

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 July 2006 (20.07.2006)

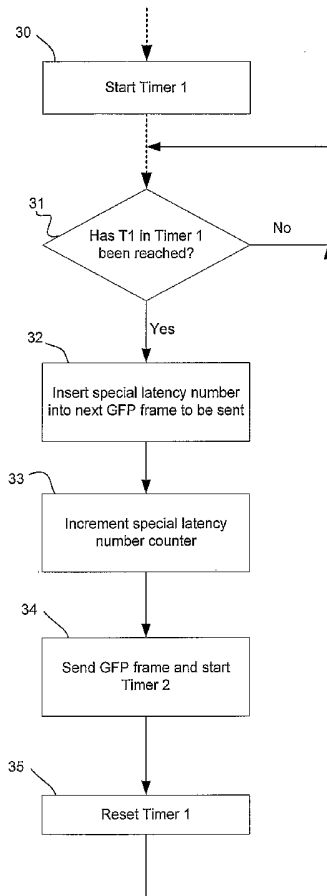
PCT

(10) International Publication Number
WO 2006/076652 A2

- (51) International Patent Classification:
H04L 12/56 (2006.01) H04J 3/06 (2006.01)
- (21) International Application Number:
PCT/US2006/001363
- (22) International Filing Date: 12 January 2006 (12.01.2006)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
11/036,596 14 January 2005 (14.01.2005) US
- (71) Applicant (for all designated States except US): CISCO TECHNOLOGY INC. [US/US]; 170 West Tasman Drive, San Jose, California 95134-1706 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): SUNDARAM, Ganesh [IN/US]; 7475 Monet Place, Rohnert Park, California 94928 (US). DIAB, John [US/US]; 2876 Liscum Street, Santa Rosa, California 95407 (US). AMIN, Hitesh [US/US]; 380 Raven Way, Petaluma, California 94954 (US).
- (74) Agents: AKA, Gary et al.; AKA CHAN LLP, 900 Lafayette Street, Suite 710, Santa Clara, California 95050 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

[Continued on next page]

(54) Title: DYNAMIC AND INTELLIGENT BUFFER MANAGEMENT FOR SAN EXTENSION



(57) Abstract: In a network system for transporting GFP-encapsulated Fibre Channel/FICON data across a SONET/SDH transport network between two Fibre Channel/FICON ports, a transport interface for one Fibre Channel/FICON port intelligently allocates the amount of buffers for receiving Fibre Channel/FICON data from the other Fibre Channel/FICON port. The first transport interface inserts a special latency instruction message into the Fibre Channel/FICON data before encapsulation in a GFP frame. After transport across the SONET/SDH network, the receiving second transport interface immediately sends the special latency instruction message back across the SONET/SDH transport network to the first transport interface which times the return of the special latency instruction message. From the time interval, the first transport interface can determine the latency of the SONET/SDH transport network and allocates the amount of buffers appropriately.



WO 2006/076652 A2



FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT,
RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA,
GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

— *without international search report and to be republished upon receipt of that report*

DYNAMIC AND INTELLIGENT BUFFER MANAGEMENT FOR SAN EXTENSION

5

BACKGROUND OF THE INVENTION

The present invention relates generally to digital communication networks, and more specifically, to methods and systems for efficiently transporting Fibre Channel/FICON client data over a SONET/SDH network path.

SONET/SDH and optical fiber have emerged as significant technologies for building large scale, high speed, IP (Internet Protocol)-based networks. SONET, an acronym for Synchronous Optical Network, and SDH, an acronym for Synchronous Digital Hierarchy, are a set of related standards for synchronous data transmission over fiber optic networks. SONET/SDH is currently used in wide area networks (WAN) and metropolitan area networks (MAN). A SONET system consists of switches, multiplexers, and repeaters, all connected by fiber. The connection between a source and destination is called a path.

One network architecture for the network interconnection of computer devices is Fibre Channel, the core standard of which is described in ANSI (American National Standards Institute) X3.230-1994. Arising out of data storage requirements, Fibre Channel currently provides for bi-directional gigabits-per-second transport over Storage Area Networks (SANs) in Fibre Channel frames that consist of standardized sets of bits used to carry data over the network system. Fibre Channel links are limited to no more than 10 kilometers. Similar to Fibre Channel is FICON, a proprietary I/O channel which was developed by IBM for the data storage requirements for main frame computers.

New standards and protocols have emerged to combine the advantages of the SONET/SDH and Fibre Channel/FICON technologies. For example, it is sometimes desirable to link two SANs, which operate with Fibre Channel or FICON protocols, over a MAN (Metropolitan Area Network), or even a WAN (Wide Area Network), which typically operate under SONET or SDH standards. This extension of SANs from 100 kilometers to over several hundred, or even thousand, kilometers, is made by mapping Fibre Channel/FICON ports to a SONET/SDH path for transport across a SONET/SDH network. One way to perform this function is to encapsulate Fibre Channel/FICON client data frames into transparent Generic Framing Protocol (GFP-T) frames and then map the

GFP-T frames into SONET/SDH frames for transport across the SONET/SDH network. In this manner two Fibre Channel/FICON ports can communicate with each other over a SONET/SDH network as though the intervening network links are part of a Fibre
5 Channel/FICON network. The Fibre Channel/FICON ports remain “unaware” of the SONET/SDH transport path. For example, see U.S. Patent Application No. 10/390,813, entitled, “Method and System for Emulating a Fibre Channel Link Over a Sonet/SDH Path,” filed March 18, 2003 and assigned to the present assignee.

For the effective movement of data across SAN networks, these network systems
10 have two types of flow control: 1) end-to-end, and 2) buffer-to-buffer credit. In both types of flow control, two Fibre Channel/FICON ports report to each other how many frames is available at the reporting port’s buffer to receive Fibre Channel/FICON frames from the other port. In end-to-end flow control, the source and destination ports are the two ports and the ports signal each other the reception of a transmitted frame by an ACK
15 Link Control frame. In buffer-to-buffer credit, the two ports on opposite sides of a link are the two ports and the ports communicate the reception of a transmitted frame with an R_Rdy Primitive signal. But flow control remains within the SAN network and is based on counting Fibre Channel/FICON frames which can vary. Flow control may also be extended across SONET/SDH transport networks which connect frame-based protocol
20 networks, such as Fibre Channel/FICON and gigabit Ethernet. See, for example, U.S. Patent Application No. 10/613,426, entitled, “Method and System For Efficient Flow Control For Client Data Frames Over GFP Across a SONET/SDH Transport Path,” filed July 3, 2003 and assigned to the present assignee.

Nonetheless, for SAN extensions, i.e., interconnecting SANs by SONET/SDH
25 transport networks, the SAN extension devices (the Fibre Channel/FICON ports communicating over a SONET/SDH network) usually provide a large amount of buffering in order to maintain a 100% throughput over very long distances. Because of the large number of buffers in the SAN extension devices, a great deal of latency can be created for the frames passing through the devices. It is possible that sometimes the latency
30 introduced by extra buffering can be a significant portion of the total latency, even compared to the latency of the long distance communication.

The present invention addresses this problem of inappropriate buffering with buffer management which is dynamic and intelligently selective for the particular SAN extension.

SUMMARY OF THE INVENTION

5 The present invention provides for a method of operating a transport interface for at least one local Fibre Channel/FICON port, the transport interface having buffers for Fibre Channel/FICON data encapsulated in GFP frames transported over a SONET/SDH network from a remote Fibre Channel/FICON port. The method has the steps of inserting a special latency instruction message into Fibre Channel/FICON data to be encapsulated in a GFP frame for transmission to the remote Fibre Channel/FICON port; sending the GFP
10 frame over the SONET/SDH transport network to the remote Fibre Channel/FICON port; timing a return of the special latency number over the said SONET/SDH transport network; determining an appropriate amount of buffers in the transport interface for GFP frames from the remote Fibre Channel/FICON port from the timing step; and allocating the appropriate amount of buffers in the transport interface for GFP frames from the
15 remote Fibre Channel/FICON port; whereby sufficient buffering is ensured in the transport interface to provide maximum throughput over the SONET/SDH network and any additional latency due to buffering in the transport interface is reduced.

Furthermore, the inserting, sending, timing, determining and allocating steps are repeated periodically so that the amount of allocated buffers is adjusted even if the latency
20 of GFP frames transported over said SONET/SDH network between the local and remote Fibre Channel/FICON ports changes. A period of about 1 second is used for the described embodiment of the present invention. The special latency instruction message, which is inserted in a Client Payload Information field of the Payload Area of the GFP frame, includes a latency sequence number to identify one sequence of inserting, sending, timing,
25 determining and allocating steps from another sequence of inserting, sending, timing, determining and allocating steps; a special character encoded in a 4-bit mapping of the 64B/65B control characters as Fh; and a command to a transport interface for the remote Fibre Channel/FICON port to resend said special latency instruction message back to the
30 transport interface for the at least one local Fibre Channel/FICON port upon receiving the special latency instruction message.

In a network system for transporting GFP-encapsulated Fibre Channel/FICON data across a SONET/SDH transport network between first and second Fibre Channel/FICON ports, the first Fibre Channel/FICON port connected to the SONET/SDH transport network through a first transport interface and the second Fibre Channel/FICON port

connected to the SONET/SDH transport network through a second transport interface, the present invention also provides for the first transport interface which has at least one integrated circuit adapted to insert a special latency instruction message into Fibre Channel/FICON data from the first Fibre Channel/FICON port and to encapsulate the Fibre Channel/FICON data in a GFP frame, to send the GFP frame over the SONET/SDH transport network to the second transport interface of the second Fibre Channel/FICON port, to time a return of the special latency instruction message over the SONET/SDH transport network, to determine an appropriate amount of buffers in the first transport interface from a time interval of the special latency instruction message to return, and to allocate the appropriate amount of buffers in the first transport interface for GFP frames from the second Fibre Channel/FICON port so that sufficient buffering is ensured in the first transport interface to provide maximum throughput over the SONET/SDH network and any additional latency due to buffering in the transport interface is reduced.

The at least one integrated circuit is further adapted to insert the special latency instruction message, to encapsulate the Fibre Channel/FICON data in a GFP frame, to send the GFP frame, to time the return of the special latency instruction message, to determine the appropriate amount of buffers and to allocate the appropriate amount of buffers periodically so that the amount of allocated buffers is adjusted as the actual latency of GFP frames transported over the SONET/SDH network between the first and second Fibre Channel/FICON ports changes.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a diagram illustrating an exemplary network employing the present invention;

5 Fig. 2A is a flow chart of operations of a local transport interface, in the exemplary network of Fig. 1 in the transmission of a special latency instruction message to a remote transport interface, according to one embodiment of the present invention; Fig. 2B is a flow chart of operations of the local transport interface after receiving a response from the remote transport interface;

10 Fig. 3A is a representative diagram of a GFP frame; Fig. 3B illustrates the special latency instruction message with latency sequence number which is inserted into a GFP frame and sent by the local transport interface in Fig. 2A, according to one embodiment of the present invention; and

15 Fig. 4 is a block diagram of a portion of a port card of Fig. 1, according to one embodiment of the present invention.

Corresponding reference characters indicate corresponding parts throughout the several views of the drawings.

DESCRIPTION OF SPECIFIC EMBODIMENTS

The following description is presented to enable one of ordinary skill in the art to make and use the invention. Descriptions of specific embodiments and applications are provided only as examples and various modifications will be readily apparent to those skilled in the art. The general principles described herein may be applied to other embodiments and applications without departing from the scope of the invention. Thus, the present invention is not to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features described herein. For purpose of clarity, details relating to technical material that is known in the technical fields related to the invention have not been described in detail.

Fig. 1 illustrates an exemplary network of Fiber Channel/FICON ports are connected over a SONET/SDH transport network 10 in which an embodiment of the present invention can operate. In the present example, it is assumed that the ports operate under Fibre Channel or FICON protocols, though the ports may also operate under other frame-based protocols, such as gigabit Ethernet, in accordance with the present invention.

In the exemplary network Fibre Channel/FICON ports 16 and 18 are connected by Fibre Channel/FICON links 15 and 17 respectively to a multi-port Fibre Channel/FICON card 14. Likewise, a second Fibre Channel/FICON port card 24 is connected by Fibre Channel/FICON links 25 and 27 to Fibre Channel/FICON ports 26 and 28 respectively. The Fibre Channel/FICON ports 16, 18, 26 and 28 are associated with elements which are interconnected by Fibre Channel/FICON protocols in SANs. These elements include data storage elements, including disk drive arrays, RAIDs, disk farms, or possibly Fibre Channel network elements, such as routers, switches, or other Fibre Channel network elements. In Fig.1 each Fibre Channel/FICON port card 14 and 24 is connected to a pair of Fibre Channel/FICON ports for purposes of illustration, and more ports may be connected to each Fibre Channel/FICON port card.

The SONET/SDH network 10 provides a transport path to connect the Fibre Channel/FICON ports 16 and 18 with the Fibre Channel ports 26 and 28 so that Fibre Channel/FICON client data can be transferred between the ports 16, 18 and 26, 28. Optical transport platforms 12 and 22, such as ONS 15454 (available from Cisco Systems, Inc. of San Jose, California), provide the interface between the Fibre Channel/FICON and SONET/SDH networks. The Fibre Channel/FICON ports 16 and 18 are connected to the

multi-port Fibre Channel/FICON card 14 which is adapted to fit into the optical transport platform 12; the Fibre Channel/FICON ports 26 and 28 are connected to the multi-port Fibre Channel/FICON card 24 which adapted to fit into the optical transport platform 22.

5 Through the Fibre Channel/FICON port cards 14 and 24, which function as transport interfaces with the platforms 12 and 22 respectively, the Fibre Channel/FICON ports 16 and 18 are interconnected to the Fibre Channel/FICON ports 26 and 28 across the SONET/SDH network transport path. The result is that there are two virtual wires for the connection between a representative Fibre Channel/FICON port at one end of the

10 SONET/SDH network 10, say, port 18, and a representative Fibre Channel port at the other end, say, port 28. As explained above, GFP-T, transparent Generic Framing Procedure, is conventionally used as the framing protocol for such a network to encapsulate the Fibre Channel/FICON payloads at one end of the SONET/SDH network 10 for transmission across the SONET/SDH network and to decapsulate the Fibre

15 Channel/FICON data at the other end. By GFP-T protocol, the GFP-T frames have fixed lengths.

While the port cards 14 and 24 and their respective optical platforms 12 and 22 are the transport interfaces for the exemplary network of Fig. 1, the transport interfaces can be considered to be located in the port cards 14 and 24 for the described embodiment of the

20 present invention. The cards 14 and 24 each have FIFO (First-In First-Out) buffers to hold the GFP frames received from the SONET/SDH transport network 10 before the encapsulated Fibre Channel/FICON frames of the described embodiment of the present invention, are stripped out of the GFP-encapsulation frames and passed on to their Fibre Channel port destinations.

25 The port cards 14 and 24, which extend the SANs so that they are interconnected, operate as intermediate transparent devices on a SAN network. Heretofore, such SAN extension devices typically have a configuration mechanism by which the user could select the number of FIFO buffers for the frames transported across the SONET/SDH network. The mechanism helps the user choose the number of buffers required for the

30 SAN extension over a long distance and in order to maintain a 100% throughput over the long distances of SONET/SDH transport network 10, the mechanism typically selects a large amount of buffering usually provided in the SAN extension devices. However, a large amount of latency is added for the frames passing through the devices, because of the large number of buffers in the SAN extension devices.

These configuration mechanisms may not be accurate nor appropriate for the particular SONET/SDH transport path. Also, changes can occur in a transport network, e.g., an increased path delay because of a SONET/SDH switchover, thus changing the buffering requirements for the SAN extension devices.

To address these problems, the present invention accurately determines the round trip delay (a measure of distance) from one SAN extension device across a SONET/SDH transport network to another SAN extension device and back. Once the latency is accurately determined, the number of buffers required in the first SAN extension device is calculated and programmed into the hardware of the device. Since the number of required buffers are configured for the current distance between the two SAN extension devices, any additional latency due to extra buffering is avoided. Only the required number of buffers on the SAN extension devices is allocated to reduce latency. For example, 1G (base clock rate of 1.0625 GHz for Fibre Channel/FICON data transfer) Fibre Channel/FICON client data sent over a 1200Km (one-way) transport path requires 600 (2Kbyte) buffers for a sustained 100% throughput. However, if the same 600 buffers are used for a 200Km circuit, the extra 500 buffers add an unwanted latency of about 5ms, thereby making the solution unsuitable for certain applications. Due to the inherent bursty nature of Fibre Channel/FICON traffic, the extra buffering can be filled with an additional 500 frames to add undesired latency. By limiting the number of buffers used, traffic is backpressured all the way to the Fibre Channel/FICON source and thereby reduces unwanted latency on all traffic.

It should be noted that although a 1G Fibre Channel/FICON client operation speed is mentioned above, the present invention works effectively with Fibre Channel/FICON clients operating at 2G (double base clock rate or 2.125 GHz) or any other Fibre Channel/FICON speed.

Also, with the present invention any SONET/SDH switchover or protection event which leads to a new SONET/SDH path and new distance is automatically detected and the amount of buffers is adjusted accordingly. The buffer adjustments are performed without any hits or errors to the SAN traffic.

In accordance with the present invention, a special latency instruction message with an incrementing latency sequence number is periodically inserted into the GFP Client Payload Information field of the GFP-T frames encapsulating the Fibre Channel/FICON payload frames that are to be transported across SONET/SDH transport

path. The special latency instruction message with latency sequence number includes a special K character that is not used in the Fibre Channel/FICON protocol and is never forwarded to the Fibre Channel/FICON client. It is only used between the Fibre Channel/FICON-Over-SONET/SDH equipment, such as the transport interfaces, i.e., the port cards 14 and 24, with the interconnecting SONET/SDH transport network 10 in the Fig. 1 network, for example. Upon sending the Fibre Channel/FICON frames encapsulated in the GFP frame with the special latency number, the local transport interface, i.e., the GFP transmitter, starts a timer.

10 At the remote or receiving transport interface, i.e., the GFP receiver, immediately responds to the special latency number by sending it back to the GFP transmitter across the SONET/SDH transport network.

Upon receiving the special latency instruction message and latency sequence number, the local transport interface reads its timer and has an accurate determination of the latency in sending frames across the SONET/SDH network to and from the GFP receiver. From the latency determination, the number of buffers required in the local transport interface, the port card 14 in this embodiment, is calculated and programmed into the hardware of the device. The local transport interface monitors the latency in the SONET/SDH transport path to the remote transport interface continually by repeating the procedure described above periodically. In the described embodiment this period is 1 second.

Figs. 2A and 2B are flow charts which illustrate the steps of operation of an exemplary local transport interface, the port card 14 in this case, in transmitting and receiving encapsulating GFP frames to and from a remote transport interface, the port card 24, in accordance with the present invention. In this manner the local and remote transport interfaces effectively extend their respective SANs to each other's SAN.

As shown in Fig. 2A, after the initialization of the local and remote transport interfaces and initial communication is established over the SONET/SDH transport path, as indicated by a dotted arrow, a Timer 1 in the port card 14 is started in step 30. After engaging in different operations not directly related to the present operations, the port card 14 reaches a decision step 31. Has the Timer 1 reached a value T1, in this example, one second? If not, other operations between steps 30 and 31 not directly related to the present invention are resumed. For example, the transfer of GFP superblock frames may be

resumed. If so, on the other hand, then a special latency sequence number from a special latency counter in the port card 14 is inserted into the next GFP frame by step 32. In step 33 the special latency counter is incremented and the GFP frame is sent by step 34 from the port card 14 to the port card 24 across the SONET/SDH transport network 10 and at the same time, the port card 14 also starts a Timer 2. In step 36 the Timer 1 is reset and the process returns back to the operations between steps 30 and 31.

Fig. 3A illustrates an exemplary GFP-T frame 50 with the special latency instruction message and latency sequence number which is assembled and sent by the local port card 14. The special instruction message and latency sequence number is placed in the GFP Client Payload Information field 54, such as found in the exemplary GFP frame 50 with its component Core Header 51 and Payload Area 52. Within the Payload Area 52 is a Payload Header field 53 and the Client Payload Information field 54. Fig. 3B illustrates the 36-bit special latency instruction message and latency sequence number which is inserted into the 10B/8B client data stream from the Fibre Channel/FICON port 18 at the Fibre Channel/FICON port card 14. In accordance with GFP-T procedures (see clause 8 of the ITU-T Generic Framing Procedure standard G.7041/Y.1303, for example), the decoded 10B/8B client data is mapped into 65B/64B block code and then into 65B/64B superblocks for placement into a GFP Client Payload Information field. In the special latency instruction message, the 4-bit "8h" (8 in hexadecimal) defines the data[31:24] as a control word; the 8-bit "44h" within the control word is a special control character to the Fibre Channel/FICON receiving port card, the port card 24 in this example. The 20-bit incrementing latency sequence number identifies each latency determining operation. It should be noted that, as other Fibre Channel/FICON control characters, the special latency control character "44h" is encoded in a 4-bit mapping of 65B/64B control characters. The special latency control character is mapped as Fh ("1111"), with the associated command in the following 4 bits to instruct the receiving port card to send back the special latency instruction message and back over the SONET/SDH transport network to the initial transmitter port card.

Thus, across the SONET/SDH transport network 10, the remote port card 24 upon receipt of the GFP frame sent by the port card 14 by step 35, immediately sends a GFP frame with the special latency instruction message with its latency sequence number back across the SONET/SDH transport network 10 to the port card 14. Upon receiving the GFP frame in step 38 shown in Fig. 2B, the port card 14 determines by step 39 whether the GFP frame contains the special latency sequence number. If not, the process returns to

handle the GFP frame as a transmission of data from the remote port card 24 to the local port card 14. If the GFP frame does have the special latency sequence number, then by step 40 T2 from Timer 2 indicates the time required for a GFP frame from the local port card 14 to traverse the SONET/SDH network 10 to the remote port card 24 and back again.

The appropriate number of buffers for transmission of GFP frames are calculated from T2 time interval in step 41 and in step 42 that number of buffers is allocated in the local port card 14 for the GFP frames from the transmitting port card 24. For a 100% throughput and a minimum transmission latency, it has been found that the allocation of a buffer of 2Kbytes of memory for every 2Km is effective. Thus as stated previously, for a determined time T corresponding to a 1200Km one-way trip across the SONET/SDH transport network 10, it is determined that 600 buffers, each buffer having a memory capacity of 2Kbytes, is appropriate. For a 200Km circuit, 100 buffers are appropriate based on the speed of light and a maximum Fibre Channel/FICON frame size of 2148 bytes. This calculation in step 41 is straightforward to those skilled in the card and can be substituted with a simple look-up table. Finally, Timer 2 is reset to be restarted by step 35 shown in Fig. 2A.

The distances and hence the roundtrip time T2 may change with failovers in the SONET/SDH network 10. The individual links in the network 10 may change with various failures in the links of the network 10 and the path rerouted. To accommodate these changes, the Timer 1 assures that the local port card 14 constantly monitors the latency between the local port card 14 and the remote port card 24 with a period T1, one second in this embodiment. As the distance between the port cards 14 and 24 change, measured in roundtrip time, the appropriate number of buffers is continually evaluated and set in the local port card 14. In a similar fashion, the remote port 24 sets the appropriate number of buffers for GFP frames from the local port card 14.

The embodiment of the present invention described above is best implemented in the port cards 14 and 24 in the exemplary network of Fig. 1. The operations described above require a timer and counter, besides logic. A hardware implementation in an ASIC (Application Specific Integrated Circuit) or an FPGA (Field Programmable Gate Array) is preferred for a high-speed implementation of the present invention for optimal transmission of the client data frames across the SONET/SDH transport network.

Where throughput is not necessarily paramount, the present invention might be implemented in firmware, such as the ROM (Read-Only Memory) of a microcontroller, or in software which offers certain advantages. For instance, the processor unit instructed by the software might also perform operations other than those described, or upgrades can be made easily in software. Fig. 4 shows a block diagram of a representative computer system 60 that may be used to execute the software of an embodiment of the invention. The computer system 60 includes memory 62 which can store and retrieve software programs incorporating computer code that implements aspects of the invention, data for use with the invention, and the like. Exemplary computer readable storage media include CD-ROM, floppy disk, tape, flash memory, semiconductor system memory, and hard drive. The computer system 60 further includes subsystems such as a central processor 61, fixed storage 64 (e.g., hard drive), removable storage 66 (e.g., CD-ROM drive), and one or more network interfaces 67, all connected by a system bus 68. Other computer systems suitable for use with the invention may include additional or fewer subsystems. For example, computer system 60 may include more than one processor 61 (i.e., a multi-processor system) or a cache memory. The computer system 60 may also include a display, keyboard, and mouse (not shown) for use as a host. Therefore, while the description above provides a full and complete disclosure of the preferred embodiments of the present invention, various modifications, alternate constructions, and equivalents will be obvious to those with skill in the art. Thus, the scope of the present invention is limited solely by the metes and bounds of the appended claims.

CLAIMS

1. A method of operating a transport interface for at least one local Fibre Channel/FICON port, said transport interface providing buffers for Fibre Channel/FICON data encapsulated in GFP frames transported over a SONET/SDH network from a remote Fibre Channel/FICON port, said method comprising:
- 5
- inserting a special latency instruction message into Fibre Channel/FICON data from said at least one local Fibre Channel/FICON port, said Fibre Channel/FICON data encapsulated in a GFP frame for transmission to said remote Fibre Channel/FICON port;
- 10
- sending said GFP frame over said SONET/SDH transport network to said remote Fibre Channel/FICON port;
- timing a return of said special latency instruction message over said SONET/SDH transport network;
- determining an appropriate amount of buffers in said transport interface for Fibre Channel/FICON data encapsulated in GFP frames from said remote Fibre Channel/FICON port from said timing step; and
- 15
- allocating said appropriate amount of buffers in said transport interface for GFP frames from said remote Fibre Channel/FICON port;
- whereby sufficient buffering is ensured in said transport interface to provide maximum throughput over said SONET/SDH network and any additional latency due to buffering in said transport interface is reduced.
- 20
2. The method of claim 1 wherein said inserting, sending, timing, determining and allocating steps are repeated periodically so that said amount of allocated buffers is adjusted as the latency of GFP frames transported over said SONET/SDH network between said local and remote Fibre Channel/FICON ports changes.
- 25
3. The method of claim 2 wherein said inserting, sending, timing, determining and allocating steps are repeated with a period of about 1 second.
4. The method of claim 1 wherein said special latency instruction message is inserted in a Client Payload Information field of a Payload Area of said GFP frame.
- 30
5. The method of claim 4 wherein said special latency instruction message includes a latency sequence number to identify one sequence of said inserting, sending, timing, determining and allocating steps from another sequence of said inserting, sending, timing, determining and allocating steps.

6. The method of claim 4 wherein said special latency instruction message includes a special character encoded in a 4-bit mapping of the 64B/65B control characters as Fh.

5 7. The method of claim 4 wherein said special latency instruction message includes a command to a transport interface for said remote Fibre Channel/FICON port to resend said special latency instruction message back to said transport interface for said at least one local Fibre Channel/FICON port upon receiving said special latency instruction message.

10 8. The method of claim 1 wherein said determining step comprises looking up said appropriate amount of buffers corresponding to a time interval from said timing step from a table.

 9. The method of claim 1 wherein said determining step comprises calculating said appropriate amount of buffers from a time interval from said timing step.

15 10. In a network system for transporting GFP-encapsulated Fibre Channel/FICON data across a SONET/SDH transport network between first and second Fibre Channel/FICON ports, said first Fibre Channel/FICON port connected to said SONET/SDH transport network through a first transport interface and said second Fibre Channel/FICON port connected to said SONET/SDH transport network through a second transport interface, said first transport interface comprising:

20 at least one integrated circuit adapted to insert a special latency instruction message into Fibre Channel/FICON data from said first Fibre Channel/FICON port and to encapsulate said Fibre Channel/FICON data in a GFP frame, to send said GFP frame over said SONET/SDH transport network to second transport interface of said second Fibre
25 Channel/FICON port, to time a return of said special latency instruction message over said SONET/SDH transport network, to determine an appropriate amount of buffers in said first transport interface from a time interval of said special latency instruction message to return, and to allocate said appropriate amount of buffers in said first transport interface for GFP frames from said second Fibre Channel/FICON port so that sufficient buffering is
30 ensured in said first transport interface to provide maximum throughput over said SONET/SDH network and any additional latency due to buffering in said transport interface is reduced.

11. The first transport interface of claim 10 wherein said at least one integrated circuit is adapted to insert said special latency instruction message, to encapsulate said Fibre Channel/FICON data in a GFP frame, to send said GFP frame, to time said return of
5 said special latency instruction message, to determine said appropriate amount of buffers and to allocate said appropriate amount of buffers periodically so that said amount of allocated buffers is adjusted as the actual latency of GFP frames transported over said SONET/SDH network between said first and second Fibre Channel/FICON ports changes.

12. The first transport interface of claim 11 wherein said at least one integrated
10 circuit is adapted to insert said special instruction, to encapsulate said Fibre Channel/FICON data in a GFP frame, to send said GFP frame, to time said return of said special latency instruction message, to determine said appropriate amount of buffers and to allocate said appropriate amount of buffers repeatedly with a period of about 1 second.

13. The first transport interface of claim 10 wherein said at least one integrated
15 circuit is adapted to insert said special latency instruction message in a Client Payload Information field of a Payload Area of said GFP frame.

14. The first transport interface of claim 13 wherein said special instruction
message includes a latency sequence number to identify one sequence of said inserting,
sending, timing, determining and allocating steps from another sequence of said inserting,
20 sending, timing, determining and allocating steps.

15. The first transport interface of claim 13 wherein said special instruction
message includes a special character encoded in a 4-bit mapping of the 64B/65B control
characters as Fh.

16. The first transport interface of claim 13 wherein said special latency
25 instruction message includes a command to said second transport interface for said second Fibre Channel/FICON port to resend said special latency instruction message back to said first transport interface for said first Fibre Channel/FICON port upon receiving said special latency instruction message.

17. The first transport interface of claim 10 wherein said at least one integrated
30 circuit is adapted to determine said appropriate amount of buffers in said first transport interface from said time interval of return of said special latency instruction message by

looking up said appropriate amount of buffers corresponding to said time interval of return of said special latency instruction message.

18. The first transport interface of claim 10 wherein said at least one integrated
5 circuit is adapted to determine said appropriate amount of buffers in said first transport interface from said time interval of return of said special latency instruction message by calculating said appropriate amount of buffers from said time interval of return of said special latency instruction message.

19. In a network system for transporting GFP-encapsulated Fibre Channel/FICON
10 data across a SONET/SDH transport network between first and second Fibre Channel/FICON ports, said first Fibre Channel/FICON port connected to said SONET/SDH transport network through a first transport interface and said second Fibre Channel/FICON port connected to said SONET/SDH transport network through a second transport interface, said first transport interface comprising:

15 means for inserting a special latency instruction message into said Fibre Channel/FICON data from said first Fibre Channel/FICON port and for encapsulating said Fibre Channel/FICON data in a GFP frame for transmission to said second Fibre Channel/FICON port;

20 means for sending said GFP frame over said SONET/SDH transport network to said second Fibre Channel/FICON port;

means for timing a return of said special latency instruction message over said SONET/SDH transport network;

25 means for determining an appropriate amount of buffers in said first transport interface for Fibre Channel/FICON data encapsulated in GFP frames from said second Fibre Channel/FICON port from said timing means; and

means for allocating said appropriate amount of buffers in said transport interface for GFP frames from said second Fibre Channel/FICON port;

30 whereby sufficient buffering is ensured in said first transport interface to provide maximum throughput over said SONET/SDH network and any additional latency due to buffering in said transport interface is reduced.

20. The transport interface of claim 19 wherein said inserting, sending, timing, determining and allocating means operate periodically so that said amount of allocated buffers is adjusted as the latency of GFP frames transported over said SONET/SDH network between said local and remote Fibre Channel/FICON ports changes

21. The transport interface of claim 20 wherein said inserting, sending, timing, determining and allocating means operate repeatedly with a period of about 1 second.

22. The transport interface of claim 19 wherein said special latency instruction message is inserted in a Client Payload Information field of a Payload Area of said GFP frame.

23. The transport interface of claim 22 wherein said special latency instruction message includes a latency sequence number to identify one sequence of said inserting, sending, timing, determining and allocating operations from another sequence of said inserting, sending, timing, determining and allocating operations.

24. The transport interface of claim 22 wherein said special latency instruction message includes a special character encoded in a 4-bit mapping of the 64B/65B control characters as Fh.

25. The transport interface of claim 22 wherein said special latency instruction message includes a command to said second transport interface for said second Fibre Channel/FICON port to resend said special latency instruction message back to said first transport interface for said first Fibre Channel/FICON port upon receiving said special latency instruction message




Fig. 1

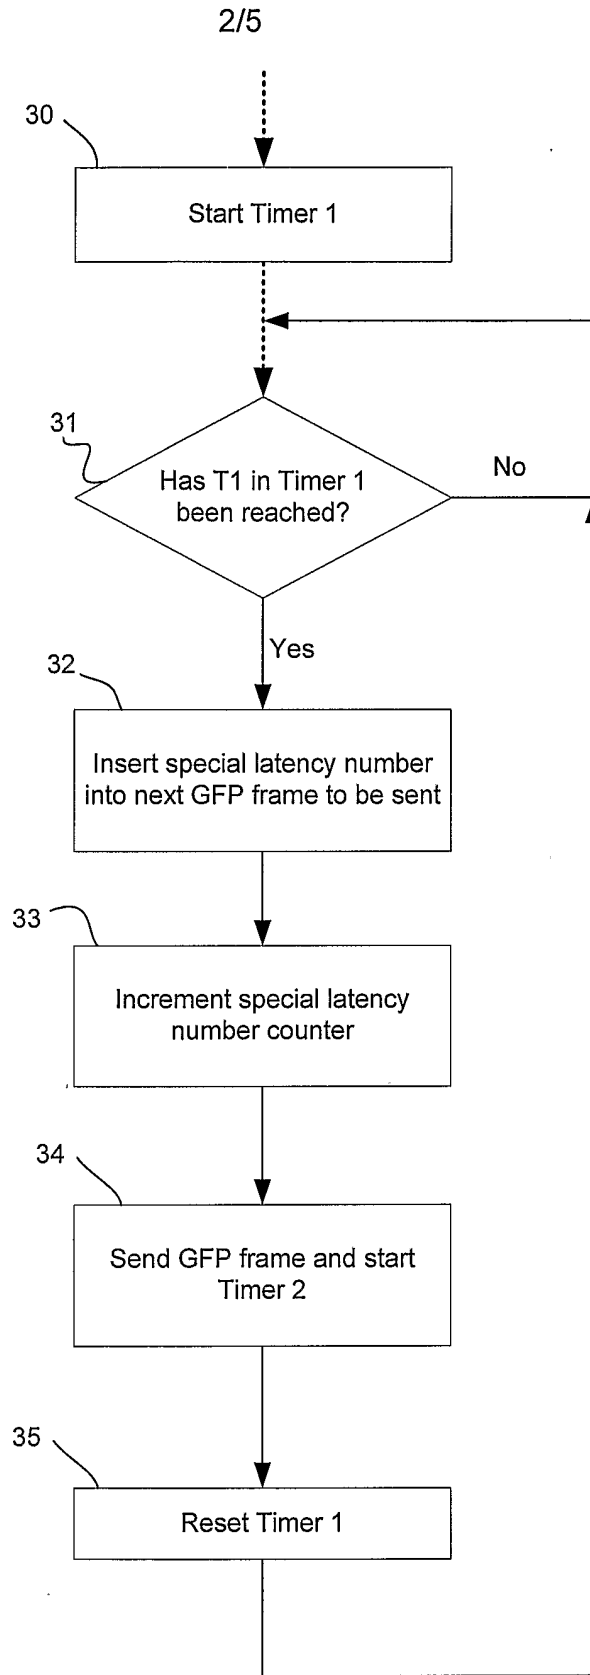


Fig. 2A

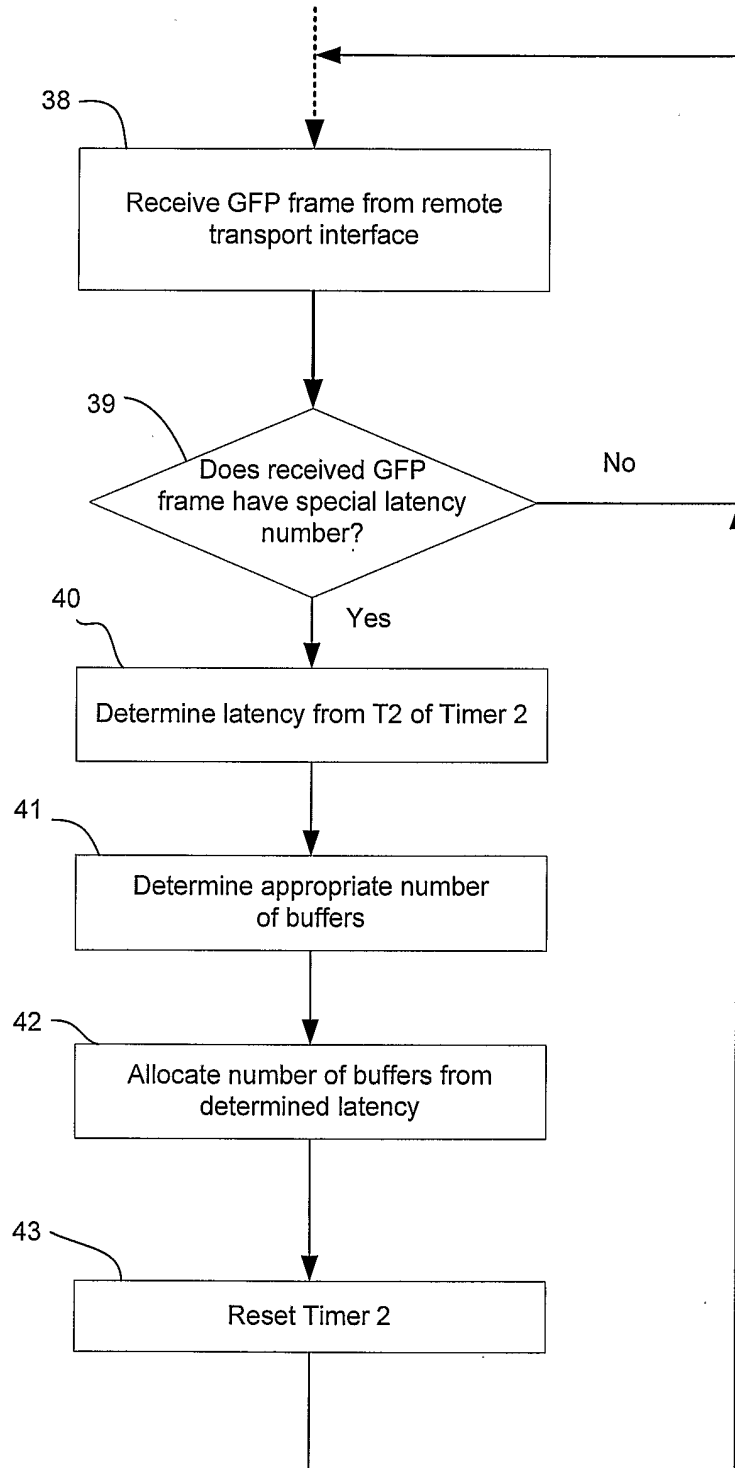


Fig. 2B




Fig. 3A

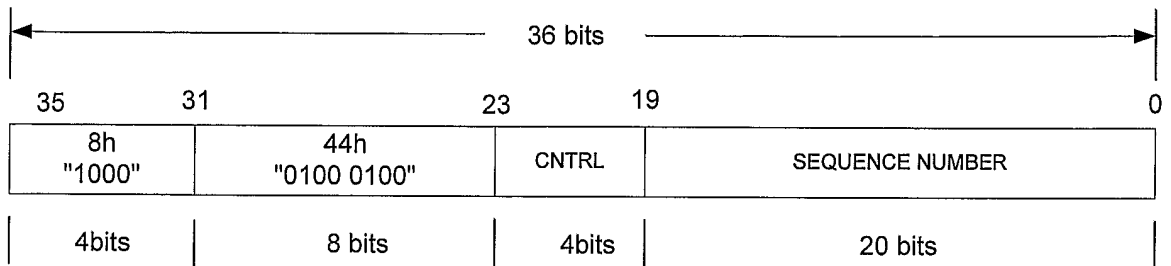


FIG. 3B

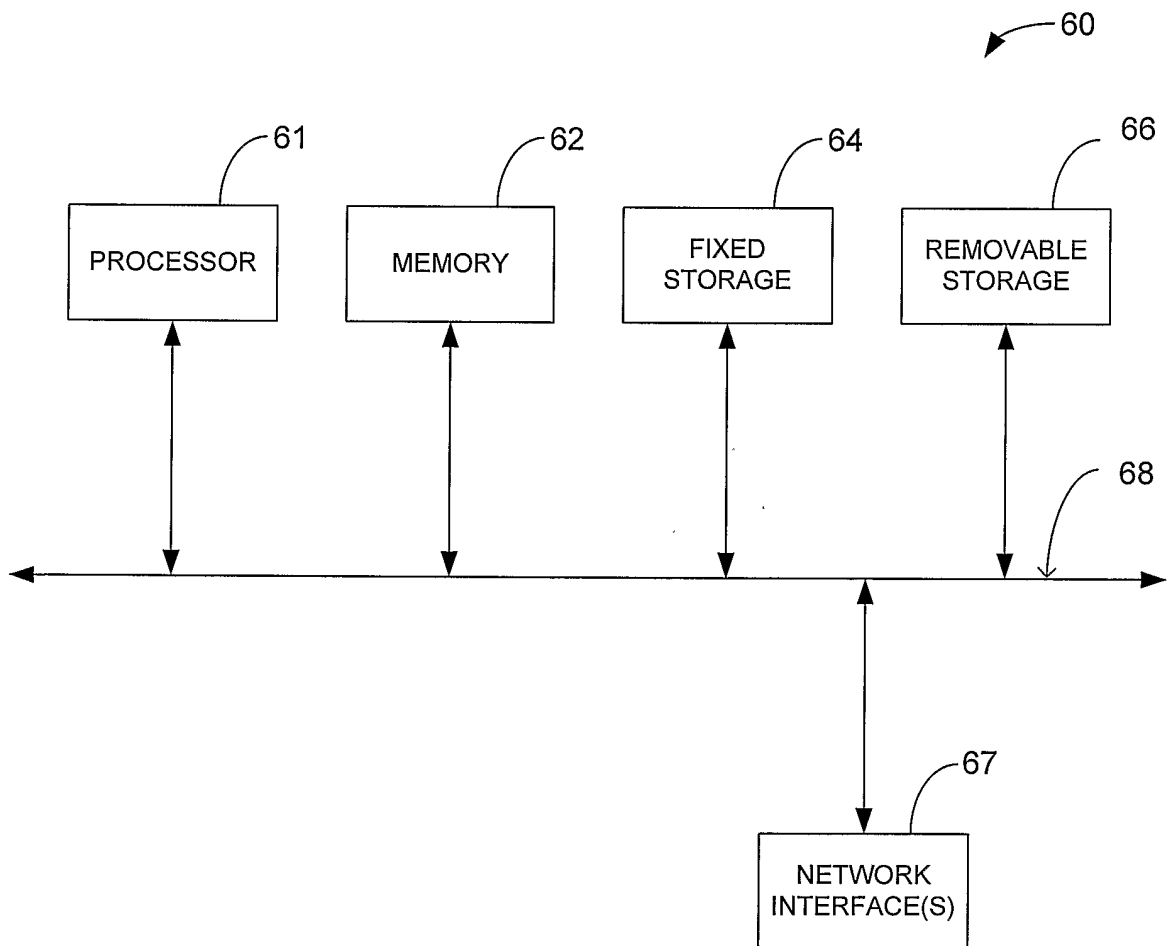


FIG. 4