

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 December 2002 (12.12.2002)

PCT

(10) International Publication Number
WO 02/098821 A2

(51) International Patent Classification⁷: **C04Q**

(21) International Application Number: PCT/US02/17057

(22) International Filing Date: 31 May 2002 (31.05.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/295,151 1 June 2001 (01.06.2001) US

(71) Applicant: **FUJITSU NETWORK COMMUNICATIONS, INC.** [US/US]; 2801 Telecom Parkway, Richardson, TX 75082 (US).

(72) Inventors: **MEZEUL, Michael J.**,; 1420 West McDermott Drive, Apt. 1836, Allen, TX 75013 (US). **YASEEN, Nimer**,; 204 Fairfax Drive, Allen, TX 75013 (US). **DHILLON, Abinder S.**,; 7101 Chase Oak Boulevard, Apt. 414, Plano, TX 75025 (US). **STEVENS, Charlie**,; 508 Cumberland Drive, Hurst, TX 76054 (US).

(74) Agents: **HUBBARD, Marc A.**, et al.,; Munsch Hardt Kopf & Harr, P.C., 4000 Fountain Place, 1445 Ross Avenue, Dallas, TX 75202-2790 (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, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, 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, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that 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: METHOD OF ADDRESSING AND ROUTING DATA



(57) Abstract: A system and method for maximum utilization of bandwidth of a resource by incrementally varying in real time the allocated bandwidth per source queue per class of service based at least in part on the current status of a destination virtual output queue fill level is disclosed. By varying the permissible bandwidth incrementally not only is the available bandwidth utilized optimally but also a desired class of service is maintained.

WO 02/098821 A2

METHOD OF ADDRESSING AND ROUTING DATA

TECHNICAL FIELD OF THE INVENTION

This invention relates to telecommunication network and equipment, and more particularly, to a method of addressing and routing data.

5 BACKGROUND OF THE INVENTION

Multi-protocol label switching or MPLS is a framework specified by Internet Engineering Task Force (IETF) that provides efficient routing, forwarding and switching traffic flows through a telecommunication network. MPLS allows a mapping of Internet protocol (IP) addresses to fixed length labels used by routers and other equipment to route and switch packets. In MPLS, packets are transmitted
10 on label-switched paths (LSPs), which are specified by a sequence of labels at each node along the path from the source to the destination. LSPs are established prior to packet transmission. The labels are distributed by protocols such as border gateway protocol (BGP). Each data packet encapsulates and carries the labels between the source and the destination. The devices that operate using MPLS can be classified into label edge routers (LERs) and label switching routers (LSRs), which route and switch the
15 data packets at the edge and in the core of the network. Therefore, the MPLS label governs inter-node packet transmission.

SUMMARY OF THE INVENTION

Although MPLS labels are used to forward packets between nodes in a network, a need has arisen
20 for a means to route the packets inside each network node. Furthermore, there is a need to determine the source port of a packet inside each network node so that a reply packet may be sent back to the source port for transmission to the source of the packet.

In accordance with an embodiment of the present invention, a method of routing packets in a system includes the steps of receiving a plurality of packets, each packet including a routing label and a
25 packet type, routing each packet to a destination within the system specified in the routing label in response to the packet type being indicative of a data packet, and routing each packet to a processor within the system and sending a reply packet to a sender specified in the routing label in response to the packet type being indicative of a control packet.

In accordance with another embodiment of the present invention, a method of routing packets
30 internally within a telecommunication system includes the steps of receiving a packet including a routing label and a packet type, routing the packet to a destination within the system specified in the routing label in response to the packet type being indicative of a data packet, and routing the packet to a processor within the system and sending a reply packet to a sender specified in the routing label in response to the packet type being indicative of a control packet.

In accordance with yet another embodiment of the present invention, a method of routing packets internally within a network node includes the steps of receiving a packet including a routing label, routing the packet to a destination within the system specified by a shelf identifier, a slot identifier and a link identifier in the routing label in response to the packet being a data packet, and routing the packet to a processor within the system and sending a reply packet to a sender specified by a shelf identifier, a slot identifier and a link identifier in the routing label in response to the packet being a control packet.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, the objects and advantages thereof, reference is now made to the following descriptions taken in connection with the accompanying drawings in which:

FIGURE 1 is a block diagram of an exemplary system which may implement the addressing and routing scheme of the present invention;

FIGURE 2 is a diagram of an embodiment of a format of a routing label according to the teachings of the present invention;

FIGURE 3 is a diagram of an embodiment of a routing table according to the teachings of the present invention.

DETAILED DESCRIPTION OF THE DRAWINGS

The preferred embodiment of the present invention and its advantages are best understood by referring to FIGURES 1 through 3 of the drawings, like numerals being used for like and corresponding parts of the various drawings.

FIGURE 1 is a block diagram of an exemplary system 10, which may implement the addressing and routing scheme of the present invention. Like many telecommunications equipment, system 10 consists of a plurality of circuit cards plugged into a plurality of slots located on one or more shelves. System 10 includes a plurality of service cards such as Ethernet or SONET optical service cards 14 and 16 residing on one or more shelves 12, each with one or more components 18-24 such as network processors or other addressable circuits. Each component 18-24 includes one or more service ports 26 which are coupled to customer premises equipment (CPE) and a first set of ports 28 of a switch fabric 32. Service cards 14 and 16 may be interface cards such as 10/100 BaseT Ethernet, Gigabit Ethernet, OC-N service cards. Switch fabric 32 resides on a shelf 30 along with a data processing unit 34, which is also coupled to a port of the switch fabric. Data processing unit 34 possesses an address mapping table or label information base (in multi-protocol label switching terminology) 36 in memory or another form of data storage that may be accessed and updated. A second set of switch fabric ports 38 is coupled to a plurality of channel interface cards 44-50 residing on one or more shelves 40-42. Channel interface cards 44-50 are further coupled to the telecommunication network via interface ports 52 and 54. Although not

shown explicitly, other circuit cards performing other functionality may be disposed or coupled between switch fabric 32 and channel interface cards 44-50.

In operation, data coming into system 10 via service cards 14-16 are routed via switch fabric 32 to the proper port 38 for transport to the telecommunication network via interface cards 44-50. Switch fabric 32 also routes data messages to data processing unit 34 for processing, and routing reply messages therefrom back to the senders. The addressing and routing scheme of the present invention provides a way to identify the origin and destination of the packet. In general, data transmission occurs on label-switched paths (LSPs) pursuant to multi-protocol label switching (MPLS). LSPs are a sequence of labels at each and every node along the path from the source to the destination. The labels may be distributed using protocols such as label distribution protocol (LDP), border gateway protocol (BGP), or other suitable protocols.

FIGURE 2 is a diagram of an embodiment of a format of a routing label 60 according to the teachings of the present invention. Routing label 60, also called a shelf-interface-channel (SIC) label, includes a shelf field 62 and a slot field 63 that are used to contain a shelf number and a slot number assigned to the shelf and the slot, respectively, of the sender of the packet if the packet is destined for data processing unit 34. If the packet is not destined for data processing unit 34, then the shelf and slot fields contains the shelf and slot numbers of the shelf and slot of the destination circuit card. A link field 64 is used to contain the representation of the source physical link terminated on the slot if the packet is intended for the data processing unit 34, otherwise the link field contains the destination link. A channel field 65 is used to contain the channel number of a source channel multiplexed on the physical link if the packet is destined for data processing unit 34, otherwise the channel field contains the destination channel. Flow 66 is a field in the routing table to indicate the source virtual flow. A virtual flow is a logical representation or grouping of system components. A packet type field 67 is used to indicate whether the packet is a data packet, a traffic manager packet, a network processor packet, or a switch fabric interface (SFI) packet. In general, the packet type indicates the intended destination of the packet – a data packet is destined for an output port on system 10; a traffic manager packet is destined for a traffic manager (not explicitly shown) that resides in a circuit card specified by the shelf and slot fields; a network processor packet is destined for the network processor residing in a card specified by the shelf and slot fields; and a SFI packet is destined for data processing unit 34 coupled to a switch fabric port. A prototype field 68 is used to indicate whether the packet is an MPLS packet or a non-MPLS packet such as a packet for point-to-point protocol (PPP) negotiation. All fields are of an appropriate length in order to provide adequate information for each field.

FIGURE 3 is a diagram of an embodiment of a routing table or label information base 36 according to the teachings of the present invention. Label information base 36 includes a field 70 for indicating whether the entry is valid. A switch fabric (SF) port field 71 contains a specification of the port number on the switch fabric the packet originated. A stack field 72 of label information base 36

contains the operation to be performed. For example, if the stack bit in field 72 is a 0, label switch routing (LSR) function is performed; if the stack bit in field 72 is a 1, label edge routing (LER) function is performed. LSR and LER are functions performed by the appropriate routers either within the core of an MPLS network or at the edge of the access network, respectively, as known in the art. A U/M field 73 is used to indicate whether the packet is for unicast or multicast. Label information base 36 also includes an MPLS label field 74. As known in the art, an MPLS label field 74 is used to contain the MPLS label, which includes a next hop value that functions as an index to an entry in label information base 36. MPLS label field 74 also includes an EXP field, which carries the class of service (CoS) if border gateway protocol is used to distribute the MPLS label or drop precedence if interior gateway protocol (IGP) is used. MPLS label field 74 also includes a stack bit, which is valid in the interior gateway protocol label to decide which label information base entry to index. A time-to-live field is also in the MPLS label, which is used in the interior gateway protocol to be decremented when a packet is received on the line. Finally, label information base 36 also includes SIC label 60 as described above and shown in FIGURE 2.

In operation, SIC label 60 allows a packet to bear the shelf, slot, link and channel information of the source or the destination. If the packet type is SFI, the packet is destined for the control plane such as data processing unit 34. Because all packets destined for data processing unit 34 arrive from the same port on switch fabric 32 and could have originated from any other port of the switch fabric, the source of the packet has to be identified in order for a reply packet to be sent and routed. Therefore, SIC label 60 provides the shelf, slot, link and channel information that reflect the source of the packet, so that data processing unit 34 may send a reply packet back to the source for each received packet. If the packet type is not SFI, then the SIC label including the shelf, slot, link and channel numbers provides an identification of the destination of the packet. The SIC label, with the link and channel designations, provides a fine granularity of resolution so that a packet may be routed based on a specific port, wavelength, slot, and shelf of the destination. The SIC label is only valid within a piece of telecommunication equipment, such as a label switching router or a label edge router. In other words, the SIC label is only used for internal routing within each node located along the label switched paths.

The manner in which label information base is used is as in MPLS routing as known in the art. In general, a label stack mechanism is used to allow for hierarchical operation. Therefore, the MPLS label is used for routing between routers or some other equipment and the SIC label is used for routing within each router. Each level in a label stack pertains to the hierarchical level. Before any traffic begins the routers binds a label to a specific forward equivalence class (FEC) and builds the label information base tables. Downstream routers initiate the distribution of the labels and the label/FEC binding. On receipt of label bindings each LSR creates entries in the label information base. The contents of the label information base map the label to an FEC. The MPLS label in the label information base is used to find the next hop in the path, while the SIC label, according to the teachings of the present invention, is used

to find the destination circuit card, link and channel within the router. Further, a fine granularity of resolution of the destination is provided by the present invention that is heretofore not available.

While the invention has been particularly shown and described by the foregoing detailed description, it will be understood by those skilled in the art that various changes, alterations, 5 modifications, mutations and derivations in form and detail may be made without departing from the spirit and scope of the invention.

WHAT IS CLAIMED IS:

1. A method of routing packets in a system, comprising:

receiving a plurality of packets, each packet including a routing label and a packet type;

routing each packet to a destination within the system specified in the routing label in response to
5 the packet type being indicative of a data packet; and

routing each packet to a processor within the system and sending a reply packet to a sender
specified in the routing label in response to the packet type being indicative of a control packet.

2. The method, as set forth in claim 1, wherein routing each packet to a destination
10 comprises routing each packet to the destination specified by shelf and slot numbers of the destination
included in the routing label.

3. The method, as set forth in claim 1, wherein routing each packet to a destination
comprises routing each packet to the destination specified by a shelf identifier, a slot identifier, a link
15 identifier, and a channel identifier included in the routing label for transporting the packet.

4. The method, as set forth in claim 1, wherein sending the reply packet to a sender
comprises routing the reply packet to the sender specified by shelf and slot numbers of the sender
included in the routing label.

20 5. The method, as set forth in claim 1, wherein sending the reply packet to a sender
comprises routing the reply packet to the sender specified by a shelf identifier, a slot identifier, a link
identifier, and a channel identifier included in the routing label for transporting the reply packet.

25 6. The method, as set forth in claim 1, wherein routing each packet to a destination
comprises:

receiving each packet at a switch; and

switching each packet to the destination coupled to a predetermined port of the switch as
specified by shelf and slot numbers of the destination included in the routing label.

30 7. The method, as set forth in claim 1, wherein routing each packet to a destination
comprises:

receiving each packet at a switch; and

routing each packet to the destination coupled to a predetermined port of the switch as specified
35 by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label
for transporting the packet.

8. The method, as set forth in claim 1, wherein routing each packet to a processor and sending the reply packet to a sender comprises:

receiving each packet at a switch;

switching the packet to a predetermined port of the switch coupled to the processor; and

switching the reply packet to the sender coupled to a second predetermined port of the switch specified by shelf and slot numbers of the sender included in the routing label.

9. The method, as set forth in claim 1, wherein routing each packet to a processor and sending the reply packet to a sender comprises:

receiving each packet at a switch;

switching the packet to a predetermined port of the switch coupled to the processor; and

switching the reply packet to the sender coupled to a second predetermined port of the switch specified by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label for transporting the reply packet.

10. A method of routing packets internally within a telecommunication system, comprising: receiving a packet including a routing label and a packet type; routing the packet to a destination within the system specified in the routing label in response to the packet type being indicative of a data packet; and

routing the packet to a processor within the system and sending a reply packet to a sender specified in the routing label in response to the packet type being indicative of a control packet.

11. The method, as set forth in claim 10, further comprising pushing the routing label onto a label information table stack after receiving the packet.

12. The method, as set forth in claim 11, further comprising popping the routing label onto the label information table stack after receiving the packet at the destination within the system.

13. The method, as set forth in claim 11, further comprising popping the routing label from the label information table stack after receiving the packet at the processor within the system.

14. The method, as set forth in claim 10, wherein routing a packet to a destination comprises routing each packet to the destination specified by shelf and slot numbers of the destination included in the routing label.

15. The method, as set forth in claim 10, wherein routing a packet to a destination comprises routing each packet to the destination specified by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label for transporting the packet.

5 16. The method, as set forth in claim 10, wherein sending the reply packet to a sender comprises routing the reply packet to the sender specified by shelf and slot numbers of the sender included in the routing label.

10 17. The method, as set forth in claim 10, wherein sending the reply packet to a sender comprises routing the reply packet to the sender specified by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label for transporting the reply packet.

18. The method, as set forth in claim 10, wherein routing a packet to a destination comprises:
receiving the packet at a switch; and
15 switching the packet to the destination coupled to a predetermined port of the switch as specified by shelf and slot numbers of the destination included in the routing label.

19. The method, as set forth in claim 10, wherein routing a packet to a destination comprises:
receiving the packet at a switch; and
20 switching the packet to the destination coupled to a predetermined port of the switch as specified by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label for transporting the packet.

20. The method, as set forth in claim 10, wherein routing a packet to a processor and sending
25 the reply packet to a sender comprises:
receiving the packet at a switch;
switching the packet to a predetermined port of the switch coupled to the processor; and
switching the reply packet to the sender coupled to a second predetermined port of the switch
specified by shelf and slot identifiers of the sender included in the routing label.

30

21. The method, as set forth in claim 10, wherein routing a packet to a processor and sending the reply packet to a sender comprises:

receiving the packet at a switch;

switching the packet to a predetermined port of the switch coupled to the processor; and

5 switching the reply packet to the sender coupled to a second predetermined port of the switch specified by a shelf identifier, a slot identifier, a link identifier, and a channel identifier included in the routing label for transporting the reply packet.

22. A method of routing packets internally within a network node, comprising:

10 receiving a packet including a routing label;

routing the packet to a destination within the system specified by a shelf identifier, a slot identifier and a link identifier in the routing label in response to the packet being a data packet; and

15 routing the packet to a processor within the system and sending a reply packet to a sender specified by a shelf identifier, a slot identifier and a link identifier in the routing label in response to the packet being a control packet.

23. The method, as set forth in claim 22, further comprising pushing the routing label onto a label information table stack after receiving the packet.

20 24. The method, as set forth in claim 23, further comprising:

accessing the routing label in the label information table stack; and

popping the routing label from the label information table stack after receiving the packet at the destination within the system.

25 25. The method, as set forth in claim 23, further comprising popping the routing label from the label information table stack after receiving the packet at the processor within the system.

26. The method, as set forth in claim 22, wherein routing a packet to a destination further comprises routing each packet to the destination further specified by a channel identifier included in the
30 routing label for transporting the packet.

27. The method, as set forth in claim 22, wherein sending the reply packet to a sender comprises routing the reply packet to the sender specified by a channel identifier included in the routing label for transporting the reply packet.

35

28. The method, as set forth in claim 22, wherein routing a packet to a destination comprises:
receiving the packet at a switch; and
switching the packet to the destination coupled to a predetermined port of the switch as specified
by the shelf, slot and link identifiers of the destination included in the routing label.

5

29. The method, as set forth in claim 22, wherein routing a packet to a destination comprises:
receiving the packet at a switch; and
switching the packet to the destination coupled to a predetermined port of the switch as specified
by the shelf identifier, slot identifier, link identifier, and a channel identifier included in the routing label
for transporting the packet.

10

30. The method, as set forth in claim 22, wherein routing a packet to a processor and sending
the reply packet to a sender comprises:
receiving the packet at a switch;
switching the packet to a predetermined port of the switch coupled to the processor; and
switching the reply packet to the sender coupled to a second predetermined port of the switch
specified by the shelf, slot and link identifiers of the sender included in the routing label.

15

31. The method, as set forth in claim 22, wherein routing a packet to a processor and sending
the reply packet to a sender comprises:
receiving the packet at a switch;
switching the packet to a predetermined port of the switch coupled to the processor; and
switching the reply packet to the sender coupled to a second predetermined port of the switch
specified by the shelf identifier, slot identifier, link identifier, and a channel identifier included in the
routing label for transporting the reply packet.

20

25

1/2

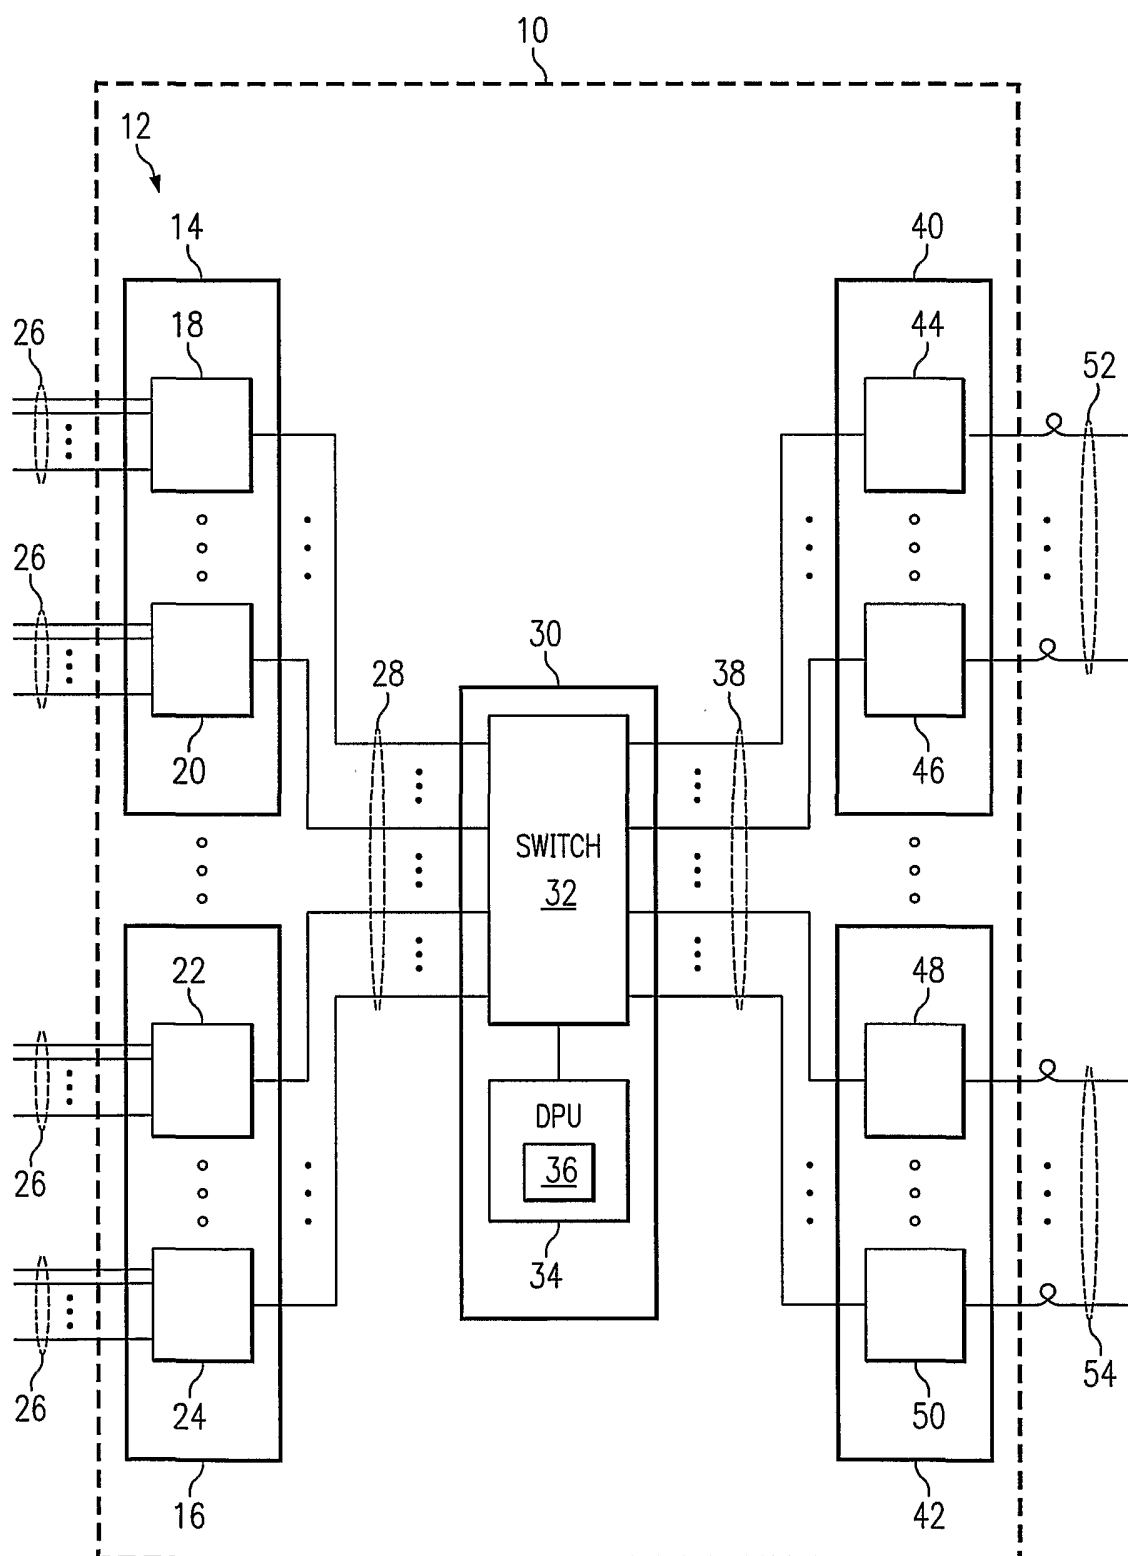


FIG. 1




FIG. 2




FIG. 3