




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification⁶ : G06F 19/00</p>	<p>A1</p>	<p>(11) International Publication Number: WO 98/12657 (43) International Publication Date: 26 March 1998 (26.03.98)</p>
<p>(21) International Application Number: PCT/US97/16218 (22) International Filing Date: 16 September 1997 (16.09.97) (30) Priority Data: 08/710,404 17 September 1996 (17.09.96) US (71) Applicant: MARATHON TECHNOLOGIES CORPORATION [US/US]; 1300 Massachusetts Avenue, Boxboro, MA 01719 (US). (72) Inventors: BISSETT, Thomas, D.; 78 Sunset Drive, Northborough, MA 01532 (US). FITZGERALD, Martin, J., V.; 5 Holbrook Street, Medway, MA 02053 (US). LEVEILLE, Paul, A.; 12 Stratton Road, Grafton, MA 01519 (US). MCCOLLUM, James, D.; 275 Swift Road, Whitinsville, MA 01588 (US). MUENCH, Erik; 677 Salem Street, Groveland, MA 01834 (US). TREMBLAY, Glenn, A.; 139 South Street, Upton, MA 01568 (US). (74) Agent: WALPERT, Gary, A.; Fish & Richardson, P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).</p>		<p>(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>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.</i></p>

(54) Title: FAULT RESILIENT/FAULT TOLERANT COMPUTING

(57) Abstract

Data transfer to computing elements (14a, 14b) is synchronized in a computer system (10) that includes the computing elements and controllers (12) that provide data from data sources to the computing elements. A request for data from a computing element (14a, 14b) is intercepted and transmitted to the controllers (12). At least a first controller responds by transmitting requested data to the computing element and by indicating how a second controller will respond to the intercepted request.



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.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LJ	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

FAULT RESILIENT/FAULT TOLERANT COMPUTINGBackground of the Invention

The invention relates to fault resilient and fault
5 tolerant computing.

Fault resilient computer systems can continue to
function in the presence of hardware failures. These
systems operate in either an availability mode or an
integrity mode, but not both. A system is "available"
10 when a hardware failure does not cause unacceptable
delays in user access. Accordingly, a system operating
in an availability mode is configured to remain online,
if possible, when faced with a hardware error. A system
has data integrity when a hardware failure causes no data
15 loss or corruption. Accordingly, a system operating in
an integrity mode is configured to avoid data loss or
corruption, even if the system must go offline to do so.

Fault tolerant systems stress both availability
and integrity. A fault tolerant system remains available
20 and retains data integrity when faced with a single
hardware failure, and, under some circumstances, when
faced with multiple hardware failures.

Disaster tolerant systems go one step beyond fault
tolerant systems and require that loss of a computing
25 site due to a natural or man-made disaster will not
interrupt system availability or corrupt or lose data.

Typically, fault resilient/fault tolerant systems
include several processors that may function as computing
elements or controllers, or may serve other roles. In
30 many instances, it is important to synchronize operation
of the processors or the transmission of data between the
processors.

Summary of the Invention

In one aspect, generally, the invention features
35 synchronizing data transfer to a computing element in a

- 2 -

computer system including the computing element and controllers that provide data from data sources to the computing element. A request for data made by the computing element is intercepted and transmitted to the controllers. Controllers respond to the request and at least one controller responds by transmitting requested data to the computing element and by indicating how another controller will respond to the intercepted request.

10 Embodiments of the invention may include one or more of the following features. A controller may respond to the intercepted request by indicating that the controller has no data corresponding to the intercepted request and by indicating that another controller will respond to the intercepted request by transmitting data to the computing element. Each response to the intercepted request by a controller may include an indication as to how each other controller will respond to the intercepted request.

20 The computing element may compare the responses to the intercepted request for consistency. When each response includes an indication as to how each other controller will respond to the intercepted request, the comparison may include comparing the indications for consistency. When responses of two or more controllers include requested data, the comparison may include comparing the data for consistency. The computing element may notify the controllers of the outcome of the comparison and that responses have been received from all of the controllers.

30 A controller may be disabled when the responses are not consistent. In addition, an error condition may be generated if the computing element does not receive responses from all of the controllers within a predetermined time period.

- 3 -

A data source may be associated with a controller, and the controller may obtain the requested data from the data source in response to the intercepted request.

A controller may maintain a record of a status of another controller, and may use the record when
5 indicating how the other controller will respond to the intercepted request. When a data source is associated with the other controller, the record may include the status of the data source. Each controller may maintain
10 records of statuses of all other controllers and may use the records to indicate how the other controllers will respond to the intercepted request. When each controller is associated with a data source, each controller may maintain records of statuses of data sources associated
15 with all other controllers.

When a status of a data source associated with a controller changes, the controller may transmit to the computing element an instruction to discard responses from other controllers to the intercepted request. The
20 computing element may respond to the instruction by discarding responses from other controllers to the intercepted request and by transmitting to the controllers a notification that the responses have been discarded. A controller may respond to the notification
25 by updating a record of the status of the data source. After updating the record, the controller may retransmit the requested data to the computing element and indicate how the other controller will respond to the intercepted request.

30 When a data source is associated with each controller, each controller may respond to the intercepted request by determining whether an associated data source is expected to process the request, and when the associated data source is expected to process the
35 request, transmitting the request to the associated data

- 4 -

source, receiving results of the request from the associated data source, and forwarding the results of the request to the computing element. When the associated data source is not expected to process the request, the controller may respond by informing the computing element that no data will be provided in response to the request.

In another aspect, generally, the invention features maintaining synchronization between computing elements processing identical instruction streams in a computer system including the computing elements and controllers that provide data from data sources to the computing elements, with the controllers operating asynchronously to the computing element. Computing elements processing identical instruction streams each stop processing of the instruction stream at a common point in the instruction stream. Each computing element then generates a freeze request message and transmits the freeze request message to the controllers. A controller receives a freeze request message from a computing element, waits for a freeze request message from other computing elements, and, upon receiving a freeze request message from each computing element processing an identical instruction stream, generates a freeze response message and transmits the freeze response message to the computing elements. Each computing element, upon receiving a freeze response message from a controller, waits for freeze response messages from other controllers to which a freeze request message was transmitted, and, upon receiving a freeze response message from each controller, generates a freeze release message, transmits the freeze release message to the controllers, and resumes processing of the instruction stream.

Embodiments of the invention may include one or more of the following features. The common point in the instruction stream may correspond to an I/O operation,

- 5 -

the occurrence of a predetermined number of instructions without an I/O operation, or both.

A controller may include a time update in the freeze response message, and a computing element, upon receiving a freeze response message from each controller to which a freeze request message was transmitted, may update a system time using the time update from a freeze response message. The computing element may use the time update from a freeze response message generated by a particular controller.

Upon receiving a freeze response message from each controller to which a freeze request message was transmitted, a computing element may process data received from a controller prior to receipt of freeze response messages from the controllers.

In another aspect, generally, the invention features handling faults in a computer system including error reporting elements and error processing elements. An error reporting element detects an error condition and transmits information about the error condition as an error message to error processing elements connected to the error reporting element. At least one error processing element retransmits the error message to other error processing elements connected to the error processing element.

In another aspect, generally, the invention features handling faults in a computer system including error reporting elements and error processing elements. Error reporting element detect an error condition and transmit information about the error condition as error messages to error processing elements connected to the error reporting elements. At least one error processing element combines information from related error messages from multiple error reporting elements and uses the

- 6 -

combined information in identifying a source of the error condition.

The error processing element may use a state table to combine information from related error messages. The error processing element may represent an error message using an error identifier that identifies a particular error, an error target that identifies a subcomponent that caused the error represented by the error message, and a reporting source that identifies an error reporting element that generated the error message and a path over which the error message was received. The error processing element may determine whether error messages are related by comparing a received error message against states representing previously received error messages.

Other features and advantages will become apparent from the following description, including the drawings, and from the claims.

Brief Description of the Drawings

Fig. 1 is a block diagram of a partially fault resilient system.

Fig. 2 is a block diagram of system software of the system of Fig. 1.

Fig. 3 is a flowchart of a procedure used by an IOP Monitor of the system software of Fig. 2.

Fig. 4 is a block diagram of an IPI module of the system of Fig. 1.

Fig. 5 is a state transition table for the system of Fig. 1.

Fig. 6 is a block diagram of a fault resilient system.

Fig. 7 is a block diagram of a distributed fault resilient system.

- 7 -

Fig. 8 is a block diagram of a fault tolerant system.

Fig. 9 is flowchart of a fault diagnosis procedure used by IOPs of the system of Fig. 8.

5 Fig. 10 is a block diagram of a disaster tolerant system.

Fig. 11 is a block diagram of software components of a computer system.

10 Figs. 12A and 12B are flow diagrams illustrating information transferred between CEs and IOPs.

Fig. 13 is a flowchart of a flush procedure implemented by the CE Transport.

Figs. 14A and 14B are block diagrams illustrating physical and logical system configurations.

15 Fig. 15 is a flowchart of a flush procedure.

Figs. 16 and 16B are tables of microcode state transitions.

Fig. 17 is a block diagram of a system.

20 Fig. 18 is a flow chart of a procedure for error processing.

Fig. 19A is a syndrome source table.

Fig. 19B is a state transition table corresponding to the syndrome source table of Fig. 19A.

Fig. 20 is a callout hierarchy diagram.

25 Figs. 21A, 21B and 22 are tables listing callout elements.

Fig. 23 is a flowchart of a procedure for identifying a faulty component.

Description of the Preferred Embodiments

30 Fig. 1 illustrates a fault resilient system 10 that includes an I/O processor ("IOP") 12 and two computing elements ("CEs") 14a, 14b (collectively referred to as CEs 14). System 10 includes only a single

- 8 -

IOP 12 and therefore cannot recover from a failure in IOP 12. As such, system 10 is not entirely fault resilient.

IOP 12 includes two inter-processor interconnect ("IPI") modules 16a, 16b that are connected,
5 respectively, to corresponding IPI modules 18a, 18b of CEs 14 by cables 20a, 20b. IOP 12 also includes a processor 22, a memory system 24, two hard disk drives 26, 28, and a power supply 30. Similarly, each CE 14 includes a processor 32, a memory system 34, and a power
10 supply 36. Separate power supplies 36 are used to ensure fault resilience in the event of a power supply failure. Processors 32a, 32b are "identical" to each other in that, for every instruction, the number of cycles required for processor 32a to perform an instruction is
15 identical to the number of cycles required for processor 32b to perform the same instruction. In the illustrated embodiment, system 10 has been implemented using standard Intel 486 based motherboards for processors 22, 32 and four megabytes of memory for each of memory systems 24,
20 34.

IOP 12 and CEs 14 of system 10 run unmodified operating system and applications software, with hard drive 26 being used as the boot disk for the IOP and hard drive 28 being used as the boot disk for CEs 14. In
25 truly fault resilient or fault tolerant systems that include at least two IOPs, each hard drive would also be duplicated.

In the illustrated embodiment, the operating system for IOP 12 and CEs 14 is DOS. However, other
30 operating systems can also be used. Moreover, IOP 12 can run a different operating system from the one run by CEs 14. For example, IOP 12 could run Unix while CEs 14 run DOS. This approach is advantageous because it allows CEs 14 to access peripherals from operating systems that do
35 not support the peripherals. For example, if CEs 14 were

- 9 -

running an operating system that did not support CD-ROM drives, and IOP 12 were running one that did, CEs 14 could access the CD-ROM drive by issuing I/O requests identical to those used to, for example, access a hard drive. IOP 12 would then handle the translation of the I/O request to one suitable for accessing the CD-ROM drive.

Fig. 2 provides an overview of specialized system software 40 used by system 10 to control the booting and synchronization of CEs 14, to disable local time in CEs 14, to redirect all I/O requests from CEs 14 to IOP 12 for execution, and to return the results of the I/O requests, if any, from IOP 12 to CEs 14.

System software 40 includes two sets of IPI BIOS 42 that are ROM-based and are each located in the IPI module 18 of a CE 14. IPI BIOS 42 are used in bootup and synchronization activities. When a CE 14 is booted, IPI BIOS 42 replaces the I/O interrupt addresses in the system BIOS interrupt table with addresses that are controlled by CE Drivers 44. The interrupt addresses that are replaced include those corresponding to video services, fixed disk services, serial communications services, keyboard services, and time of day services.

CE Drivers 44 are stored on CE boot disk 28 and are run by CEs 14. CE Drivers 44 intercept I/O requests to the system BIOS and redirect them through IPI modules 18 to IOP 12 for execution. CE Drivers 44 also respond to interrupt requests from IPI modules 18, disable the system clock, and, based on information supplied by IOP Monitor 48, control the time of day of CEs 14.

An IOP Driver 46 that is located on IOP boot disk 26 and is run by IOP 12 handles I/O requests from CEs 14 by redirecting them to an IOP Monitor 48 for processing. Thereafter, IOP Driver 46 transmits the results of the

- 10 -

requests from IOP Monitor 48 to CEs 14. IOP Driver 46 communicates with CE drivers 44 using a packet protocol.

IOP Monitor 48 is located on IOP boot disk 26 and is run by IOP 12. IOP Monitor 48 controls system 10 and
5 performs the actual I/O requests to produce the results that are transmitted by IOP Driver 46 to CEs 14.

System software 40 also includes console software 49 that runs on IOP 12 and permits user control of system 10. Using console software 49, a user can reset, boot,
10 or synchronize a CE 14. The user can also set one or both of CEs 14 to automatically boot (autoboot) and/or automatically synchronize (autosync) after being reset or upon startup. The ability to control each CE 14 is useful both during normal operation and for test
15 purposes. Using console software 49, the user can also place system 10 into either an integrity mode in which IOP Monitor 48 shuts down both CEs 14 when faced with a miscompare error, a first availability mode in which IOP Monitor 48 disables CE 14a when faced with a miscompare
20 error, or a second availability mode in which IOP Monitor 48 disables CE 14b when faced with a miscompare error. Finally, console software 49 allows the user to request the status of system 10. In an alternative embodiment, console software 49 is implemented using a separate
25 processor that communicates with IOP 12.

Each CE 14 runs a copy of the same application and the same operating system as that run by the other CE 14. Moreover, the contents of memory systems 34a and 34b are the same, and the operating context of CEs 14 are the
30 same at each synchronization time. Thus, IOP Monitor 48 should receive identical sequences of I/O requests from CEs 14.

As shown in Fig. 3, IOP Monitor 48 processes and monitors I/O requests according to a procedure 100.
35 Initially, IOP Monitor 48 waits for an I/O request from

- 11 -

one of CEs 14 (step 102). Upon receiving an I/O request packet from, for example, CE 14b, IOP Monitor 48 waits for either an I/O request from CE 14a or for the expiration of a timeout period (step 104). Because
5 system 10 uses the DOS operating system, which halts execution of an application while an I/O request is being processed, IOP Monitor 48 is guaranteed not to receive an I/O request from CE 14b while waiting (step 104) for the I/O request from the CE 14a.

10 Next, IOP Monitor 48 checks to determine whether the timeout period has expired (step 106). If not (that is, if an I/O request packet from CE 14a has arrived), IOP Monitor 48 compares the checksums of the packets (step 108). If the checksums are equal, IOP Monitor 48
15 processes the I/O request (step 110). After processing the I/O request, IOP Monitor 48 issues a request to the system BIOS of IOP 12 for the current time of day (step 112).

After receiving the time of day, IOP Monitor 48
20 assembles an IPI packet that includes the time of day and the results, if any, of the I/O request (step 114) and sends the IPI packet to IOP Driver 46 (step 116) for transmission to CEs 14. When CEs 14 receive the IPI packet, they use the transmitted time of day to update
25 their local clocks which, as already noted, are otherwise disabled.

As required by DOS, execution in CEs 14 is suspended until IOP Monitor 48 returns the results of the I/O request through IOP Driver 46. Before execution is
30 resumed, the times of day of both CEs 14 are updated to a common value corresponding to the transmitted time of day from the IPI packet. Accordingly, the CEs 14 are kept in time synchronization. The transmitted time of day is designated as the meta time. If a multitasking operating
35 system were employed, execution in CEs 14 would not be

- 12 -

suspended while IOP Monitor 48 performed the I/O request. Instead, processing in CEs 14 would be suspended only until receipt of an acknowledgement indicating that IOP Monitor 48 had begun processing the I/O request (step
5 110). The acknowledgement would include the time of day and would be used by CEs 14 to update the local clocks.

After sending the IPI packet to IOP Driver 46, IOP Monitor 48 verifies that both CEs 14 are online (step 118), and, if so, waits for an I/O request from one of
10 the CEs 14 (step 102).

If the timeout period has expired (step 106), IOP Monitor 48 disables the CE 14 that failed to respond (step 119) and processes the I/O request (step 110).

If there is a miscompare between the checksums of
15 the packets from CEs 14 (step 108), IOP Monitor 48 checks to see if system 10 is operating in an availability mode or an integrity mode (step 120). If system 10 is operating in an availability mode, IOP Monitor 48 disables the appropriate CE 14 based on the selected
20 availability mode (step 122), and processes the I/O request (steps 110-116). Thereafter, IOP Monitor 48 determines whether the disabled CE 14 has been repaired and reactivated (step 118). If not, IOP Monitor 48 waits for an I/O request from the online CE 14 (step 124).
25 With one of the CEs 14 disabled, system 10 is no longer fault resilient and IOP Monitor 48 immediately processes a received I/O request (step 110).

If system 10 is operating in an integrity mode when a miscompare is detected, IOP Monitor 48 disables
30 both CEs 14 (step 126) and stops processing (step 128).

Referring again to Figs. 1 and 2, when the application or the operating system of, for example, CE 14a makes a non-I/O call to the system BIOS, the system BIOS executes the request and returns the results to the
35 application without invoking system software 40.

- 13 -

However, if the application or the operating system makes an I/O BIOS call, CE Driver 44a intercepts the I/O request. After intercepting the I/O request, CE Driver 44a packages the I/O request into an IPI packet and
5 transmits the IPI packet to IOP 12.

When IPI module 16a of IOP 12 detects transmission of an IPI packet from CE 14a, IPI module 16a generates an interrupt to IOP Driver 46. IOP Driver 46 then reads the IPI packet.

10 As discussed above, IOP Monitor 48 responds to the IPI packet from CE 14a according to procedure 100. As also discussed, assuming that there are no hardware faults, IOP Driver 46 eventually transmits an IPI packet that contains the results of the I/O request and the time
15 of day to CEs 14.

IPI modules 18 of CEs 14 receive the IPI packet from IOP 12. CE Drivers 44 unpack the IPI packet, update the time of day of CEs 14, and return control of CEs 14 to the application or the operating system running on CEs
20 14.

If no I/O requests are issued within a given time interval, the IPI module 18 of a CE 14 generates a so-called quantum interrupt that invokes the CE Driver 44 of the CE 14. In response, the CE Driver 44 creates a
25 quantum interrupt IPI packet and transmits it to IOP 12. IOP Monitor 48 treats the quantum interrupt IPI packet as an IPI packet without an I/O request. Thus, IOP Monitor 48 detects the incoming quantum interrupt IPI packet (step 102 of Fig. 3) and, if a matching quantum interrupt
30 IPI packet is received from the other CE 14 (steps 104, 106, and 108 of Fig. 3), issues a request to the system BIOS of IOP 12 for the current time of day (step 112 of Fig. 3). IOP Monitor 48 then packages the current time of day into a quantum response IPI packet (step 114 of
35 Fig. 3) that IOP Driver 46 then sends to CEs 14 (step 116

- 14 -

of Fig. 3). CE Drivers 44 respond to the quantum response IPI packet by updating the time of day and returning control of CEs 14 to the application or the operating system running on CEs 14.

5 If IOP Monitor 48 does not receive a quantum interrupt IPI package from the other CE 14 within a predefined timeout period (step 106 of Fig. 3), IOP Monitor 48 responds by disabling the non-responding CE 14.

10 As shown in Fig. 1, IPI modules 16, 18 and cables 20 provide all of the hardware necessary to produce a fault resilient system from the standard Intel 486 based motherboards used to implement processors 22, 32. An IPI module 16 and an IPI module 18, which are implemented
15 using identical boards, each perform similar functions.

As illustrated in Fig. 4, an IPI module 18 includes a control logic 50 that communicates I/O requests and responses between the system bus of a processor 32 of a CE 14 and a parallel interface 52 of
20 IPI module 18. Parallel interface 52, in turn, communicates with the parallel interface of an IPI module 16 through a cable 20. Parallel interface 52 includes a sixteen bit data output port 54, a sixteen bit data input port 56, and a control port 58. Cable 20 is configured
25 so that data output port 54 is connected to the data input port of the IPI module 16, data input port 56 is connected to the data output port of the IPI module 16, and control port 58 is connected to the control port of the IPI module 16. Control port 58 implements a
30 handshaking protocol between IPI module 18 and the IPI module 16.

Control logic 50 is also connected to an IPI BIOS ROM 60. At startup, control logic 50 transfers IPI BIOS 42 (Fig. 2), the contents of IPI BIOS ROM 60, to
35 processor 32 through the system bus of processor 32.

- 15 -

A QI counter 62, also located on IPI module 18, generates quantum interrupts as discussed above. QI counter 62 includes a clock input 64 that is connected to the system clock of processor 32 and a gate input 66 that is connected to control logic 50. Gate input 66 is used to activate and reset the counter value of QI counter 62. When activated, QI counter 62 decrements the counter value by one during each cycle of the system clock of processor 32. When the counter value reaches zero, QI counter 62 generates a quantum interrupt that, as discussed above, activates CE Driver 44 (Fig. 2).

CE Driver 44 deactivates QI counter 62 at the beginning of each I/O transaction. CE Driver 44 deactivates QI counter 62 by requesting an I/O write at a first address, known as the QI deactivation address. Control logic 50 detects the I/O write request and deactivates QI counter 62 through gate input 66. Because this particular I/O write is for control purposes only, control logic 50 does not pass the I/O write to parallel interface 52. At the conclusion of each I/O transaction, CE Driver 44 resets and activates QI counter 62 by requesting an I/O write to a second address, known as the QI activation address. Control logic 50 responds by resetting and activating QI counter 62.

In an alternative approach, quantum interrupts are generated through use of debugging or other features available in processor 32. Some commonly available processors include debugging or trap instructions that trap errors by transferring control of the processor to a designated program after the completion of a selected number of instructions following the trap instruction. In this approach, each time that CE Driver 44 returns control of processor 32 to the application or operating system, CE Driver 44 issues a trap instruction to indicate that control of processor 32 should be given to

- 16 -

CE Driver 44 upon completion of, for example, 300 instructions. After processor 32 completes the indicated 300 instructions, the trap instruction causes control of processor 32 to be returned to CE Driver 44. In the event that an I/O request activates CE Driver 44 prior to completion of the indicated number of instructions, CE Driver 44 issues an instruction that cancels the trap instruction.

IPI Module 18 is also used in activating an offline CE 14. As discussed below, before an offline CE 14 is activated, the contents of the memory system 34 of the active CE 14 are copied into the memory system 34 of the offline CE 14. To minimize the effects of this copying on the active CE 14, the processor 32 of the active CE 14 is permitted to continue processing and the memory is copied only during cycles in which the system bus of the processor 32 of the active CE 14 is not in use.

To enable processor 32 to continue processing while the memory is being copied, IPI module 18 accounts for memory writes by the processor 32 to addresses that have already been copied to the offline CE 14. To do so, control logic 50 monitors the system bus and, when the processor 32 writes to a memory address that has already been copied, stores the address in a FIFO 68. When the memory transfer is complete, or when FIFO 68 is full, the contents of memory locations associated with the memory addresses stored in FIFO 68 are copied to the offline CE 14 and FIFO 68 is emptied. In other approaches, FIFO 68 is modified to store both memory addresses and the contents of memory locations associated with the addresses, or to store the block addresses of memory blocks to which memory addresses being written belong.

IPI module 18 also handles non-BIOS I/O requests. In some computer systems, the BIOS is too slow to

- 17 -

effectively perform I/O operations such as video display. As a result, some less structured or less disciplined operating systems, such as DOS or UNIX, allow applications to circumvent the BIOS and make non-BIOS I/O requests by directly reading from or writing to the addresses associated with I/O devices. These non-BIOS I/O requests, which cannot be intercepted by changing the system interrupt table, as is done in connection with, for example, I/O disk reads and writes, are problematic for a system in which synchronization requires tight control of the I/O interface.

To remedy this problem, and to assure that even non-BIOS I/O requests can be isolated and managed by IOP 12, IPI module 18 includes virtual I/O devices that mimic the hardware interfaces of physical I/O devices. These virtual I/O devices include a virtual display 70 and a virtual keyboard 72. As needed, other virtual I/O devices such as a virtual mouse or virtual serial and parallel ports could also be used.

In practice, control logic 50 monitors the system bus for read or write operations directed to addresses associated with non-BIOS I/O requests to system I/O devices. When control logic 50 detects such an operation, control logic 50 stores the information necessary to reconstruct the operation in the appropriate virtual device. Thus, for example, when control logic 50 detects a write operation directed to an address associated with the display, control logic 50 stores the information necessary to reconstruct the operation in virtual display 70. Each time that a BIOS I/O request or a quantum interrupt occurs, CE Driver 44 scans the virtual I/O devices and, if the virtual devices are not empty, assembles the information stored in the virtual devices into an IPI packet and transmits the IPI packet to IOP 12. IOP 12 treats the packet like a BIOS I/O

- 18 -

request using procedure 100 discussed above. When control logic 50 detects a read addressed to a virtual I/O device, control logic 50 assembles the read request into an IPI packet for handling by IOP 12. IOP 12 treats the IPI packet like a standard BIOS I/O request.

Referring to Fig. 5, each CE 14 always operates in one of eight states. Because there are only a limited number of permissible state combinations, system 10 always operates in one of fourteen states. The major CE operating states are OFFLINE, RTB (ready to boot), BOOTING, ACTIVE, RTS (ready to sync), WAITING, M_SYNC, (synchronizing as master), and S_SYNC (synchronizing as slave). IOP Monitor 48 changes the operating states of CEs 14 based on the state of system 10 and user commands from console software 49. Through console software 49, a user can reset a CE 14 at any time. Whenever the user resets a CE 14, or a fault occurs in the CE 14, IOP Monitor 48 changes the state of the CE 14 to OFFLINE.

At startup, system 10 is operating with both CEs 14 OFFLINE (state 150). System 10 operates in the upper states of Fig. 5 (states 152-162) when CE 14a becomes operational before CE 14b and in the lower states (states 166-176) when CE 14b is the first to become operational. If CEs 14 become operational simultaneously, the first operational CE 14 to be recognized by IOP Monitor 48 is treated as the first to become operational.

When a CE 14 indicates that it is ready to boot by issuing a boot request, the state of the CE 14 changes to RTB if the CE 14 is not set to autoboot or to BOOTING if the CE 14 is set to autoboot. For example, if CE 14a issues a boot request when both CEs 14 are OFFLINE, and CE 14a is not set to autoboot, then the state of CE 14a changes to RTB (state 152). Thereafter, IOP Monitor 48 waits for the user, through console software 49, to boot CE 14a. When the user boots CE 14a, the state of CE 14a

- 19 -

changes to BOOTING (state 154). If the user resets CE 14a, the state of CE 14a changes to OFFLINE (state 150).

If both CEs 14 are OFFLINE when CE 14a issues a boot request, and CE 14a is set to autoboot, the state of CE 14a changes to BOOTING (state 154). If CE 14a boots successfully, the state of CE 14a changes to ACTIVE (state 156).

When CE 14a is ACTIVE, and CE 14b issues a boot request, or if CE 14b had issued a boot request while the state of CE 14a was transitioning from OFFLINE to ACTIVE (states 152-156), the state of CE 14b changes to RTS (state 158) if CE 14b is set to autosync and otherwise to WAITING (state 160). If the state of CE 14b changes to RTS (state 158), IOP Monitor waits for the user to issue a synchronize command to CE 14b. When the user issues such a command, the state of CE 14b changes to WAITING (state 160).

Once CE 14b is WAITING, IOP Monitor 48 copies the contents of memory system 34a of CE 14a into memory system 34b of CE 14b. Once the memory transfer is complete, IOP Monitor 48 waits for CE 14a to transmit a quantum interrupt or I/O request IPI packet. Upon receipt of such a packet, IOP Monitor 48 changes the state of CE 14a to M_SYNC and the state of CE 14b to S_SYNC (state 162), and synchronizes the CEs 14. This synchronization includes responding to any memory changes that occurred while IOP Monitor 48 was waiting for CE 14a to transmit a quantum interrupt or I/O request IPI packet. Upon completion of the synchronization, the states of the CEs 14 both change to ACTIVE (state 164) and system 10 is deemed to be fully operational.

In an alternative implementation, IOP Monitor 48 does not wait for memory transfer to complete before changing the state of CE 14a to M_SYNC and the state of CE 14b to S_SYNC (state 162). Instead, IOP Monitor 48

- 20 -

makes this state change upon receipt of an IPI packet from CE 14a and performs the memory transfer as part of the synchronization process.

Similar state transitions occur when CE 14b is the first CE 14 to issue a boot request. Thus, assuming that CE 14b is not set to autoboot, CE 14b transitions from OFFLINE (state 150) to RTC (state 166) to BOOTING (state 168) to ACTIVE (state 170). Similarly, once CE 14b is ACTIVE, and assuming that CE 14a is not set to autosync, CE 14a transitions from OFFLINE (state 170) to RTS (state 172) to WAITING (state 174) to S_SYNC (state 176) to ACTIVE (state 164).

In other embodiments of the invention, for example, referring to Fig. 6, a fault resilient system 200 includes two IOPs 202 and two CEs 204. Each CE 204 is connected, through an IPI card 206 and a cable 208, to an IPI card 210 of each IOP 202. IOPs 202 are redundantly connected to each other through IPI cards 210 and cables 212. Because every component of system 200 has a redundant backup component, system 200 is entirely fault resilient. In an alternative approach, cables 208 and 210 could be replaced by a pair of local area networks to which each IOP 202 and CE 204 would be connected. Indeed, local area networks can always be substituted for cable connections.

System 200 is operating system and application software independent in that it does not require modifications of the operating system or the application software to operate. Any single piece of hardware can be upgraded or repaired in system 200 with no service interruption. Therefore, by sequentially replacing each piece of hardware and allowing system 200 to resynchronize after each replacement, the hardware of system 200 can be replaced in its entirety without service interruption. Similarly, software on system 200

- 21 -

can be upgraded with minimal service interruption (that is, during the software upgrade, the application will become unavailable for an acceptable period of time such as two seconds). Also, disaster tolerance for purposes
5 of availability can be obtained by placing each IOP/CE pair in a separate location and connecting the pairs through a communications link.

Referring to Fig. 7, a distributed, high performance, fault resilient system 220 includes two
10 systems 200, the IOPs 202 of which are connected to each other, through IPI modules, by cables 222. System 220 uses distributed computing environment software to achieve high performance by running separate portions of an application on each system 200. System 220 is fault
15 tolerant and offers the ability to perform both hardware and software upgrades without service interruption.

Referring to Fig. 8, a fault tolerant system 230 includes three IOPs (232, 234, and 236) and three CEs (238, 240, and 242). Through IPI modules 244 and cables
20 246, each IOP is connected to an IPI module 244 of each of the other IOPs. Through IPI modules 248 and cables 250, each CE is connected to an IPI module 244 of two of the IOPs, with CE 238 being connected to IOPs 232 and 234, CE 240 being connected to IOPs 232 and 236, and CE
25 242 being connected to IOPs 234 and 236. Like system 200, system 230 allows for hardware upgrades without service interruption and software upgrades with only minimal service interruption.

As can be seen from a comparison of Figs. 7 and 8,
30 the CEs and IOPs of systems 200 and 230 are identically configured. As a result, upgrading a fault resilient system 200 to a fault tolerant system 230 does not require any replacement of existing hardware and entails the simple procedure of adding an additional CE/IOP pair,
35 connecting the cables, and making appropriate changes to

- 22 -

the system software. This modularity is an important feature of the paired modular redundant architecture of the invention.

Because the components of system 230 are triply
5 redundant, system 230 is more capable of identifying the source of a hardware fault than is system 10. Thus, while system 10 simply disables one or both of CEs 14 when an error is detected, system 230 offers a higher degree of fault diagnosis.

10 In one approach to fault diagnosis, as shown in Fig. 9, each IOP (232, 234, 236) of system 230 performs fault diagnosis according to a procedure 300. Initially, each IOP (232, 234, 236) checks for major faults such as power loss, broken cables, and nonfunctional CEs or IOPs
15 using well known techniques such as power sensing, cable sensing, and protocol timeouts (step 302). When such a fault is detected, each IOP disables the faulty device or, if necessary, the entire system.

After checking for major faults, each IOP waits to
20 receive IPI packets (that is, quantum interrupts or I/O requests) from the two CEs to which the IOP is connected (step 304). Thus, for example, IOP 232 waits to receive IPI packets from CEs 238 and 240. After receiving IPI packets from both connected CEs, each IOP transmits the
25 checksums ("CRCs") of those IPI packets to the other two IOPs and waits for receipt of CRCs from the other two IOPs (step 306).

After receiving the CRCs from the other two IOPs, each IOP may generate a three by three matrix in which
30 each column corresponds to a CE, each row corresponds to an IOP, and each entry is the CRC received from the

- 23 -

column's CE by the row's IOP (step 308). Thus, for example, IOP 232 generates the following matrix:

		<u>CE 238</u>	<u>CE 240</u>	<u>CE 242</u>
5	IOP 232	CRC	CRC	X
	IOP 234	CRC	X	CRC
	IOP 236	X	CRC	CRC

After generating the matrix, IOP 232 sums the entries in each row and each column of the matrix. If the three row sums are equal and the three column sums are equal (step 10 310), then there is no fault and IOP 232 checks again for major faults (step 302).

If either the three rows' sums or the three columns' sums are unequal (step 310), then IOP 232 compares the CRC entries in each of the columns of the 15 matrix. If the two CRC entries in each column match (step 312), then IOP 232 diagnoses that a CE failure has occurred and disables the CE corresponding to the column for which the sum does not equal the sums of the other columns (step 314).

20 If the CRC entries in one or more of the matrix columns do not match (step 312), then IOP 232 determines how many of the columns include mismatched entries. If the matrix includes only one column with mismatched entries (step 315), then IOP 232 diagnoses that the path 25 between the IOP corresponding to the matrix row sum that is unequal to the other matrix row sums and the CE corresponding to the column having mismatched entries has failed and disables that path (step 316). For purposes of the diagnosis, the path includes the IPI module 244 in 30 the IOP, the IPI module 248 in the CE, and the cable 250.

If the matrix includes more than one column with mismatched entries (step 314), then IOP 232 confirms that one matrix row sum is unequal to the other matrix row sums, diagnoses an IOP failure, and disables the IOP 35 corresponding to the matrix row sum that is unequal to the other matrix row sums (step 318).

- 24 -

If, after diagnosing and accounting for a CE failure (step 314), path failure (step 316), or IOP failure (step 318), IOP 232 determines that system 300 still includes sufficient non-faulty hardware to remain operational, IOP 232 checks again for major faults (step 5 302). Because system 230 is triply redundant, system 230 can continue to operate even after several components have failed. For example, to remain operating in an availability mode, system 230 only needs to have a single 10 functional CE, a single functional IOP, and a functional path between the two.

Using procedure 300, each IOP (232, 234, 236) can correctly diagnose any single failure in a fully operational system 230 or in a system 230 in which one 15 element (that is, a CE, an IOP, or a path) has previously been disabled. In a system 230 in which an element has been disabled, each IOP accounts for CRCs that are not received because of the disabled element by using values that appear to be correct in comparison to actually 20 received CRCs.

Procedure 300 is not dependent on the particular arrangement of interconnections between the CEs and IOPs. To operate properly, procedure 300 only requires that the output of each CE be directly monitored by at least two 25 IOPs. Thus, procedure 300 could be implemented in a system using any interconnect mechanism and does not require point to point connections between the CEs and IOPs. For example, the CEs and IOPs could be connected to at least two local area networks. In an alternative 30 approach, instead of summing the CRC values in the rows and columns of the matrix, these values can be compared and those rows or columns in which the entries do not match can be marked with a match/mismatch indicator.

A simplified version of procedure 300 can be 35 implemented for use in a system 200. In this procedure,

- 25 -

each IOP 202 of system 200 generates a two by two matrix in which each column corresponds to a CE 204 and each row corresponds to a IOP 202:

		CE 204	CE 204
5	IOP 202	CRC	CRC
	IOP 202	CRC	CRC

After generating the matrix, each IOP 202 attaches a mismatch indicator to each row or column in which the two entries are mismatched.

10 If there are no mismatch indicators, then system 200 is operating correctly.

 If neither row and both columns have mismatch indicators, then an IOP 202 has faulted. Depending on the operating mode of system 200, an IOP 202 either
15 disables another IOP 202 or shuts down system 200. The IOP 202 to be disabled is selected based on user supplied parameters similar to the two availability modes used in system 10.

 If both rows and neither column have mismatch
20 indicators, then a CE 204 has faulted. In this case, IOPs 202 respond by disabling a CE 204 if system 200 is operating in an availability mode or, if system 200 is operating in an integrity mode, shutting down system 200. If both rows and one column have mismatch indicators,
25 then one of the paths between the IOPs 202 and the CE 204 corresponding to the mismatched column has failed. Depending on the operating mode of system 200, IOPs 202 either disable the CE 204 having the failed path or shut down system 200. If both rows and both column have
30 mismatch indicators, then multiple faults exist and IOPs 202 shut down system 200.

 If one row and both columns have mismatch indicators, then the IOP 202 corresponding to the mismatched row has faulted. Depending on the operating
35 mode of system 200, the other IOP 202 either disables the faulty IOP 202 or shuts down system 200. If one row and

- 26 -

one column have mismatch indicators, then the path between the IOP 202 corresponding to the mismatched row and the CE 204 corresponding to the mismatched column has failed. Depending on the operating mode of system 200, 5 IOPs 202 either account for the failed path in future processing or shut down system 200.

Referring to Fig. 10, one embodiment of a disaster tolerant system 260 includes two fault tolerant systems 230 located in remote locations and connected by 10 communications link 262, such as Ethernet or fiber, and operating in meta time lockstep with each other. To obtain meta time lockstep, all IPI packets are transmitted between fault tolerant systems 230. Like system 220, system 260 allows for hardware and software 15 upgrades without service interruption.

As shown, the paired modular redundant architecture of the invention allows for varying levels of fault resilience and fault tolerance through use of CEs that operate asynchronously in real time and are 20 controlled by IOPs to operate synchronously in meta time. This architecture is simple and cost-effective, and can be expanded or upgraded with minimal difficulty.

Fig. 11 depicts the components of the system software architecture of an alternative embodiment of the 25 invention that includes multiple CEs and multiple IOPs. Each CE 1100 includes a set of one or more Physical Device Redirectors 1105, a CE Transport layer 1110, and an IPI Driver 1115. The Physical Device Redirectors 1105 intercept I/O requests directed to peripheral devices, 30 package the requests, and send the requests to the CE Transport 1110. The Physical Device Redirectors 1105 also receive responses to the requests from the CE Transport 1110, unpackage the responses, and return the responses to the operating system or applications

- 27 -

software that initially made the I/O requests to the peripheral devices.

The system software architecture depicted in Fig. 11 supports a multi-threaded processing environment. In such an environment, each thread is a separate stream of instructions that may be processed by the computing element. When a Physical Device Redirector 1105 intercepts an I/O request by a particular thread, processing of that thread stops until the Physical Device Redirector 1105 returns the response to the request.

The CE Transport 1110 communicates I/O requests between the CEs and the IOPs. The CE Transport also keeps track of responses expected and received from the IOPs on a per request basis. The CE Transport searches for completed requests (i.e., requests to which all IOPs have responded) as IOP responses to redirected requests are received and sends the resultant data to the Physical Device Redirector 1105.

Each IOP 1150 includes an IPI driver 1155, a state control program (SCP) 1160, a fault handler 1165, an IOP Transport layer 1170, a Device Synchronization Layer (DSL) 1175, and a set of one or more Physical Device Providers 1180. The IPI Drivers 1115, 1155 of both the CEs and the IOPs control the actual transmission and reception of data along the interconnect paths 1185 between the CEs and the IOPs. The state control program 1160 initiates and responds to state transitions involving the IOPs and the CEs. The fault handler 1165 responds to reports of detected faults by deconfiguring appropriate hardware components. The IOP Transport 1170 transports data between the IPI Driver and the other software components of the IOPs. The Device Synchronization Layer 1175 is responsible for synchronizing all redirected requests and responses between IOP-based peripherals. This synchronization

- 28 -

includes the reprocessing and restructuring of the requests and responses due to state transitions of the IOPs, the CEs, or the peripheral devices. The Physical Device Providers 1180 unpackage requests received from the DSL 1175, process the requests, package responses, and send the responses to the DSL 1175. Each IOP also may include an ethernet driver 1190 or other mechanism that permits direct communication between the IOPs.

As noted above, the architecture of the described embodiments requires redirection of all I/O requests from the CEs to one or more IOPs. Responses to these redirected requests must be synchronized. Accordingly, all IOPs responding to a redirected request must respond to the request in a known way and must describe how all other IOPs will respond. Moreover, an IOP must respond to a redirected request even if the IOP does not have a device capable of processing the request. These requirements permit software running on the CEs to compare the responses of the IOPs for consistency.

Figs. 12A and 12B provide an illustrative example of appropriate IOP responses (Fig. 12B) to requests by a CE (Fig. 12A). As shown, each disk of a two-disk shadow set is associated with a different IOP, with the disk 1200 associated with the first IOP 1205 functioning normally and the disk 1210 associated with the second IOP 1215 being offline. A request from the CE 1100 to the shadow set for data is redirected to the IOPs, and both IOPs respond. (The second IOP 1215 receives all requests directed to the shadow set even though its associated disk 1210 is offline.) Each IOP indicates that the first IOP 1205 is responding with the expected data and that the second IOP 1215 is responding with no data.

The CE Transport 1110 (Fig. 11) monitors the responses provided by the IOPs. The CE Transport compares the responses for consistency once all IOPs have

- 29 -

responded to a request. The responses must be consistent in indicating which IOPs responded to the request with actual data and which IOPs did not. Moreover, if more than one IOP responds with data, then the data provided
5 by the IOPs must be consistent. If the responses pass the consistency checks, then the CE Transport provides a response containing actual data to the appropriate Physical Device Redirector 1105 in the CE 1100 and discards the no-data response. Thereafter, the CE
10 Transport informs the IOPs that all IOPs have responded to the request and further informs the IOPs of the outcome of the comparison. If the responses do not pass the consistency checks, then one of the IOPs is disabled and, in effect, removed from the architecture.

15 The CE Transport 1110 implements the request-response architecture employed by the system of Fig. 11 using the procedure 1300 illustrated in Fig. 13. The CE Transport first instructs the IPI Driver 1115 to send a request from a Physical Device Redirector 1105 to
20 all IOPs 1150 (step 1305). At the same time, the CE Transport 1110 initializes a timer. If the timer has not expired (step 1310) before all responses are received (step 1315), then the CE Transport compares the responses to determine whether they are consistent (step 1320). If
25 the responses are consistent and there are multiple responses that include data (step 1325), then the CE transport compares the data responses (step 1330). If the data responses are consistent, or if there is only a single data response, then the CE Transport sends the
30 data to the appropriate Physical Device Redirector 1105 (step 1335). Finally, the CE Transport instructs the IPI driver to send a Response Complete message to all of the IOPs (step 1340).

 If the timer expires (step 1310) before all
35 responses are received, or if inconsistent responses are

- 30 -

received (steps 1320, 1330), then the CE Transport reports the occurrence of an error (step 1345) and instructs the IPI driver to send a Response Complete message to all of the IOPs (step 1340).

5 In the IOPs, the Device Synchronization Layer (DSL) 1175 makes all IOP-based physical devices appear to the CEs as logical or virtual devices. In addition, the DSL combines device and state information from all IOPs and uses this information to project a single logical IOP
10 to the Physical Device Redirectors of the CEs. Thus, a system configured as illustrated in Fig. 14A would appear to the CE-based Physical Device Redirectors as having the logical representation illustrated in Fig. 14B.

 The DSL 1175 represents devices logically so that
15 the CE-based redirectors have no knowledge of the locations and characteristics of physical devices. For example, a SCSI redirector would not know that Disk C: (1400, 1405) and Disk D: (1410, 1415) are shadowed drives and would simply treat them as though they were each a
20 single drive (1420, 1425). Also, an Ethernet redirector would not know that multiple Ethernet controllers provide a primary controller (1430) and a secondary controller (1435) and instead would be aware only of a single Ethernet controller (1440).

25 The DSL 1175 receives all redirected requests from the IOP Transport 1170. The DSL then decides whether to send a request to a Physical Device Provider 1180 or to send a no-data response. A critical function of the DSL is that it only sends a request to a device provider when
30 that provider is expected to process the request. For example, if two IOPs control a shadow set of disks, with the first IOP controlling an offline disk and the second IOP controlling an active disk, the DSL of the first IOP would not send a request to its disk, and would instead
35 respond with a no-data response. The DSL on the second

- 31 -

IOP would send the request to its disk, and would eventually respond with the data from its disk along with an indication that the first IOP would be providing a no-data response.

5 With every response, the DSL 1175 indicates how every other IOP will respond. The DSL accomplishes this by maintaining a record of the status of every other IOP and associated device in an internal device state table.

10 The DSL 1175 also accounts for the effects of state changes in the CEs, the IOPs or the peripheral devices. For example, the DSL accounts for the situation that arises when the state of a peripheral device changes after an IOP indicates how the other IOPs will respond to a request, but before the other IOPs actually respond to
15 the request. If the other IOPs were to respond after the state change, their responses could differ from the response expected by the first IOP. The DSL solves this problem by disabling all response transmissions from the affected device and initiating a flush sequence that
20 causes the CE Transport to discard all previously received, incomplete responses for the particular device. (As previously noted, the CE Transport does not consider a request to be complete until responses to the request have been received from all IOPs.)

25 Upon receipt of a flush completion indication from the CE Transport 1110, the DSL 1175 of each IOP 1150 updates its internal device state table and reenables transmission of responses for the particular device. Finally, the DSL reprocesses any incomplete responses
30 that were submitted to the CE Transport prior to the flush.

 As noted above, the CE Transport 1110 sends an indication to the DSL 1175 of each IOP upon completion of each response. As such, the DSL always knows which
35 responses are complete. This permits the DSL to keep

- 32 -

track of incomplete responses. When a device state change occurs and an ensuing flush is issued, the DSL can, upon receipt of the flush complete indication, reset its internal knowledge of the device state and re-issue any affected incomplete requests. This functionality is critical, for example, to the implementation of standby-primary processing since the DSL may have responded to a particular request on behalf of a standby device with a no-data response. If the primary device failed and was unable to process the request, the DSL would initiate the flush sequence and re-issue the request. This time the standby device would be considered the primary device (since the primary device had failed) and would receive the request.

The internal state table of the DSL 1175 keeps track of all flushes that can affect the I/O devices and the IOPs. The DSL, through use of the state table, permits multiple flushes to affect one or more devices, and permits multiple state transitions to occur at a single time. This permits seamless handling of multiple related or unrelated component state changes.

The DSL 1175 also mandates that all device state changes originate from the IOP 1150 that owns the device. This permits the DSLs of different IOPs to have a different simultaneous idea of the state of a device without risk that inconsistency problems will occur at the CE Transport 1110. This also permits the DSL to be fully distributed since there is no need to freeze request or response queues while a master software entity determines whether a steady state operation has been achieved.

In most circumstances, the Physical Device Providers 1180 need not consider the state of a peripheral device or an IOP because the DSL 1175 only sends a request to a device provider when the device

- 33 -

provider is expected to process the request. Similarly, the DSL does not consider the I/O policy associated with a particular device. For example, the DSL does not consider whether a disk device has a shadowed or a
5 single-ended I/O policy. However, the DSL does use I/O policies to determine which providers on which IOPs will receive a particular request to process. This permits the DSL to arbitrarily treat any device as shadowed, singled-ended, virtual, or primary/standby even though
10 all combinations of I/O policies and device types do not necessarily make sense.

The DSL 1175 handles all device state transitions, including device failure, device activation, device merge, and manual device enable/disable. In addition,
15 the DSL transparently handles all IOP state transitions, including IOP joining, IOP activation, IOP removal, and IOP graceful shutdown, as these transitions relate to device states. The DSL also responds automatically to requests that cannot be satisfied on behalf of the device
20 providers. The DSL provides a full featured application program interface (API) set that is useable by developers.

The DSL 1175 provides automatic request timeout support. In this regard, the DSL starts a recovery
25 process if a CE-originated request does not complete within a specified period of time. During the recovery process, the DSL determines which IOP 1150 has stalled and notifies the fault handler.

The DSL 1175 can apply any I/O policy to any
30 physical device. For example, the DSL can configure a hard disk as a single-ended device instead of a shadowed device. Similarly, the DSL can configure, for example, a CD-ROM player or a serial port as a standby/active device instead of as a single-ended device.

- 34 -

As discussed above, the DSL 1175 initiates a flush sequence in the CE Transport 1110 of one or more CEs 1100 in response to a state change in a peripheral device or an IOP 1150. The flush sequence causes the CE Transport
5 to flush all outstanding activity in the messaging pipeline of the corresponding CE with respect to indicated devices or Physical Device Providers 1180 and to provide notification of completed requests as a result of the flushing.

10 The flush mechanism provides varying granularity of request-response synchronization as required by the DSL. Thus, a system-wide flush can be implemented for all devices (or Physical Device Providers), or a flush can be implemented for a class of devices or a specific
15 device.

Upon completion of the flush sequence, the DSL of each IOP knows exactly which request-response pairs have been processed and completed. The DSL uses this knowledge to reissue, re-execute or re-transmit any
20 necessary request-response pairs to permit recovery from the transition (or stimulus) within the system that disrupted the steady state operation of one or more devices. This allows the DSL to react to changes in the system that affect steady state operation of the devices.
25 Once a flush sequence is initiated and completed, the DSL can determine exactly which requests or responses must be reprocessed or redirected to other devices for completion.

Each CE Transport 1110 maintains a database of
30 outstanding requests. The database includes a list of all outstanding I/O requests in the system, each identified by a unique identifier called a XRN (Transport Reference Number). The flush sequence is carried out according to the procedure 1500 illustrated in Fig. 15.
35 First, the DSL 1175 of each IOP 1150 initiates a flush

- 35 -

sequence by sending a flush request to all of the CEs 1100 (step 1505). The DSL then suspends all request/response processing activity for the devices that are involved in the flush sequence until the flush
5 sequence is completed (step 1510). Each CE receives a flush request and waits for matching flush requests from every IOP in the system (step 1515). The CE enters the flush request into the database and tracks the flush request in the same way that the CE tracks operating
10 system requests. When a flush request has been received from all IOPs, the CE indicates that the first portion of the flush sequence is completed. The completion of this portion of the flush sequence signifies that all activity that was in the IOP-to-CE message pipelines has been
15 flushed out and processed by the CEs.

Before acknowledging completion of the flush sequence, the CE first sends acknowledgments to the IOPs as to which outstanding requests have been completed (step 1520). In some cases, I/O requests may be
20 completed by flushing responses through the IOP-to-CE pipelines. The CE sends a SWTACK (Software Transaction Acknowledgment) for each request that was completed and removes the completed request from the database. Each SWTACK contains the request's original XRN. The XRN
25 allows the IOPs to associate the SWTACK with the proper completed request.

After sending a SWTACK for each completed request, the CE sends a SWTACK for the flush sequence (step 1530). Messages are delivered in order through the CE-to-IOP
30 message pipelines (or are reordered to reflect their transmission sequence). Accordingly, the flush SWTACK serves to flush the request completion notification SWTACKs through the CE-to-IOP message pipelines. Thus, when the DSL receives the SWTACK for the flush sequence
35 (step 1535), the DSL has already received and processed

- 36 -

all SWTACKs for requests that were originally in the message pipelines and have since completed (step 1525). Upon receiving the SWTACK for the flush sequence, the DSL knows the state of all request/response activity in the system. The DSL responds to this information by updating the state of the system (step 1540). Thereafter, the DSL resumes request-response activity for affected devices (step 1545), and resends any incomplete affected requests (step 1550). This re-synchronizes the devices against the new state to achieve a steady state of operation for the device or devices involved in the flush sequence.

Since each flush sequence is uniquely tagged with its own XRN, more than one flush can be in progress at a time. The flush processing software in the CE Transport and the DSL abides by certain rules as to how outstanding flush sequences are processed and in what order they are acknowledged. This allows the DSL to preempt or override previously-issued flush sequences due to subsequent or secondary transitions that might occur within the system.

The system software implements a freeze protocol to ensure that IOP-to-CE communications will not affect meta time synchronization of the CEs. As previously discussed, the IOPs operate asynchronously to the CEs and to other IOPs due to the asynchrony inherent in I/O devices. For this reason, communication between the CEs and the IOPs needs to occur in a way that will not disturb the meta time synchronization of the CEs. CE-to-IOP communication is synchronous to the CE instruction stream and will not affect CE lockstep as long as sufficient buffering is provided. However, IOP-to-CE communication is by nature asynchronous to the instruction stream of each CE. Accordingly, IOP-to-CE communication, if handled improperly, could affect each CE differently and result in divergence of the CE instruction streams. The freeze protocol serves to delay

- 37 -

processing of asynchronous data from the IOPs until the data can be handled synchronously by all CEs. In particular, the freeze protocol serves to delay implementation of the procedure 1300 by the CE Transports 5 1110 until all of the CE Transports 1110 are ready to implement the procedure 1300.

The freeze protocol provides four primary features: synchronized processing of input data streams from the asynchronous IOPs across the loosely 10 synchronized CEs; synchronized time updates to the CEs; a deep-freeze mechanism that allows an IOP to synchronously hold all CEs in a captive state for an extended period of time; and fault detection/diagnosis with respect to the communication paths between the CEs and the IOPs.

15 The freeze protocol provides CE synchronization using a so-called freeze cycle. The CE Transport 1110 of a CE may initiate a freeze cycle each time that it is activated by a Physical Device Redirector 1105 to service a redirected I/O operation. However, to prevent the use 20 of excessive bandwidth in performing freeze cycles, implementations of the software may initiate a freeze cycle every time that a certain number of I/O requests occurs or a certain number of instructions are processed without an I/O request. For example, the CE Transport 25 1110 may implement a freeze cycle with every fifth I/O request or every ten thousand instructions.

The CE Transport 1110 initiates a freeze cycle by transmitting a high-priority freeze request message to all active IOPs and waiting for freeze response messages 30 from all active IOPs. Since all of the CEs are processing the same instruction stream, the CE transport 1110 of each active CE will transmit a freeze request message to all active IOPs. Each IOP receives the freeze request messages from the CEs. When an IOP has received 35 a freeze request message from all active CEs, this

- 38 -

indicates that the CEs are at the same point in their instruction streams (i.e., in synchronization) and that it is permissible for the CEs to process received data using the procedure 1300. Accordingly, the IOP responds
5 by sending a freeze response message to all active CEs.

The CEs 1100 receive the freeze response messages from the IOPs and place the messages in the normal-priority message queues of the IPI Driver 1115 to provide an inter-CE synchronization point for the CE Transport
10 1110. Finally, after receiving freeze response messages from all active IOPs, the CEs terminate the freeze cycle by transmitting freeze release messages to the IOPs.

The CE Transport 1110 invokes the freeze cycle through the IPI Driver 1115. The IPI Driver responds by
15 initiating the freeze cycle (i.e., sending the freeze request message) and returning control to the CE Transport 1110. This permits the CE Transport 1110 to continue processing in parallel with the IPI Driver's handling of the freeze protocol messages. The CE
20 Transport performs whatever useful work it can, such as transmitting normal priority messages, or draining and processing messages from the receive queue of the CE 1100. The CE Transport does not return control to the operating system until the CE Transport has encountered a
25 freeze response message in the normal-priority receive queue corresponding to each active IOP.

The IPI Driver 1115 attempts to complete the freeze cycle as quickly as possible. To this purpose, the IPI Driver composes a priority message and sends the
30 message to all active IOPs 1150. The message carries no data to the IOPs and merely serves as an indication that the freeze cycle has started. The IPI Driver starts a freeze response timeout counter after sending the priority message to detect failure of any IOP to respond.
35 Typically, the length of this counter is on the order of

- 39 -

twice the worst-case message transmit time. This provides each IOP with sufficient time to complete transmission of any normal priority message that it may be transmitting upon receipt of the freeze request message before the IOP transmits a freeze response message. Typically, normal priority messages have a size limit of, for example, 64 kilobytes, to ensure that the worst-case message transmit time will have a reasonable value.

10 The IPI Driver 1155 of each IOP 1150 places a timestamp update in the freeze response message. This timestamp corresponds to the number of clock ticks since the last freeze response message transmitted by the IOP. Thus, a CE 1100 may receive different timestamp updates from different IOPs. Accordingly, the CEs designate one IOP 1150 as a meta-time server to ensure that all CEs will update their local time clock identically.

 After receiving a freeze request message from all active CEs, the IPI Driver 1155 of the IOP 1150 sends the freeze response message as a microcode-formatted priority message. Upon sending the message, the IOP initiates a freeze release timeout counter to detect failure of a CE to respond to the freeze response.

25 The IPI Driver 1115 of each CE is interrupted to service the incoming freeze response message, and responds by reformatting the message into a normal-priority message and placing the message at the bottom of the normal-priority message queue. Placement of the message in the message queue completes the time-critical portion of the freeze cycle, since it permits the CE Transport 1110 to return control to the operating system.

 Once the CE's IPI Driver 1115 has received and processed the freeze response from the last active IOP, 35 the IPI Driver 1115 broadcasts a freeze release message

- 40 -

to the active IOPs. This completes the freeze cycle for the CE 1100. The IOP 1150 receives the release message and cancels the freeze release timer upon receipt of a release message from every active CE 1100.

5 As noted above, time updates are provided to the operating system by including time increments in the freeze response packet. The IOP IPI Drivers 1155 maintain the time increments using an internal 100 μ s timer interrupt. Only one of the IOPs is designated as
10 the time provider, and the time increment from that IOP is used by all of the CEs. The IOP's IPI Driver 1155 transmits the delta time since the last freeze response packet was transmitted as the time increment. The CE Transport 1110 uses this value to update the time-of-day
15 clock of the operating system after a freeze response packet has been processed for all receive queues.

A deep freeze protocol is a variation of the normal freeze protocol and serves to suspend normal activity between the CEs and the IOPs to allow major
20 system state transitions to occur. The deep freeze state is invoked by the IOP software, which uses the IOP Transport 1170 to command the IOP IPI Driver 1155 to replace the next freeze response message with a deep freeze response message. The format of the deep freeze
25 response message is identical to that of the normal freeze response message with the exception that different opcodes are used to designate the two types of messages. The deep freeze response carries a meta-time update like a normal freeze response. In addition, the deep freeze
30 response causes the initiating IOP to disable its transmitter without initiating a timeout counter.

The IOP's IPI Driver 1155 sends the deep freeze response to all active CEs to inform them that a deep freeze state is requested. The CEs respond by converting
35 the deep freeze response to a normal-priority response

- 41 -

message and adding the message to the normal-priority message queue, with an indication that the message is a deep freeze response instead of a normal freeze response. The CEs continue to process normal freeze response
5 messages from the other IOPs in the normal manner.

After receiving a normal freeze response or a deep freeze response from all active IOPs, the CE's IPI Driver 1115 sends a deep freeze request message to the IOPs that are not yet in the deep freeze state and restarts the
10 freeze response timer.

Receipt of a deep freeze request informs the IOPs that another IOP has injected a deep freeze cycle into the current freeze cycle. (Normally, the IOP IPIs would have received a freeze release message.) Each IOP
15 responds to the freeze request by cancelling the freeze release timeout counter, turning off the IOP's transmitter path, sending to all active CEs a deep freeze response message with a meta-time update since the previous freeze response, and restarting the freeze
20 release timeout counter.

The IPI Drivers 1155 of the CEs receive the deep freeze responses and insert them into the appropriate receive queues. Upon receipt of a deep freeze response from each active IOP, the CE Transport 1110 cancels the
25 response timer and issues the normal freeze release message to the IOPs.

The CE Transport 1110, in the mean time, has encountered a combination of normal freeze response messages and deep freeze response messages in the various
30 receive queues. Detection of a single deep freeze response causes the CE transport to process beyond the normal freeze response to the deep freeze response. Hence, the CE IPI Driver 1115 must ensure that both the freeze responses and the deep freeze responses make their
35 way into the receive queues in the proper order.

- 42 -

Only priority messages can be exchanged between components when the system is in the deep freeze state. In addition, no new freeze cycles will be initiated by the CE Transport while the system is in the deep freeze
5 state.

Once the major state transition requiring the deep freeze cycle has completed, the deep freeze state is terminated by initiating a deep freeze termination cycle with the issuance of a deep freeze termination request.
10 The deep freeze termination cycle typically is originated by the IOP Transport 1170 that invoked the deep freeze cycle, but any IOP Transport 1170 can invoke a deep freeze termination cycle with the same result.

The deep freeze termination request is a
15 register-level request to the IPI Driver 1155 that causes the driver to broadcast a deep freeze termination priority message to all active CEs. The IPI Driver 1115 of each CE receives this message and echoes the message to all active IOPs. Upon receiving the deep freeze
20 termination message, the IOPs activate their transmission paths and exit the deep freeze state.

The freeze response timers and the freeze release timers account for the occurrence of errors during the freeze protocol. If a CE's freeze response timer
25 expires, the CE generates a high priority system error (SYSERR) packet containing the CE's freeze status virtual register and sends it to all active IOPs. The IOPs supplement the SYSERR packet with relevant state information of their own and forward the SYSERR packet to
30 the Fault Handlers 1165.

If an IOP's freeze release timer expires, the IOP generates a local SYSERR packet containing only state information from the detecting IOP and passes this SYSERR packet to the Fault Handler 1165. The detecting IOP then
35 sends a similar SYSERR packet to all active CEs. The CEs

- 43 -

supplement the SYSERR packet with their own state information and echo the SYSERR packet back to all active IOPs. The IOPs store additional state information in the echoed SYSERR and forward the message to the Fault
5 Handler 1165.

After generating a SYSERR packer, the IPI Driver 1115, 1155 of the CE 1100 or the IOP 1150 waits for the Fault Handler 1165 residing on the IOP 1150 to resolve the error condition. The Fault Handler determines which
10 path is at fault and commands the IOP IPI Driver 1155 to disable the faulty path. After disabling the faulty path, the IOP IPI Driver evaluates the state of the freeze protocol and resumes normal processing once the requirements for such processing are met.

15 The IOPs do not apply a timeout for the reception of freeze requests from all active CEs. If one or more CEs fail to send a freeze request message, or if a freeze request message is not received for some other reason, the other CEs will eventually generate a freeze response
20 timeout SYSERR packet.

The CE IPI Driver 1115 state transitions required to service the freeze protocol, including the deep freeze extension, are illustrated in Fig. 16A. The IOP IPI
Driver 1155 state transitions required to service the
25 freeze protocol, including the deep freeze extension, are illustrated in Fig. 16B.

The system implements a message based fault notification and reporting environment using Fault
Handler 1165. From an error-processing perspective, the
30 system can be viewed of as including error reporting elements and error processing elements. The error reporting elements are any components in the system that detect error conditions or determine status conditions and transmit the information to an error processing
35 element. The error processing elements, or fault

- 44 -

handlers, receive error information from the reporting elements.

The collection of error messages that result from a single fault are referred to as a fault event. A fault handler uses the error messages associated with a fault event to identify a particular system component that has failed and caused the fault event. The identified component is referred to as a callout. A fault handler may also take or initiate action to resolve the fault and restore normal, if degraded, system operation.

Error reporting elements may be either hardware or software entities. The only requirement is that they must be capable of transmitting error information to the error processing elements or causing such information to be transmitted.

An error reporting element that detects a system error encapsulates the system error into a uniformly formatted packet referred to as a SYSERR. The error reporting element then sends the SYSERR (or causes the SYSERR to be sent) to all error processing elements that have connectivity to the reporting element. The architecture of the physical system is configured so that any error condition causes generation of at least one SYSERR packet. The physical system is also configured so that, ideally, each error reporting element is connected to every error processing element. This permits the lack of an error indication when one was expected to be used as diagnostic information.

A single SYSERR packet may not unambiguously identify the source of a fault event in many instances. When this situation arises, the fault handlers rely on the diagnostic information provided by SYSERR packets from multiple sources to unambiguously identify the source of the fault.

- 45 -

To provide connectivity between the error reporting elements and the fault handlers, some transport components are capable of echoing SYSERRs generated by a reporting element to other fault handlers in the system
5 that are directly connected to the transport components. Thus, fault handlers that are not directly connected to a reporting element can still obtain error information from that element. For example, as shown in Fig. 17, a fault handler 1700 on an IOP 1705 cannot directly receive a
10 SYSERR packet generated by an IOP 1710. To account for this, the IPI adapter 1715 on the CE 1720 serves as a SYSERR reflector and echoes a SYSERR produced by IOP 1710 to IOP 1700.

A fault handler 1165 may also be able to probe
15 other system components to obtain error information. In addition, separate fault handlers may communicate with each other to probe the viability of the system components on which they reside, to test the communication paths between the system components, and,
20 assuming that the communication paths are intact, to ensure that each fault handler reaches the same diagnosis in response to a fault event.

A fault handler 1165 groups encountered errors into sets of co-related errors, referred to as syndromes.
25 Syndromes generally indicate a faulty component, or a suspected list of faulty components, with more specificity than individual errors in the syndrome are able to provide.

Each fault handler uses a state table to parse the
30 incoming errors into specific syndromes. Each syndrome represents a state in the table. If possible, the fault handler uses an incoming error to transition the state table to a new state.

The fault handler processes errors according to
35 the procedure 1800 illustrated in Fig. 18. Initially,

- 46 -

the fault handler represents each error as a canonical error by converting the error to a normal form that uniquely identifies the error (step 1805). For example, the fault handler might convert the error to a triplet
5 that includes an error identifier that identifies a particular error, an error target that identifies the subcomponent about which the error is complaining, and a reporting source that identifies the subcomponent that reported the error and the path over which the error was
10 received.

The fault handler then processes the canonical error. First, the fault handler compares the error against states represented by previously established syndromes (step 1810) to see if the error will transition
15 the state of a previously established syndrome (step 1815). If such a transition can be performed, then the fault handler performs the transition (step 1820) and concludes state processing for the error (step 1825). If the error cannot transition any existing syndrome, the
20 fault handler creates a new syndrome at an INITIAL state (step 1830) and determines whether the error can transition the syndrome to a starting state of any syndrome (step 1835). If so, the fault handler performs the transition (step 1820) and concludes state processing
25 for the error (step 1825). If the error is not anticipated by a starting state of a syndrome, the fault handler converts the error to an unsyndromed error (step 1840) and concludes state processing for the error (step 1845). An unsyndromed error is a catch-all for errors
30 that are unanticipated, misreported, or incompletely transmitted or received. For example, if canonical errors are in the triplet form described above, then a canonical unsyndromed error might consist of an identifier for "UNSYNDROMED", an error target
35 corresponding to the reporting source of the bogus error,

- 47 -

and a reporting source of the component containing the fault handler.

A set of potential callouts is associated with each state in the state table. When a syndrome
5 transitions to a particular state, the set of callouts is referred to as the callout list of the syndrome.

The state table is a tree-structured list of canonical errors. The tag on each node of the tree is an identifier for a canonical error. Each node points to a
10 list of other errors, or to a list of callouts, or to both.

The state table can be created from a source document that consists of groups of errors (syndromes). The syndrome text syntactically indicates whether one
15 error should occur before another in a syndrome. If no such indication is provided, then the errors are assumed to be unordered. A list of callouts is associated with each syndrome.

The state table is created by permutating each
20 syndrome's errors, converting each error to canonical form, and mapping the syndrome into the table. The terminal node of each permutation points to the callout list for the syndrome. For example, suppose the source table contains the syndromes with associated callouts
25 illustrated in Fig. 19A, and the system topology is assumed to be that of Fig. 17. The "Error" column of Fig. 19A identifies different errors reported to the fault handler in a SYSERR message. In particular, a NAK error indicates excessive retries on a transmit path, an
30 EDC error indicates a low-level protocol failure on a receive path, and on a NAK (echoed) error is a NAK error transmitted by one component and echoed by another. The "Path" column identifies the path on which an error was detected, and the "Received from" column identifies the
35 machine that reported the SYSERR (M1 indicates IOP 1705

- 48 -

(machine 1) and M2 indicates CE 1720 (machine 2)).
Finally, the "Callouts" column provides a list of possible failures in the system that could cause the collection of errors listed in the "Error" column.

5 Using the source table of Fig. 19A for syndrome #1 to be detected by the fault handler, a NAK error reported against path M1-M2 by M1, a NAK error reported against path M1-M2 by M1 echoed from M2, and an EDC error reported against path M1-M2 by M2 must all occur. Then
10 the resulting callout is all of the entries in the "Callout" column for syndrome #1.

 The state transition table resulting from the above source would then appear as illustrated in Fig. 19B, where the canonical form error designations are in
15 the triplet form: error(error_path, received_from). The above state transition table is really a tree structure. The root is the "Initial state". Starting from the initial state there are six ways of arriving at a Syndrome #1 callout. There are three required error
20 messages to arrive at a syndrome #1 callout and they can arrive in any possible order (3 factorial = 6 possible orderings). Syndrome #2 callouts require two errors in any order (2 factorial = 2 possible orderings). Thus a NAK (M1-M2, M1 v. M2) and an EDC (M1-M2, M2) arriving in
25 any order without the third error NAK (M1-M2, M1) produces a syndrome #2 callout. In this case, the existence of the NAK (M1-M2, M1) uniquely identifies syndrome #1.

 The state table is constructed by permutating sets of errors and, therefore, can become very large. This is
30 particularly true if the syndromes comprising the table are particularly complex or if there are a large number of system components generating errors. The table size may be reduced by logical to physical mapping. In a system with redundant components, errors reported against
35 one component are indistinguishable (up to unit

- 49 -

identification) from errors reported against the
redundant set that includes that component. Accordingly,
the table size can be reduced by identifying errors and
callouts in logical form so that the table is, in effect,
5 reflective of errors corresponding to only one component
of a redundant set. With this approach, a mapping of
physical to logical identifiers is maintained for each
syndrome as the syndrome is constructed. When a
canonical error is compared against the errors in a state
10 of a syndrome, the error must be further transformed into
a logical canonical form relative to that syndrome. The
logical to physical mapping may vary from syndrome to
syndrome.

The size of the state table may also be reduced by
15 subtree folding. Many parts of the lower structure of
the tree-structured state table are identical to other
parts. Identical subtrees can be collapsed into a single
copy, even if they originate from different syndromes.
Performing this optimization tends to mitigate the
20 explosive growth of the state table as syndrome
complexity increases: larger syndromes generate more
duplicate subtrees that can be eliminated by subtree
folding.

Each syndrome indicates zero, one, or more
25 potential faulty components, or callouts, in the system.
The fault handler produces a single diagnosis by
combining these indications into an event callout list.
Ideally, the callout list includes only a single element
that unambiguously identifies a system component or
30 function that failed and precipitated the error event.
The callout list may, however, contain more than one
element.

The final callout list is formed by taking the
intersection of the most-likely callouts associated with
35 each syndrome. Some of the faults indicated by a

- 50 -

syndrome are more likely than others. For example, a transmission error may result from a transient point-to-point error, a common transmitter or common receiver fault, cable integrity problems, or incipient
5 power failure on one end. Incipient power failure could explain almost any erroneous behavior. However, without corroborating evidence, a transient fault is a more likely explanation for a single transmission failure than is an incipient power failure. In addition, some faults
10 identify components less specifically and can be presumed to identify more specific components subsumed therein. For example, the callout for one syndrome may specify an interconnect failure between IOP 1705 and CE 1720. This callout would include, for example, the IPI adapters 1715
15 at both ends of the interconnection as well as the cabling 1725 between them. Another syndrome may indicate a non-specific error on CE 1720 (that is, all hardware and software components of CE 1720, including the IPI adapter 1715 of CE 1720). The combination of the callout
20 lists of these syndromes clearly indicates that the IPI adapter 1715 of CE 1720 is faulty.

Two callouts are combined by selecting a system component that is common to the two callouts, but least encompassing. For any two callouts, there is at least
25 one other callout that contains an intersection of the system components covered by the first two. The product of the combination of callouts is such a component that covers the least part of the system. For example, suppose that in the loosely-connected system shown in
30 Fig. 17 there is a callout against the interconnection of IOP 1705 and CE 1720 and also against the interconnection of IOP 1710 and CE 1720. These interconnections intersect at the IPI adapter 1715 in CE 1720. They also could be considered to intersect at CE 1720, in general.
35 However, the adapter is the callout chosen since this is

- 51 -

the least-encompassing component that can be identified as the intersection of the faulty interconnections.

The approach described above for combining callouts results in there being only one product callout for each combination of two other callouts. Accordingly, a "multiplication table" could be established for use by the fault handler in quickly establishing a new callout from any two others. The multiplication table is formed by creating a callout hierarchy diagram such as is illustrated in Fig. 20. Each box in the diagram represents a callout, and the product of two callouts is defined as the lowest callout (in the hierarchy) common between the two. The exception to this rule is that a callout combined with itself is itself (i.e., all entries are returned instead of the lowest entry). Thus, for example:

$CXUux * CXUvx = RX(CEx),$
 $RX(CEx) * TX(CEx) = IPI(CEx),$
 $RX(IOPu) * Capability(IOPu) = Power(IOPu),$ and
 $Device(IOPu) * Capability(CEx) = Tuple Power,$

while

$CXUux * CXUux = CXUux.$

NOTE: In the discussions that follow, the abbreviations being used are:

CNux-unidirectional inconnection between Machine u and Machine x
 CBux-bidirectional inconnection between Machine u and Machine x

The multiplication table formed by this method can be very large, especially in systems with large numbers of redundant components. Fortunately, it is not necessary to actually form the multiplication table. Rather, the multiplication can be performed by following the same rules for two callouts as would be required in computing the table. All that is required is an instantiation of

- 52 -

the directed graph comprising the callout hierarchy diagram. Furthermore, redundant callouts in the system may be eliminated by representing the callouts in logical form. (However, the logical callouts must be mapped to
5 physical callouts when the multiplication is performed.)

The fault handler never combines callouts within a list. Instead, the fault handler combines two callout lists by "cross-multiplying" the lists to form a third list, where the third list contains all unique
10 combinations of the original lists. For example, cross-multiplication of the list 1 and list 2 elements of Fig. 21A would result in the product elements illustrated in Fig. 21B. Eliminating duplicate callouts results in a product list of Common Machine 1/2 Power Supply; Machine
15 1 Power; Machine 2 Power; Machine 1 Adapter; Machine 2 Adapter; and 1 ↔ 2 Bidirectional Interconnect.

Callout lists from independent sources can be combined in this same manner. For example, suppose that a fault handler 1700 exists on each of IOP 1705 and 1710
20 in the system shown in Fig. 17. If communication exists between the two fault handlers, then the callout lists from the two fault handlers can be combined into a callout list for the entire system. This combined list will have equal or greater specificity than the callout
25 list from each fault handler standing taken independently.

After all syndrome callout lists are combined, nonsensical callouts are removed from the list. This is done by ranking each callout according to the scope of
30 compromise to system operation that it implies. The final callout list (i.e., the diagnosis) is formed by selecting callouts of only the lowest rank (lower rank is arbitrarily chosen to apply to callouts of lesser scope). For example, with the callout ranking illustrated in Fig.
35 22, the callout list of Fig. 21B can be reduced to a

- 53 -

single callout: Machine 1 ↔ Machine 2 Bidirectional Interconnect.

The ranking procedure involves a probabilistic assumption. When all else is equal, callouts of lower rank are more probably the cause of a fault event than are callouts of a higher rank. In the above example, it is possible that an incipient power failure in the common power supply between IOP 1700 (machine 1) and CE 1710 (machine 2) evoked the syndromes, but without corroborating evidence (in the form of another syndrome, or another error that would modify one of the existing syndromes) it would be misleading to report this as a callout.

Callouts represent the finest granularity of diagnosis that the fault handler is capable of producing. Generally, callouts are too specific to be of use to service personnel. However, they can be mapped to Field Replaceable Units (or FRUs) that represent the subcomponents of the system that can be identified for service or replacement. For instance, the fault handler may be capable of identifying a common receive port on an interconnect adapter. The FRU corresponding to this callout would be the adapter card.

In summary, a Fault Handler 1165 identifies a faulty FRU according to the procedure 2300 illustrated in Fig. 23. Starting with the callout list from a first syndrome (step 2305), the Fault Handler 1165 maps the logical callout to a physical list (step 2310). The Fault Handler 1165 then identifies this physical list with the designation CURRENT (step 2315).

If there are more syndromes to process (step 2320), the Fault Handler 1165 converts the logical callout list for the new callout to physical callouts (step 2325), cross-multiplies the list against CURRENT

- 54 -

(step 2330) and stores the resulting product as CURRENT (step 2335).

Once all syndromes have been processed (step 2320), the Fault Handler 1165 reduces CURRENT to the most
5 specific callout or callouts (step 2340), maps the callout or callouts to one or more FRUs (step 2345) and concludes the procedure (step 2350).

Other embodiments are within the scope of the following claims.

- 55 -

What is claimed is:

1. A method of synchronizing data transfer to a computing element in a computer system, the computer system including the computing element and controllers
5 that provide data from data sources to the computing element, said method comprising the steps of:
intercepting a request for data made by a computing element;
transmitting the intercepted request to the
10 controllers; and
responding, through the controllers, to the intercepted request, wherein at least a first controller responds by transmitting requested data to the computing element and by indicating how a second controller will
15 respond to the intercepted request.
2. The method of claim 1, further comprising the step of responding, through the second controller, to the intercepted request by indicating that the second controller has no data corresponding to the intercepted
20 request and by indicating that the first controller will respond to the intercepted request by transmitting data to the computing element.
3. The method of claim 1, further comprising the step of including, in each response to the intercepted
25 request by a controller, an indication as to how each other controller will respond to the intercepted request.
4. The method of claim 1, further comprising the step of comparing the responses to the intercepted request for consistency.
- 30 5. The method of claim 4, further comprising the step of including in each response an indication as to

- 56 -

how each other controller will respond to the intercepted request,

wherein the step of comparing comprises comparing the indications for consistency.

5 6. The method of claim 5, wherein, when responses of multiple controllers include requested data, the step of comparing further comprises comparing the data for consistency.

10 7. The method of claim 4, further comprising the steps of:

 comparing, through the computing element, the responses for consistency after the computing element has received responses from all of the controllers, and
 notifying the controllers, through the computing
15 element, of the outcome of the comparison and that responses have been received from all of the controllers.

 8. The method of claim 4, further comprising the step of disabling one of the controllers when the responses are not consistent.

20 9. The method of claim 1, further comprising the step of generating an error condition if the computing element does not receive responses from all of the controllers within a predetermined time period.

 10. The method of claim 1, wherein a data source
25 is associated with the first controller, the method further comprising the step of obtaining, in response to the intercepted request and through the first controller, the requested data from the data source.

- 57 -

11. The method of claim 1, further comprising the steps of:

maintaining, through the first controller, a record of a status of the second controller, and

5 using the record when indicating how the second controller will respond to the intercepted request.

12. The method of claim 11, wherein a data source is associated with the second controller, and further comprising the steps of:

10 maintaining, through the first controller, a record of a status of the data source, and

using the record when indicating how the second controller will respond to the intercepted request.

13. The method of claim 11, further comprising
15 maintaining, through each controller, records of statuses of all other controllers and having controllers use the records when indicating how the other controllers will respond to the intercepted request.

14. The method of claim 13, further comprising:
20 associating each controller with a data source;
maintaining, through each controller, records of statuses of data sources associated with all other controllers; and

having controllers use the records when indicating
25 how the other controllers will respond to the intercepted request.

15. The method of claim 12, further comprising,
when a status of the data source associated with the second controller changes, transmitting, through the
30 second controller, an instruction to the computing

- 58 -

element to discard responses to the intercepted request from other controllers.

16. The method of claim 15, further comprising responding, through the computing element, to the
5 instruction by discarding responses from other controllers to the intercepted request and by transmitting to the controllers a notification that the responses have been discarded.

17. The method of claim 16, further comprising
10 responding, through the first controller, to the notification by updating the record of the status of the data source associated with the second controller.

18. The method of claim 17, further comprising, after updating the record, retransmitting, through the
15 first controller, the requested data to the computing element and indicating, through the first controller, how the second controller will respond to the intercepted request.

19. The method of claim 1, wherein a data source
20 is associated with each controller and each of the controllers responds to the intercepted request by:

determining whether an associated data source is expected to process the request;

when the associated data source is expected to
25 process the request, transmitting the request to the associated data source, receiving results of the request from the associated data source, and forwarding the results of the request to the computing element; and

when the associated data source is not expected to
30 process the request, informing the computing element that no data will be provided in response to the request.

- 59 -

20. A method of synchronizing data transfer to computing elements in a computer system, the computer system including the computing elements, data sources, and controllers that provide data from the data sources
5 to the computing elements, said method comprising the steps of:

- intercepting a request for data made by a computing element;
- transmitting the intercepted request to the
10 controllers; and
- responding to the intercepted request, through each of the controllers, by:
 - determining whether an associated data source is expected to process the request;
 - 15 when the associated data source is expected to process the request, transmitting the request to the associated data source, receiving results of the request from the associated data source, and forwarding the results of the request to the computing element; and
 - 20 when the associated data source is not expected to process the request, informing the computing element that no data will be provided in response to the request.

21. A method of maintaining synchronization between computing elements processing identical
25 instruction streams in a computer system, the computer system including the computing elements and controllers that provide data from data sources to the computing elements, wherein the controllers operate asynchronously to the computing elements, said method comprising the
30 steps of:

- at computing elements processing identical instruction streams, each computing element: stopping processing of the instruction stream at a common point in

- 60 -

the instruction stream, generating a freeze request message, and transmitting the freeze request message to controllers;

at a controller: receiving a freeze request
5 message from a computing element, waiting for a freeze request message from other computing elements, and, upon receiving a freeze request message from each computing element processing an identical instruction stream, generating a freeze response message and transmitting the
10 freeze response message to the computing elements; and

at the computing elements processing identical instruction streams, each computing element: upon receiving a freeze response message from a controller, waiting for freeze response messages from other
15 controllers to which a freeze request message was transmitted, and, upon receiving a freeze response message from each said controller, generating a freeze release message, transmitting the freeze release message to controllers, and resuming processing of the
20 instruction stream.

22. The method of claim 21, wherein the common point in the instruction stream corresponds to an I/O operation.

23. The method of claim 21, wherein the common
25 point in the instruction stream corresponds to occurrence of a predetermined number of instructions without an I/O operation.

24. The method of claim 21, wherein the step of
generating a freeze response message comprises including
30 a time update in the freeze response message, the method further comprising having a computing element, upon receiving a freeze response message from each controller

- 61 -

to which a freeze request message was transmitted, updating a system time using the time update from a freeze response message.

25. The method of claim 24, wherein the step of
5 updating a system time comprises using the time update from a freeze response message generated by a particular controller.

26. The method of claim 21, further comprising
10 having a computing element, upon receiving a freeze response message from each controller to which a freeze request message was transmitted, processing data received from a controller prior to receipt of a freeze response message.

27. A method of handling faults in a computer
15 system, the computer system including error reporting elements and error processing elements, the method comprising:

detecting, through an error reporting element, an error condition and transmitting information about the
20 error condition as an error message to error processing elements connected to the error reporting element, and retransmitting the error message, through at least one error processing element, to other error processing elements connected to the at least one error processing
25 element.

28. A method of handling faults in a computer system, the computer system including error reporting elements and error processing elements, the method comprising:

30 detecting, through error reporting elements, an error condition and transmitting information about the

- 62 -

error condition as error messages to error processing elements connected to the error reporting elements, and combining, through at least one error processing element, information from related error messages from multiple error reporting elements and using the combined information in identifying a source of the error condition.

29. The method of claim 28, wherein the at least one error processing element uses a state table to combine information from related error messages.

30. The method of claim 29, wherein the at least one error processing element represents an error message using an error identifier that identifies a particular error, an error target that identifies a subcomponent that caused the error represented by the error message, and a reporting source that identifies an error reporting element that generated the error message and a path over which the error message was received.

31. The method of claim 29, wherein the error processing element determines whether error messages are related by comparing a received error message against states representing previously received error messages.

32. A computer system including:
a computing element,
data sources, and
controllers that provide data from the data sources to the computing element, wherein:
the computing element is configured to intercept a request for data made by software running on the computing element and to transmit the intercepted request to the controllers; and

- 63 -

at least a first controller is configured to respond to the intercepted request by transmitting requested data to the computing element and by indicating how a second controller will respond to the intercepted
5 request.

33. The system of claim 32, wherein the second controller is configured to respond to the intercepted request by indicating that the second controller has no data corresponding to the intercepted request and by
10 indicating that the first controller will respond to the intercepted request by transmitting data to the computing element.

34. The system of claim 32, wherein each controller is configured to include in response to the
15 intercepted request an indication as to how each other controller will respond to the intercepted request.

35. The system of claim 32, wherein the computing element is configured to compare the responses to the intercepted request for consistency.

20 36. The system of claim 35, wherein the computing element is configured to compare the responses for consistency after the computing element has received responses from all of the controllers and to notify the controllers of the outcome of the comparison and that
25 responses have been received from all of the controllers.

37. The system of claim 32, wherein the computing element is configured to generate an error condition if the computing element does not receive responses from all of the controllers within a predetermined time period.

- 64 -

38. The system of claim 32, wherein:
a first data source is associated with the first
controller, and
the first controller is configured to obtain the
5 data requested in the intercepted request from the first
data source.

39. The system of claim 32, wherein the first
controller is configured to maintain a record of a status
of the second controller and to use the record when
10 indicating how the second controller will respond to the
intercepted request.

40. The system of claim 39, wherein:
a first data source is associated with the second
controller, and
15 the first controller is configured to maintain a
record of a status of the first data source and to use
the record when indicating how the second controller will
respond to the intercepted request.

41. The system of claim 40, wherein the second
20 controller is configured to transmit an instruction to
the computing element to discard responses from other
controllers to the intercepted request when a status of
the first data source changes.

42. The system of claim 41, wherein the computing
25 element is configured to respond to the instruction by
discarding responses from other controllers to the
intercepted request and by transmitting to the
controllers a notification that the responses have been
discarded.

- 65 -

43. The system of claim 42, wherein the first controller is configured to respond to the notification by updating the record of the status of the first data source.

5 44. The system of claim 43, wherein the first controller is configured, after updating the record, to retransmit the requested data to the computing element and to indicate how the second controller will respond to the intercepted request.

10 45. The system of claim 32, wherein:
 a data source is associated with each controller,
 and
 each controller is configured to respond to the intercepted request by:
15 determining whether an associated data source is expected to process the request;
 when the associated data source is expected to process the request, transmitting the request to the associated data source, receiving results of the request
20 from the associated data source, and forwarding the results of the request to the computing element; and
 when the associated data source is not expected to process the request, informing the computing element that no data will be provided in response to the
25 request.

 46. A computer system including:
 a computing element,
 data sources, and
 controllers that provide data from the data
30 sources to the computing element, wherein:
 the computing element is configured to intercept a request for data made by software running on the

- 66 -

computing element and to transmit the intercepted request to the controllers; and

each controller is configured to respond to the intercepted request, by:

5 determining whether an associated data source is expected to process the request;

 when the associated data source is expected to process the request, transmitting the request to the associated data source, receiving results of the request
10 from the associated data source, and forwarding the results of the request to the computing element; and

 when the associated data source is not expected to process the request, informing the computing element that no data will be provided in response to the
15 request.

47. A computer system including:

computing elements configured to process identical instruction streams,

data sources, and

20 controllers configured to provide data from the data sources to the computing elements and to operate asynchronously to the computing elements, wherein:

each of the computing elements is further configured to:

25 stop processing of the instruction stream at a common point in the instruction stream,

generate a freeze request message, and

transmit the freeze request message to the controllers;

30 a controller is further configured to:

receive a freeze request message from a computing element,

wait for a freeze request message from other computing elements, and,

- 67 -

upon receiving a freeze request message from each computing element, generate a freeze response message and transmit the freeze response message to the computing elements; and

5 each of the computing elements is further configured to:

upon receiving a freeze response message from a controller, to wait for freeze response messages from other controllers to which a freeze request message was
10 transmitted, and,

upon receiving a freeze response message from each said controller, to generate a freeze release message, transmit the freeze release message to the controllers, and resume processing of the instruction
15 stream.

48. A computer system including:

error reporting elements, and
error processing elements, wherein:

an error reporting element is configured to detect
20 an error condition and transmit information about the error condition as an error message to error processing elements connected to the error reporting element, and
at least one error processing element is
configured to retransmit the error message to other error
25 processing elements connected to the at least one error processing element.

49. A computer system including:

error reporting elements, and
error processing elements, wherein:

30 error reporting elements are configured to detect an error condition and transmit information about the error condition as error messages to error processing elements connected to the error reporting elements, and

- 68 -

at least one error processing element is
configured to combine information from related error
messages from multiple error reporting elements and use
the combined information in identifying a source of the
5 error condition.




FIG. 1


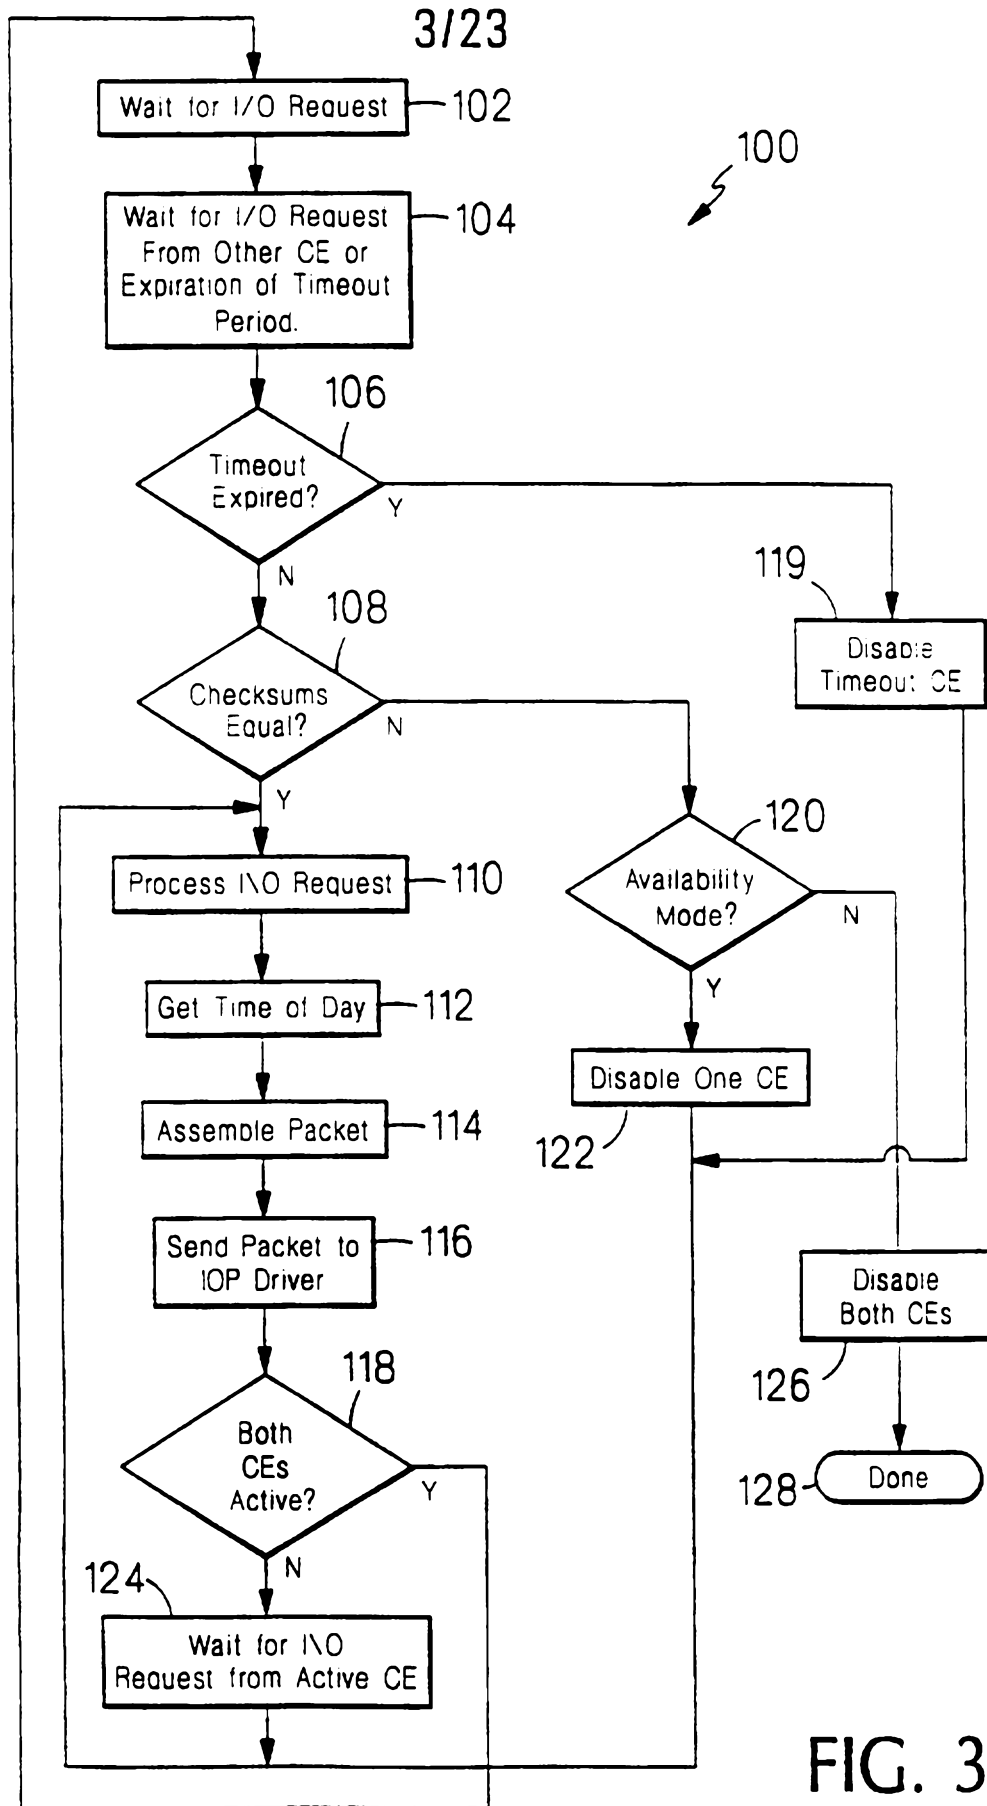


FIG. 2



4/23




FIG. 4




FIG. 5

6/23




FIG. 6




FIG. 7

7/23




FIG. 8

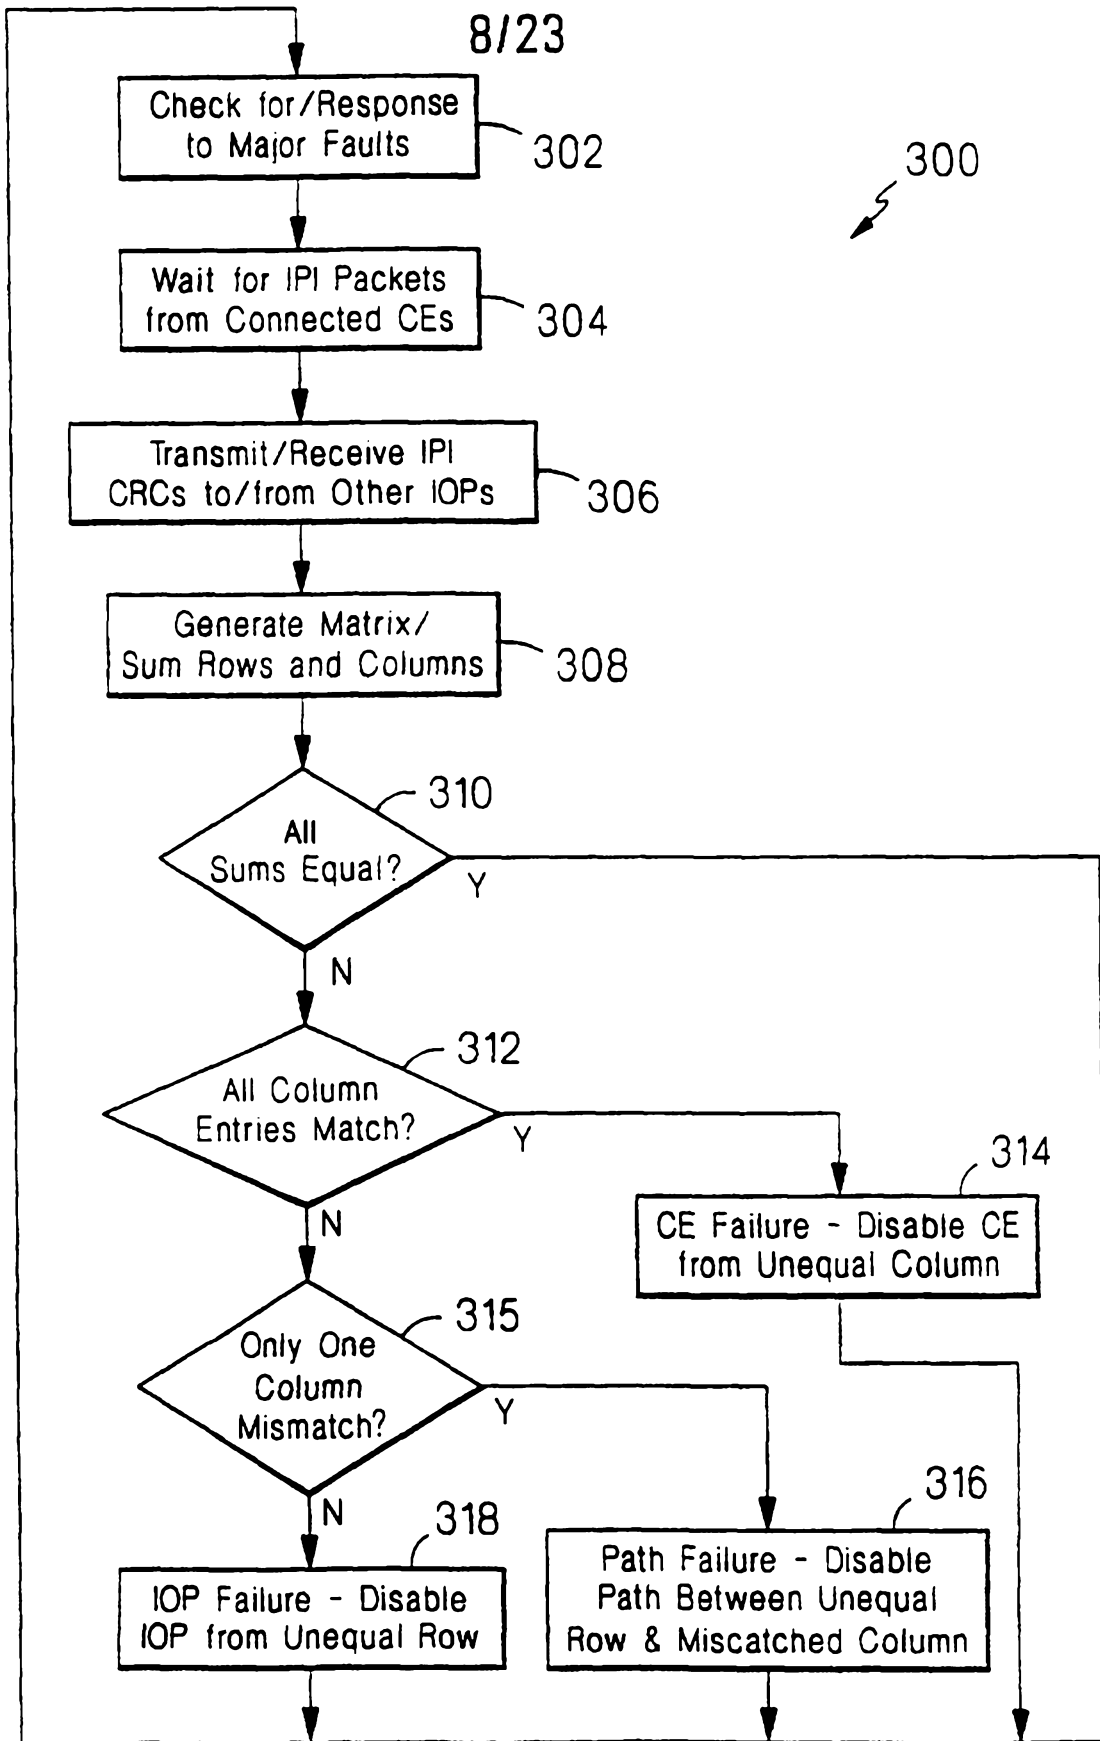


FIG. 9

9/23




FIG. 10




FIG. 11

10/23

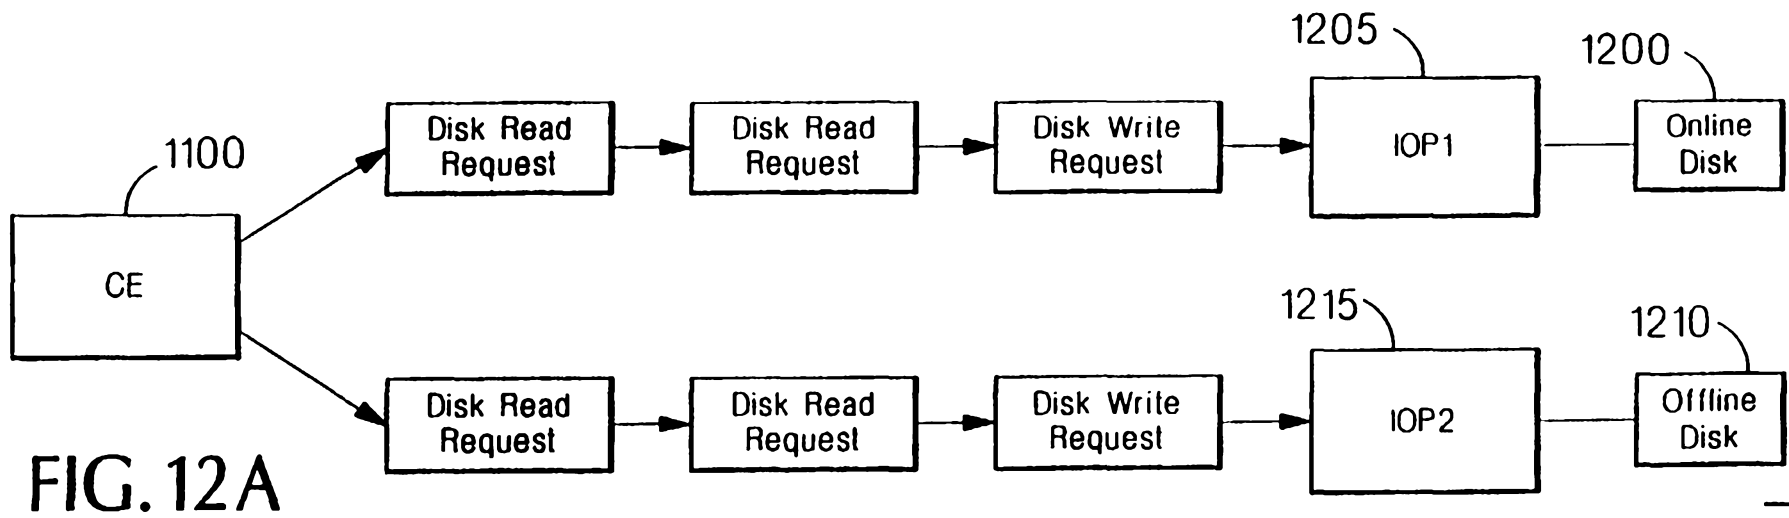


FIG. 12A

11/23

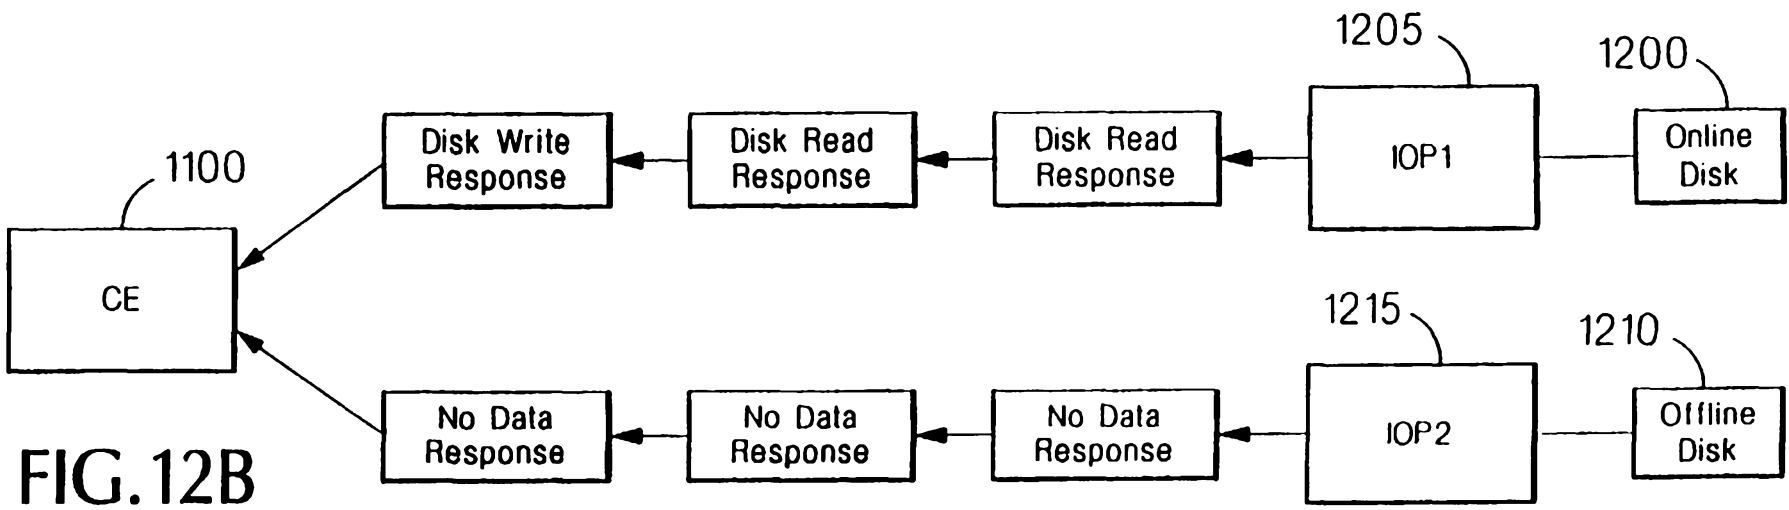


FIG. 12B

12/23

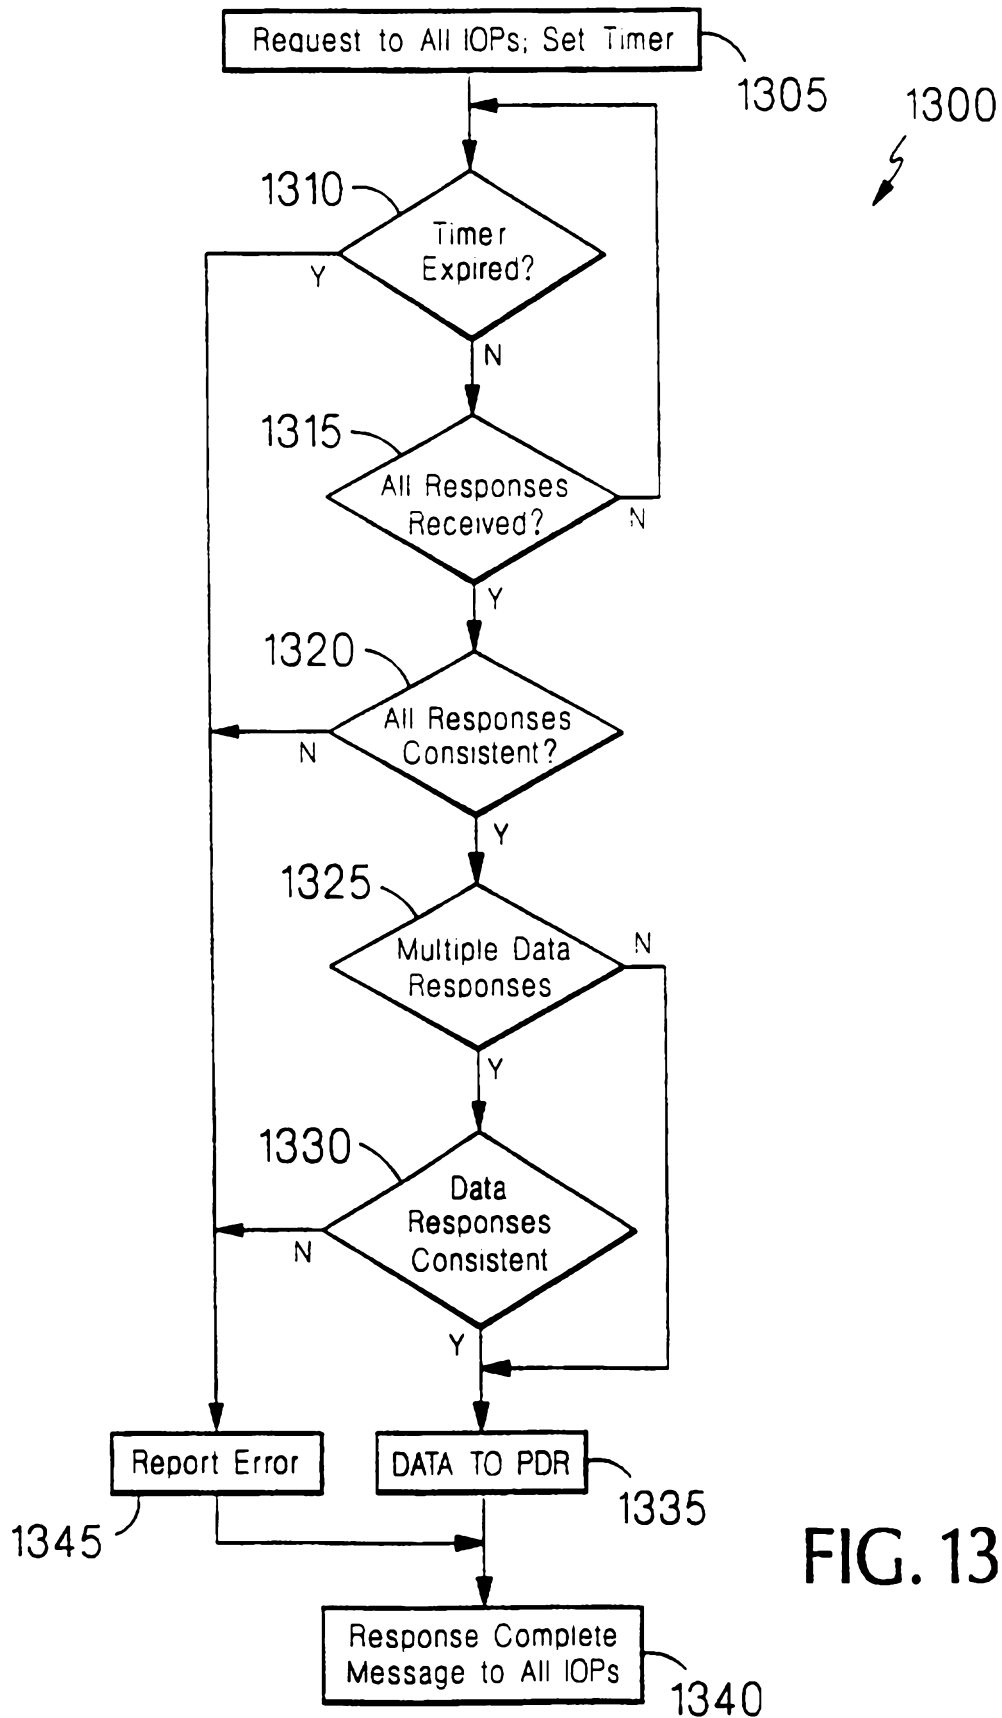


FIG. 13

13/23

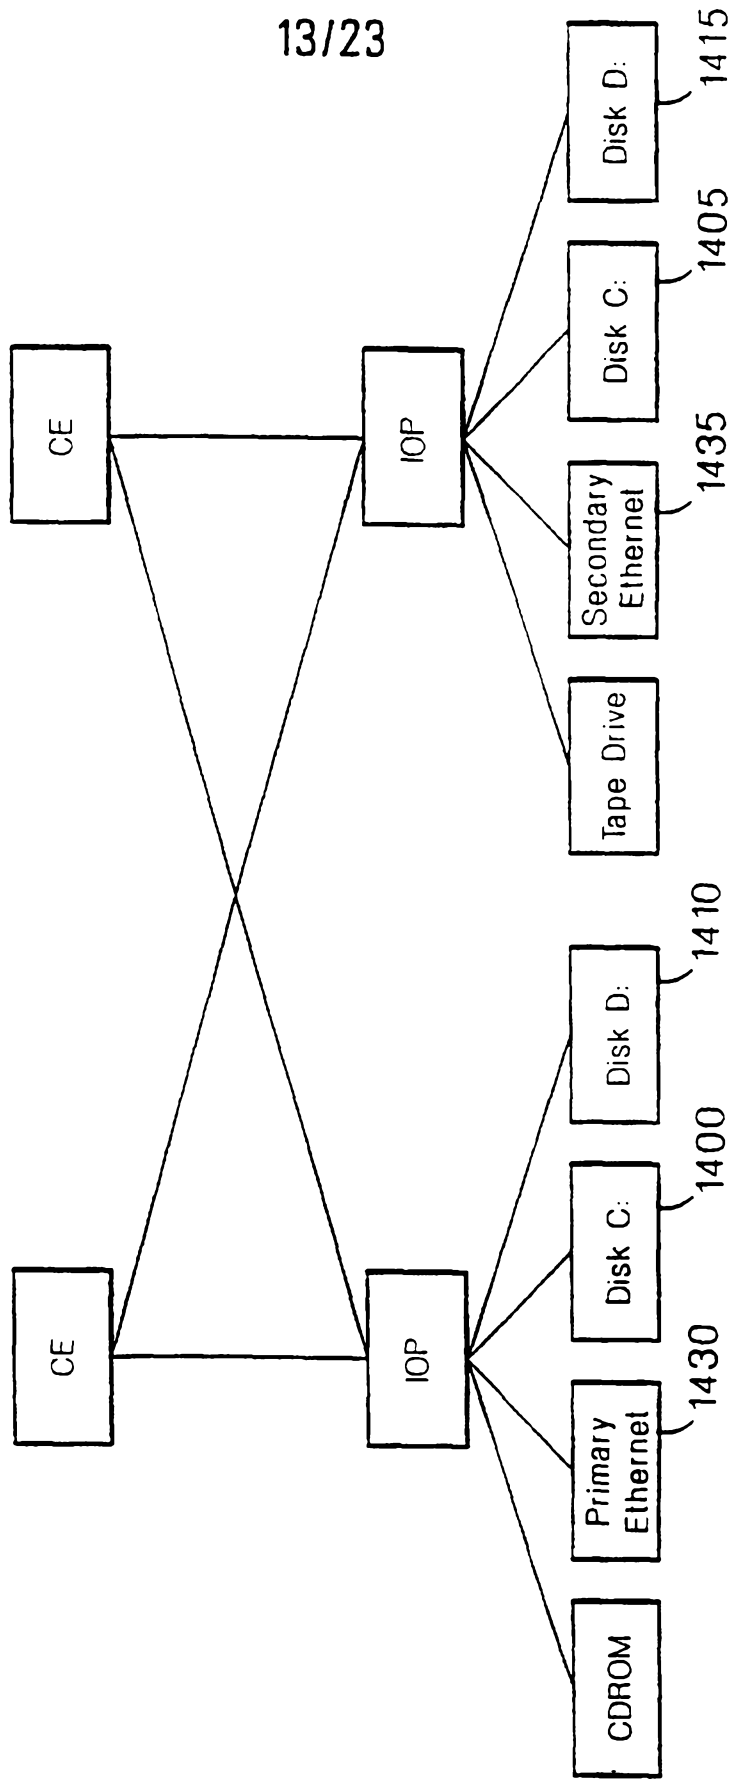


FIG.14A Physical Configuration

14/23




FIG.14B Logical Configuration

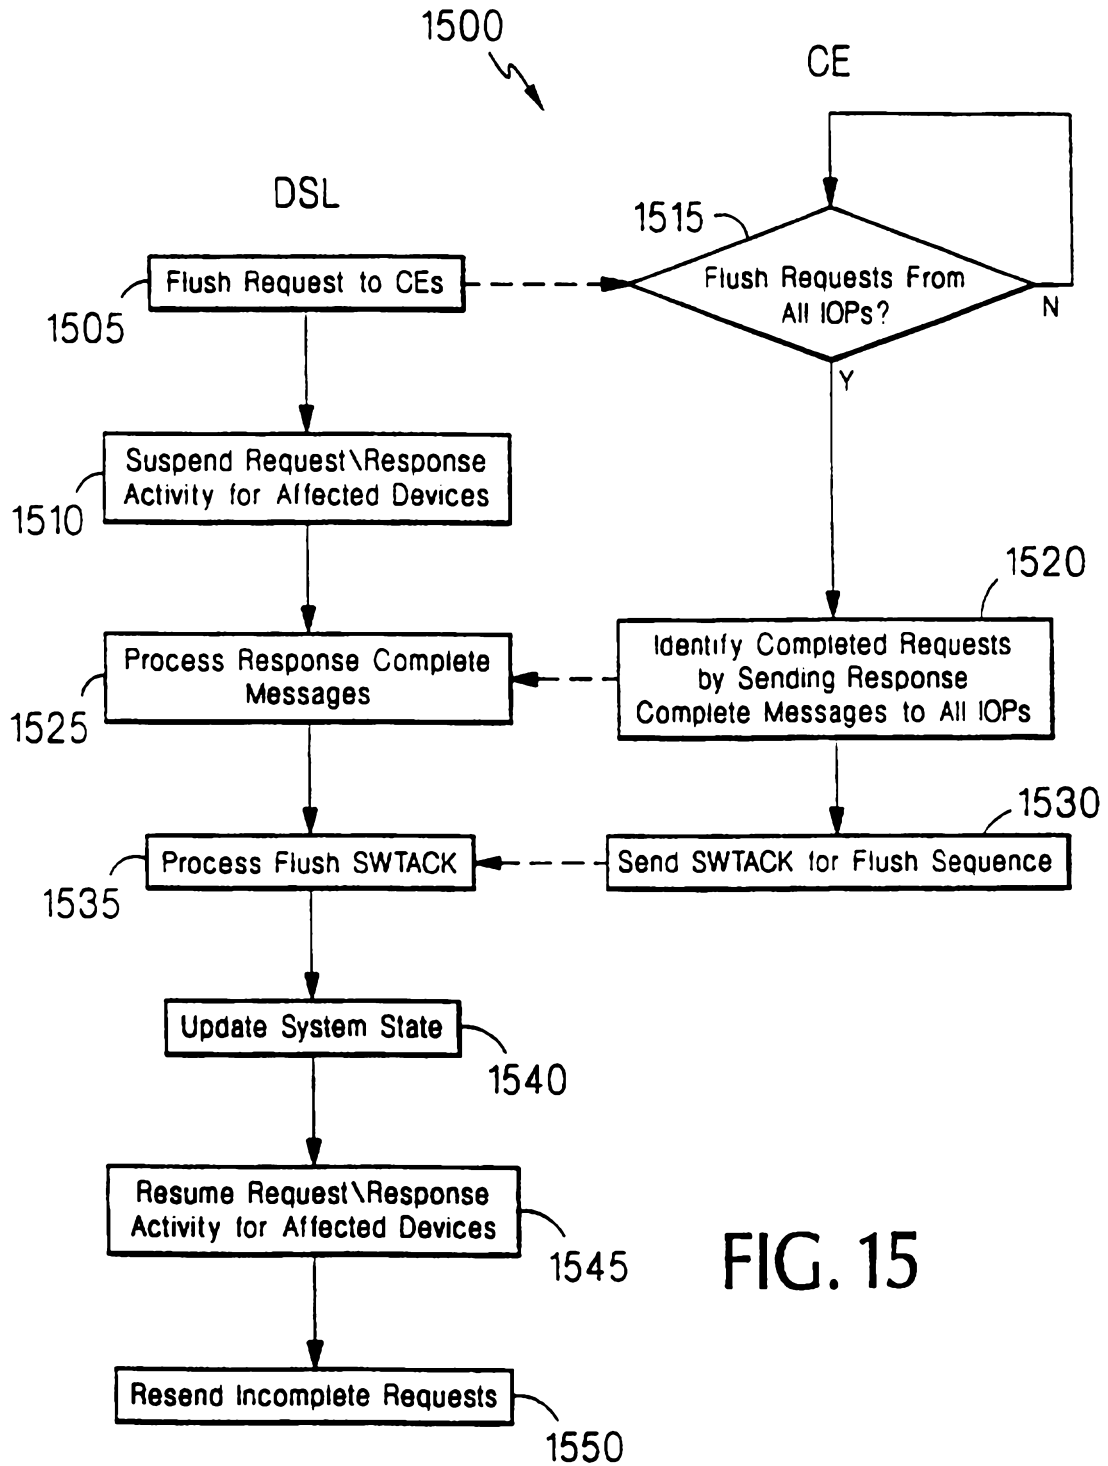


FIG. 15

CURRENT STATE	NEW STATE						
IDLE	REQ_SND CEX Entry	RSP_WAIT	DREQ_SND	DRSP_WAIT	REL_SND	RSP_TMO	
REQ_SND	Send retry	Req sent					
RSP_WAIT		Another Rsp	All rsp rcvd		All rsp rcvd	Rsp timeout	
DREQ_SND			Send retry	DReq sent			
DRSP_WAIT				Another	All DRsp	DRsp	
REL_SND	Rel sent				Sent retry		
RSP_TMO	IPI Reset		Port state			Port state chg	

16/23

FIG. 16A

CURRENT STATE	NEW STATE						
IDLE	REQ_WAIT Rev rcvd	RSP_SND	REL_WAIT	DRSP_SND	DREL_WAIT	REL_TMO	
REQ_WAIT	Another Req	All Req rcvd					
RSP_SND		Send retry	Rsp sent				
REL_WAIT	All Rel rcvd		Another Rel	DReq rcvd	DRsp sent	Rel timeout	
DRSP_SND				Send retry	Another DRel	DRel timeout	
DREL_WAIT	All DRel					Port state chg	
REL_TMO	IPI Reset, or	port state chg					

FIG. 16B

17/23




FIG. 17

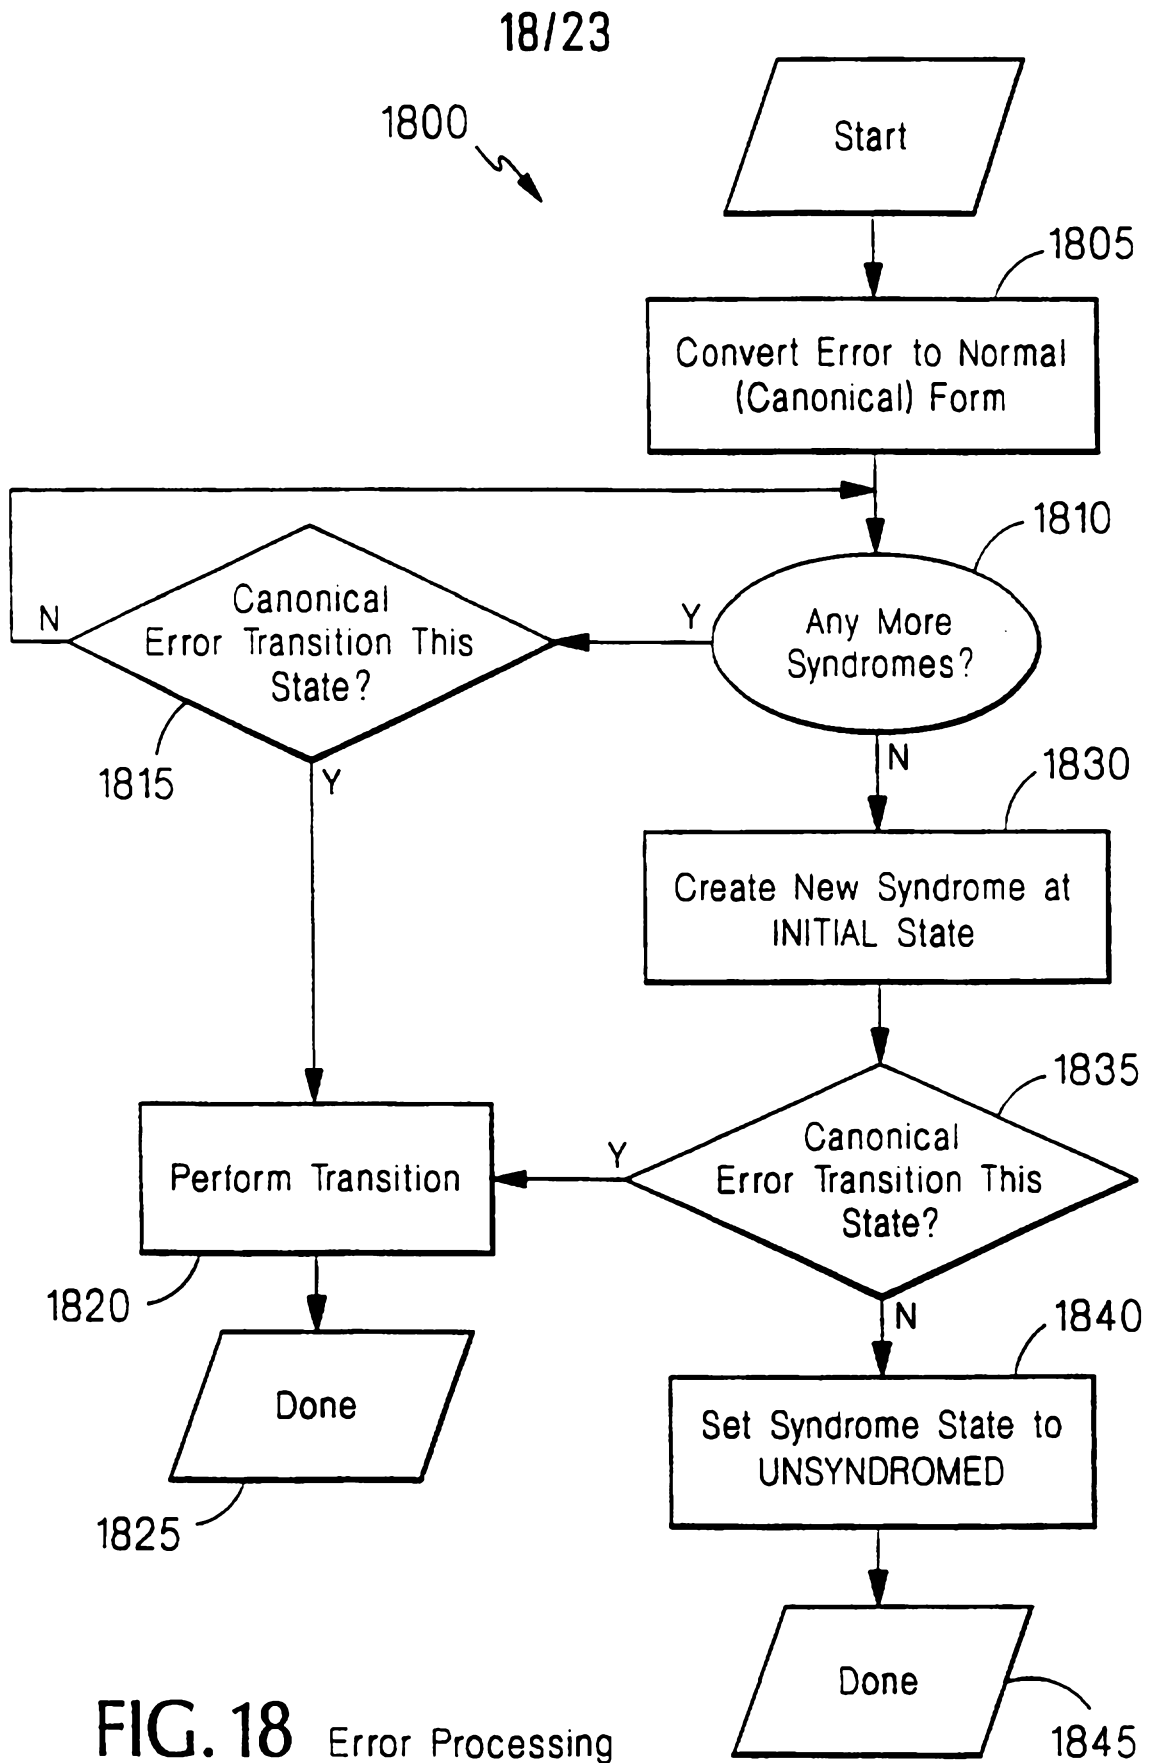


FIG. 18 Error Processing

Syndrome			
#	Error	Path	Received from
1	NAK	M1 → M2	M1
	NAK (echoed)	M1 → M2	M1 (echoed by M2)
	EDC	M1 → M2	M2
2	NAK (echoed)	M1 → M2	M1 (via M2 echo)
	EDC	M1 → M2	M2

FIG. 19A




FIG. 19B

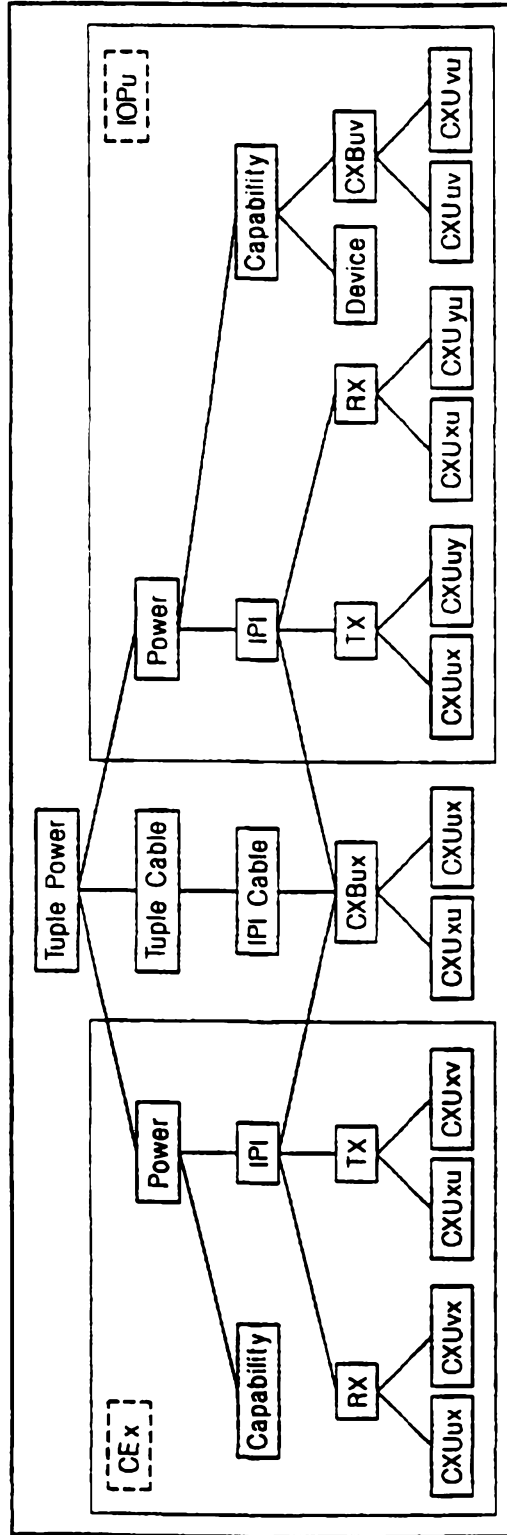


FIG. 20

List 1	List 2
Machine 1 Power	Machine 2 Power
Machine 1 Interconnect Adapter	Machine 2 Interconnect Adapter
Machine 1 ↔ Machine 2 Unidirectional InterCnx	Machine 2 → Machine 1 Unidirectional InterCnx

FIG. 21A

List 1 Element	List 2 Element	Product
Machine 1 Power	Machine 2 Power	Common Machine 1/2 Power Supply
Machine 1 Power	Machine 2 Adapter	Common Machine 1/2 Power Supply
Machine 1 Power	2 → 1 Interconnect	Machine 1 Power
Machine 1 Adapter	Machine 2 Power	Common Machine 1/2 Power Supply
Machine 1 Adapter	Machine 2 Adapter	Common Machine 1/2 Power Supply
Machine 1 Adapter	2 → 1 Interconnect	Machine 1 Adapter
1 → 2 Interconnect	Machine 2 Power	Machine 2 Power
1 → 2 Interconnect	Machine 2 Adapter	Machine 2 Adapter
1 → 2 Interconnect	2 → 1 Interconnect	1 ↔ 2 Bidirectional Interconnect

FIG. 21B

22/23

Callout	Rank
Common machine 1/2 Power Supply	5
Machine 1 Power	4
Machine 2 Power	4
Machine 1 Adapter	3
Machine 2 Adapter	3
Machine 1 ↔ Machine 2 Bidirectional Interconnect	2
Machine 1 → Machine 2 Unidirectional Interconnect	1
Machine 1 → Machine 1 Unidirectional Interconnect	1

FIG. 22

23/23

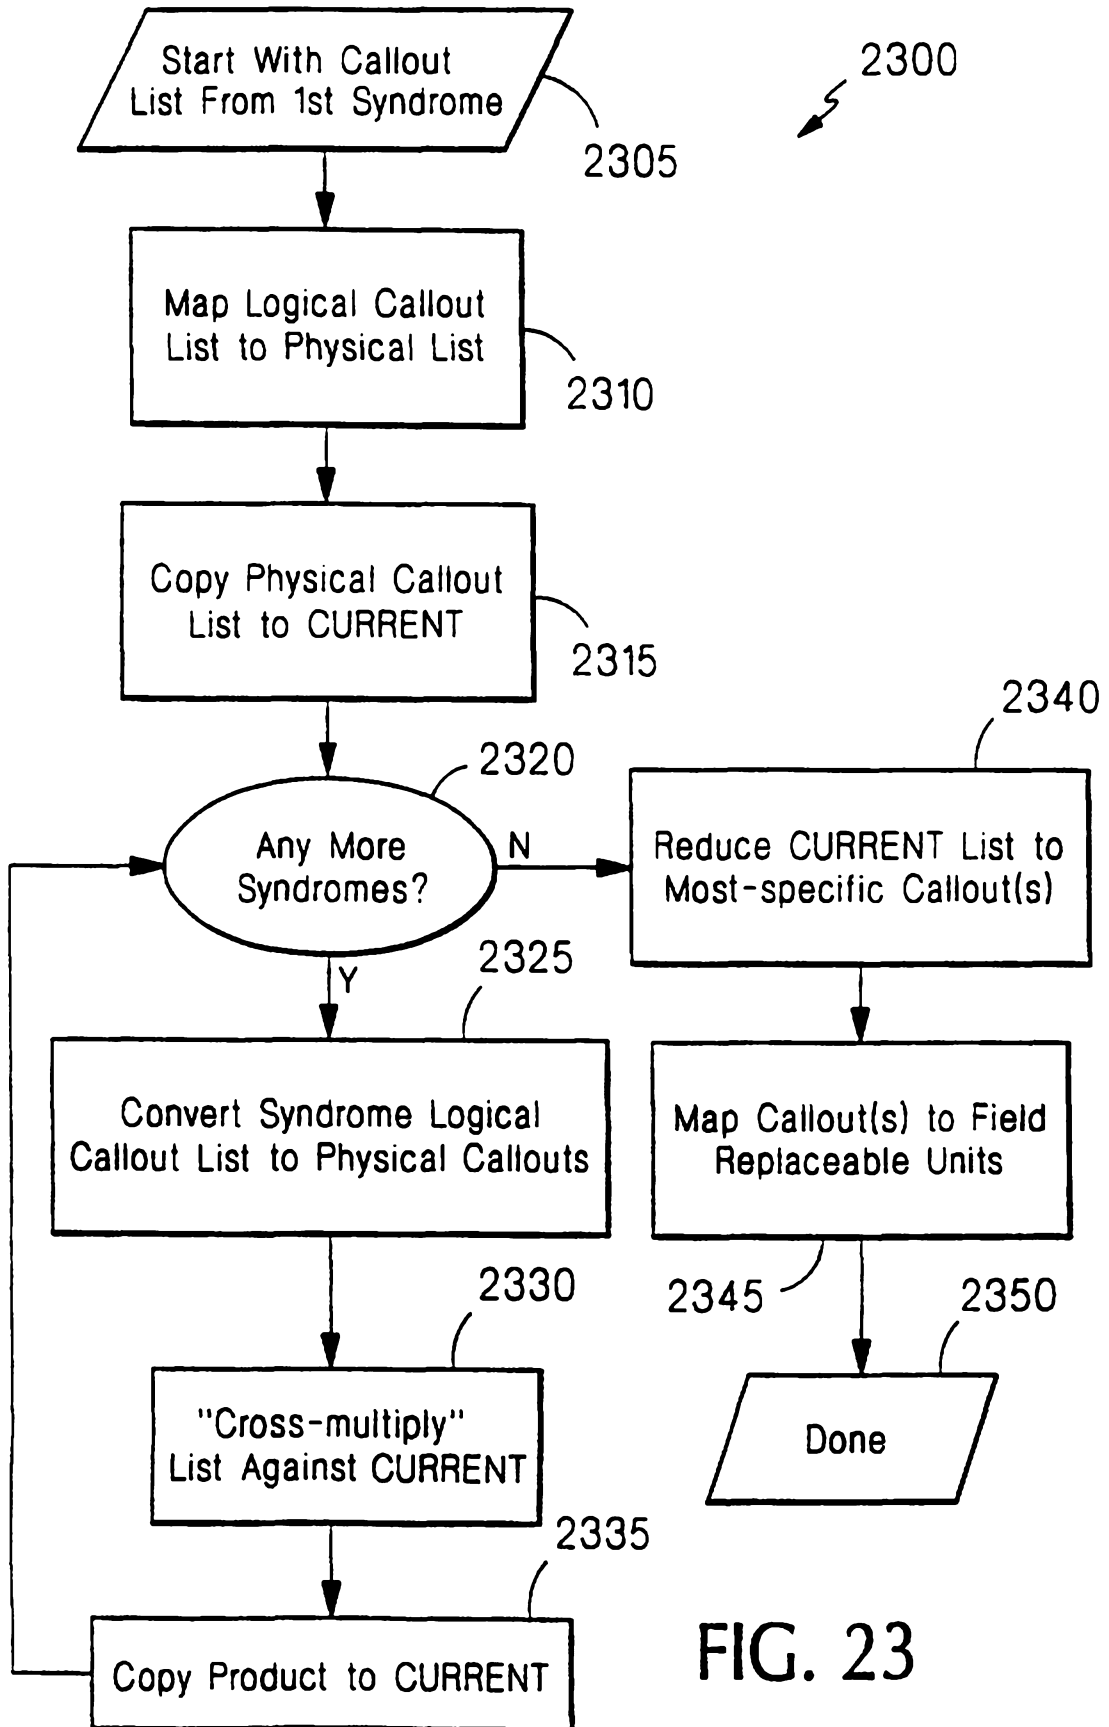


FIG. 23

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US97/16218

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 19/00

US CL :364/131; 395/182.1

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 364/131-136, 184-187; 395/182.09, 182.1, 185.08, 185.1, 200.78, 881; 371/47.1,68.1

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS, search terms: intercept, request, synchronize, controllers, error

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 5,088,021 A (McLAUGHLIN et al) 11 February 1992, col. 4, line 62 to col. 5, line 30	1-3, 10-14, 32-34, and 38-40 ----- 4-9 and 35-37
Y	US 5,261,092 A (McLAUGHLIN et al) 09 November 1993, col. 10, line 26 to col. 11, line 24.	4-9 and 35-37
A	US 5,142,470 A (BRISTOW et al) 25 August 1992, see the entire document.	1-26 and 32-47
A	US 5,491,625 A (PRESSNALL et al) 13 February 1996, see the entire document.	1-26 and 32-47

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	"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
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier document published on or after the international filing date	"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
"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)	"A" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

06 NOVEMBER 1997

Date of mailing of the international search report

26 JAN 1998

Name and mailing address of the ISA/US

Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

B. F. Ruggiero
JOSEPH F. RUGGERO

Telephone No. (703) 305-3800

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/16218

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,555,372 A (TETREULT et al) 10 September 1996, col. 3, lines 8-43 and col. 8, lines 6-42.	27 and 48
X	US 5,423,025 A (GOLDMAN et al) 06 June 1995, col. 4, line 61 to col. 6, line 33.	28, 29, and 49
A	US 5,448,722 A (LYNNE et al) 05 September 1995, see the entire document.	27-31, 48, and 49

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/16218

Box I Observations where certain claims were found unsearchable (Continuation of Item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:

3. Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of Item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

Please See Extra Sheet.

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US97/16218

BOX II. OBSERVATIONS WHERE UNITY OF INVENTION WAS LACKING

This ISA found multiple inventions as follows:

This application contains the following inventions or groups of inventions which are not so linked as to form a single inventive concept under PCT Rule 13.1. In order for all inventions to be searched, the appropriate additional search fees must be paid.

Group I, claim(s) 1-26 and 32-47, drawn to a method and system for synchronizing data transfer.

Group II, claim(s) 27-31, 48, and 49, drawn to a method and system for fault handling and error reporting.

The inventions listed as Groups I and II do not relate to a single inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons:

The data transfer system and method of Group I is not limited to fault or error conditions and the fault handling system and method of Group II does not require the data transfer techniques of Group I.