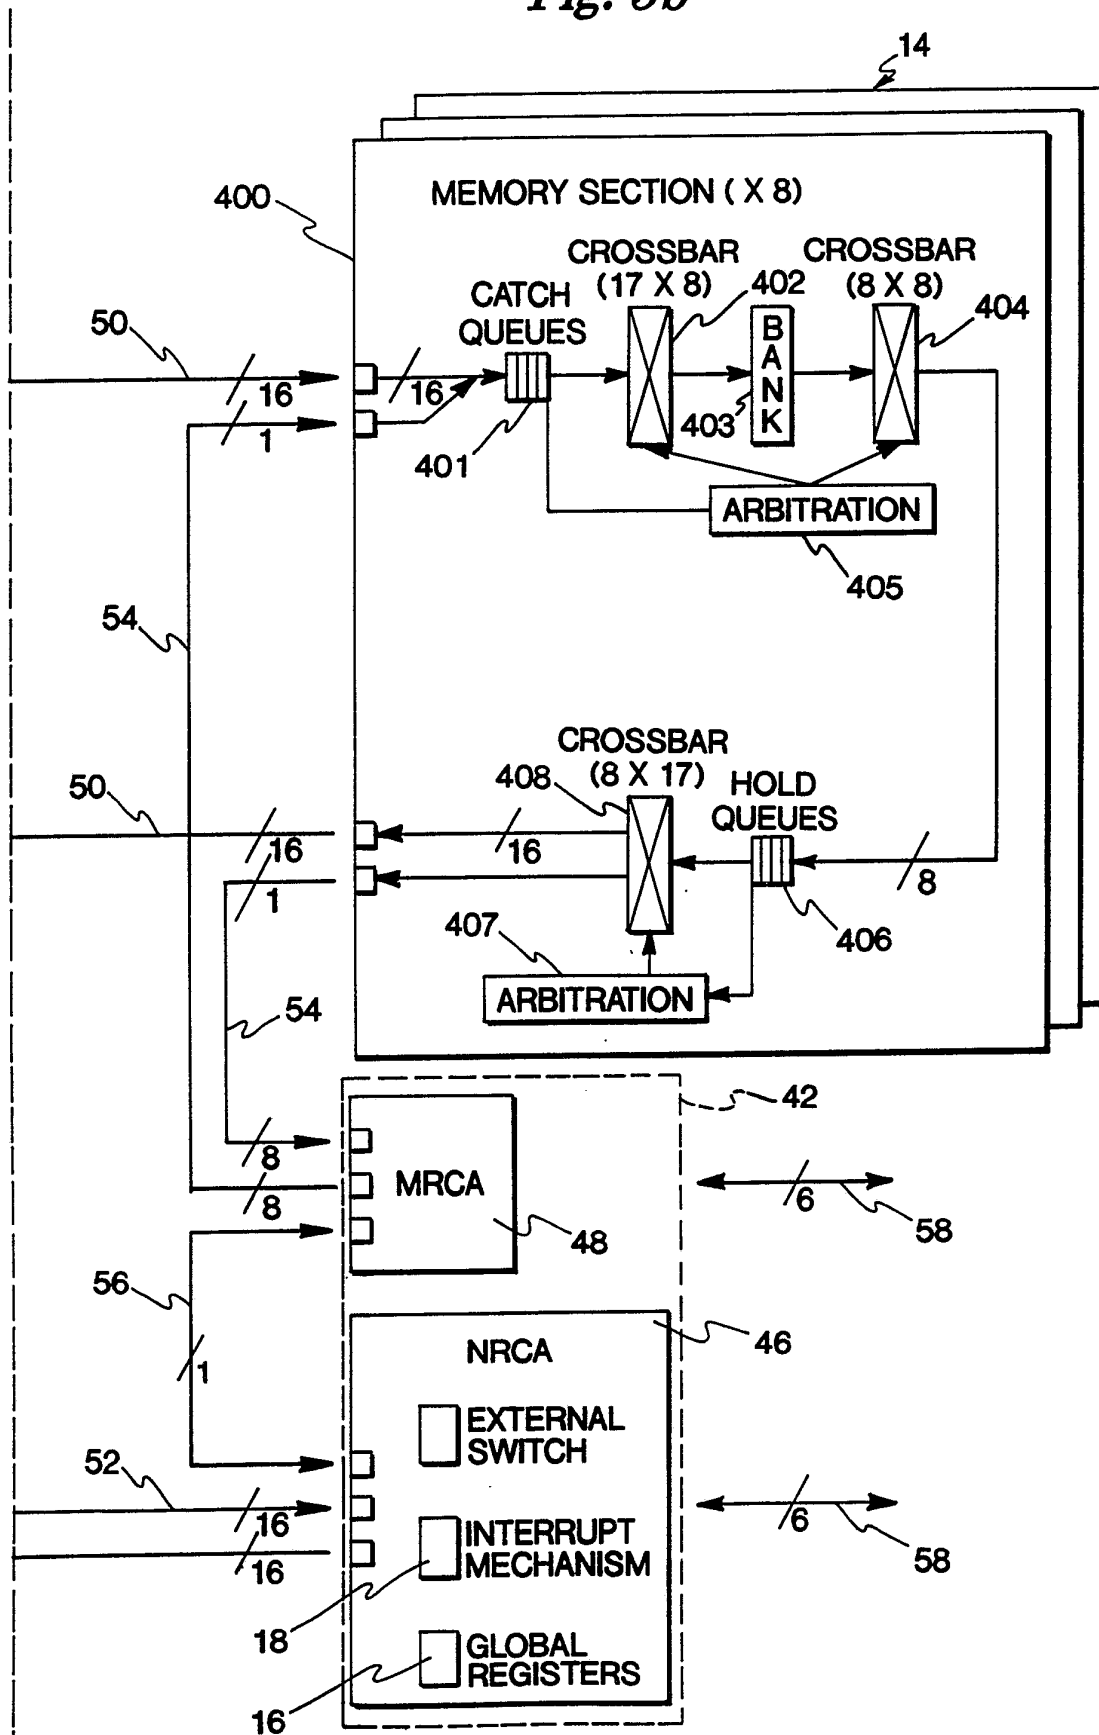
*Fig. 5b*

Fig. 6a

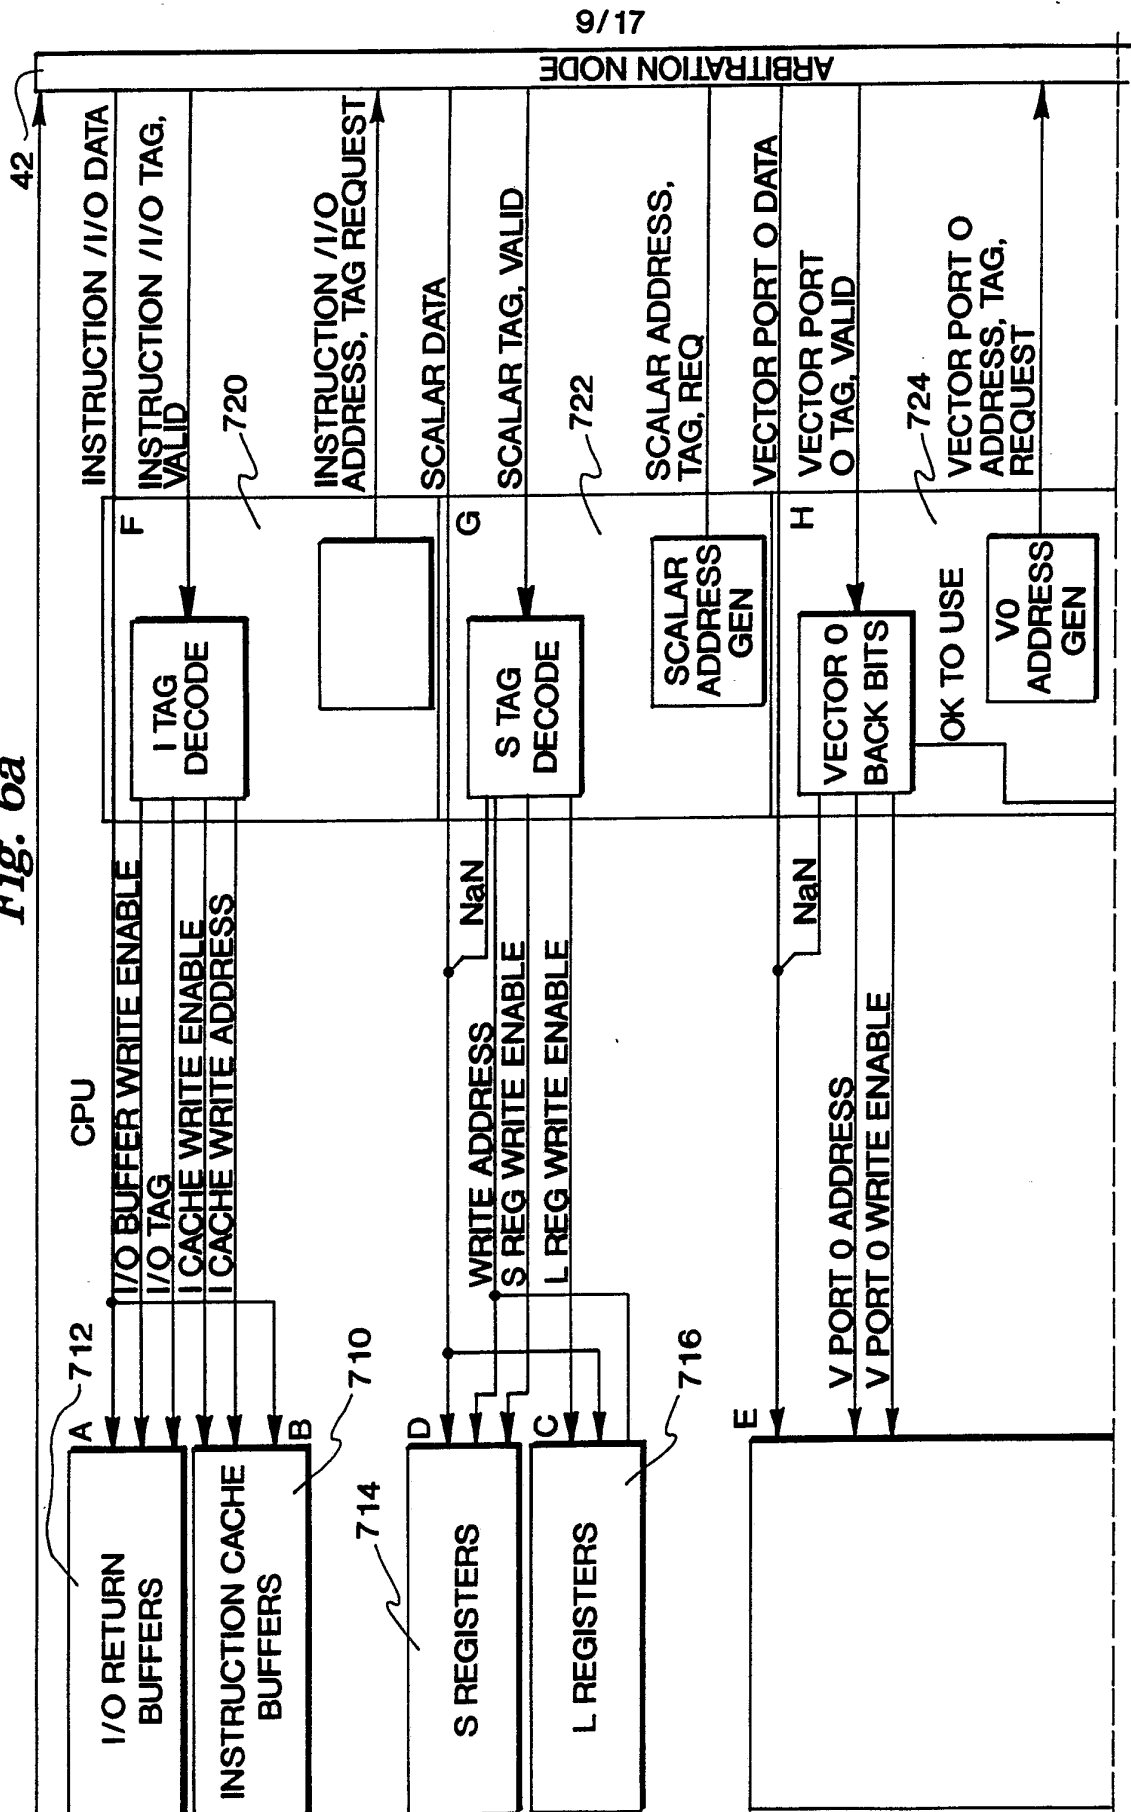
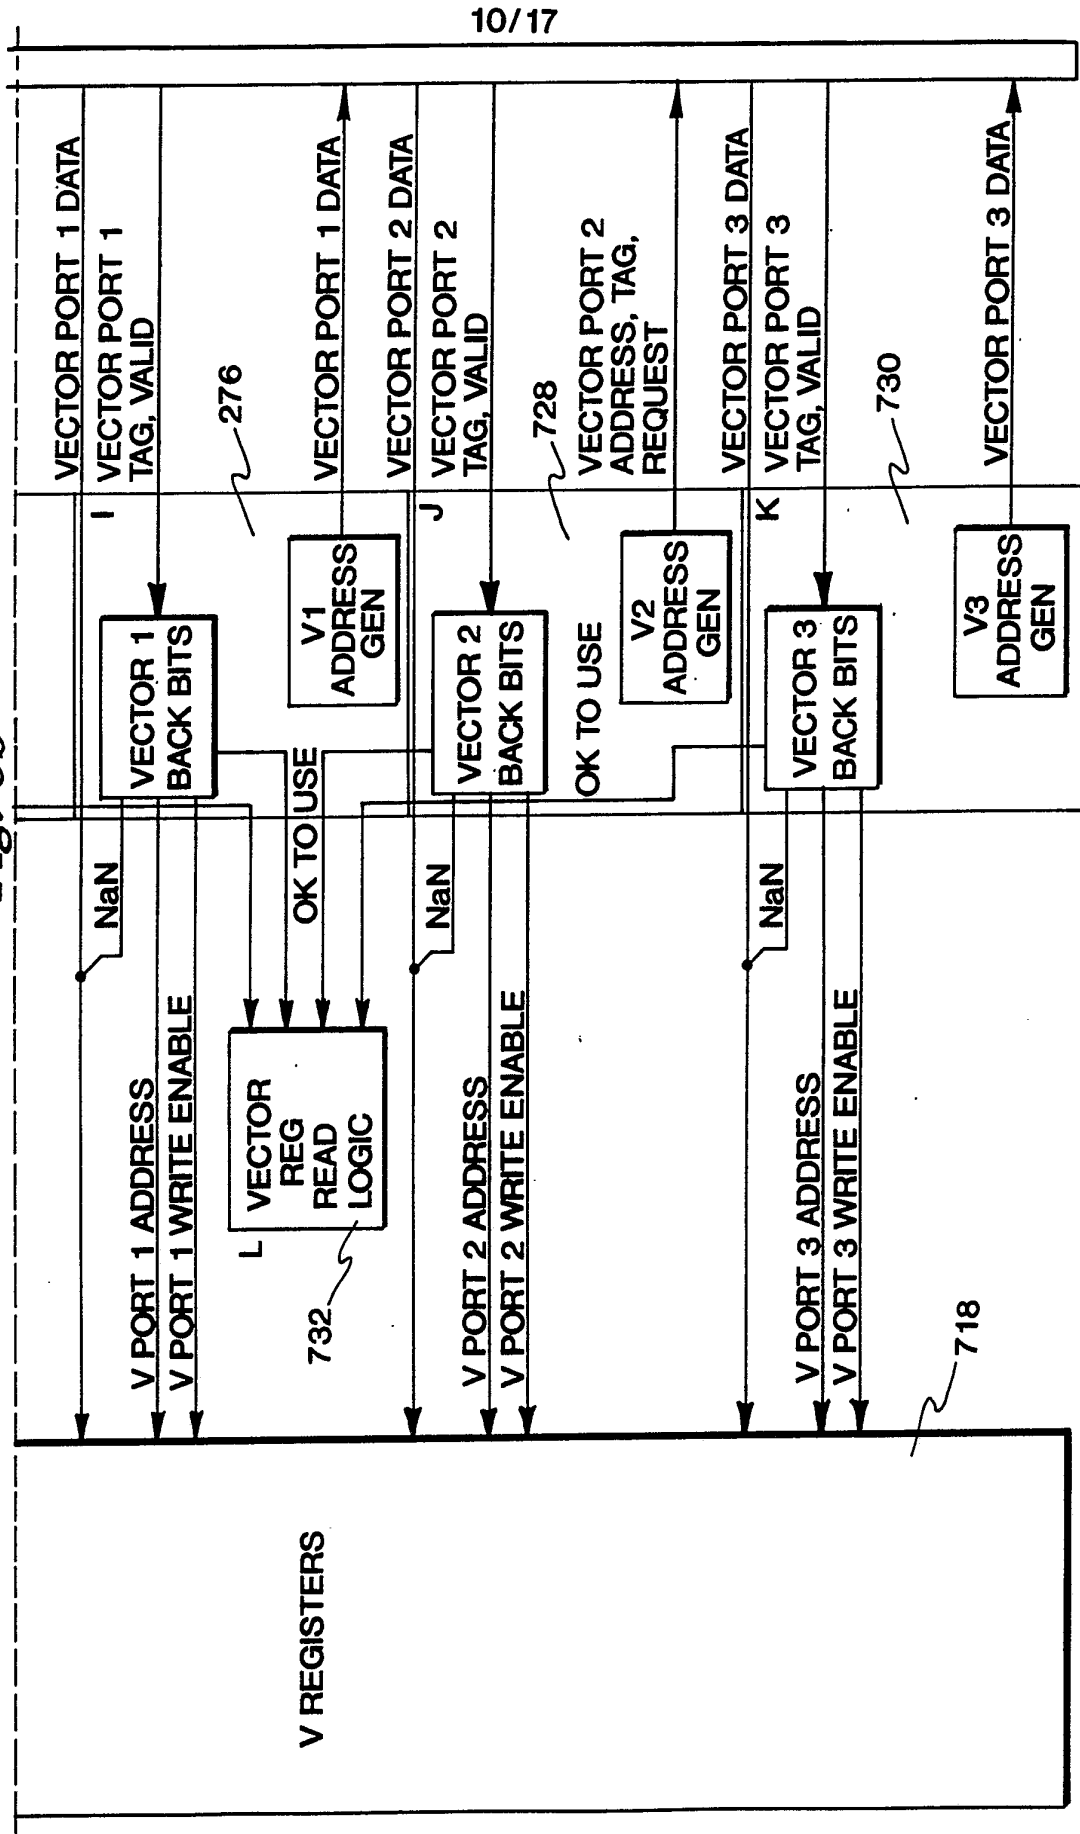


Fig. 6b




11/17

*Fig. 7a*

VECTOR PORT TAGS

THE 8-BIT TAG IS PARTITIONED AS FOLLOWS:




WHERE THE CAN BIT INDICATES THAT A NaN SHOULD BE SUBSTITUTED FOR THE RETURN DATA. THE L/R BIT INDICATES WHICH OF THE TWO QUEUES (1 PER POSSIBLE OUTSTANDING LOAD) TO USE FOR REORDERING. THE 6 LSBS OF THE TAG DIRECT WHICH VECTOR REGISTER ELEMENT NUMBER TO UPDATE WITH THE INCOMING DATA.

*Fig. 7b*

SCALAR PORT TAGS

THE 11-BIT TAG IS PARTITIONED AS FOLLOWS:



WHERE THE CAN BIT INDICATES THAT A NaN SHOULD BE SUBSTITUTED FOR THE RETURN DATA. THE L/S BIT INDICATES WHETHER TO WRITE AN L OR S REGISTER WHOSE NUMBER IS INDICATED BY BITS EIGHT THROUGH ZERO.

Fig. 7c

INSTRUCTION /IO PORT TAGS


THE 10-BIT TAG IS PARTITIONED AS FOLLOWS FOR INSTRUCTION FETCHES:



WHERE BNUM INDICATES ONE OF TWO BUFFER NUMBERS, AND E4 THROUGH E0 INDICATE THE ENTRY NUMBER IN THE BUFFER.

Fig. 7d

AND AS FOLLOWS FOR I/O FETCHES:




WHERE T9 THROUGH T0 ARE STORED IN THE I/O FIFO ALONG WITH THE RETURN DATA.



Fig. 7e

THE 21-BIT CLUSTER TAG IS PARTITIONED AS FOLLOWS:



WHERE:


- A3:A0 ARBITRATION NODE IDENTIFICATION NUMBER
- AP PARITY OVER ARBITRATION NODE IDENTIFICATION NUMBER FIELD
- SC STORE CLASS BIT
- P3:P0 PORT IDENTIFICATION NUMBER
- C9:C ORIGINAL ARBITRATION NODE PORT TAG
- CP PARITY OVER STORE CLASS BIT FIELD, PORT IDENTIFICATION NUMBER FIELD, AND ORIGINAL ARBITRATION NODE PORT TAG FIELD

Fig. 9


COMMAND	BLOCK WORD	TAG (BINARY)	FUNCTION
1	1X11000000		TRANSFER STATUS
2	1X11000001		SMTC COMMAND
3	1X11000010		SM FIRST BLOCK ADDRESS
4	1X11000011		BLOCK COUNT
5	1X11000100		SM BLOCK ADDRESS INCREMENT
6	1X11000101		MAIN MEMORY FIRST WORD ADDRESS
7	1X11000110		SIGNAL DEVICE SELECT

WHERE X = 0 FOR SMTc0 AND X = 1 FOR SMTc1 IN THE TAG FIELD.

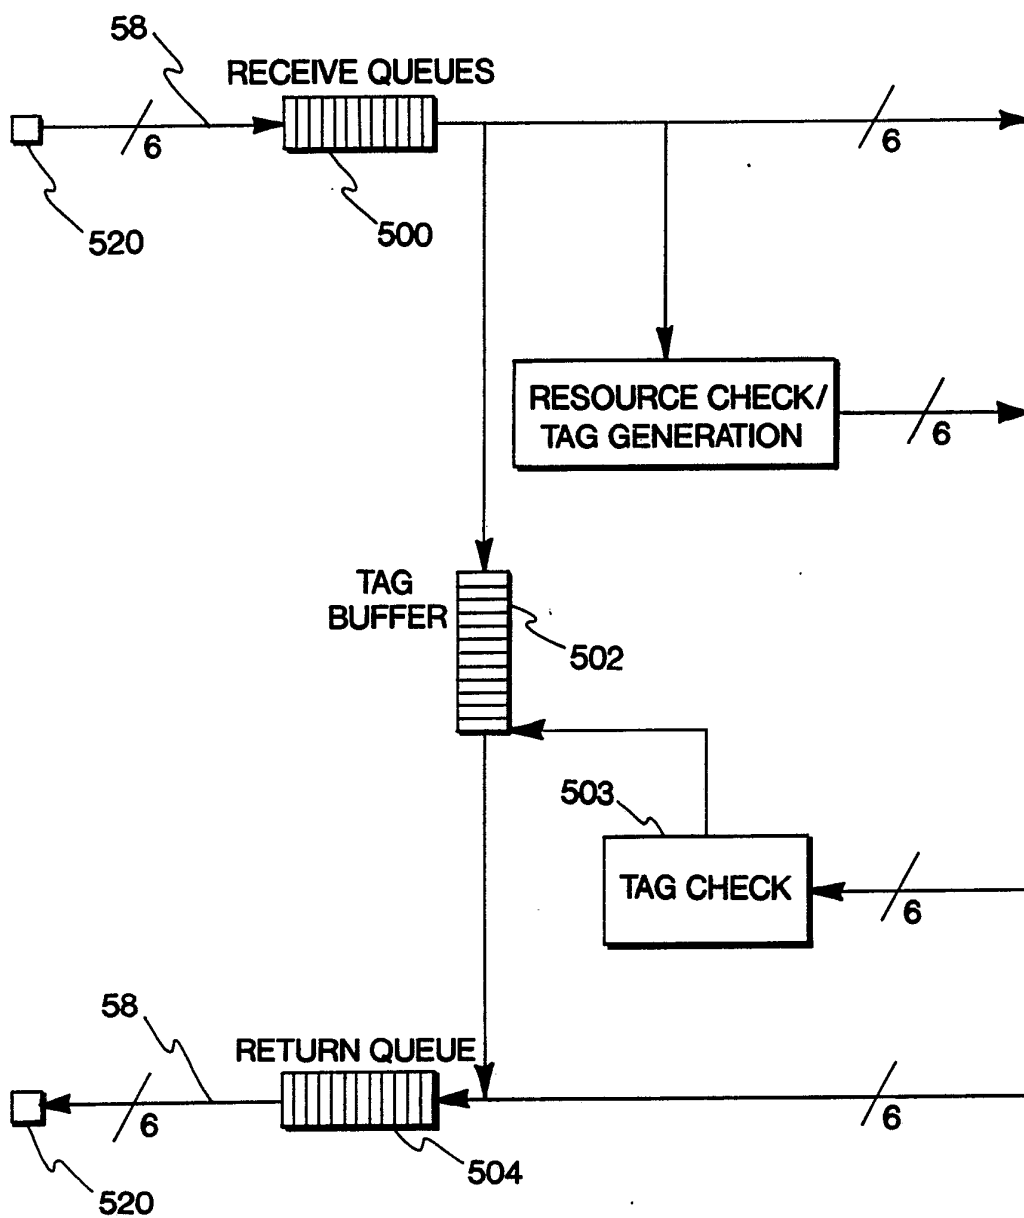
14/17




15/17

*Fig. 10*

16/17

*Fig. 11a*

17/17

*Fig. 11b*

# INTERNATIONAL SEARCH REPORT

International Application No. PCT/US91/04054

<b>I. CLASSIFICATION OF SUBJECT MATTER</b> (if several classification symbols apply, indicate all) <sup>6</sup> According to International Patent Classification (IPC) or to both National Classification and IPC IPC(5): G06F 13/00, 12/00, 15/16																													
<b>II. FIELDS SEARCHED</b> <div style="text-align: center; margin-top: 10px;">Minimum Documentation Searched <sup>7</sup></div> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 25%; padding: 5px;">Classification System</td> <td style="padding: 5px;">Classification Symbols</td> </tr> <tr> <td style="padding: 5px;">U.S. CL.</td> <td style="padding: 5px;">395/650</td> </tr> </table> <div style="text-align: center; margin-top: 10px;">Documentation Searched other than Minimum Documentation to the Extent that such Documents are Included in the Fields Searched <sup>8</sup></div>			Classification System	Classification Symbols	U.S. CL.	395/650																							
Classification System	Classification Symbols																												
U.S. CL.	395/650																												
<b>III. DOCUMENTS CONSIDERED TO BE RELEVANT <sup>9</sup></b> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="width: 10%; padding: 5px;">Category <sup>*</sup></th> <th style="width: 60%; padding: 5px;">Citation of Document, <sup>11</sup> with indication, where appropriate, of the relevant passages <sup>12</sup></th> <th style="width: 30%; padding: 5px;">Relevant to Claim No. <sup>13</sup></th> </tr> </thead> <tbody> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">Y</td> <td style="padding: 5px;">US,A 4,807,111 (COHEN et al.) 21 FEBRUARY 1989 see the entire document</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">Y</td> <td style="padding: 5px;">US,A 4,805,107 (KIECKHAFFER et al.) 14 FEBRUARY 1989 see column 2</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">Y</td> <td style="padding: 5px;">US,A 4,644,461 (JENNINGS) 17 FEBRUARY 1987 see columns 7-50</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">Y</td> <td style="padding: 5px;">US,A 4,663,706 (ALLEN et al.) 05 MAY 1987 see entire document</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A</td> <td style="padding: 5px;">US,A 4,779,194 (JENNINGS et al.) 18 OCTOBER 1988 see the entire document</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A,E</td> <td style="padding: 5px;">US,A 5,031,089 (LIU et al.) 09 JULY 1991 see entire document</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A</td> <td style="padding: 5px;">US,A 4,615,001 (HUDGINS Jr.) 30 SEPTEMBER 1986 see the entire document</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> <tr> <td style="text-align: center; vertical-align: top; padding: 5px;">A</td> <td style="padding: 5px;">US,A 4,320,451 (BACHMAN et al.) 16 MARCH 1982 see columns 6-35</td> <td style="text-align: center; vertical-align: top; padding: 5px;">1-20</td> </tr> </tbody> </table>			Category <sup>*</sup>	Citation of Document, <sup>11</sup> with indication, where appropriate, of the relevant passages <sup>12</sup>	Relevant to Claim No. <sup>13</sup>	Y	US,A 4,807,111 (COHEN et al.) 21 FEBRUARY 1989 see the entire document	1-20	Y	US,A 4,805,107 (KIECKHAFFER et al.) 14 FEBRUARY 1989 see column 2	1-20	Y	US,A 4,644,461 (JENNINGS) 17 FEBRUARY 1987 see columns 7-50	1-20	Y	US,A 4,663,706 (ALLEN et al.) 05 MAY 1987 see entire document	1-20	A	US,A 4,779,194 (JENNINGS et al.) 18 OCTOBER 1988 see the entire document	1-20	A,E	US,A 5,031,089 (LIU et al.) 09 JULY 1991 see entire document	1-20	A	US,A 4,615,001 (HUDGINS Jr.) 30 SEPTEMBER 1986 see the entire document	1-20	A	US,A 4,320,451 (BACHMAN et al.) 16 MARCH 1982 see columns 6-35	1-20
Category <sup>*</sup>	Citation of Document, <sup>11</sup> with indication, where appropriate, of the relevant passages <sup>12</sup>	Relevant to Claim No. <sup>13</sup>																											
Y	US,A 4,807,111 (COHEN et al.) 21 FEBRUARY 1989 see the entire document	1-20																											
Y	US,A 4,805,107 (KIECKHAFFER et al.) 14 FEBRUARY 1989 see column 2	1-20																											
Y	US,A 4,644,461 (JENNINGS) 17 FEBRUARY 1987 see columns 7-50	1-20																											
Y	US,A 4,663,706 (ALLEN et al.) 05 MAY 1987 see entire document	1-20																											
A	US,A 4,779,194 (JENNINGS et al.) 18 OCTOBER 1988 see the entire document	1-20																											
A,E	US,A 5,031,089 (LIU et al.) 09 JULY 1991 see entire document	1-20																											
A	US,A 4,615,001 (HUDGINS Jr.) 30 SEPTEMBER 1986 see the entire document	1-20																											
A	US,A 4,320,451 (BACHMAN et al.) 16 MARCH 1982 see columns 6-35	1-20																											
<div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p><sup>*</sup> Special categories of cited documents: <sup>10</sup></p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> </div> <div style="width: 45%;"> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance: the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&amp;" document member of the same patent family</p> </div> </div>																													
<b>IV. CERTIFICATION</b> <table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="width: 50%; padding: 5px;">           Date of the Actual Completion of the International Search             12 SEPTEMBER 1991         </td> <td style="width: 50%; padding: 5px;">           Date of Mailing of this International Search Report   <div style="font-size: 1.5em; font-weight: bold;">17 OCT 1991</div> </td> </tr> <tr> <td style="padding: 5px;">           International Searching Authority             ISA/US         </td> <td style="padding: 5px;">           Signature of Authorized Person  <div style="text-align: center;"> <b>NGUYEN NGOC-HO</b>  <b>INTERNATIONAL DIVISION</b>  <i>For ROBERT S. HAUSER</i> </div> </td> </tr> </table>			Date of the Actual Completion of the International Search  12 SEPTEMBER 1991	Date of Mailing of this International Search Report  <div style="font-size: 1.5em; font-weight: bold;">17 OCT 1991</div>	International Searching Authority  ISA/US	Signature of Authorized Person <div style="text-align: center;"> <b>NGUYEN NGOC-HO</b>  <b>INTERNATIONAL DIVISION</b>  <i>For ROBERT S. HAUSER</i> </div>																							
Date of the Actual Completion of the International Search  12 SEPTEMBER 1991	Date of Mailing of this International Search Report  <div style="font-size: 1.5em; font-weight: bold;">17 OCT 1991</div>																												
International Searching Authority  ISA/US	Signature of Authorized Person <div style="text-align: center;"> <b>NGUYEN NGOC-HO</b>  <b>INTERNATIONAL DIVISION</b>  <i>For ROBERT S. HAUSER</i> </div>																												