



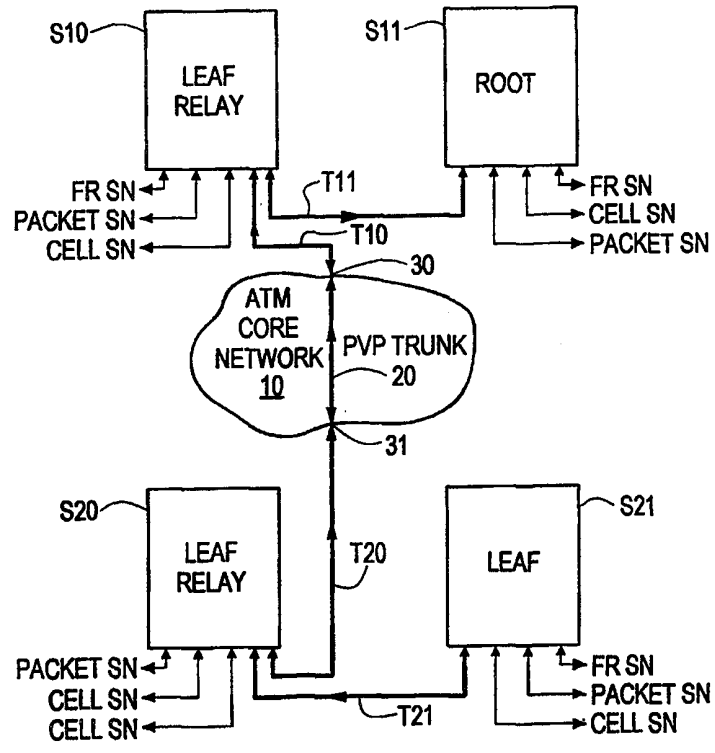
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : H04L 12/28, H04J 3/16, 3/17</p>	<p>A1</p>	<p>(11) International Publication Number: WO 99/14899 (43) International Publication Date: 25 March 1999 (25.03.99)</p>
<p>(21) International Application Number: PCT/US98/19303 (22) International Filing Date: 16 September 1998 (16.09.98) (30) Priority Data: 60/059,245 18 September 1997 (18.09.97) US 09/074,364 7 May 1998 (07.05.98) US (71) Applicant: ASCEND COMMUNICATIONS, INC. [US/US]; 1 Robbins Road, Westford, MA 01886 (US). (72) Inventors: PUNTAMBEKAR, Arvind, S.; 6 Jack Rabbit Lane, Westford, MA 01886 (US). SHELLY, Sushil; 1818 Farmington Drive, Shrewsbury, MA 01545 (US). (74) Agents: LEBOVICI, Victor, B. et al.; Weingarten, Schurgin, Gagnebin & Hayes LLP, Ten Post Office Square, Boston, MA 02109 (US).</p>	<p>(81) Designated States: AU, CA, JP, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: VIRTUAL PATH MERGING IN A MULTIPOINT-TO-POINT NETWORK TUNNELING PROTOCOL

(57) Abstract

A computer network includes frame or packet-based subnetworks (S10, S11, S20 and S21) connected by switches, the switches being interconnected by high capacity trunks (T10, T11, T20 and T21) using a connection based data transfer protocol similar to Asynchronous Transfer Mode (ATM). Some of the trunks include a Permanent Virtual Path (PVP) trunk crossing an ATM core network (10), the PVP trunk including one or more bidirectional PVPs. A multi-point-to-point (MPT) protocol is used among the switches to transfer packets as groups to cells directly from "leaf", or source, switches to "root", or destination, switches without required significant routing-related processing during cell transmission. The switches allocated virtual path identifiers in a conserving manner such that (i) MPT paths from multiple leaf switches are merged to one path with a single virtual path identifier terminating at a root switch; (ii) on the PVP trunks, a virtual path identifier already allocated for an outgoing connection is allocated to an incoming connection ahead of any virtual path identifiers that are completely unallocated; and (iii) a range of virtual path identifiers is pre-provisioned at the core network access points, so that a switch connected to an access point virtual path identifiers from the pre-provisioned range on behalf of upstream switches to extend MPTs across the core network.



FOR THE PURPOSES OF INFORMATION ONLY

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

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

-1-

TITLE OF THE INVENTION

**Virtual Path Merging in a Multipoint-to-Point
Network Tunneling Protocol**

5

CROSS REFERENCE TO RELATED APPLICATIONS

10

This application claims priority from provisional patent application Serial No. 60/059,245, filed September 18, 1997, entitled "Virtual Path Utilization in a Connection Based Network"; and patent application Serial No. 09/074,364, filed May 7, 1998, entitled "Virtual Path Merging in a Multipoint-to-Point Network Tunneling Protocol".

15

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR
DEVELOPMENT

--Not Applicable--

BACKGROUND OF THE INVENTION

20

The present invention is related to the field of computer networks, and more particularly to networks including Asynchronous Transfer Mode (ATM) switches employing a connection-based multipoint-to-point tunneling protocol to transfer connectionless data traffic, such as data traffic carried by Internet Protocol (IP) packets.

25

30

35

Many computer networks employ connectionless protocols for transferring data among nodes. In a connectionless protocol, data is transferred as a series of one or more datagrams, each transmitted along a network segment when no higher-priority datagrams are being transmitted on the segment. A well-known example of such a connectionless protocol is the Internet Protocol (IP). IP datagrams, or packets, are forwarded by devices known as routers that determine the network segment on which the packet is to be forwarded based on a destination address included in the packet, and then forward the packet over the respective network segment.

-2-

5 Connectionless protocols differ from connection-oriented protocols, in which data traffic is sent over pre-established connections between sending and receiving nodes, much like a telephone call. For example, datagrams may become lost or suffer extensive delay in an error-prone or congested network. One source of delay is the need to dynamically determine and implement the routing of the datagram. In a connection-oriented network, traffic is routed along a previously established and allocated route, and thus routing
10 in a connection-oriented network is generally simpler to implement and enables higher-speed packet forwarding.

15 Connection-oriented data protocols more closely resemble protocols used in standard audio telephony, and also better support streaming transmissions such as digital video. Thus as the need for transmission of voice, video, and data over a common network has increased, there has been a trend toward increasing use of connection-oriented protocols. Asynchronous Transfer Mode (ATM) is one example of a connection-oriented protocol that is receiving increased use.
20 In fact, the use of ATM switches within the core of the public data communications network has become more common, and thus the ATM protocol has become an important industry-standard protocol.

25 Connection-oriented networks like ATM networks employ switches rather than routers to route traffic. Connections through a switch are established prior to the beginning of data transmission; the connections correspond to a path through multiple switch elements linking source and destination nodes. Once a path is established, it remains
30 in place until specifically torn down, regardless of whether data traffic is being transmitted at any given time. Some connections may be long-lived; in fact, there is in some networks a notion of a "permanent" path that might be dedicated, for example, to carrying large volumes of traffic
35 between specified sites.

-3-

5 Connection-oriented networks must use some means for identifying connections over which data is to be forwarded. ATM employs a 2-tier switching technique that uses two separate connection identifiers. An ATM data cell includes an 8-bit Virtual Path Identifier (VPI) as well as a 16-bit Virtual Channel Identifier (VCI). This technique allows network elements such as switches to make switching decisions based on either the VPI or the VCI. Although other arrangements are possible, networks commonly employ "VPI switching", in which VCIs identify an individual connection between a source and a destination over a trunk-like path carrying numerous connections, and the VPIs are used to identify virtual paths within the network. Many virtual paths may be employed at a given physical port of a network element such as a switch.

10 In a large network that includes an ATM core network, the 8-bit limitation on VPI space (i.e., 256 paths) requires that the space be well managed to maximize the number of usable paths. For example, a large network could not afford to reserve some VPI bits for a separate signalling function, allocating only the remaining bits for virtual path identification, because such an allocation would result in too few allocatable virtual paths. Thus there is a general need to manage relatively small connection identifier spaces like the VPI space in a network using the ATM protocol.

15 Prior switches have used a connection-oriented protocol like ATM on inter-switch trunks that carry high-volume inter-switch data traffic. At one end of such a trunk, a switch concentrates data traffic originating at multiple input ports onto the trunk, while at the other end a switch demultiplexes the traffic for distribution to its output ports. The switches have had interfaces to cell-based subnetworks like ATM, and also interfaces to frame- or packet-based subnetworks, an example of a packet subnetwork being the Internet Protocol (IP) network mentioned above. Thus the switches have been designed to forward both connection-

-4-

oriented and connectionless data traffic over the connection-oriented inter-switch trunks.

5 The flow of connectionless data traffic in the connection-oriented inter-switch subnetwork is known as "tunneling". Cells arriving at an intermediate switch along an inter-switch path are simply switched from one incoming virtual path to a corresponding outgoing virtual path, the correspondence having been previously established by a separate connection-establishment procedure. This operation contrasts with operation in a set of interconnected routers, 10 in which each router is required to examine an address or other routing-related information in an arriving a packet, determine the next hop for the packet, and then transmit the packet over the next hop. Switches employing tunneling also retain router-like functionality, because such functionality 15 is needed to establish virtual paths at connection-establishment time. However, the routing function is normally bypassed at data-transmission time.

20 One particular tunneling technique is known as multipoint-to-point tunneling, or MPT. Each switch in an MPT environment is the "root" of a set of converging paths emanating from the other switches, known as "leafs". When a switch receives a datagram at a port, it determines which switch is connected to the subnetwork containing the destination node, and then forwards the datagram as a group 25 of cells on the MPT path for which the destination switch is the root. Intermediate switches simply switch the cells from an incoming port to an outgoing port in accordance with the previously-established MPT path definition. The destination switch accumulates the cells, reconstructs the datagram, 30 determines which output port the subnetwork containing the destination node is connected to, and sends the datagram over that port.

-5-

BRIEF SUMMARY OF THE INVENTION

In accordance with the present invention, a network of switches is disclosed in which the switches send connectionless data traffic such as IP packets between subnetworks via a connection-based network protocol, and do so in a fashion that makes efficient use of network connection identifiers.

In the disclosed network operating method, multipoint-to-point (MPT) traffic destined for a switch is merged to a single virtual path and thus uses only one VPI, even if the traffic may originate in any of a number of source switches. When a connection to a target switch over a trunk is being established, the root switch first determines whether a connection to any switches reachable via the trunk already exists. If not, the switch allocates a previously-unallocated VPI to be used with the new connection. The switch then signals the VPI to be used for the new connection to the target switch. The signalled VPI is the VPI of the existing connection, if any, and the allocated VPI otherwise. Thus where possible MPT data traffic from multiple leaf switches is merged to one virtual path, resulting in the allocation of fewer VPIs in the network than would otherwise occur.

In another technique used with the above-described merging technique, a range of VPIs is allocated at trunks connected to the ATM core network. A switch connected to the trunk allocates these VPIs only for MPT connections to itself or to upstream switches, that is, switches that can reach the core network only through the VPI-allocating switch. This technique ensures that upstream switches are able to establish MPTs on the core network despite the limited VPI space.

Other aspects, features, and advantages of the present invention are disclosed in the detailed description which follows.

-6-

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

Figure 1 is a block diagram of a network embodying the principles of the present invention;

5 Figure 2 is a schematic diagram of a data cell used in the network of Figure 1;

Figure 3 is a schematic diagram of a Virtual Channel Identifier (VCI) field that is part of the data cell of Figure 2;

10 Figure 4 is a block diagram of the network of Figure 1 showing relationships among the switches for a first multipoint-to-point tunnel (MPT) established therein;

Figures 5 through 8 are schematic diagrams depicting various data structures and their relationships in the switches of Figure 1 after the first MPT is established;

15 Figure 9 is a block diagram of the network of Figure 1 showing relationships among the switches for a second MPT established therein; and

20 Figures 10 through 13 are schematic diagrams depicting data structures and their relationships in the switches of Figure 1 after both the first and second MPTs have been established.

DETAILED DESCRIPTION OF THE INVENTION

25 The disclosure of provisional patent application Serial No. 60/059,245, filed September 18, 1997, entitled "Virtual Path Utilization in a Connection Based Network" is incorporated by reference herein.

30 Figure 1 shows a network in which four network switches S10, S11, S20, and S21 are interconnected by trunks T10, T11, T20, and T21. Each switch contains a switch processor SP and various interface cards. Each interface card is connected to a different subnetwork, and different card types are used for connections to different types of subnetworks. Switch S10, for example, includes two Frame cards and two Cell cards. One Frame card is attached to Frame Relay subnetwork (FR SN), the other to a packet-based subnetwork (PACKET SN)

35

-7-

such as an Internet Protocol (IP) subnetwork. One of the Cell cards is connected to a cell-based subnetwork (CELL SN) such as Asynchronous Transfer Mode (ATM). The other switches in Figure 1 are shown with similar configurations. Network nodes attached to subnetworks connected to a common switch communicate with each other via local connections made between interface cards in the switch; nodes attached to subnetworks connected to different switches communicate with each other via connections made in each switch creating a path between the subnetworks over the trunk or trunks interconnecting the end switches.

The connection between switches S10 and S11 over trunk T11 is a direct connection, as is the connection between switches S20 and S21 over trunk T21. The connection between switches S10 and S20, however, is through an ATM core network 10. This connection includes three sub-connections, shown in Figure 1 as trunk T10, trunk T20, and Permanent Virtual Path (PVP) Trunk 20. Trunk T10 connects switch S10 to one end of the PVP Trunk 20 at a first core network access point 30, and trunk T20 connects switch S20 to the other end of the PVP Trunk at a second core network access point 31. The access points 30, 31 are for example user ports of ATM switches used in the core network 10. The PVP Trunk 20 forms a connection through the ATM core network 10 between the two network access points 30 and 31, thus completing the connection between switches S10 and S20.

The PVP Trunk 20 is a set of one or more PVPs that have been administratively pre-provisioned by a network manager for MPT use by the switches S10, S11, S20, and S21. The PVPs are identified at each access point by a corresponding range of VPIS, as described in greater detail below. The size of the range of VPIS at either access point is the same; however, the starting point of the range in general is different. For example, the PVPs may be identified by VPIS 10-19 at access point 30, and by VPIS 20-29 at access point 31. The switches S10 and S11 are informed of the starting

-8-

and ending VPIs that bracket the range of VPIs that designate a PVP when the network is initialized, prior to the establishment of a connection through the core network 10.

5 The core network 10 performs the bookkeeping required to switch those cells arriving at access point 31, having one VPI, toward access point 30, and also to replace the VPI with the corresponding VPI identifying the connection at access point 30 when the cells exit the core network. The core network 10 generally contains numerous other connections among other pairs of access points (not shown in Figure 1).
10 Some of these may also be pre-provisioned PVP Trunks for use by other switches (not shown) operating in the manner disclosed herein; others however are created and terminated on demand in the usual fashion for an ATM network.

15 The connections to a Cell card carry network traffic in fixed-length data units known as "cells", such as ATM cells. The connections to a Frame card carry network traffic in variable-length data units called "frames". The term "frame" as used herein includes what is commonly known as a packet.
20 Thus the frame cards provide an interface to, among other things, Internet Protocol (IP) routers exchanging data and messages using IP packets.

Figure 2 shows the general structure of an ATM cell. It includes a 5-byte header and a 48-byte payload. The header includes the 8-bit VPI 32, the 16-bit VCI 33, and other header fields 34. The payload is shown as a series of data bytes 35. The VPI field 32 and the VCI field 33 together uniquely identify a connection. In the system of Figure 1, the switches S10, S11, S20 and S21 and the core network 10 employ the VPI switching, so that switching decisions are made based on the VPI field 32 and not the VCI field 33. The VCI field 33 conveys information between source and destination switches and thus is not disturbed by either the core network 10 or any switches acting as
30 intermediate or relay switches.
35

-9-

Figure 3 shows the structure of the VCI field 33 as used in the system of Figure 1. The VCI field 33 includes a 5-bit Forwarding Engine Identifier (FE ID) field 36 and an 11-bit Reassembly Identifier (RE ID) field 37. The use of these fields is described in greater detail below.

MPT Setup

A method by which MPTs are established is described below. The method includes functionality in a routing program used in the network of Figure 1 known as Open Shortest Path First (OSPF). OSPF is responsible for knowing whether and how network elements are connected together, including the switches and subnets shown in Figure 1. This functionality is largely conventional, and thus is not described herein. Extensions to OSPF needed to support the disclosed method are noted where appropriate.

Figure 4 shows a first MPT, hereinafter referred to as MPT(1), via which switch S10 receives data from the other switches S11, S20 and S21. MPT(1) includes a set of unidirectional virtual path on each of the trunks T11, T10, T20 and T21, as well as on the PVP Trunk 20. These paths are indicated by the arrowheads on the trunks in Figure 4. For MPT(1), switch S10 is a "root" switch, and the other three switches are "leaf" switches. Switch S20 is also referred to as a "relay" switch because it performs the function of relaying cells from leaf switch S21 toward root switch S10. MPT(1) is established by switch S10 using MPT signalling in a manner described below.

In the following description, the terms "upstream" and "downstream" are used to denote direction in an MPT as follows: "upstream" means "toward the root switch", and "downstream" means "away from the root switch". Also, the terms "ingress" and "egress" are used to refer to the Cell cards at either end of a trunk with respect to a particular MPT. "Egress" means the Cell card that sends data cells upstream on the trunk, and "ingress" means the Cell card that

-10-

receives data cells sent by an egress Cell card. Note that this definition refers only to the direction of DATA flow for a particular MPT. The direction of signalling is irrelevant, as is the fact that a Cell card of one type with respect to
 5 a given MPT can also be a Cell card of the other type with respect to another MPT.

The switches of Figure 1 use a data structure known as the VC Entry Data Structure in conjunction with the method disclosed herein. This data structure is a large collection
 10 of virtual channel entries (VC Entries). Each VC Entry includes one or more fields that identifies the entry as a certain type, depending on the functions for which the VC Entry is used. The types are introduced and described in more detail below. In the illustrated embodiment, there are
 15 2048 VC Entries allocated for use by MPT. Other VC Entries not discussed herein may be used for other purposes in the switches of Figure 1.

The following types of VC Entries are used:

Ref	Name	Location
R	Root	SP of root switch
D	Default Conn.	Cell & Frame cards of root switch
V	VP Termination	Trunk port of root switch
L	Leaf	SP of leaf switch
P	Parent	Upstream trunk port on leaf switch
C	Child	Downstream trunk port on leaf switch
F	FE Array	SP of leaf switch
RI	Reassembly Identifier	Frame card of root switch

In addition to the VC Entry Data Structure, another data
 30 structure called "INCircuit" is used. This structure is shown in the Figures as IN; it is used on the Frame cards of leaf switches. Each INCircuit has an array of 32 connection

-11-

identifiers associated with it, used to map a VPI/VCI pair to an internal path through the switching fabric within a switch.

5 The method by which switch S10 establishes MPT(1) is now described in conjunction with Figures 4 through 8. It is assumed that switch S10 connects to switches S11, S20, and S21 in that order. First, switch S10 determines whether there is an existing MPT to switch S11. The MPT system keeps track of MPTs as they are created, and thus is aware at any
10 given time whether an MPT to a given switch exists. In this case, it is assumed that no prior MPT exists. Thus the SP in switch S10 allocates a Root VC Entry R, and then allocates an 11-bit Reassembly Identifier (RE ID). The Root VC Entry R signifies the termination point for cells being transmitted
15 on the MPT being created. The RE ID is stored by each Frame card in switch S10 for later use, and is also included in a Call message created by switch S10 to be addressed and sent to switch S11.

20 The ingress Cell card 50 on switch S10 allocates a first VP Termination VC Entry (V) 52, which has associated therewith a VPI reserved for use by MPT. This VPI is to be included in the VPI field 32 of all data cells sent to switch S10 by switch S11 to identify MPT(1). The ingress Cell card 50 inserts this VPI into the Call message as a data field and
25 sends the message to switch S11 over trunk T11. The Call message is sent as a group of cells each using a pre-established value in the VPI field 32 that is reserved for signalling.

30 The egress Cell card 60 on switch S11 recognizes the signalling VPI and re-assembles the message. Recognizing the message as the first Call received at the port for trunk T11, the egress Cell card creates a Parent VC Entry (P) 52 which is used to handle subsequent MPT data traffic. The Cell card also determines that the target of the Call message is switch
35 S11, and thus passes the Call on to the SP. The SP allocates a Leaf VC Entry (L) 64 having an FE Array VC Entry (F) 66

-12-

associated with it. The FE Array VC Entry 66 identifies "forwarding engines" (FEs) residing on the leaf switch (not shown in the Figures). An FE is an independent controller on a frame card that is responsible for the frame-to-MPT interface. On an egress switch, an FE handles frame-to-cell conversion and initiating the sending of cells on an MPT; on an ingress switch, an FE receives cells from an MPT, and handles cell-to-frame conversion and delivery of frames to the correct subnetwork. In one embodiment, there may be up to sixteen Frame cards in a switch, and up to two FEs on a Frame card. Thus the 5 bits in the FE ID field 36 uniquely identify one of a possible 32 FEs at the destination switch.

At the same time that the Leaf VC Entry 64 is allocated, the routing program OSPF is informed that the leaf switch is being added as a leaf of an MPT. OSPF stores the information associating destination routing addresses with the MPT, for use in a manner described below.

Having established itself as a leaf switch on MPT(1), switch S11 returns a CONFIRM message to switch S10 indicating that the MPT(1) connection has been established, and including a bit map indicating which FEs exist at the leaf switch S11. The root switch S10 responds by issuing a Call message to each FE on switch S11, each one including a different RE ID allocated by the SP in S10. The Parent VC Entry 62 on switch S11 forwards the Calls to the Frame cards, each of which responds by allocating an InCircuit structure IN in which the RE ID accompanying the Call is stored, and then returning a CONFIRM message. Once the root switch S10 has connected to each FE on switch S11, switch S11 is fully established as a leaf switch on MPT(1).

Having thus connected to leaf switch S11, root switch S10 proceeds to connect to switch S20. This process differs slightly from the above-described process for connecting to switch S11, because switches S10 and S20 are connected via the ATM core network 10. Again in this case there is no existing MPT to switch S10. The SP is aware through

-13-

configuration information that trunk T10 is a PVP trunk. Thus the VP Termination VC Entry (V) 54 allocated for the connection to switch S20 is one associated with one of the pre-provisioned VPIs identifying a PVP on the PVP trunk 20 at access point 30. The Call message sent by switch S10 signals this VPI by sending an index value having a known relationship to the allocated VPI. This index value is referred to as the VPCI, for Virtual Path Connection Index. The VPCI is used because as described above switch S20 in general uses a different VPI to identify a PVP than does switch S10. What is important is that both switches understand which particular PVP within the PVP trunk 20 is allocated for MPT(1); the use of the VPCI enables such understanding. One straightforward technique for arriving at the VPCI is calculating the difference between the allocated VPI and the starting VPI in the range pre-provisioned for the PVP Trunk at access point 30.

Upon receiving the Call message, switch S20 (Fig. 7) determines the VPI to be used for MPT(1) from the received VPCI. Switch S20 knows to do this because it is aware through configuration information that trunk T20 is a PVP trunk. If the VPI signalling technique is as described above wherein the VPCI is a difference value, switch S20 adds the received VPCI to the starting VPI in the range pre-provisioned for the PVP Trunk at access point 31 to obtain the VPI to be allocated.

An example is presented assuming that VPIs 10-19 have been pre-provisioned at access point 30, and VPIs 20-29 have been pre-provisioned at access point 31. Assuming that switch S10 allocates VPI 10, it therefore sends a VPCI of 0 to switch S20, indicating that switch S20 should allocate the 0th VPI after the starting VPI (i.e., the starting VPI) in the range pre-provisioned at access point 31. Switch S20 thus allocates VPI 20 to MPT(1).

Once the two switches S10 and S20 know which PVP to use on the PVP trunk 20, the remainder of the processing for

-14-

establishing S20 and its FEs as leafs on MPT(1) is the same as discussed above for switch S11. In the case of switch S20, only one InCircuit structure (IN) 70 is created, because the switch has only one Frame card.

5 Switch S10 then proceeds to add switch S21 to MPT(1). In this case, MPT(1) already exists to switch S20. Therefore, a new MPT is not created; rather, the existing one is extended as follows: Switch S10 allocates a new RE ID and includes it in a Call message addressed to switch S21
10 including the same VPCI as used when calling S20. The egress Cell card 72 on switch S20 realizes that the Call is to be forwarded over trunk T21. The SP on switch S20 determines whether an MPT to switch S21 exists. In this case, no MPT exists yet, so the ingress Cell card 72 on switch S20
15 allocates a Child VC Entry (C) 74 associated with the existing Parent VC Entry (P) 76 and also with a VPI to be used on trunk T21. The ingress Cell card 72 modifies the Call message to signal the VPI associated with the Child VC Entry 72, and then forwards the Call message to switch S21
20 using the signalling VPI. In this manner switch S20 acts as an intermediate or relay switch between switches S10 and S21 for MPT(1).

From this point the signalling between switch S10 and switch S21 is the same as that between switch S10 and the
25 other two switches S11 and S20, with the Parent-Child connection in the Cell card 72 on switch S20 providing the necessary bridge between trunks T21 and T20, and with no further allocation of VPIs by the Cell card 72 in switch S20. The VPI used for MPT(1) on trunk T21 may be different from
30 the VPI used on trunk T20; thus the ingress Cell card 72 on switch S20 is responsible for maintaining the necessary mapping between these two VPIs.

Figure 9 illustrates a second MPT (MPT(2)), established by switch S11 as the root switch. The processing by S11 to
35 establish MPT(2) is like that described above for switch S10. In this case, switch S11 first establishes switch S10 as a

-15-

leaf first, then establishes switch S20 via switch S10, and finally establishes switch S21 via switches S10 and S20. Unlike the process used by switch S10, however, S11 is not constrained to use pre-provisioned PVPs or the VPCI signalling technique described above, because its single trunk connection T11 is direct rather than through the core network 10.

During the establishment of MPT(2), switch S10 plays the role of leaf and relay switch rather than root. Switch S10 behaves slightly differently as a relay than does switch S20, because in this case the downstream trunk for switch S10 is a PVP trunk, whereas for switch S20 the MPT(1) downstream trunk is the direct trunk T21. When the Child VC Entry (C) 100 is created on switch S10, the VPI allocated is one in the range of VPIs pre-provisioned for use by switch S10 on PVP Trunk 20, and a corresponding VPCI is calculated and signalled to switch S20 in the same manner described above for MPT(1). This need to enable a switch connected to the core network 10 to act as a relay for upstream switches is in fact the purpose for pre-provisioning a range of VPIs at an access point of the core network 10; it guarantees that at least some pre-determined number of VPIs will be available for use by switches upstream of switch S10 to establish their MPTs.

Figures 10 through 13 respectively show the results at each switch S11, S10, S20, and S21 after MPT(2) is established. It can be seen that switch S11 as root acquires a VP Termination VC Entry (V) 110, default connections (D) 112, and Reassembly Identifier VC Entries (RI) 114. Switch S10 acquires a Parent VC Entry (P) 102 at trunk T11, the Child VC Entry 100 at trunk T10, a Leaf VC Entry (L) 104, and two InCircuit structures (IN) 106. Switches S20 and S21 acquire another set of VC Entries exactly like those for MPT(1).

In the four-switch system shown in Figures 1, 4, and 9, two additional MPTs would also be established, one each for

-16-

switches S20 and S21. These have been omitted for simplicity of description. These additional MPTs would be established in the same manner as described above for switches S10 and S11.

5 The MPTs existing on the direct trunks T11 and T21 are unidirectional connections, that is, each VPI used on either trunk is associated with a single MPT, and therefore with data flow in only one direction on the trunk. This features simplifies VPI allocation at direct trunks; VPIs are simply
10 drawn from a pool at either end of a trunk without regard to VPI allocation at the other end. In the ATM core network 10, however, each PVP is bidirectional, that is, the VPI used for the PVP from access point 30 to access point 31 is also used for the PVP from access point 31 to access point 30. Thus
15 switches S10 and S20 allocate VPIs on the PVP trunk 20 differently than on the direct trunks T11 and T21. Specifically, the switches S10 and S20 allocate a VPI that is already allocated to an MPT in the opposite direction, if any exist, before allocating any unused VPIs. While this
20 complicates VPI allocation somewhat, it helps preserve VPI space in the ATM core network 10. To accomplish the allocation, the VC Entries (V, C) used at a trunk ingress store information indicating whether the trunk is a PVP trunk, and also whether the associated VPI is being used for
25 an MPT in the egress direction.

 An example of the above is given assuming that switch S20 establishes an MPT(3) after MPT(1) and MPT(2) have been established in the preceding manner. When switch S20
30 allocates a VPI at trunk T20, it allocates VPI 0 again, because VPI 0 was already allocated in the opposite direction. If switch S20 were not constrained to use such VPIs, it might have allocated, for example, VPI 5 to MPT(3) (assuming that VPIs are allocated sequentially and that VPIs 1-4 were already in use). Thus in such a case VPI 5 has been
35 spared from non-essential use, leaving it available for allocation to another connection.

-17-

MPT Data Flow

Data cell transmission on an MPT involves a leaf switch as the source, a root switch as the destination, and possibly one or more intermediate or relay switches. The processing at these switches is described in turn.

At a leaf switch, one of the FEs on a frame card receives a frame at a port. The FE consults OSPF to determine which MPT to use to forward the frame. OSPF maintains a mapping from network addresses to InCircuit structures IN, which are provided to OSPF as they are created during the MPT setup processes described above. OSPF also maintains the network addresses of all the FEs in the system as assigned by a network manager. Thus the leaf switch determines the following from the destination address in the frame: (1) the MPT to send the data cells on to reach the destination (root) switch, and (2) the connection ID to be included with the data cells in order to reach the correct FE and Reassembly Identifier VC Entry RI on the destination switch.

The frame is then segmented into data cells, each one including in its header the VPI allocated to the MPT to be used, and a VCI including the 11-bit connection ID and the 5-bit FE identifier. These cells are then passed to the Parent VC Entry P associated with the VPI of the MPT, and sent to the next switch over the trunk at which the Parent VC Entry exists.

The switch at the other end of the trunk receives the incoming data cells, and determines its next action based on the type of VC Entry associated with the incoming VPI. If the VC Entry is a Child VC Entry C, the switch is a relay switch, and it thus switches the data cells to the outgoing trunk at the corresponding Parent VC Entry P, replacing the VPI with the VPI associated with the Parent VC Entry P. If the VC Entry associated with the incoming VPI is a VP Termination VC Entry V, the switch is a root switch. A root switch examines the FE ID field 36 of the VCI 33 to determine

-18-

which Default connection D to forward the cells over. The destination FE uses the RE ID field 37 of the VCI 33 to index into a reassembly table on the frame card that points to the appropriate Reassembly Identifier VC Entry RI to reassemble the cells on. The FE accumulates data cells on this Reassembly Identifier VC Entry RI until a complete frame has been received, at which time the forwarding logic on the frame card forwards the frame to one of the card's ports as indicated by the destination address.

An MPT tunneling protocol has been described using virtual path merging and other techniques to conserve VPI space in an ATM core network. The techniques used are applicable more generally to the management of connection identifier space in connection-oriented networks. Additionally, several of the specific features of the illustrated embodiment may be achieved by other means without departing from the present invention. For example, the trunks connected to a given switch may be connected to different Cell cards, rather than to the same card as shown. The MPTs need not be established in the order given, nor is it necessary for an intermediate switch to have been established as a leaf before a switch downstream of the intermediate switch becomes a leaf.

It will be apparent to those skilled in the art that modifications to and variations of the above-described methods and apparatus are possible without departing from the inventive concepts disclosed herein. Accordingly, the invention should be viewed as limited solely by the scope and spirit of the appended claims.

-19-

CLAIMS

We claim:

- 5 1. In a switch having a trunk connected thereto via which other switches are reachable, a method of establishing a new connection to a target one of the other switches to receive data cells therefrom, the method comprising the steps of:
- 10 determining whether a connection to any of the other switches over which data cells are to be received already exists on the trunk;
- if such a connection is determined to exist, sending a connection-establishment message over the trunk, the message being addressed to the target switch and signalling a virtual path identifier previously allocated for use with the
- 15 existing connection; and
- if no such connection is determined to exist, performing the steps of:
- allocating an unallocated virtual path identifier
- 20 to be used with the new connection; and
- sending a connection-establishment message over the trunk, the message being addressed to the target switch and signalling the allocated virtual path identifier.
- 25 2. A method according to claim 1, further comprising the step of informing a network routing program of the existence of the new connection upon establishment thereof.
- 30 3. A method according to claim 1, further comprising the step of including an index value in a signalling field in the connection-establishment message, the index value signalling the virtual path identifier to be used with the new connection by identifying the position of the virtual path identifier being signalled in a predetermined ordered set of
- 35 virtual path identifiers.

-20-

4. A method according to claim 3, further comprising the step of determining the difference between the value of the virtual path identifier to be signalled and a base member of the set of virtual path identifiers, and wherein the virtual path identifier is signalled by setting the index value equal to the determined difference.

5. A method according to claim 1, wherein (i) the switch is an intermediate switch, the trunk is a downstream trunk, and the reachable switches are downstream switches, (ii) the intermediate switch has an upstream trunk connected thereto via which the intermediate switch can be reached by an upstream switch, (iii) the intermediate switch has a pre-provisioned set of virtual path identifiers exclusively available for use with connections on the downstream trunk between the downstream switches and either the intermediate switch or the upstream switch, and (iv) the virtual path identifier signalled to the target downstream switch is selected from the pre-provisioned set of virtual path identifiers and upon allocation unambiguously identifies the intermediate switch as the recipient of data cells received over the downstream trunk, and further comprising the steps of:

receiving a connection-establishment message from the upstream switch over the upstream trunk;

determining whether one of the downstream switches is the target of the received connection-establishment message; and

if one of the downstream switches is determined to be the target of the received connection-establishment message, performing the steps of:

determining whether a connection to any of the downstream switches over which the switch is to receive data cells destined for the upstream switch already exists on the downstream trunk, such a connection being identified by a corresponding one of the pre-provisioned

-21-

virtual path identifiers other than the virtual path identifier signalled in claim 1;

if such a connection is determined to exist, sending a connection-establishment message over the downstream trunk, the message being addressed to the target switch and signalling a virtual path identifier previously allocated for use with the existing connection; and

if no such connection is determined to exist, performing the steps of:

allocating one of the pre-provisioned virtual path identifiers other than the virtual path identifier signalled in claim 1 for use with a new connection over which the intermediate switch is to receive data cells from the downstream switches destined for the upstream switch; and

sending a connection-establishment message over the trunk, the message being addressed to the target switch and signalling the allocated virtual path identifier.

6. A method according to claim 5, wherein each virtual path identifier in the pre-provisioned set of virtual path identifiers identifies a bidirectional connection on the downstream trunk, such that each virtual path identifier in the set is available for allocation to both an incoming connection over which the intermediate switch receives data cells from downstream switches and to an outgoing connection over which the intermediate switch sends data cells to downstream switches, and wherein the allocating step comprises the step of allocating a virtual path identifier previously allocated to an outgoing connection, if any, before allocating a virtual path identifier not previously allocated to an outgoing connection.

-22-

7. A method according to claim 1, wherein the signalled virtual path identifier is a selected one from a set of virtual path identifiers each of which identifies a bidirectional connection on the trunk, such that each virtual path identifier in the set is available for allocation to both an incoming connection over which the switch receives data cells from the other switches as well as to an outgoing connection over which the switch sends data cells to the other switches, and wherein the allocating step comprises the step of allocating a virtual path identifier previously allocated to an outgoing connection, if any, before allocating a virtual path identifier not previously allocated to an outgoing connection.

8. In a switch having a trunk connected thereto via which other switches are reachable, the switch including a set of virtual path identifiers each of which identifies a bidirectional connection on the trunk, such that each virtual path identifier in the set is available for allocation to both an incoming connection over which the switch receives data cells from the other switches as well as to an outgoing connection over which the switch sends data cells to the other switches, a method of establishing a new connection to a target one of the other switches to receive data cells therefrom, the method comprising the steps of:

determining whether a virtual path identifier allocated to an outgoing connection but not allocated to an incoming connection exists;

if such a virtual path identifier is determined to exist, then allocating the existing such virtual path identifier for use with the new connection;

If such a virtual path identifier is determined not to exist, then allocating a virtual path identifier not allocated to either an incoming or outgoing connection for use with the new connection; and

-23-

sending a connection-establishment message over the trunk, the message being addressed to the target switch and signalling the virtual path identifier allocated for use with the new connection.

5

9. A method according to claim 8, further comprising the step of determining, prior to allocating a virtual path identifier, whether a connection to any of the other switches over which data cells are to be received therefrom already exists on the downstream trunk, and wherein the allocating steps are performed only if such a connection is determined not to already exist.

10




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6

S20




FIG. 7

S21




FIG. 8

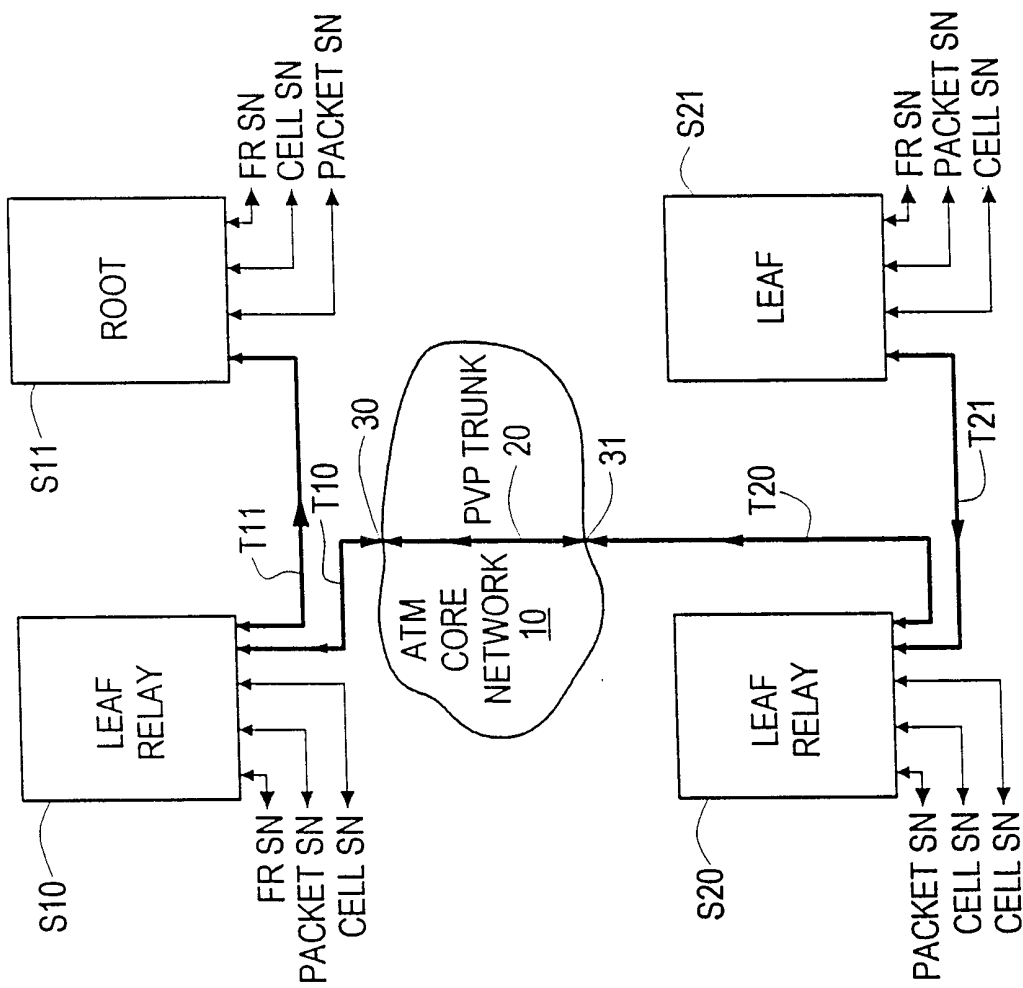


FIG. 9




FIG. 10




FIG. 11




FIG. 12

S21




FIG. 13

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/19303

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04L 12/28; H04J 3/16, 3/17
US CL :370/392, 395, 397, 398, 399, 401, 409, 433, 437

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/392, 395, 397, 398, 399, 401, 409, 433, 437

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

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

ATM, trunk, tunneling, virtual path identifier, VPI, multipoint to point, new connection, set up

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,479,401 A (BITZ et al) 26 December 1995, col. 5, lines 33-68)	1-9
Y	US 5,570,361 A (NORIZUKI et al) 29 October 1996, col 11, line 1-36)	1-9
X	US 5,583,868 A (RASHID et al) 10 December 1996, col. 5, lines 64- col. 6, lines 1-40)	1-9
X,P	US 5,717,689 A (AYANOGLU) 10 Febuary 1998, col. 18, lines 6-col, 19, lines 1-46)	1-9

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	*T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X* document of particular relevance: the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*&* document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means	
P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search
19 NOVEMBER 1998

Date of mailing of the international search report
26 JAN 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Authorized officer

ANDREW LEE

Joni Hill

Facsimile No. (703) 305-3230

Telephone No. (703) 305-4700