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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
16 October 2003 (16.10.2003)

PCT

(10) International Publication Number
WO 03/085900 A1

(51) International Patent Classification[7]: H04L 12/46

(21) International Application Number: PCT/US03/09442

(22) International Filing Date: 25 March 2003 (25.03.2003)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
10/114,394 1 April 2002 (01.04.2002) US

(71) Applicant *(for all designated States except US)*: **ANDIAMO SYSTEMS, INC.** [US/US]; 375 East Tasman Drive, San Jose, CA 95134 (US).

(72) Inventors; and
(75) Inventors/Applicants *(for US only)*: **LEE, Scott, S.** [US/US]; 409 Calado Avenue, Campbell, CA 95008 (US). **DUTT, Dinesh, G.** [IN/US]; 1176 Corral Avenue, Sunnyvale, CA 94086 (US). **EDSALL, Thomas, J.** [US/US]; 13208 Peacock Court, Cupertino, CA 95014 (US).

(74) Agent: **WEAVER, Jeffrey, K.**; Beyer Weaver & Thomas, LLP, P.O. Box 778, Berkeley, CA 94704-0778 (US).

(81) Designated States *(national)*: AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States *(regional)*: ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:
— *with international search report*

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.

(54) Title: LABEL SWITCHING IN FIBRE CHANNEL NETWORKS

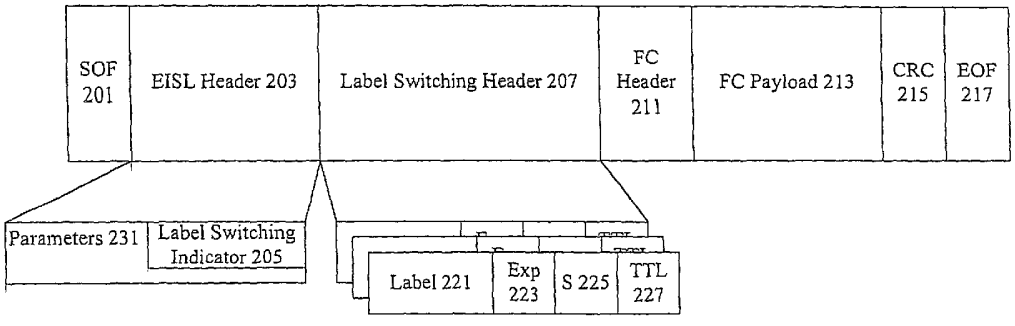

(57) Abstract: Methods and apparatus are provided for label switched routing in fibre channel networks. Techniques are provided for implementing label switching based on particular characteristics of fibre channel networks. By using label switching, mechanisms such as traffic engineering, security, and tunneling through networks that do not support fibre channel frames can be implemented.

WO 03/085900 A1

LABEL SWITCHING IN FIBRE CHANNEL NETWORKS

Cross-reference to Related Applications

The present application is related to concurrently filed U.S. Patent Application No. 10/114,568 (Attorney Docket No. ANDIP008) by Maurilio Cometto and Scott S. Lee and titled Methods and Apparatus For Fibre Channel Frame Delivery, the entirety of which is incorporated by reference for all purposes. The present application is also related to U.S. Patent Application No. 10/034,160 by Tom Edsall, Dinesh Dutt, and Silvano Gai and titled Extended ISL Header as of filing on December 26, 2001, the entirety of which is incorporated by reference for all purposes.

Background of the Invention

1. Field of the Invention.

The present invention relates to fibre channel networks. More specifically, the present invention relates to methods and apparatus for label switching in fibre channel networks.

2. Description of Related Art

In connectionless networks such as packet-switched networks, label switching has conventionally been used to allow for various features. However, it has been difficult to extend label switching into fibre channel network because of particular characteristics of fibre channel networks.

It is therefore desirable to provide methods and apparatus for using label switching in fibre channel networks not only to allow faster access to routing table entries, but also to generally improve fibre channel frame delivery.

Summary of the Invention

Methods and apparatus are provided for label switched routing in fibre channel networks. Techniques are provided for implementing label switching based on particular characteristics of fibre channel networks. By using label switching, mechanisms such as traffic engineering, security, and tunneling through networks that do not support fibre channel frames can be implemented.

According to various embodiments, a method for routing fibre channel frames in a fibre channel fabric is provided. The method includes receiving a fibre channel frame having a first stack of labels at a fibre channel switch and referencing

an entry in a label information base at the fibre channel switch based on the stack of incoming labels. The method also includes removing the first stack of labels from the fibre channel frame and forwarding the fibre channel frame.

According to various embodiments, a method for tunneling fibre channel frames is provided. The method includes receiving a fibre channel frame at a gateway between a first network supporting fibre channel and a second network not supporting fibre channel. The method also includes identifying an incoming label associated with the fibre channel frame, the incoming label determined using fibre channel routing mechanisms. The incoming label associated with the fibre channel frame is swapped with an outgoing label, the outgoing label determined by referencing an entry in the label information base associated with the gateway. The method also includes inserting additional labels to the fibre channel frame, wherein the additional labels are determined using non-fibre channel routing mechanisms. The additional labels used to forward the frame in the second network.

According to still other embodiments, a method for configuring a tunnel in a fibre channel network is provided. The method includes receiving augmented link state update information at an ingress fibre channel label switching router, selecting a route from the ingress fibre channel label switching router through a plurality of core fibre channel label switching routers to an egress fibre channel label switching router using the augmented link state update information, and generating a tunnel setup message having information identifying the plurality of core fibre channel label switching routers.

According to other embodiments, an ingress fibre channel label switching router is provided. The ingress fibre channel label switching router includes a memory and a processor. The processor is operable to receive augmented link state update information, select a route through a plurality of core fibre channel label switching routers to an egress fibre channel label switching router using the augmented link state update information, and generate a tunnel setup message having information identifying the plurality of core fibre channel label switching routers.

These and other features and advantages of the present invention will be presented in more detail in the following specification of the invention and the accompanying figures, which illustrate by way of example the principles of the invention.

Brief Description of the Drawings

The invention may best be understood by reference to the following description taken in conjunction with the accompanying drawings, which are illustrative of specific embodiments of the present invention.

Figure 1 is a diagrammatic representation of a network that can use the techniques of the present invention.

Figure 2 is a diagrammatic representation of a fibre channel frame supporting label switching.

Figure 3 is a diagrammatic representation of a routing table in a label switching router.

Figure 4A is a flow process diagram showing generation of a tunnel at an ingress label switching router.

Figure 4B is a flow process diagram showing generation of a tunnel at a core label switching router.

Figure 4C is a flow process diagram showing generation of a tunnel at an egress label switching router.

Figure 5A is a flow process diagram showing ingress label switching router operations.

Figure 5B is a flow process diagram showing core label switching router operations.

Figure 5C is a flow process diagram showing egress label switching router operations.

Figure 6 is a diagrammatic representation showing tunneling using label switching.

Figure 7 is a diagrammatic representation showing traffic engineering using label switching.

Figure 8 is a diagrammatic representation of one example of a fibre channel network that supports in order delivery.

Detailed Description of Specific Embodiments

Reference will now be made in detail to some specific embodiments of the invention including the best modes contemplated by the inventors for carrying out the invention. Examples of these specific embodiments are illustrated in the accompanying drawings. While the invention is described in conjunction with these

specific embodiments, it will be understood that it is not intended to limit the invention to the described embodiments. On the contrary, it is intended to cover alternatives, modifications, and equivalents as may be included within the spirit and scope of the invention as defined by the appended claims.

Methods and apparatus of the present invention provide for label switching of fibre channel frames. According to various embodiments, fibre channel frames include label stacks that allow fibre channel frames to be tunneled through networks that do not support fibre channel frames and rerouted around downed links. Fibre channel frames can also be delivered in order using label switching.

In a typical connectionless fibre channel network such as a class two or class three fibre channel network, a packet travels from one network entity to the next network entity based on an independent forwarding decision at each switch. The next hop for a frame is determined based on information including a destination address in the frame header. In a label switching domain, however, label switching routers make forwarding decisions based not on the destination address in the frame header but instead based on label information associated with a frame. No analysis of the packet header or frame header is needed at each hop.

Instead, the label in the packet or frame is used to index an entry in a forwarding table that contains the next hop information and a set of one or more new labels to be used as the packet is forwarded. The next hop can then use a new label to forward the frame. Although label switching was originally developed in TCP/IP networks to simplify access to routing table entries, the techniques of the present invention contemplate using label switching in fibre channel networks to enable features such has traffic engineering, tunneling, and in order delivery in addition to facilitating routing table access. Label switching as Multiprotocol Label Switching (MPLS) for IP networks is described in RFC 3031.

Several obstacles prevent the implementation of label switching in fibre channel networks. One obstacle is that some fibre channel devices require that fibre channel frames be delivered in order. Label switching used in TCP/IP networks often can deliver packets out of order. However, network entities in TCP/IP networks can handle out of order packets. Destination nodes typically reorder packets received out of sequence. Some fibre channel devices, however, can not handle out of order frames. Furthermore, fibre channel frames do not have a mechanism for carrying labels.

Figure 1 is a diagrammatic representation of the network that can use the techniques of the present invention. According to various embodiments, a label switching domain 131 includes edge label switching routers 111 and 121, and core label switching routers 113, 115, 117, and 119. An interconnected set of network entities that support label switching for forwarding frames is referred to herein as a label switching domain. A switch that is outside the label switching domain 131 can use conventional techniques for accessing routing table entries and forwarding frames. In one example, a switch 101 may send a frame without a label to a label switching router 111. Without label switching, a label switching router 111 typically uses the destination address provided in the frame to identify a next hop for the frame. The next hop then similarly accesses the destination address and determine the subsequent hop based on the destination address to continue forwarding the frame until the frame finally reaches the destination.

Using label switching however, an edge label switching router 111 adds a label to the frame from switch 101. The label can then be used at a subsequent hop, core label switching router 113, to access a routing table to again determine where to transmit the frame. Instead of using the destination address at router 113, a label is used at router 113 to access a forwarding entry. When a label switching router is connected to a label unaware switch, the label switching router is referred to herein as an edge label switching router. Any device that does not support the use of labels for forwarding decisions is referred to herein as a label unaware switch. The edge label switching router can be an ingress label switching router when it handles traffic entering into the label switching domain. The edge label switching router can be an egress label switching router when it handles traffic leaving the label switching domain. In one example where a frame is transmitted from switch 101 to switch 105, edge label switching router 111 would be the ingress label switching router while edge label switching router 121 would be the egress label switching router. It should be noted that some label switching routers could be core and edge label switching routers.

The label switching routers connected to other label switching enabled routers are referred to herein as core label switching routers. According to various embodiments, all the incoming packets or frames received at core label switching routers include labels. Consequently, core label switching routers only need to look at the incoming label in order to make the forwarding decision. According to other

embodiments, some incoming packets or frames received at a core label switching router do not include labels. Consequently, some packets and frames are switched without the use of labels.

It should be noted that a switch can be both a core label switching router and an edge label switching router. Edge label switching router 121 is connected to label switching router 117 and 119 while it is connected to label unaware router 105 and label unaware storage device 107.

Figure 2 is a diagrammatic representation of a frame that can include a label for label switching. According to various embodiments, the label switching header 207 is located between the Extended Inter-switch Links (EISL) header 203 and the fibre channel header 211. EISL is described in application number __/_____ titled Methods And Apparatus For Encapsulating A Frame For Transmission In A Storage Area Network by Thomas J. Edsall, Dinesh G. Dutt, and Silvano Gai (Attorney Docket No. ANDIP001). The EISL header 203 includes a label switching indicator 205 that provides information on whether a label is provided in the frame.

According to specific embodiments, the label switching header 207 includes a stack of 32-bit words. Each label includes a 20-bit label for accessing an entry in a routing table, a 3-bit experimental field, a 1-bit EOS field, and an 8-bit time-to-live (TTL) field. The label value is used at a label switching router to determine the next hop, the stack bit is used to indicate that the bottom of a stack of labels has been reached. It should be noted that the label switching header can include more than one label. The time-to-live field is decremented at each hop like the time-to-live field in a conventional TCP/IP packet.

Although the label switching header described above is included between an EISL header and a fibre channel header 207, it should be noted that a label switching header can be included in a variety of different fields associated with the frame. For example, a label switching header can be included in a frame that does not include an EISL header 203. In this example, the label switching indicator can be provided in a fibre channel header 211, and label switching information can be included in the fibre channel payload. However, it should be noted that a label switching indicator and label switching information can instead be included in an extended fibre channel header.

Figure 3 is a diagrammatic representation of a label information base (LIB). In the example shown, label switching information and routing table information is

maintained in the same entity. However, a routing table and label switching information can be maintained in separate entities. An entity including label switching information is referred to herein as a label information base (LIB). Some LIB entities may include only a label information base and no routing table information.

According to various embodiments, a LIB with routing table information includes entries pairing a destination identifier with a next hop. That is, the routing tables include a destination identifier column 303 and a next hop column 305. To support label switching, an in or incoming label column in 301 as well as an out or outgoing label column 307 is included. When a frame is received, a label can be used to access an entry in the routing table corresponding to the label in the frame. In one example, if the label in the frame is 2000, the switch recognizes that the next hop is switch 43 and the out label should be 3000. In this example the destination ID is not used to determine the next hop.

In still other examples, the number of labels to push or pop and a determination of which labels to insert can be made not only using the destination ID, but also by policies that are configured in the switch. Some policies include the port number and source and destination pair. It should be noted that routing tables are provided on a per virtual network basis (virtual storage area network (VSAN) or virtual local area network (VLAN)). That is, a routing table is available for each VSAN/VLAN the label switching router is a part of. A SAN (or a VSAN) may be a network such as an Infiniband network.

A variety of techniques can be used for generating a LIB. In one embodiment, a LIB is generated upon the receipt of link state update packets under the FSPF protocol. According to other embodiments, a LIB is generated upon receiving augmented link state update packets. Link state packets with additional information such as bandwidth availability allowing traffic engineering are referred to herein as augmented link state update packets. A LIB can be generated periodically or upon the identification of a change in link state. According to various embodiments, a newly generated LIB is associated with an incarnation number. A combination of all the incarnation numbers in a fibre channel fabric is herein referred to as a topology version number. Using a topology version number can allow for in order delivery of fibre channel frames. Using a topology version number to allow for in order delivery is described in concurrently filed U.S. Patent

Application No. __/_____ (Attorney Docket No. ANDIP008) by Maurilio Cometto and Scott S. Lee and titled Methods and Apparatus For Fibre Channel Frame Delivery, the entirety of which is incorporated by reference for all purposes.

In one embodiment, every time a new routing table is generated at a switch, the incarnation number is incremented by one. According to various embodiments, each label switching router in a fibre channel network not only generates new forwarding routes toward each destination, but each label switching router also generates new in labels different from the previous set of in labels.

The augmented link state update packets can be used to generate conventional routing tables and/or LIBs. Packets received at a label switching router can then be routed to a next hop by using labels instead of next hop information in a routing table. Alternatively, packets can be forwarded to a next hop using label information in place of routing table information. Accordingly, label switching routers can be implemented without routing tables entirely. However, augmented link state update packets can also be used to traffic engineer and select routes not found using conventional routing table mechanisms. In one example, a traffic engineered route may forward a received next packet to a different next hop than a routing table would. Furthermore, augmented link state update packets can be used to specify a route from a source to a destination, whereas routing table information can only specify a next hop.

According to various embodiments, a source label switching router determines the best route to a particular destination. The source then sends a frame explicitly routed to each hop between the source and the destination on the selected route. A mechanism such as Resource Reservation Protocol (RSVP-TE) can be used to configure routes. RSVP-TE is described in RFC 3209, the entirety of which is incorporated by reference for all purposes.

Figure 4A is a process flow diagram showing the configuration of route using traffic engineering. At 401, the ingress label switching router determines the best route to a destination based on information such as that provided in the augmented link state update packets. The label switching router can also apply other policies for selecting a route. In one example, the label switching router may attempt to avoid certain links. At 403, the ingress label switching router creates a tunnel set up message having information identifying the label switching routers in the tunnel. Any message for configuring label switching routers on a selected route

to forward packets along the selected route is referred to herein as a tunnel setup message. At 405, the ingress label switching router sends the message to the next label switching router in the tunnel. It should be noted that the selected route can be a variety of different types of routes and tunnels.

In one example, a tunnel established is a virtual private network or VPN tunnel. In another embodiment the tunnel established is an IP tunnel. At 407, the ingress label switching router receives a response message corresponding to the tunnel set up message sent at 405. According to various embodiments, the response to the tunnel set up message originated from the destination. At 409, the label provided in the response message is programmed into the LIB.

Figure 4B is a process flow diagram showing a core label switching router in a tunnel between the source of the destination. At 421, the core label switching router receives a tunnel set up message. At 423, the core label switching router allocates labels associated with the message. At 425, the core label switching router forwards the tunnel set up message downstream to the next label switching router in the tunnel. At 427, the core label switching router receives a response message corresponding to be tunnel set up message and programs the label into the LIB at 429. At 431, the core label switching router forwards the response upstream to the previous hop in the tunnel.

Figure 4C is a process flow diagram showing an egress label switching router. At 441, the egress label switching router receives a tunnel setup message. The destination label switching router generates label information at 443 and sends a response message in the reverse direction along the same selected route at 445.

Once a tunnel is established, label switching operations may vary based on whether a label switch is an ingress label switching router, a core label switching router, or an egress label switching router. Figure 5A is a flow process diagram showing one example of label push operations at an ingress label switching router. According to various embodiments, an ingress label switching router receives a frame at 501 from a label unaware node. At 503, the ingress label switching router classifies the frame. At 505, the ingress label switching router identifies the LIB entry corresponding to the classified frame. In one embodiment, an input or output port identified may be used to select a LIB entry.

At 507, the number of labels and the labels to be pushed onto the label stack are determined. Any mechanism for holding labels and information associated with

labels is referred to herein as a label stack. A label stack can be a stack, a linked list, an array, or any structure containing label information. The frame is then modified at 509 to include the one or more labels. Modifying the frame can include updating an EISL header to show that a label is available and placing the label information into a label header.

Figure 5B is a flow process diagram showing one example of core label switching router operations. As noted above, a core label switching router receives frames from a label switching enabled router. At 521, a core label switching router receives a frame. At 523, it is determined if an LIB entry corresponds to the incoming label associated with the frame. If no entry corresponds, the frame is dropped at 531. According to various embodiments, various error reporting and notification operations can also be performed. Although it may be possible to route the frame based on a routing table next hop, the frame is dropped in various embodiments in order to limit the chance of a loop in the network. At 525, the labels to be popped, pushed, or swapped are determined based on the LIB entry. At 527, the frame is modified to add, remove, or replace label information.

Figure 5C is a flow process diagram showing one example of egress label switching router operations. At 541, the egress label switching router receives a frame from a label switching enabled router. It is determined at 543 whether a label exists in the packet. If no label exists, the packet is forwarded based on VSAN and destination ID. If a label exists, it is determined whether an LIB entry corresponds to the incoming label at 545. If no LIB entry corresponds to the label at 545, the frame is dropped at 553. Otherwise, the number of labels to pop is determined at 547. The frame is then modified at 549.

Figure 6 is a diagrammatic representation showing tunneling through a TCP/IP network 655 the does not support fibre channel frames. The label switching router 651 provides a frame to gateway 653. The frame includes a first label 665, the fibre channel header 663, and a fibre channel payload 661. According to various embodiments, the frame may not include a first label 665. The gateway 653 recognizes that it is about to forward a labeled fibre channel frame through a TCP/IP network 655 that does not support fibre channel frames. The fibre channel gateway 653 uses a frame including a top label 675, an ethernet header 677, and a second label 673. The top label is used to pass traffic from gateway 653 to gateway 657 and the bottom label is used to forward the traffic to label switching router 659 after

the frame reaches gateway 657. The top label is determined by forwarding mechanisms in the TCP/IP network 655 while the second label is determined by forwarding mechanisms in the fibre channel network.

According to various embodiments, the label switching routers in the TCP/IP network 655 only operate on the topmost label 675 and do not need to access any other labels in the label stack. Accordingly, the entities in the TCP/IP network 655 do not need to be aware that the fibre channel frame is being tunneled through the TCP/IP network 655. Instead, the TCP/IP network entities merely forward frames through the network based on the topmost label.

Unlike conventional TCP/IP networks, fibre channel frames cannot be delivered out of sequence. Accordingly, a control word 671 including a sequence number is included after the label stack, or beneath second label 673, so that the gateway 657 can detect packets arriving out of sequence. It should be noted that a control word can be included in different fields in the fibre channel frame, such as in the fibre channel header. Any mechanism allowing a fibre channel gateway to detect out of order fibre channel frames after transmission through a non fibre channel network is referred to herein as a control word.

Also unlike conventional TCP/IP networks, fibre channel network frames are typically not allowed to be dropped for performance reasons. Techniques of the present invention contemplates a gateway 657 detecting that frames from gateway 653 were dropped and requesting retransmission in order to provide for efficient frame delivery. It should be noted that by tunneling using labels, security can also be provided by using network protocols such as Virtual Private Network or VPN.

Figure 7 is a diagrammatic representation showing fast link failover, according to various embodiments. A label switching router 783 receiving a frame from label switching router 781 can recognize that a link between router 783 and router 789 is down even though the routing table instructs router 783 to forward the frame directly to router 789 through the downed link. Instead of waiting for conventional link state update and link state record techniques to update routing tables, labels can be used to more quickly reroute traffic around the downed link. According to various embodiments, an additional label is pushed onto the label stack associated with the frame to forward the frame to a label switching router 785. The label switching router 785 uses the top label with the value of 100 to determine that the frame should be forwarded to a label switching router 787. The top label is

then replaced with the value of 200. A label switching router 787 then removes the incoming label 200 and forwards the frame based on the second label with the value of 20 to label switching router 789.

The downed link between label switching router 783 and label switching router 789 is bypassed. It should be noted that traffic engineering using label switching can be implemented in a variety of different manners. In one example, a system administrator can manually set up alternative routes at a label switching router 783. A switch 783 is manually configured to replace a label stack with a value of 10 with a label stack with a top label value of 100 and a second label value of 20. Link state information can also be passed into the network automatically.

Although the techniques of the present invention can be used to provide features such as fast failover, explicit source routing, and traffic engineering as noted above, the techniques of the present invention can also be used to provide for in order delivery.

Figure 8 is a diagrammatic representation of one example of a fibre channel network that supports in order delivery. In order delivery is described in concurrently filed U.S. Patent Application No. 10/114,568 (Attorney Docket No. ANDIP008) by Maurilio Cometto and Scott S. Lee and titled Methods and Apparatus For Fibre Channel Frame Delivery, the entirety of which is incorporated by reference for all purposes.

In addition to containing the destination address, a frame includes as a destination identifier an input label that allows a switch to quickly access an entry in a routing table. For example, a label switching router 804 can receive a frame with a destination of 2 and an in label of 420. The label switching router 804 can access its routing table 814 to recognize that the next hop is label switching router 802 and the output should be 220. According to various embodiments, the label switching router 804 replaces the frame label value of 420 corresponding to the in label in the routing table with a frame label of 220 corresponding to the out label in the routing table 814.

By replacing the label value, the label switching router 804 provides label information to the next hop router 802, to allow the label switching router 802 to similarly access a routing table entry quickly. It should be noted that although label switching can be provided for fast access of entries in a routing table, label

switching can be used for a variety of reasons. The techniques of the present invention provide that frames can be delivered in order by using labels.

When a label switching router 802 receives a frame from label switching router 804, the label switching router uses the label 220 to access an entry in the routing table 812. Using the in label 220, the label switching router 802 recognizes that the frame no longer needs to be forwarded, as the frame has actually arrived at its destination.

As described above, label switching may be performed in a variety of network devices. According to various embodiments, the switch includes a processor, network interfaces, and memory for maintaining LIBs. A variety of input and output ports, Media Access Control (MAC) blocks, and buffers can also be provided as will be appreciated by one of skill in the art.

In addition, although an exemplary switch is described, the above-described embodiments may be implemented in a variety of network devices (e.g., servers) as well as in a variety of mediums. For instance, instructions and data for implementing the above-described invention may be stored on a disk drive, a hard drive, a floppy disk, a server computer, or a remotely networked computer. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

CLAIMS

What is claimed is:

1. A method for routing fibre channel frames in a fibre channel fabric, the method comprising:

receiving a fibre channel frame having a first stack of labels at a fibre channel switch;

referencing an entry in a label information base at the fibre channel switch based on the stack of incoming labels;

removing the first stack of labels from the fibre channel frame; and

forwarding the fibre channel frame.

2. The method of claim 1, wherein the fibre channel frame is forwarded substantially without any label information.

3. The method of any of claims 1-2, further comprising:

inserting a second set of labels into the fibre channel frame.

4. The method of claim 3, wherein the fibre channel frame is forwarded to a next hop based on the entry in the label information base.

5. The method of claim 3, wherein the fibre channel switch is a core label switching router.

6. The method of any of claims 1-5, wherein the stack of labels comprises of one or more labels.

7. The method of claim 6, wherein the one or more of the labels are associated with one or more virtual storage area networks.

8. The method of any of claims 1-7, wherein the first label stack is included in the EISL header.

9. The method of any of claims 1-7, wherein the first label stack is included in the in the fibre channel frame header.

10. An apparatus for routing fibre channel frames in a fibre channel fabric, the apparatus comprising:

means for receiving a fibre channel frame having a first stack of labels at a fibre channel switch;

means for referencing an entry in a label information base at the fibre channel switch based on the stack of incoming labels;

means for removing the first stack of labels from the fibre channel frame; and

means for forwarding the fibre channel frame.

11. A method for tunneling fibre channel frames, the method comprising:

receiving a fibre channel frame at a gateway between a first network supporting fibre channel and a second network not supporting fibre channel;

identifying an incoming label associated with the fibre channel frame, the incoming label determined using fibre channel routing mechanisms;

swapping the incoming label associated with the fibre channel frame with an outgoing label, the outgoing label determined by referencing an entry in the label information base associated with the gateway;

inserting additional labels to the fibre channel frame, wherein the additional labels are determined using non-fibre channel routing mechanisms, the additional labels used to forward the frame in the second network.

12. The method of claim 11, wherein the second network is a TCP/IP network supporting label switching.

13. The method of any of claims 11-12, wherein the additional labels correspond to entries in a label information base of a second network.

14. The method of claim 13, further comprising adding a control word to the fibre channel frame.

15. The method of claim 14, wherein the control word is used at the tunnel end gateway to reorder frames received out of sequence.

16. The method of any of claims 11-15, wherein a label indicator is contained in the EISL header.

17. The method of any of claims 11-16, wherein a label indicator is contained in the fibre channel header.

18. A computer readable medium comprising computer code for configuring a fibre channel switch, the computer readable medium comprising:

computer code for receiving a fibre channel frame at a gateway between a first network supporting fibre channel and a second network not supporting fibre channel;

computer code for identifying an incoming label associated with the fibre channel frame, the incoming label determined using fibre channel routing mechanisms;

computer code for swapping the incoming label associated with the fibre channel frame with an outgoing label, the outgoing label determined by referencing an entry in the label information base associated with the gateway;

computer code for inserting additional labels to the fibre channel frame, wherein the additional labels are determined using non-fibre channel routing mechanisms, the additional labels used to forward the frame in the second network.

19. An ingress fibre channel label switching router, the router comprising:

a memory; and

a processor operable to receive augmented link state update information, select a route through a plurality of core fibre channel label switching routers to an egress fibre channel label switching router using the augmented link state update information, and generate a tunnel setup message having information identifying the plurality of core fibre channel label switching routers.

20. The router of claim 19, wherein the processor is further operable to forward the tunnel setup message along the selected route.

21. The method of claim 20, wherein the processor is further operable to receive a response corresponding to the tunnel setup message, the response including label information.

22. The method of claim 21, wherein the memory includes a label information base having label information associated with the ingress fibre channel label switching router.

Figure 1

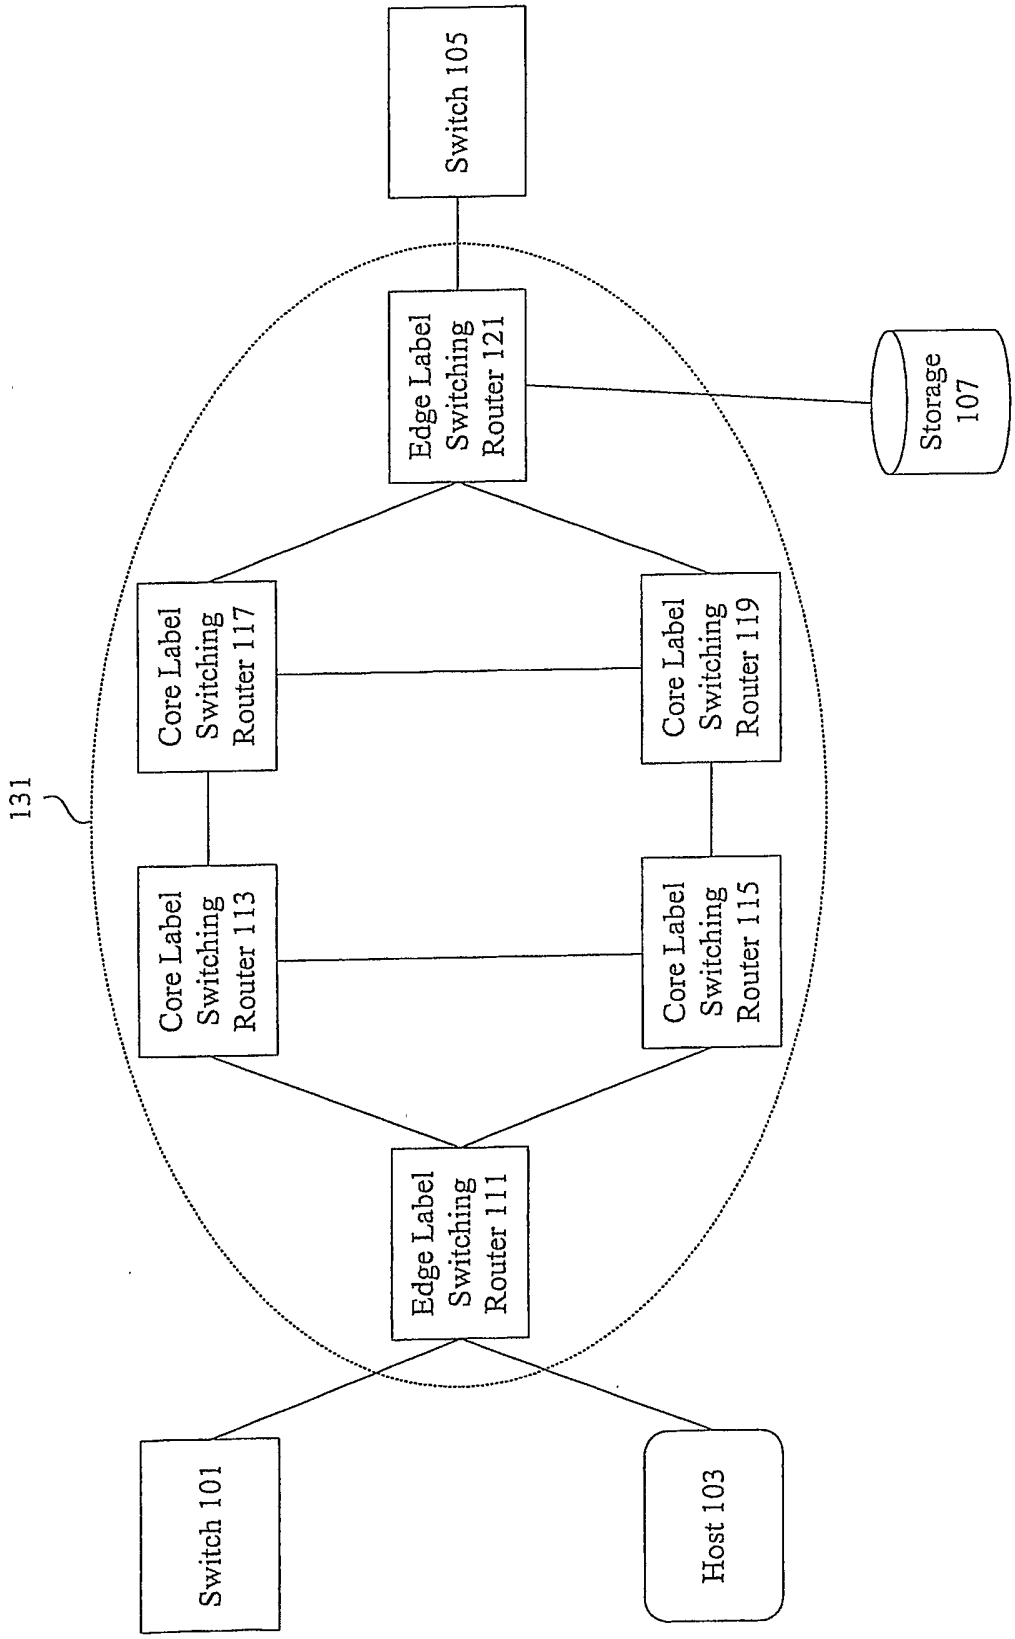
Switch 105
Edge Label Switching Router 121
Storage 107
131
Core Label Switching Router 117
Core Label Switching Router 119
Core Label Switching Router 113
Core Label Switching Router 115
Edge Label Switching Router 111
Switch 101
Host 103

Figure 2

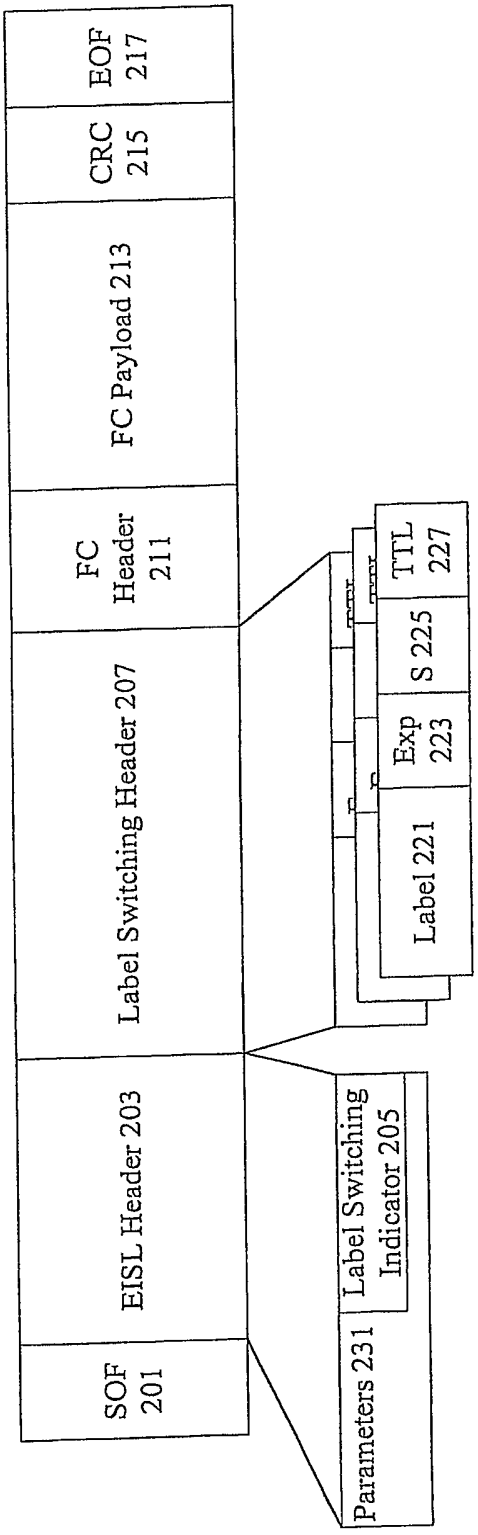
SOF 201
EISL Header 203
Label Switching Header 207
FC Header 211
FC Payload 213
CRC 215
EOF 217
Parameters 231
Label Switching Indicator 205
Label 221
Exp 223
S 225
TTL 227

Figure 3

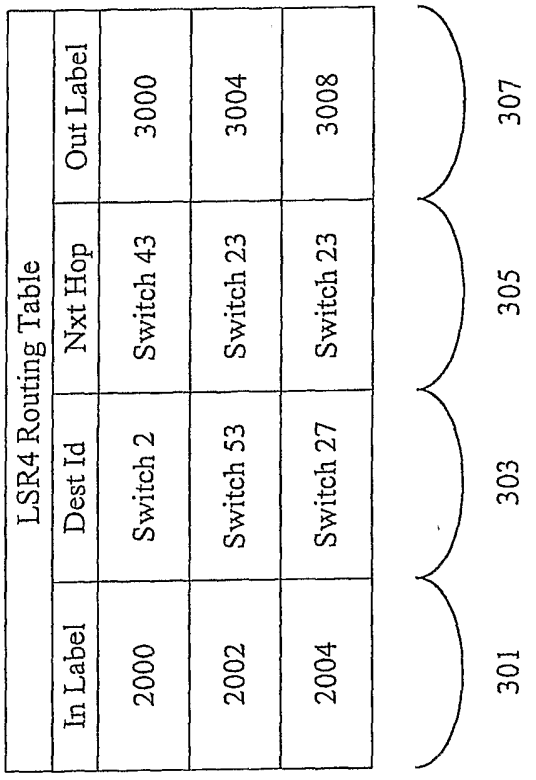
LSR4 Routing Table
In Label
Dest Id
Nxt Hop
Out Label
2000
Switch 2
Switch 43
3000
2002
Switch 53
Switch 23
3004
2004
Switch 27
Switch 23
3008
301
303
305
307

Figure 4A

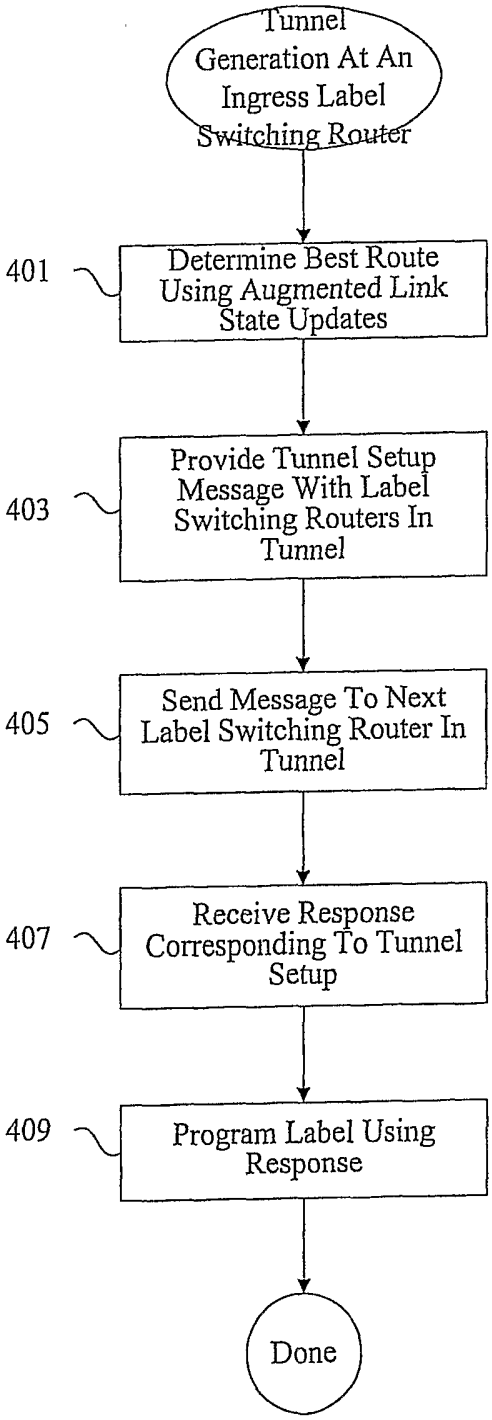
Tunnel Generation At An Ingress Label Switching Router
401
Determine Best Route Using Augmented Link State Updates
403
Provide Tunnel Setup Message With Label Switching Routers In Tunnel
405
Send Message To Next Label Switching Router In Tunnel
407
Receive Response Corresponding To Tunnel Setup
409
Program Label Using Response
Done

Figure 4B

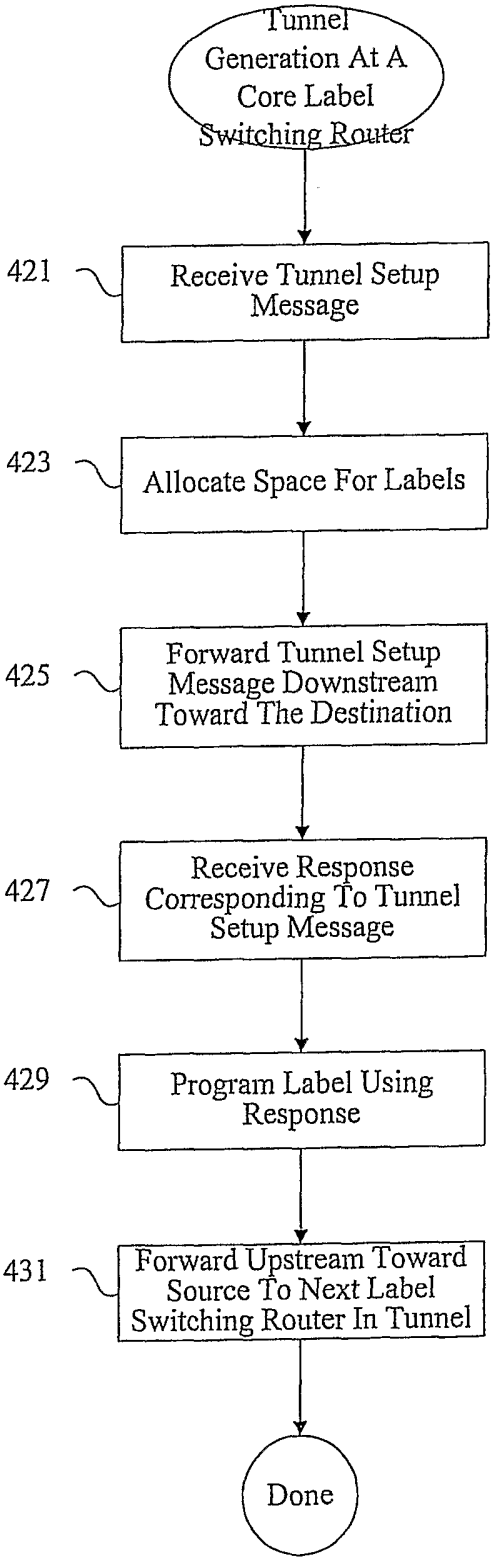
Tunnel Generation At A Core Label Switching Router
421
Receive Tunnel Setup Message
423
Allocate Space For Labels
425
Forward Tunnel Setup Message Downstream Toward The Destination
427
Receive Response Corresponding To Tunnel Setup Message
429
Program Label Using Response
431
Forward Upstream Toward Source To Next Label Switching Router In Tunnel
Done

Figure 5A

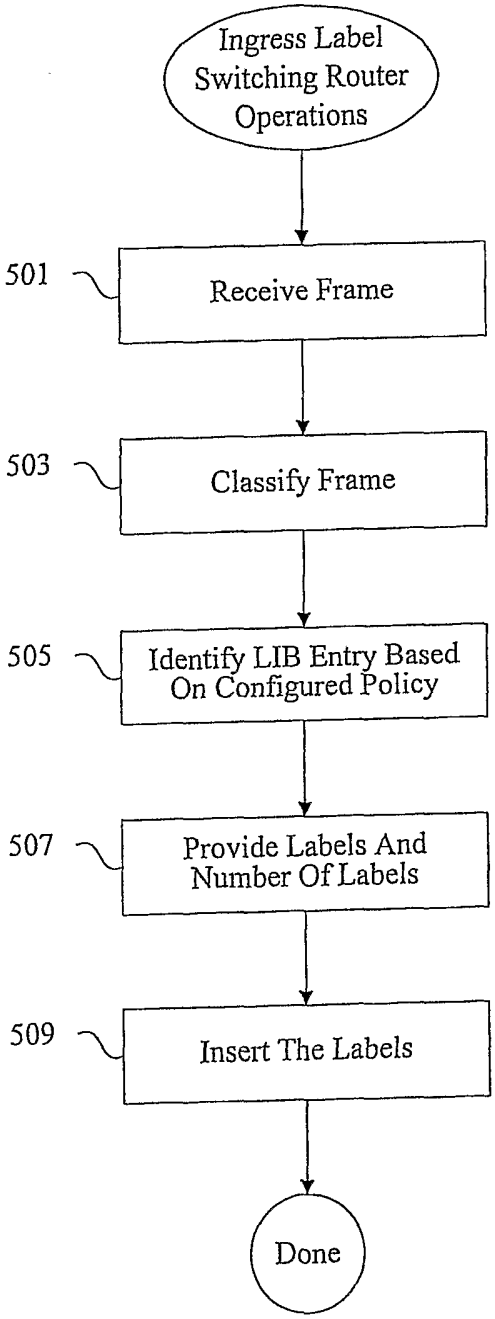
Ingress Label Switching Router Operations
501
Receive Frame
503
Classify Frame
505
Identify LIB Entry Based On Configured Policy
507
Provide Labels And Number Of Labels
509
Insert The Labels
Done

Figure 5B

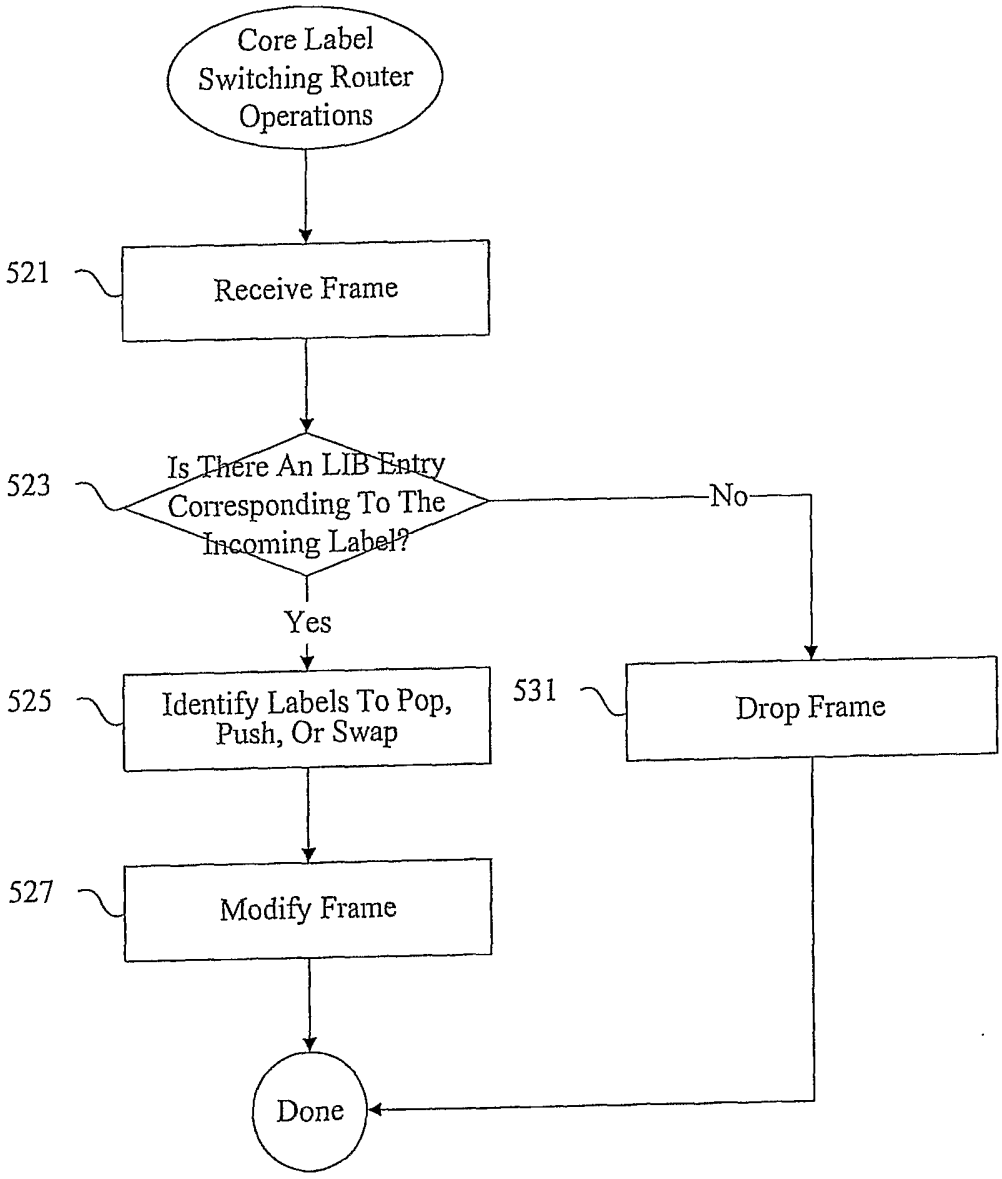
Core Label Switching Router Operations
521
Receive Frame
523
Is There An LIB Entry Corresponding To The Incoming Label?
No
Yes
525
Identify Labels To Pop, Push, Or Swap
531
Drop Frame
527
Modify Frame
Done

Figure 4C

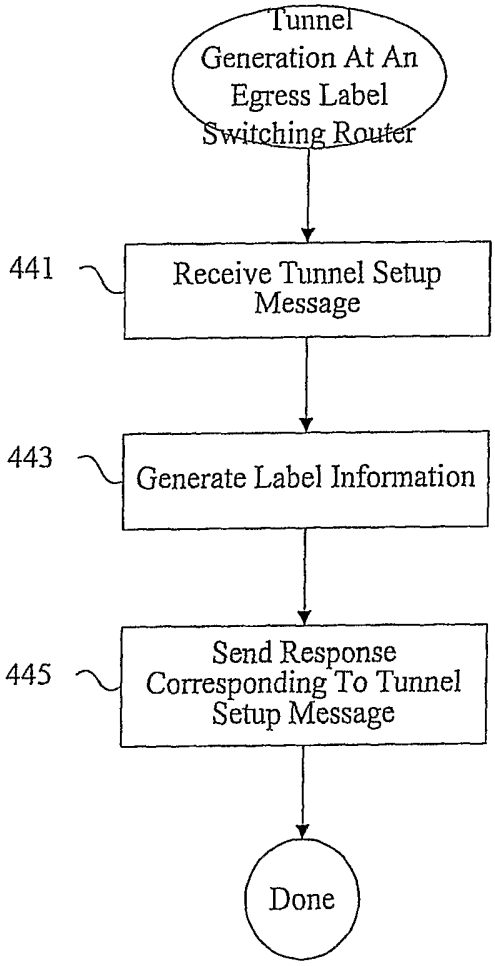
Tunnel Generation At An Egress Label Switching Router
441
Receive Tunnel Setup Message
443
Generate Label Information
445
Send Response Corresponding To Tunnel Setup Message
Done

Figure 5C

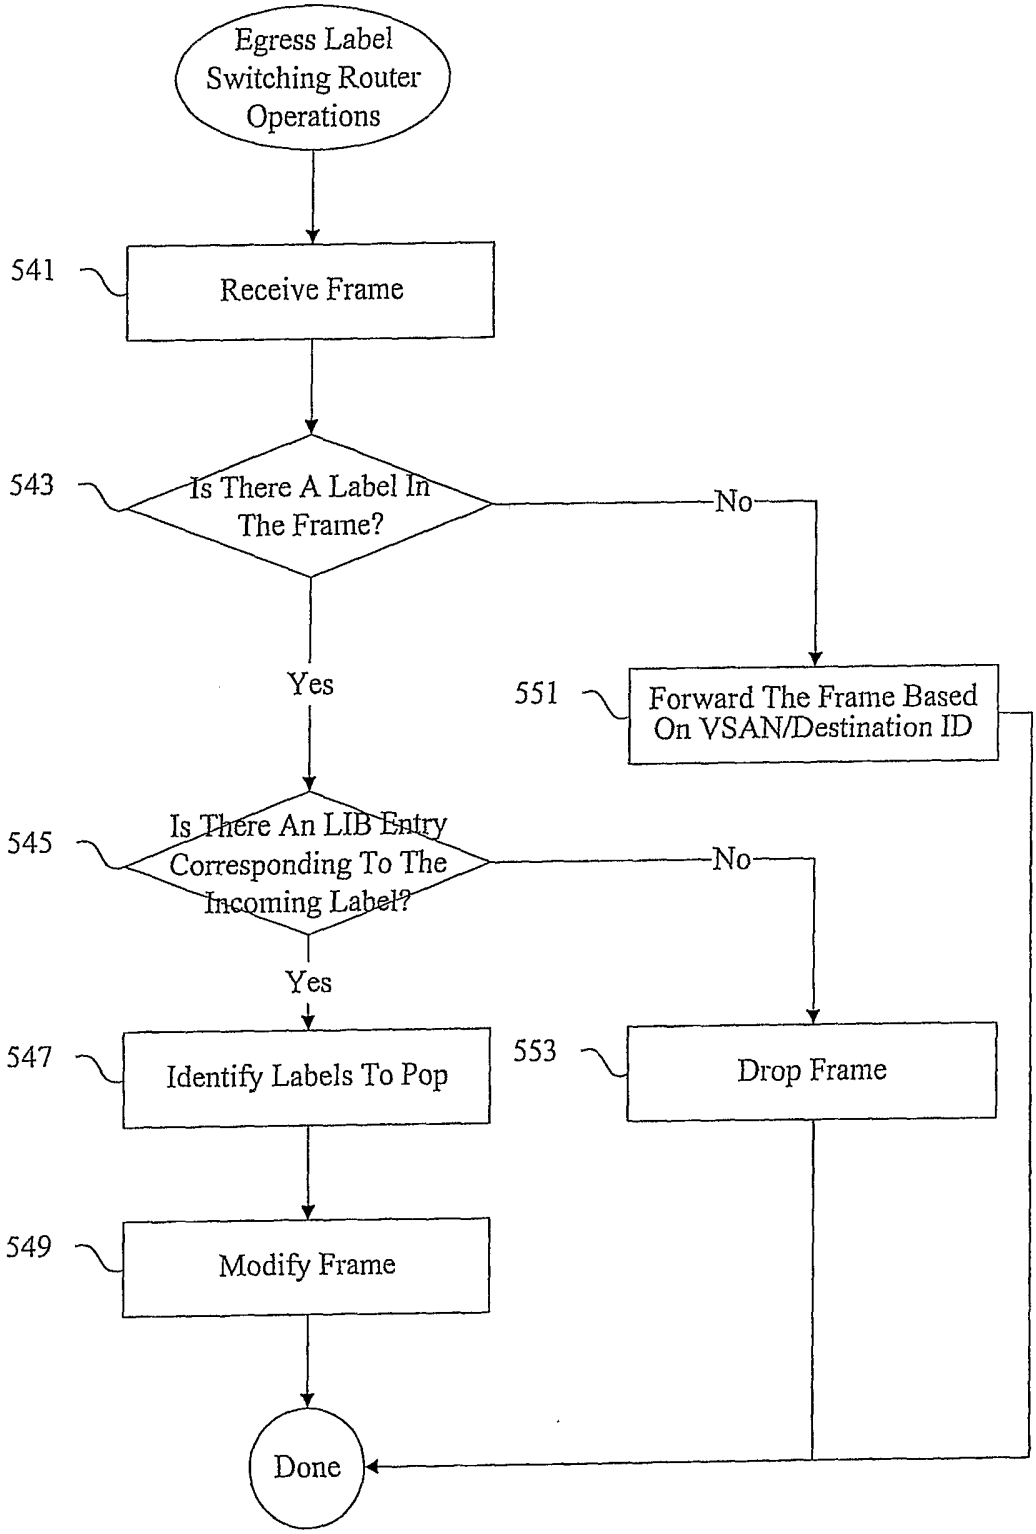
Egress Label Switching Router Operations
541
Receive Frame
543
Is There A Label In The Frame?
No
Yes
551
Forward The Frame Based On VSAN/Destination ID
545
Is There An LIB Entry Corresponding To The Incoming Label?
No
Yes
547
Identify Labels To Pop
553
Drop Frame
549
Modify Frame
Done

Figure 6

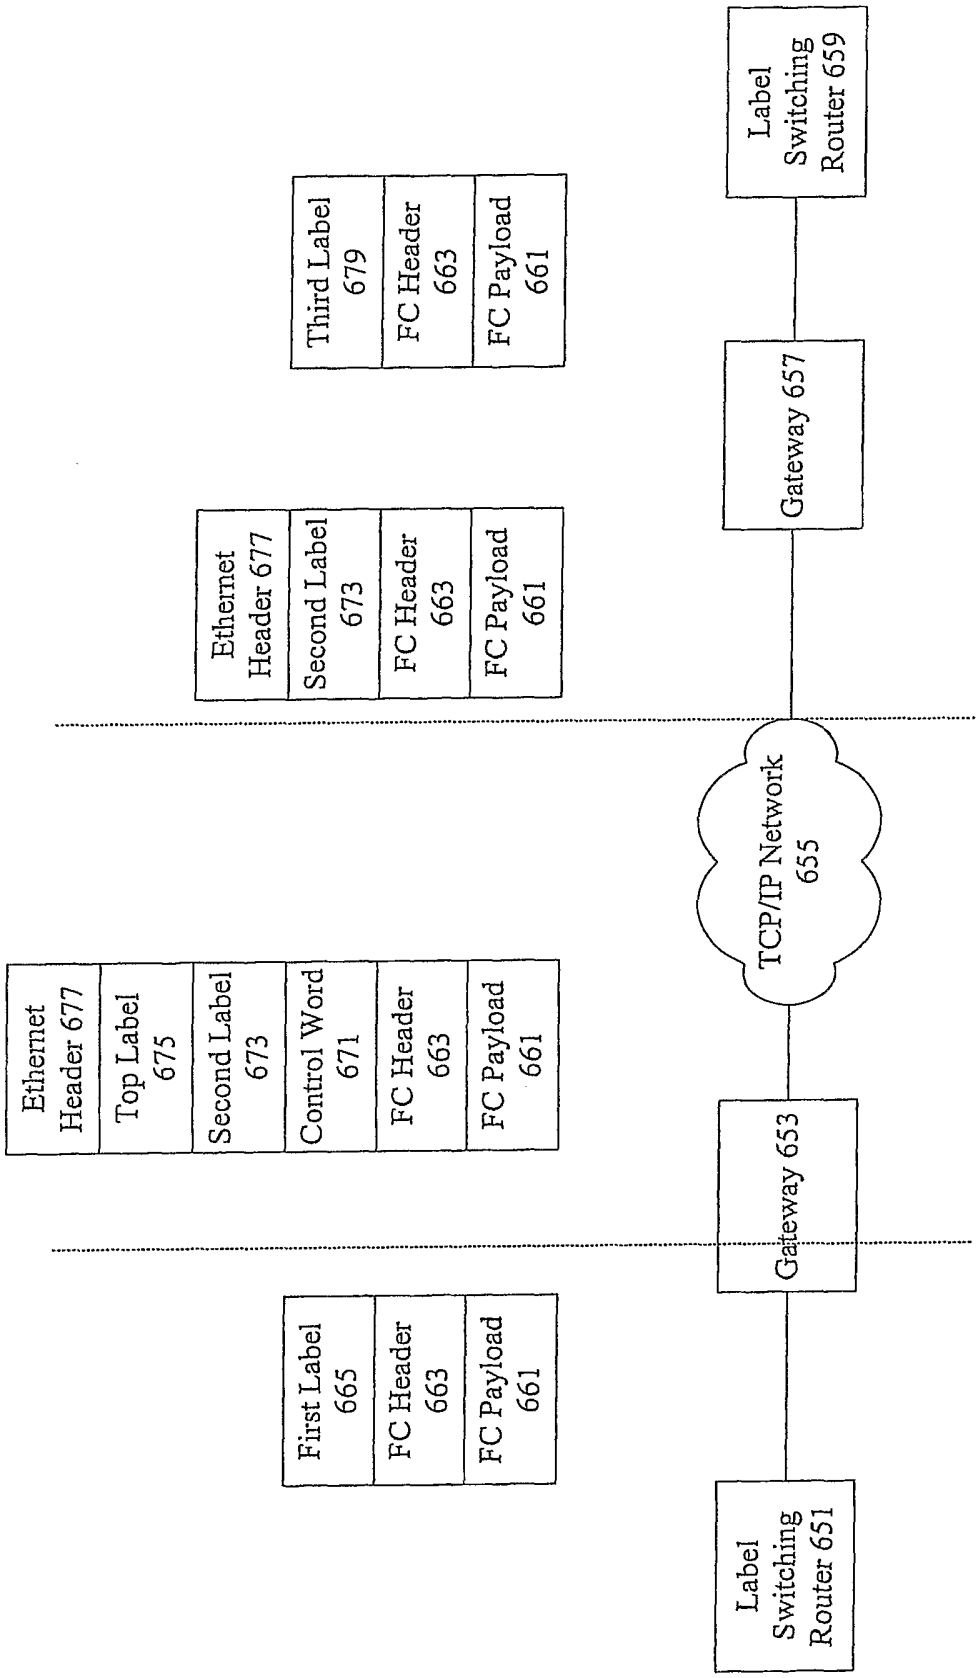
Label Switching Router 651
Gateway 653
Label Switching Router 659
Gateway 657
TCP/IP Network 655
First Label 665
FC Header 663
FC Payload 661
Ethernet Header 677
Top Label 675
Second Label 673
Control Word 671
FC Header 663
FC Payload 661
Ethernet Header 677
Second Label 673
FC Header 663
FC Payload 661
Third Label 679
FC Header 663
FC Payload 661

Figure 7

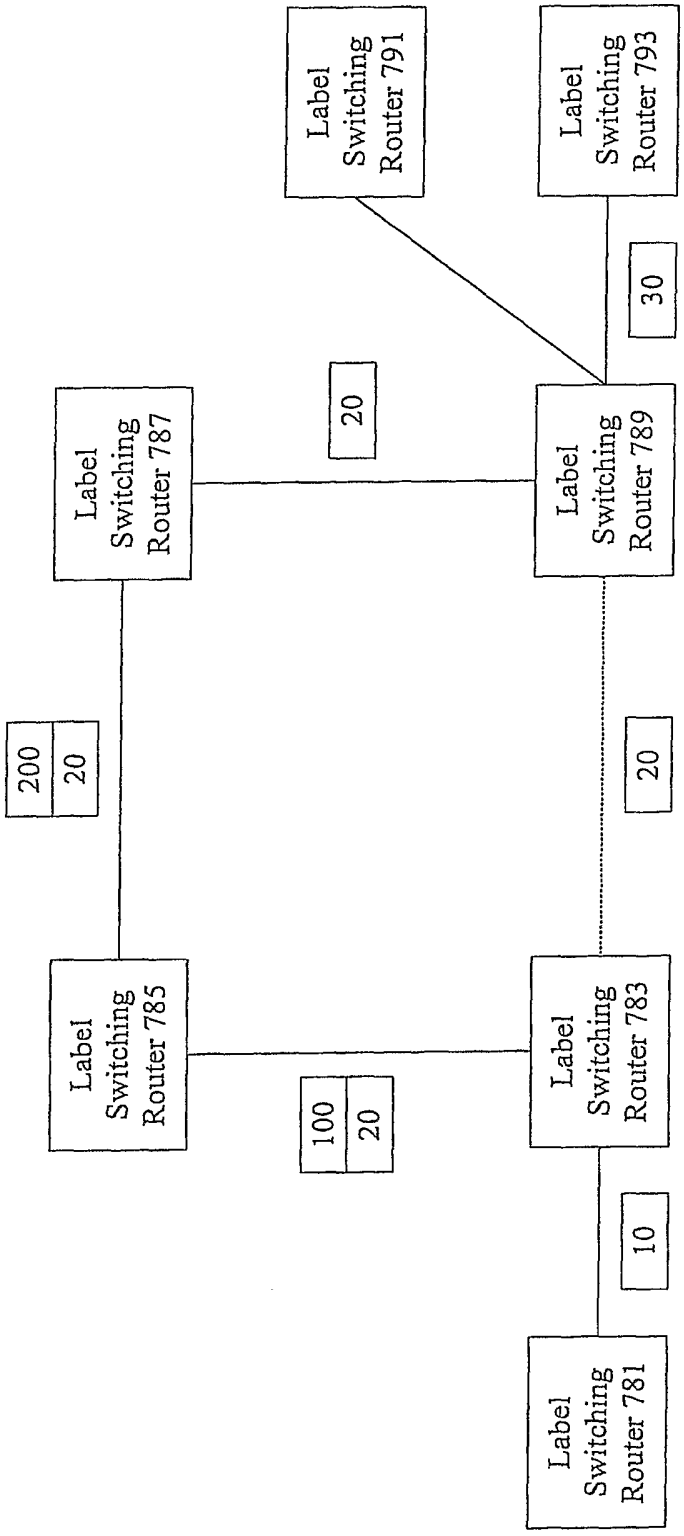
Label Switching Router 781
Label Switching Router 783
Label Switching Router 785
Label Switching Router 787
Label Switching Router 789
Label Switching Router 791
Label Switching Router 793
10
100
20
200
20
20
20
30

Figure 8

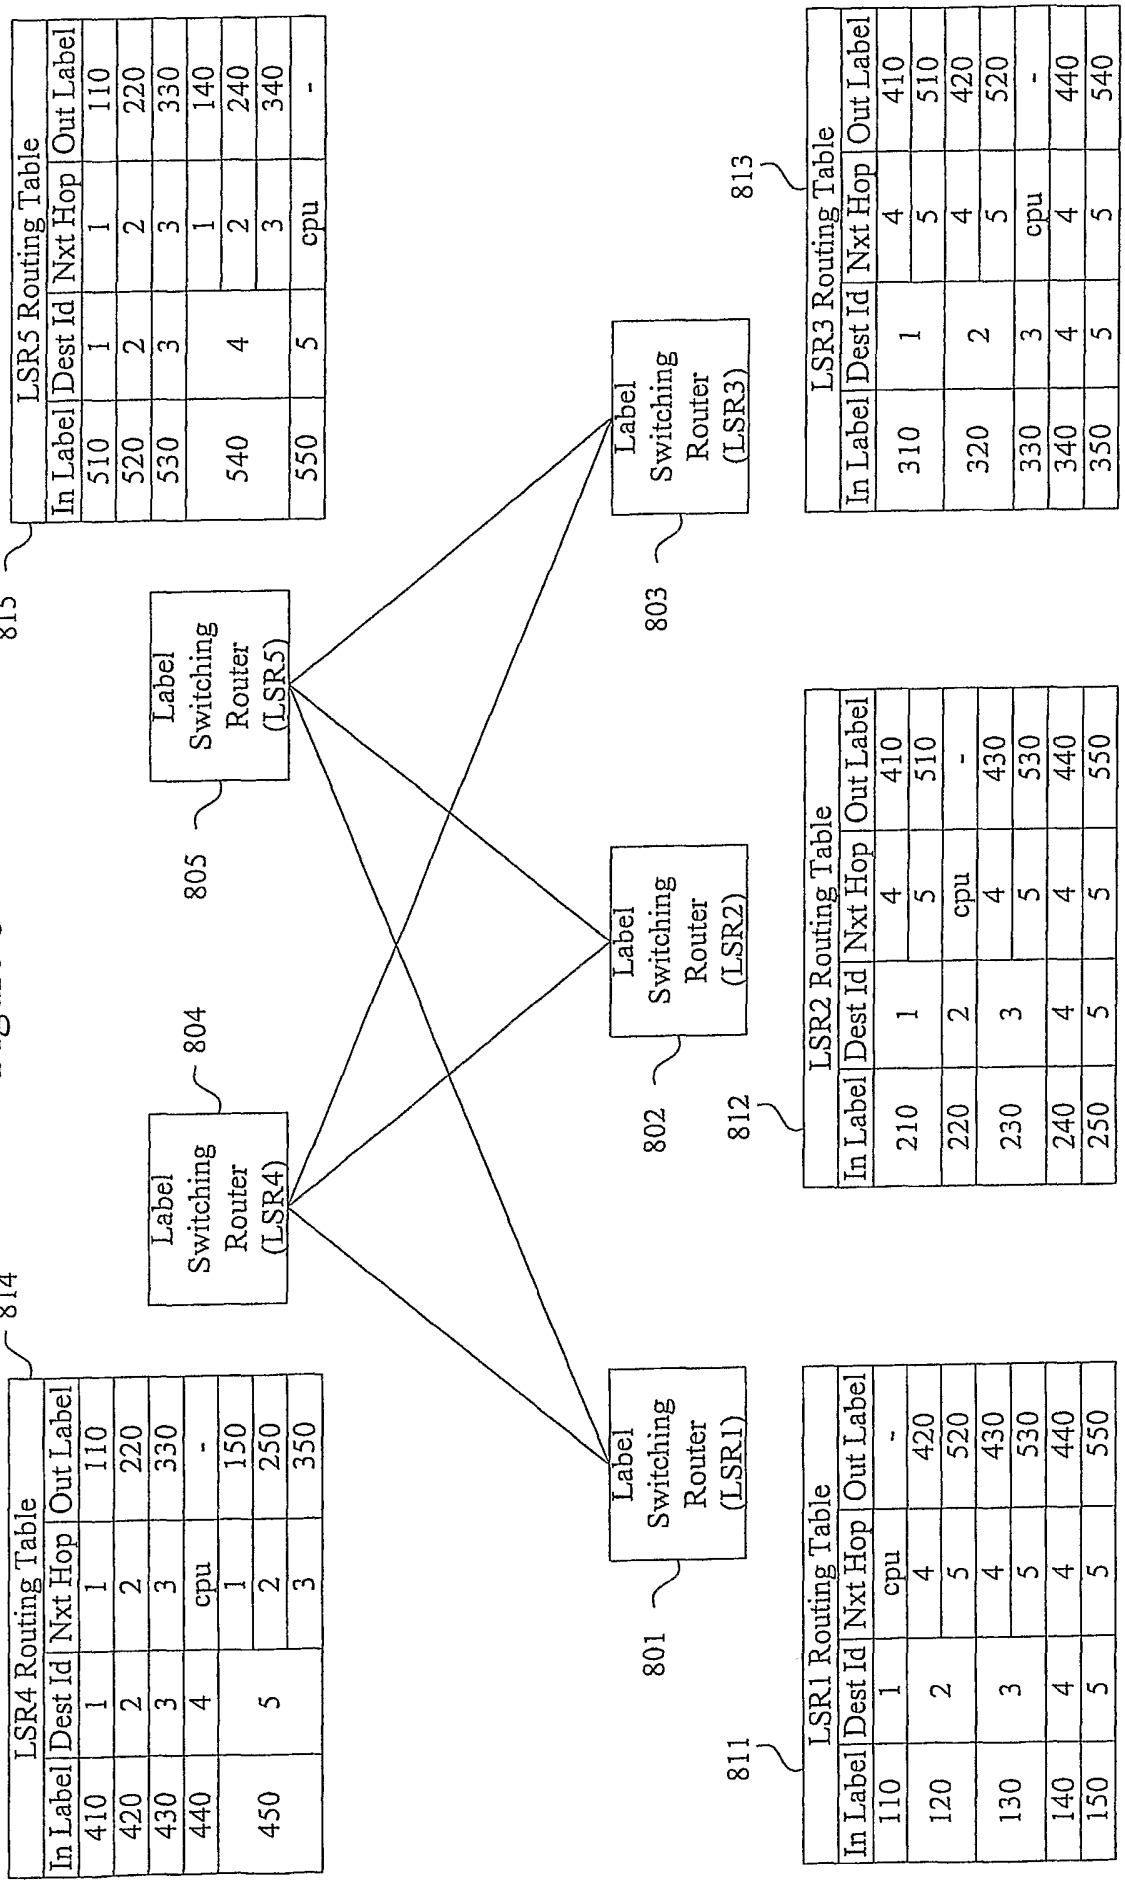

LSR5 Routing Table

In Label	Dest Id	Nxt Hop	Out Label
510	1	1	110
520	2	2	220
530	3	3	330
540	4	1	140
		2	240
		3	340
550	5	cpu	-

LSR3 Routing Table

In Label	Dest Id	Nxt Hop	Out Label
310	1	4	410
		5	510
320	2	4	420
		5	520
330	3	cpu	-
340	4	4	440
350	5	5	540

LSR2 Routing Table

In Label	Dest Id	Nxt Hop	Out Label
210	1	4	410
		5	510
220	2	cpu	-
230	3	4	430
		5	530
240	4	4	440
250	5	5	550

LSR4 Routing Table

In Label	Dest Id	Nxt Hop	Out Label
410	1	1	110
420	2	2	220
430	3	3	330
440	4	cpu	-
450	5	1	150
		2	250
		3	350

LSR1 Routing Table

In Label	Dest Id	Nxt Hop	Out Label
110	1	cpu	-
120	2	4	420
		5	520
130	3	4	430
		5	530
140	4	4	440
150	5	5	550

INTERNATIONAL SEARCH REPORT

Inte nal Application No
PCT/US 03/09442

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 H04L12/46

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)
IPC 7 H04L

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 practical, search terms used)
EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	EP 1 187 406 A (NORTEL NETWORKS LTD) 13 March 2002 (2002-03-13) abstract; claims 1,3,4,6,9,12,14,15,17,21,23; figures 1,6,9,10 ---	1,11,19
A	M. RAJAGOPAL, R. BHAGWAT, W. RICHARD: "IP and ARP over Fibre Channel" REQUEST FOR COMMENTS: 2625, 'Online! 30 June 1999 (1999-06-30), XP002246207 Retrieved from the Internet: <URL:www.ietf.org/rfc/rfc2625.txt> 'retrieved on 2003-07-02! Section 1 --- -/--	1,11,19

[X] Further documents are listed in the continuation of box C.

[X] Patent family members are listed in annex.

° Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"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)

"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

"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

"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

"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.

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
2 July 2003	15/07/2003

Name and mailing address of the ISA	Authorized officer
European Patent Office, P.B. 5818 Patentlaan 2 NL – 2280 HV Rijswijk Tel. (+31–70) 340–2040, Tx. 31 651 epo nl, Fax: (+31–70) 340–3016	Matos, N

INTERNATIONAL SEARCH REPORT

Inte ɔnal Application No
PCT/US 03/09442

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication,where appropriate, of the relevant passages	Relevant to claim No.
A	CHARLES MONIA: "iFCP - A protocol for Internet Fibre Channel Storage Networking" NISHAN SYSTEMS, 'Online! 12 December 2000 (2000-12-12), XP002246205 Retrieved from the Internet: <URL:www.ietf.org/proceedings/00dec/slides /ips-11.pdf> 'retrieved on 2003-07-02! the whole document ---	1,11,19
A	JAMES NARDI: "What's After Fibre Channel?" COMPUTERWORLD, 'Online! 15 October 2001 (2001-10-15), XP002246206 Retrieved from the Internet: <URL:www.computerworld.com/hardwaretopics/ storage/story/0,10801,64710,00.html> 'retrieved on 2003-07-02! the whole document -----	1,11,19

INTERNATIONAL SEARCH REPORT

information on patent family members

Inte onal Application No
PCT/US 03/09442

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 1187406 A	13-03-2002	CA 2351771 A1	28-12-2001
		EP 1187406 A1	13-03-2002