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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
8 August 2002 (08.08.2002)

PCT

(10) International Publication Number
WO 02/062076 A2

(51) International Patent Classification[7]: H04Q

(21) International Application Number: PCT/IL02/00084

(22) International Filing Date: 30 January 2002 (30.01.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/265,535 31 January 2001 (31.01.2001) US

(71) Applicant *(for all designated States except US)*: **TULIP NETWORKS LTD.** [IL/IL]; P.O.Box 12571, 46733 Herzliya (IL).

(72) Inventor; and

(75) Inventor/Applicant *(for US only)*: **VYHNAK, Petr** [IL/IL]; 8 D'Israeli, 34334 Haifa (IL).

(74) Agent: **MILLER - SIERADZKI ADVOCATES & PATENT ATTORNEYS**; P.O.Box 6145, 31061 Haifa (IL).

(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, US, 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

[Continued on next page]

(54) Title: DISTRIBUTED MULTICASTING FOR AN ATM NETWORK

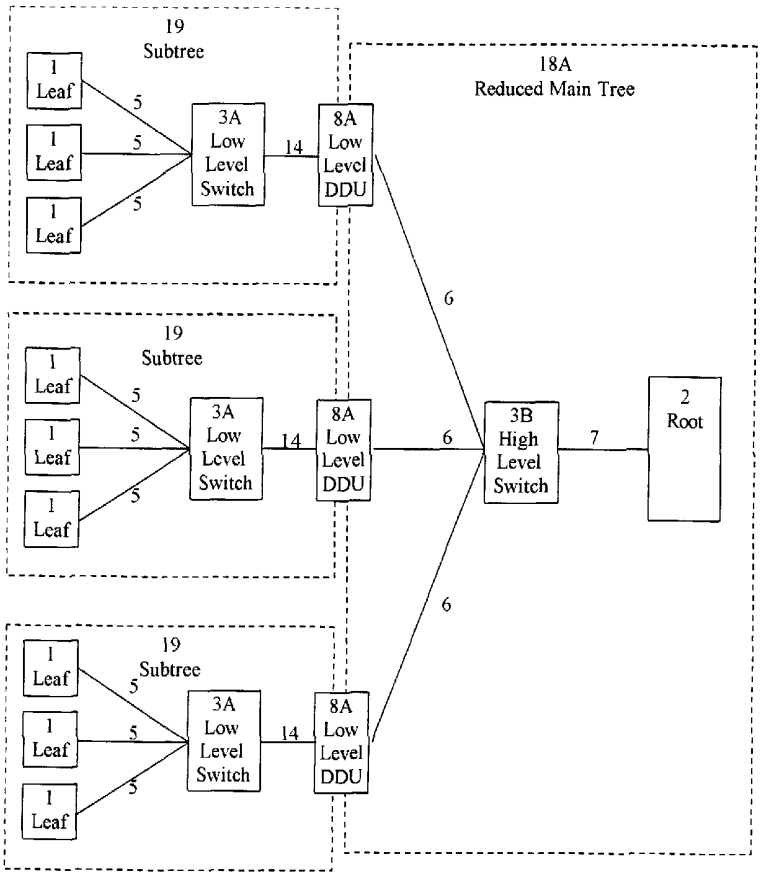

(57) Abstract: An apparatus (8) for distributed multicasting over an ATM network; the apparatus paired with a switch and comprising: a root communications interface (11); a switch communications interface (10); a substitution unit (20) that will substitute cell addresses; and a control unit (21) comprising computing means that will control the operation of the apparatus and maintain a table of cell addresses and control values. In response to a request from an endpoint (1A) to join a multicast that no other endpoint connected to the apparatus (8) is a member of, the control unit (21) joins the multicast tree as a virtual leaf and sets up a multicast subtree, acting as a virtual root. As the root (2) transmits cells over the multicast tree, the substitution unit (20) substitutes the addresses of cells from the root with the corresponding address virtual multicast.

WO 02/062076 A2

(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.

DISTRIBUTED MULTICASTING FOR AN ATM NETWORK

FIELD OF THE INVENTION

The present invention relates to multicasting over an asynchronous transfer mode network.

BACKGROUND OF THE INVENTION

Asynchronous Transfer Mode (ATM) is a network technology based on transferring data in cells or packets of a fixed size. ATM equipment is often used to transmit video, audio, and data over the same network.

FIG 1 shows a typical ATM network comprising leaves 1, root 2, and switches 3A and 3B. These network components are connected by data communication connections 5, 6, and 7. The connections can be through various media, such as fiber optic, twisted pair copper, or xDSL.

Multicasting is a unidirectional connection over a network from root **2** to one or more leaves 1. An example of multicasting is distribution of video, audio, or other multimedia over some type of digital subscriber line, (e.g., xDSL).

ATM currently supports multicasting but in an inefficient manner, placing too much load on the root. ATM forum standard UNI 3.1 does not define how a leaf joins a multicast connection. (The leaf has somehow to communicate with the root in order for the root to issue the request for the network to add the leaf to the multicast connection. How the leaf does this is not covered by the standard.)

This creates a scalability problem in the case where the multicast tree is very big and changing frequently as leaves join or quit the connection. In such a tree the need for the root to handle each leaf's joining/quitting places a very high load on the root.

ATM forum standard, UNI 4.0, specifies a new feature called leaf-initiated join (LIJ), which reduces some of the load on the root. But support for UNI 4.0 is not such widespread and still requires root intervention to implement access control. In other words, unless access to the multicast is to be unrestricted, the root must intervene in LIJ to verify if the endpoint can join the multicast. Again, this creates a scalability problem.

The solution is to logically split the multicast tree into smaller subtrees by introducing an intermediary component that can act as a root towards the leaves and as a leaf towards the root. Such distributed multicasting has the advantage of eliminating the need to send the same information multiple times over the same connection, and thereby improves bandwidth usage.

There are many applications for such an invention, one of which is multicasting TV over an ATM network to television sets (the leaves 1 of the multicast). In such an application subscribers are constantly joining and leaving the multicast. The intermediaries reduce the bandwidth and processing demands on the root by handling data transmission and managing subscriber access.

The intermediary component relieves the root of having to process leaf-join requests while maintaining security. This is superior to the leaf-initiated join of UNI 4.0 which either allows anyone to join (lacks security) or requires that the root process each leaf-join request (not scalable).

In summary, it is a main object of the present invention to provide a means for splitting a multicast tree into smaller subtrees.

It is a further object of the present invention to provide a means for relieving the root of a multicast ATM network of handling leaf-join requests while maintaining network security.

It is a further object of the present invention to provide a means for efficiently multicasting video, audio, and other high bandwidth data over an ATM network.

Still further objects and advantages will become apparent from a consideration of the ensuing description and drawings.

BRIEF DESCRIPTION OF THE INVENTION

There is thus provided in accordance with a preferred embodiment of the present invention, an apparatus for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots which transmit data in cells, at least one of a plurality of switches which switch the data to endpoints, and at least one of a plurality of endpoints which subscribe to the data from the roots; the apparatus paired with a switch on the network and comprising:

a) a root communications interface that will receive and transmit cells to and from the root;

b) a switch communications interface that will receive and transmit cells to and from the switch;

c) a substitution unit that will substitute cell addresses and which is in communication with the communications interfaces;

d) a control unit comprising computing means that will control the operation of the apparatus and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution unit;

whereby, in response to a request from an endpoint to join a multicast that no other endpoint connected to the apparatus is a member of, the control unit will issue a join-multicast request to the root; and upon the root establishing a root–apparatus multicast virtual circuit, the control unit will add the address of the circuit to the table and notify the paired switch to establish an apparatus-endpoint virtual circuit and will add the address of the apparatus-endpoint multicast virtual circuit to the table, and upon a subsequent request by any endpoint connected to the apparatus to join the multicast, the control unit will notify the paired switch to add the endpoint to the multicast without involving the root; and

whereby the substitution unit substitutes the addresses of cells from the root with the corresponding address of the apparatus-endpoint multicast virtual circuit; and

whereby the apparatus acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

Furthermore, in accordance with another preferred embodiment of the present invention, the substitution unit comprises a processor that substitutes part or all of an ATM cell header according to a substitution table.

Furthermore, in accordance with another preferred embodiment of the present invention, the substitution unit will substitute the address of an incoming cell;

Furthermore, in accordance with another preferred embodiment of the present invention, the apparatus is integrated into the switch,

Furthermore, in accordance with another preferred embodiment of the present invention, the function of the apparatus is implemented in the programming of the switch and wherein the communications interface to the switch is a logical interface and the communications interface to the root is implemented through the switch's interface to the root.

Furthermore, in accordance with another preferred embodiment of the present invention, the switch is a digital subscriber line access multiplexer ATM switch.

Furthermore, in accordance with another preferred embodiment of the present invention, the apparatus is applied to distributing audio data.

Furthermore, in accordance with another preferred embodiment of the present invention, the apparatus is applied to video distribution.

Furthermore, in accordance with another preferred embodiment of the present invention, the apparatus is applied to video distribution over digital subscriber line systems.

There is thus also provided in accordance with a preferred embodiment of the present invention, a method for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots transmitting data in cells, at least one of a plurality of switches switching the data to endpoints, and at least one of a plurality of endpoints subscribing to the data from the roots; the method applied in cooperation with a switch on the network and comprising:

a) providing a root communications means that will receive and transmit cells to and from the root;

b) providing a switch communications means that will receive and transmit cells to and from the switch;

c) providing a substitution means that will substitute cell addresses and which is in communication with the communications interfaces;

d) providing a control means that will control the process and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution means;

whereby, in response to a request from an endpoint to join a multicast that no other endpoint connected to the cooperating switch is a member of, the control means will issue a join-multicast request to the root; and upon the root establishing a root–apparatus multicast virtual circuit, the control means will add the address of the circuit to the table and notify the cooperating switch to establish an apparatus-endpoint virtual circuit and will add the address of the apparatus-endpoint multicast virtual circuit to the table, and upon a subsequent request by any endpoint connected to the cooperating switch to join the multicast, the control means will notify the cooperating switch to add the endpoint to the multicast without involving the root; and

whereby the substitution means substitutes the addresses of cells from the root with the corresponding address of the apparatus-endpoint multicast virtual circuit; and

whereby the apparatus acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

There is thus also provided in accordance with a preferred embodiment of the present invention, a method for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots transmitting data in cells, at least one of a plurality of switches switching the data to endpoints, and at least one of a plurality of endpoints subscribing to the data from the roots; the method applied in cooperation with a switch on the network and comprising:

a) providing a root communications means that will receive and transmit cells to and from the root;

b) providing a switch communications means that will receive and transmit cells to and from the switch;

c) providing a substitution means that will substitute cell addresses and which is in communication with the communications interfaces;

d) providing a control means that will control the process and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution means;

whereby, an endpoint joins a multicast as follows:

i) the endpoint establishes a virtual circuit with the control means;

ii) the control means adds a control record to the table with the circuit address received from the endpoint in the incoming address field and "leaf" in the control field;

iii) the endpoint sends a request over the virtual circuit to the control means to join a given multicast tree;

iv) the control means verifies the endpoint's right to join the multicast tree and, if this is the first endpoint under the cooperating switch to request to join this particular multicast tree, then:

A) the control means creates a virtual circuit to the root by sending a signaling request to the switch above it;

B) the switch above the control means responds with signaling information about the virtual circuit including an address for the circuit;

C) the control means adds a control record to the table with the address of the virtual circuit to the root as the incoming address and "root" in the Control field;

D) the control means sends a request to the root, asking to join the multicast as a leaf;

E) the root sends signaling information to the switch above the control means to add the control means to the multicast tree;

F) the switch above the control means sends signaling information to the control means;

G) the control means stores the signaling information in the table;

H) the control means sends signaling information to the cooperating switch, requesting to establish a multicast tree the endpoint that originally requested to join the multicast;

I) the cooperating switch responds to the control means with signaling information;

J) the control means updates the table with the incoming address received from the switch above the control means, the outgoing address received from the cooperating switch, and "none" as the Control value;

v) otherwise, the control means sends only a signaling request to the cooperating switch, telling it to add the new requesting endpoint to the distribution for the multicast tree;

vi) as multicast cells arrive from the root, the substitution means substitutes incoming address received from the switch above the control means with the outgoing address received from the cooperating switch;

whereby the cooperating switch acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

BRIEF DESCRIPTION OF THE FIGURES

The invention is described herein, by way of example only, with reference to the accompanying Figures, in which like components are designated by like reference numerals.

FIG. 1 illustrates a typical ATM network and ATM multicast tree.

FIG. 2 illustrates an ATM multicast tree split into subtrees in accordance with a preferred embodiment of the present invention.

FIG. 3 illustrates nested ATM multicast trees in accordance with a preferred embodiment of the present invention.

FIG. 4 illustrates an implementation of the data substitution unit in accordance with a preferred embodiment of the present invention.

FIG. 5 illustrates a hardware implementation of a data substitution unit in accordance with a preferred embodiment of the present invention.

FIG. 6A illustrates a data substitution unit built into an ATM switch.

FIG. 6B illustrates data distribution implemented in the programming of an ATM switch.

FIG. 7 illustrates a database for an implementation of the data substitution unit in accordance with a preferred embodiment of the present invention.

FIG. 8 is a flowchart for joining a multicast tree in accordance with a preferred embodiment of the present invention.

FIG. 9 is a flowchart for processing cells received by the data substitution unit from the direction of the root.

FIG. 10 is a flowchart for processing cells received by the data substitution unit from the direction of the leaves.

FIG. 11 is a flowchart for processing cells received by the data substitution unit from the control unit.

The following component numbers are used in the drawings:

1	Leaf	1A	Endpoint (EP)
2	Root	3	ATM Switch
3A	Low level ATM switch	3B	High level ATM switch
5	Low level switch – leaf connection	6	High level - low level connection
7	Root – high level connection	7A	Root – high level DDU connection
7B	High level DDU – high level switch connection	8	Data substitution unit (DDU)

8A	Low level data substitution unit	8B	High level data substitution unit
10	Low-side ATM interface	11	High-side ATM interface
12	Low-side virtual connection 2	13	High-side virtual connection 1
14	Low level DDU – switch interface	18	Main multicast tree
18A	Reduced main multicast tree	18B	Further reduced main multicast tree
18C	Higher level subtree	19	Multicast subtree
20	Substitution unit	21	Control unit
22	Physical interface to higher level ATM	23	Physical interface to lower level ATM
24	Data distribution processor	25	Segmentation and reassembly (SAR)
26	Microcontroller	27	Shared RAM
28	RAM	29	Flash ROM
30	Ethernet connection	31	Management interface
32	PCI bus	41	ATM switch circuitry
42	Data distribution circuitry	43	ATM switch programming
44	Data distribution programming	102	Lookup table

DETAILED DESCRIPTION OF THE INVENTION

Note: In this detailed description the terms "high level" and "low level" refer to a component or connection's relative position in the multicast tree. "Low" in this context means "closer to the leaf" and "high" means "closer to the root." A low level component can be identical in construction to its high level counterpart. For example, the data substitution unit (DDU) 8 is referred to as 8A when located lower in the multicast tree and 8B when located higher in the multicast tree.

Every multicast tree is a subset of the physical ATM network. There can be more than one multicast tree present in the network simultaneously. A given physical endpoint can participate in more than one multicast tree. There can be more than one physical endpoint acting as a root but every multicast tree must have exactly one root.

FIG 1 illustrates a simple ATM network implementation of a two-level ATM multicast tree 18. Root 2 is connected via root – high level connections 7 to high level switch 3B. High level switch 3B is connected via high level - low level connection 6 to low level switches 3A. Low level switch 3A is connected via low level switch - leaf connection 5 to leaf 1.

Connections 5, 6, and 7 are standard data communication connections, such as fiber optic, twisted pair copper, xDSL modems etc. Nodes on the network that are not subscribing to the multicast are labeled as endpoints 1A to distinguish them from leaves 1.

FIG 2 illustrates the same large multicast tree split into a reduced main tree 18A and subtrees 19 in accordance with a preferred embodiment of the present invention. This is accomplished by inserting low level data substitution unit (DDU) 8A between low level switch 3A and high level switch 3B. The connection to low level switch 3A is via a new connection: low level DDU – switch interface 14. The DDU 8A takes the place of the low level switch 3A in the original high level - low level connection 6.

The reduced main tree 18A is composed of root 2, root-high level connection 7, high level switch 3B, and high level-low level connection 6. The subtrees 19 are composed of low level DDU – switch interface 14, low level switch 3A, low level switch - leaf connection 5, and leaf 1.

The DDU 8 behaves toward the reduced main multicast tree 18A as a leaf (at the end of high level-low level connection 6) and toward the subtree 19 it behaves as a root (via low level DDU – switch interface 14 to the low level switch 3A).

FIG 3 illustrates the basic operation of the DDU 8 during a multicast. Physically, cells (data) received at high-side interface 11 are passed along through to low-side interface 10. Substitution unit 20 receives the cells as if the DDU is a leaf at the end of high-side virtual connection 13 and passes the cells along as if the DDU is a root at the base of low-side virtual connection 12.

DDU 8 functionality does not inhibit ATM cells that are not related to the multicast (for example, unicast virtual circuits implementing Internet connectivity). These cells pass through the DDU unchanged.

DDUs 8 can be inserted at multiple levels in the tree, each time before a switch 3. FIG. 3 illustrates nested DDUs, breaking up reduced main multicast tree 18A into further reduced main multicast tree 18B and higher level subtree 18C. It will be noted that the "leaves" of the higher level subtree 18C are low level DDUs 8A, which also serve as roots for subtrees 19 (not shown in figure).

FIG. 5 shows the primary components of the DDU 8, which can be divided into two primary two logical parts: substitution unit 20 and control unit 21. Substitution unit 20 primarily comprises physical interfaces 23 to ATM interfaces 10 and 11, dynamic distribution processor 24, and segmentation and reassembly (SAR) unit 25 for the control traffic (ATM signaling etc.). Shared RAM 27 stores information regarding cell headers, whereby the DDU 8 converts a cell from one multicast tree to another multicast tree.

Control unit 21 is a common computer design, comprising microcontroller 26, bus 27, RAM 28 and Flash ROM 29. Control unit 21 components are connected via PCI bus 32. Control unit 21 can also include an Ethernet interface 30 to management system 31.

FIG 5 illustrates a preferred embodiment of the DDU 8, as a standalone hardware unit. However the DDU 8 functionality could instead be built into the hardware of switch 3 (FIG. 6A) or implemented as part of the firmware/software switch 3 (FIG. 6B). In those cases the low level DDU – switch interface 14 would be implemented respectively as an internal electronic interface or software interface. A particular form of software implementation could be as part of a digital subscriber line access multiplexer ATM switch.

At system setup, a table in RAM 28 is filled with the ATM addresses for all possible roots and for all potential leaves 1 that could be part of a multicast tree involving that DDU. Optionally, RAM could also be configured with security and access information for each potential leaf.

The task of the DDU 8 is only to manage the topology of the multicast. The switch does the actual distribution. DDU network management operations are managed by the DDU's substitution unit 20, working from a lookup table in shared RAM 27. The lookup table 102 is shown in FIG. 7. The table records are indexed by the incoming cell address, the outgoing address (if there is one) to replace the incoming address with, and the control destination. Control destination can be leaf, root, or no control (meaning that this incoming address is not used for control purposes).

When a network endpoint 1A requests to join a multicast and the DDU 8 to which the endpoint is connected is not itself a leaf of that multicast, then the DDU 8 has to first join the multicast. This the DDU 8 does by using LIJ if the root 2 supports multicast switched virtual circuits. If the root 2 does not then a DDU 8B can be installed higher in the network to handle LIJ requests.

FIG. 8 illustrates the process for setting up a multicast tree 19 like the one shown as in FIG. 2. The procedure follows these steps:

60 Endpoint 1A establishes a virtual circuit with DDU 8A. (More specifically, the DDU receives a signaling request from the ATM switch 3A containing information about the virtual circuit, including an address for the circuit.)

61 DDU 8A adds a control record to the table in shared RAM 27 with the circuit address received in step 60 in the incoming address field and "leaf" in the control field. The outgoing address field is not used.

62 Endpoint 1A sends a request over the virtual circuit to DDU 8A to join a given multicast tree.

63 DDU 8A verifies endpoint's right to join the multicast tree and one of two possibilities occurs, depending on whether this is the first endpoint 1A under this DDU 8A to request to join this particular multicast tree. If it is the first such request, then steps 64 to 73 are followed. Otherwise step 74 is followed.

64 If this is the first endpoint to request to join the multicast tree (meaning that no other endpoint under this DDU 8A is currently part of that tree), then the DDU 8A creates a virtual circuit to the root by sending a signaling request to ATM switch 3B.

65 ATM switch 3B responds with signaling information about the virtual circuit including an address for the circuit.

66 DDU 8A adds a control record to the table in the shared RAM 27 containing the circuit address received in step 65 as the incoming address and "root" in the Control field. The outgoing address field is not used.

67 DDU 8A sends a request to the root 2, asking to join the multicast as a leaf.

68 If the request is approved by root 2, then root 2 sends signaling information to ATM switch 3B to add the DDU 8A to the multicast tree 18A.

69 ATM switch 3B sends signaling information to DDU 8A.

70 DDU 8A stores the signaling information in RAM 28.

71 DDU 8A sends signaling information to low level ATM switch 3A, requesting to establish a multicast tree with leaf 1 (which is the endpoint 1A that initially requested to join the multicast).

72 ATM switch 3A responds to the DDU 8A with signaling information, which is saved in RAM 28.

73 DDU 8A updates table in shared RAM 27 with the incoming address received from ATM switch 3B (step 66) where the DDU acts as a leaf in main tree 19A and the outgoing address received from ATM switch 3A (step 69) where the DDU acts as a root for the new leaf 1 in new multicast subtree 19. The Control field is set to "None" for such multicast records. This record is the only information required to establish the subtree 19 for that multicast.

74 If there is already at least one other endpoint 1A connected to the same DDU 8A that is part of the multicast, then the DDU 8A only sends a signaling request to the low level ATM switch 3A, to add the new requesting endpoint 1A to the distribution for multicast tree 19.

Once an endpoint 1A has established a virtual circuit (steps 60 and 62), it can use that circuit for all its control requests regarding all the multicasts available to that leaf. Similarly, once the DDU 8 has established a virtual circuit with the root, the root can use that virtual circuit to handle all the multicasts originating from that root.

While the above description covers a single DDU connected via a switch to the root, the same principles apply if the DDU is nested under a higher DDU. In that case the first DDU would appear as a leaf to the higher DDU.

Once the distributed multicast has been established, there are three potential origins for cells received by the DDU 8 substitution unit 20: From the direction of the root 2, from the direction of the leaves 1, or from the DDU's own control unit 21.

FIG. 9 is a flowchart for processing cells received by the data substitution unit from the direction of the root. The steps are as follows:

75 A cell is received from the direction of the root.

76 The substitution unit 20 in the DDU 8 looks up the incoming cell address in table 102.

77 If it is not in the table, the cell is sent on unchanged in the direction of the leaves.

78 If the incoming address is in the table, the substitution unit checks whether that address is associated with control.

79 If the incoming address is associated with control, the cell is sent to the DDU's control unit 21.

80 Otherwise, the substitution unit replaces the address with the associated outgoing address.

82 Send the modified cell in the direction of the leaves.

FIG. 10 is a flowchart for processing cells received by the data substitution unit from the direction of the leaves. The steps are as follows:

86 A cell is received from the direction of the leaves.

88 The substitution unit 20 in the DDU 8 looks up the incoming cell address in table 102.

90 If the record indexed by the incoming address is marked as control, the cell is sent unchanged to the control unit 21.

92 Otherwise, the cell is sent unchanged in the direction of the root.

FIG. 11 is a flowchart for processing cells received by the data substitution unit from the control unit 21. The steps are as follows:

94 A cell is received from the control unit.

96 The substitution unit 20 in the DDU 8 looks up the control field of the incoming cell address in table 102.

98 If the control value is root, the cell is sent unchanged) in the direction of the root.

92 Otherwise, the cell is sent unchanged in the direction of the leaves.

It should be clear that the description of the embodiments and attached Figures set forth in this specification serves only for a better understanding of the invention, without limiting its scope.

It should also be clear that a person skilled in the art, after reading the present specification could make adjustments or amendments to the attached Figures and above described embodiments that would still be covered by the following Claims.

CLAIMS

It is claimed:

1. An apparatus for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots which transmit data in cells, at least one of a plurality of switches which switch the data to endpoints, and at least one of a plurality of endpoints which subscribe to the data from the roots; the apparatus paired with a switch on the network and comprising:

a) a root communications interface that will receive and transmit cells to and from the root;

b) a switch communications interface that will receive and transmit cells to and from the switch;

c) a substitution unit that will substitute cell addresses and which is in communication with the communications interfaces;

d) a control unit comprising computing means that will control the operation of the apparatus and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution unit;

whereby, in response to a request from an endpoint to join a multicast that no other endpoint connected to the apparatus is a member of, the control unit will issue a join-multicast request to the root; and upon the root establishing a root–apparatus multicast virtual circuit, the control unit will add the address of the circuit to the table and notify the paired switch to establish an apparatus-endpoint virtual circuit and will add the address of the apparatus-endpoint multicast virtual circuit to the table, and upon a subsequent request by any endpoint connected to the apparatus to join the multicast, the control unit will notify the paired switch to add the endpoint to the multicast without involving the root; and

whereby the substitution unit substitutes the addresses of cells from the root with the corresponding address of the apparatus-endpoint multicast virtual circuit; and

whereby the apparatus acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

2. The apparatus of claim 1, wherein the substitution unit comprises a processor that substitutes part or all of an ATM cell header according to a substitution table.

3. The apparatus of claim 1, wherein the substitution unit will substitute the address of an incoming cell;

4. The apparatus of claim 1 integrated into the switch,

5. The apparatus of claim 1 implemented in the programming of the switch and wherein the communications interface to the switch is a logical interface and the communications interface to the root is implemented through the switch's interface to the root.

6. The apparatus of claim 1 where the switch is a digital subscriber line access multiplexer ATM switch.

7. The apparatus of claim 1 applied to distributing audio data.

8. The apparatus of claim 1 applied to video distribution.

9. The apparatus of claim 1 applied to video distribution over digital subscriber line systems.

10. An apparatus for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots which transmit data in cells, at least one of a plurality of switches which switch the data to endpoints, and at least one of a plurality of endpoints which subscribe to the data from the roots; the apparatus paired with a switch on the network and comprising:

    means for passing data from a first virtual multicast circuit where the apparatus acts as virtual leaf to a corresponding virtual multicast circuit where the apparatus acts as a virtual root.

    whereby the apparatus acts as a virtual leaf towards the root and as a virtual root towards the endpoints.

11. The apparatus of claim 10 integrated into the switch,

12. The apparatus of claim 10 implemented in the programming of the switch.

13. The apparatus of claim 10 where the switch is a digital subscriber line access multiplexer ATM switch.

14. The apparatus of claim 10 applied to video distribution.

15. The apparatus of claim 10 applied to video distribution over digital subscriber line systems.

16. A method for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots transmitting data in cells, at least one of a plurality of switches switching the data to endpoints, and at least one of a plurality of endpoints subscribing to the data from the roots; the method applied in cooperation with a switch on the network and comprising:

a) providing a root communications means that will receive and transmit cells to and from the root;

b) providing a switch communications means that will receive and transmit cells to and from the switch;

c) providing a substitution means that will substitute cell addresses and which is in communication with the communications interfaces;

d) providing a control means that will control the process and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution means;

whereby, in response to a request from an endpoint to join a multicast that no other endpoint connected to the cooperating switch is a member of, the control means will issue a join-multicast request to the root; and upon the root establishing a root–apparatus multicast virtual circuit, the control means will add the address of the circuit to the table and notify the cooperating switch to establish an apparatus-endpoint virtual circuit and will add the address of the apparatus-endpoint multicast virtual circuit to the table, and upon a subsequent request by any endpoint connected to the cooperating switch to join the multicast, the control means will notify the cooperating switch to add the endpoint to the multicast without involving the root; and

whereby the substitution means substitutes the addresses of cells from the root with the corresponding address of the apparatus-endpoint multicast virtual circuit; and

whereby the apparatus acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

17. A method for distributed multicasting over an ATM network; the multicast comprising at least one of a plurality of roots transmitting data in cells, at least one of a plurality of switches switching the data to endpoints, and at least one of a plurality of endpoints subscribing to the data from the roots; the method applied in cooperation with a switch on the network and comprising:

a) providing a root communications means that will receive and transmit cells to and from the root;

b) providing a switch communications means that will receive and transmit cells to and from the switch;

c) providing a substitution means that will substitute cell addresses and which is in communication with the communications interfaces;

d) providing a control means that will control the process and maintain a table of cell addresses and control values and which is in communication with the communications interfaces and with the substitution means;

whereby, an endpoint joins a multicast as follows:

i) the endpoint establishes a virtual circuit with the control means;

ii) the control means adds a control record to the table with the circuit address received from the endpoint in the incoming address field and "leaf" in the control field;

iii) the endpoint sends a request over the virtual circuit to the control means to join a given multicast tree;

iv) the control means verifies the endpoint's right to join the multicast tree and, if this is the first endpoint under the cooperating switch to request to join this particular multicast tree, then:

A) the control means creates a virtual circuit to the root by sending a signaling request to the switch above it;

B) the switch above the control means responds with signaling information about the virtual circuit including an address for the circuit;

C) the control means adds a control record to the table with the address of the virtual circuit to the root as the incoming address and "root" in the Control field;

D) the control means sends a request to the root, asking to join the multicast as a leaf;

E) the root sends signaling information to the switch above the control means to add the control means to the multicast tree;

F) the switch above the control means sends signaling information to the control means;

G) the control means stores the signaling information in the table;

H) the control means sends signaling information to the cooperating switch, requesting to establish a multicast tree the endpoint that originally requested to join the multicast;

I) the cooperating switch responds to the control means with signaling information;

J) the control means updates the table with the incoming address received from the switch above the control means, the outgoing address received from the cooperating switch, and "none" as the Control value;

v) otherwise, the control means sends only a signaling request to the cooperating switch, telling it to add the new requesting endpoint to the distribution for the multicast tree;

vi) as multicast cells arrive from the root, the substitution means substitutes incoming address received from the switch above the control means with the outgoing address received from the cooperating switch;

whereby the cooperating switch acts as a virtual leaf towards the root and as a virtual root towards the endpoints, thereby reducing the demands on the root by managing all leaf-join requests after the first and by distributing the multicast data.

PRIOR ART

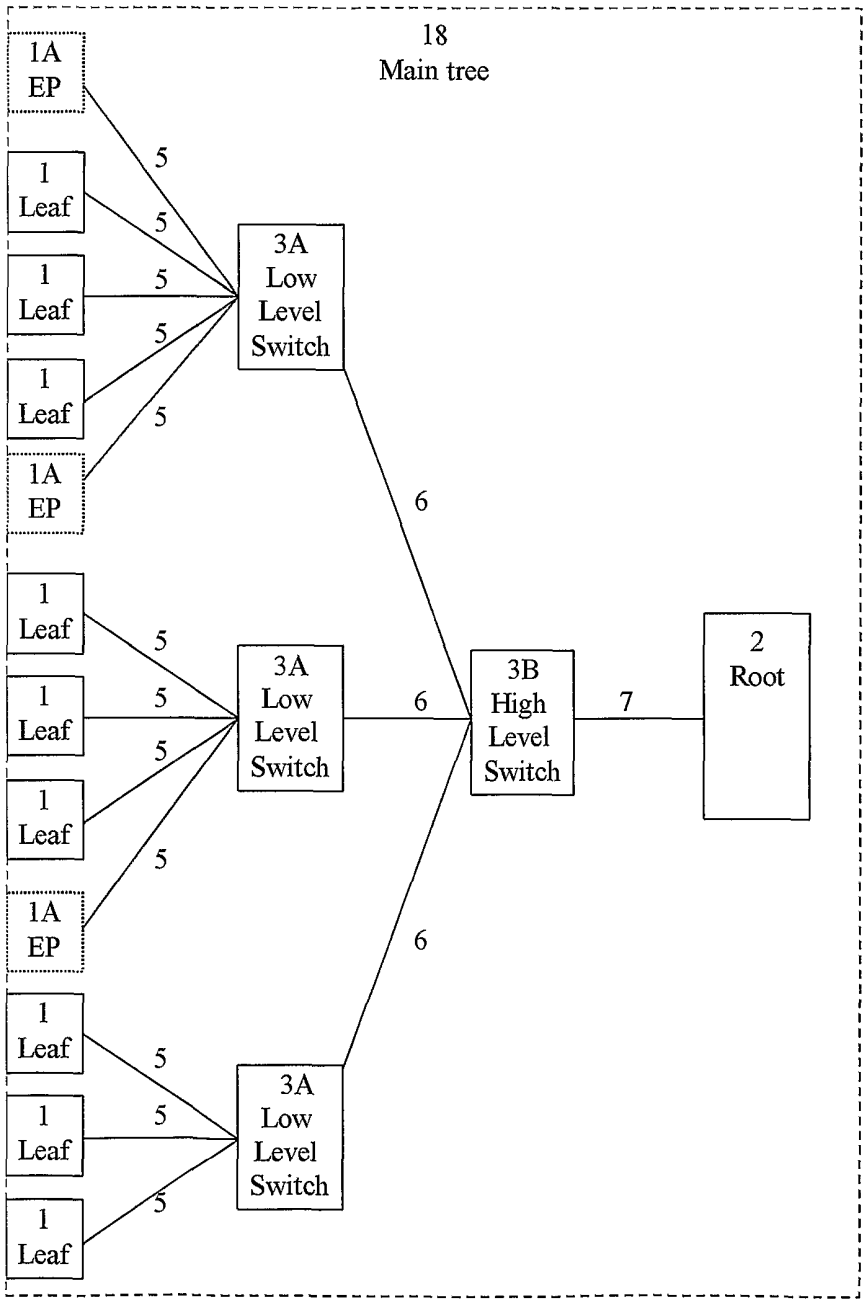

FIG. 1

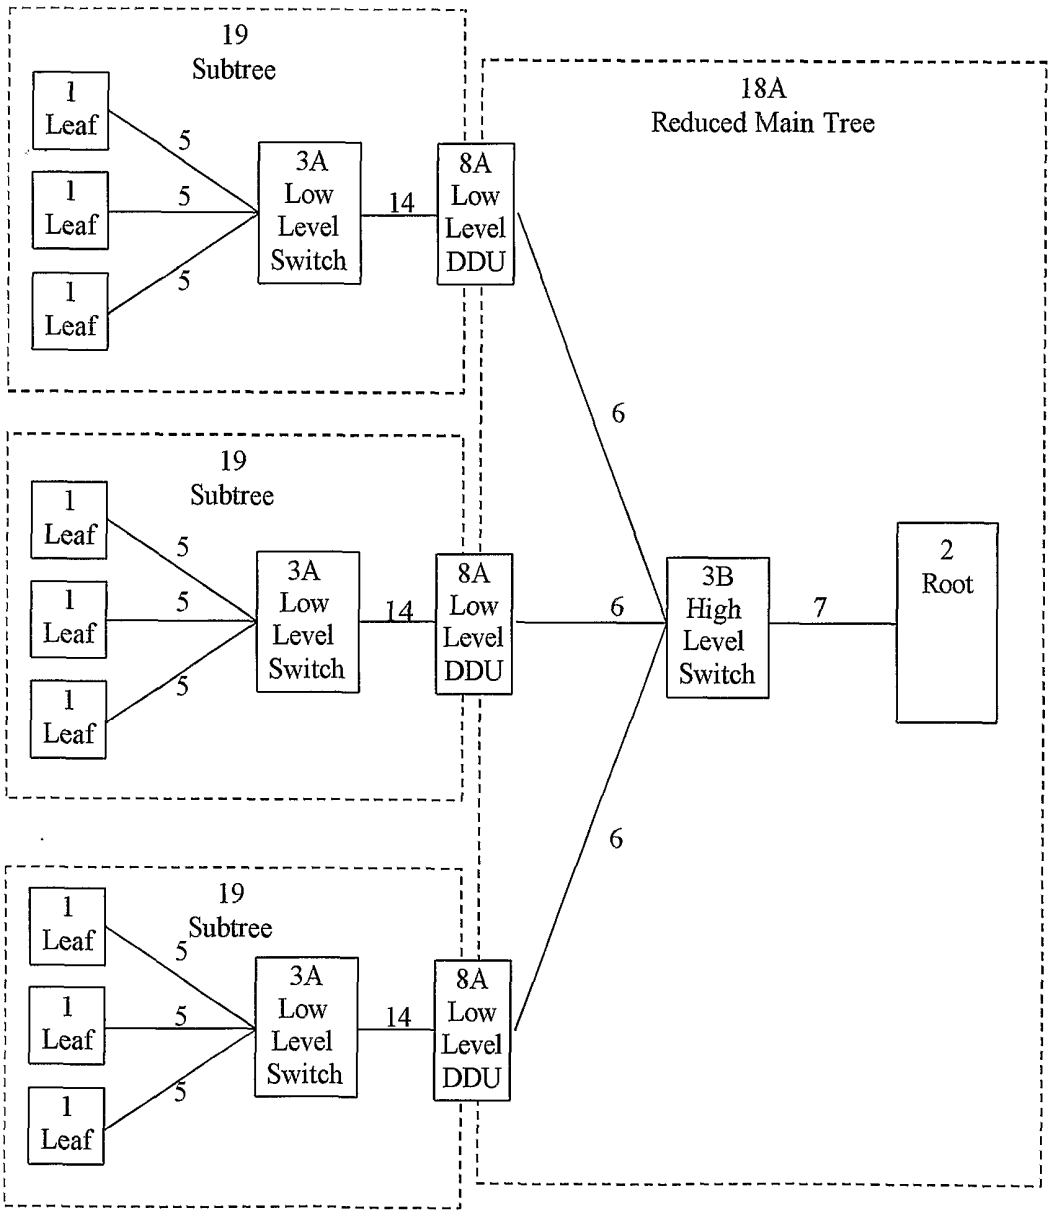
19
Subtree
1
Leaf
1
Leaf
1
Leaf
5
5
5
3A
Low
Level
Switch
14
8A
Low
Level
DDU
18A
Reduced Main Tree
6
19
Subtree
1
Leaf
1
Leaf
1
Leaf
5
5
5
3A
Low
Level
Switch
14
8A
Low
Level
DDU
6
3B
High
Level
Switch
7
2
Root
6
19
Subtree
1
Leaf
1
Leaf
1
Leaf
5
5
5
3A
Low
Level
Switch
14
8A
Low
Level
DDU

FIG. 2

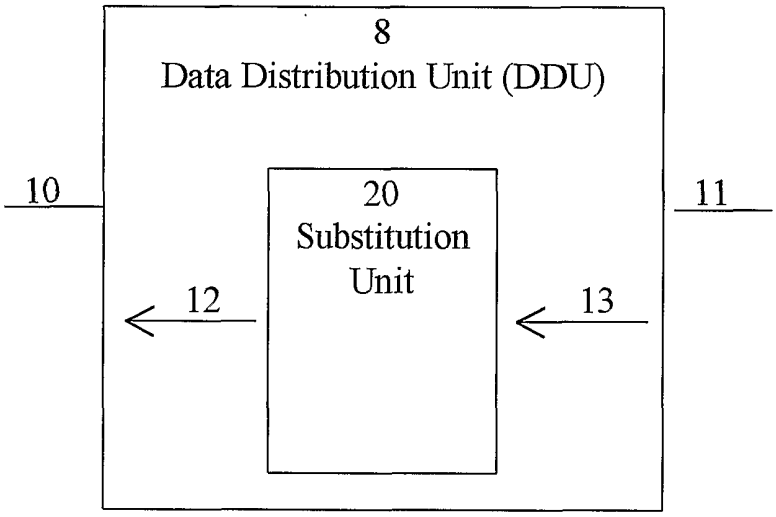
8
Data Distribution Unit (DDU)
10
20
Substitution
Unit
11
12
13

FIG. 3

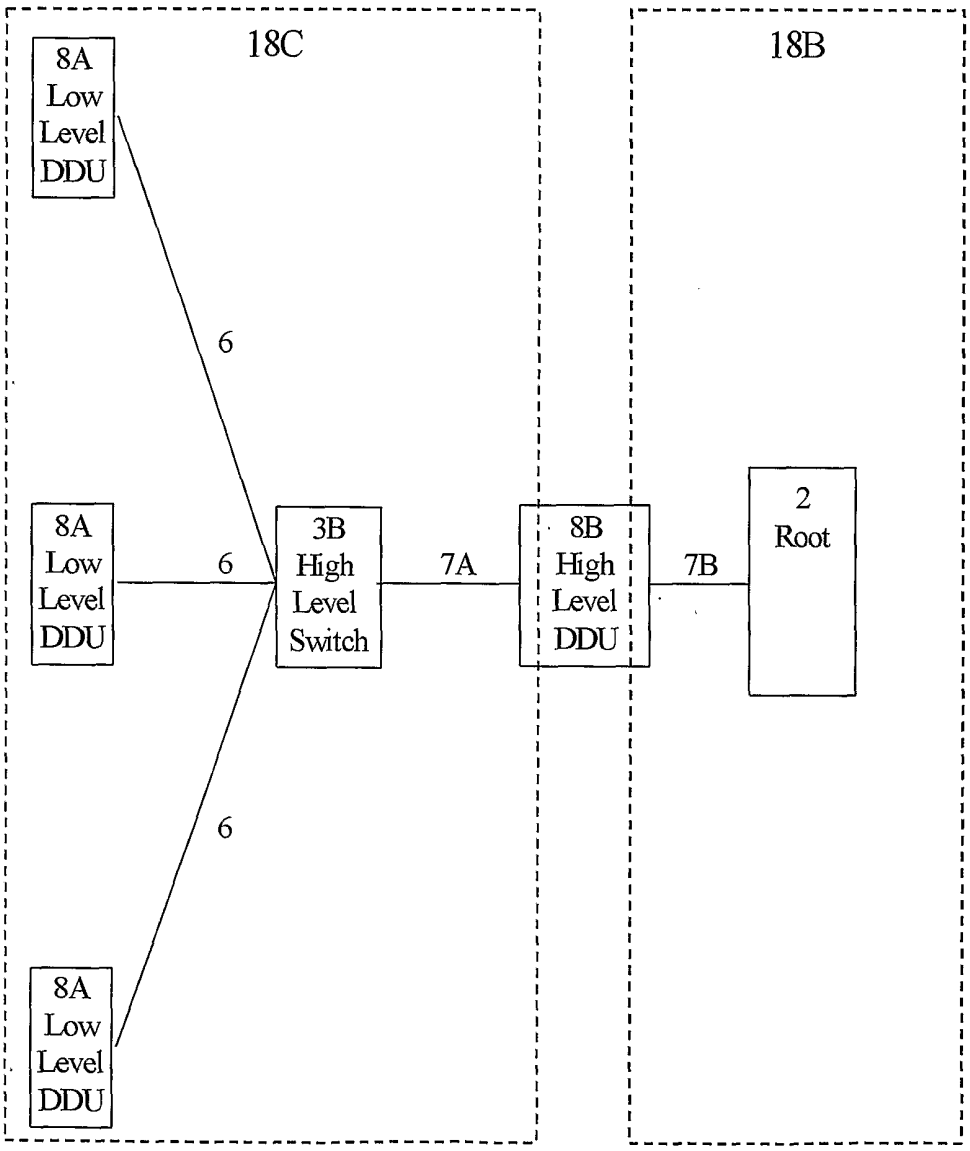
18C
18B
8A Low Level DDU
8A Low Level DDU
8A Low Level DDU
6
6
6
3B High Level Switch
7A
8B High Level DDU
7B
2 Root

FIG. 4

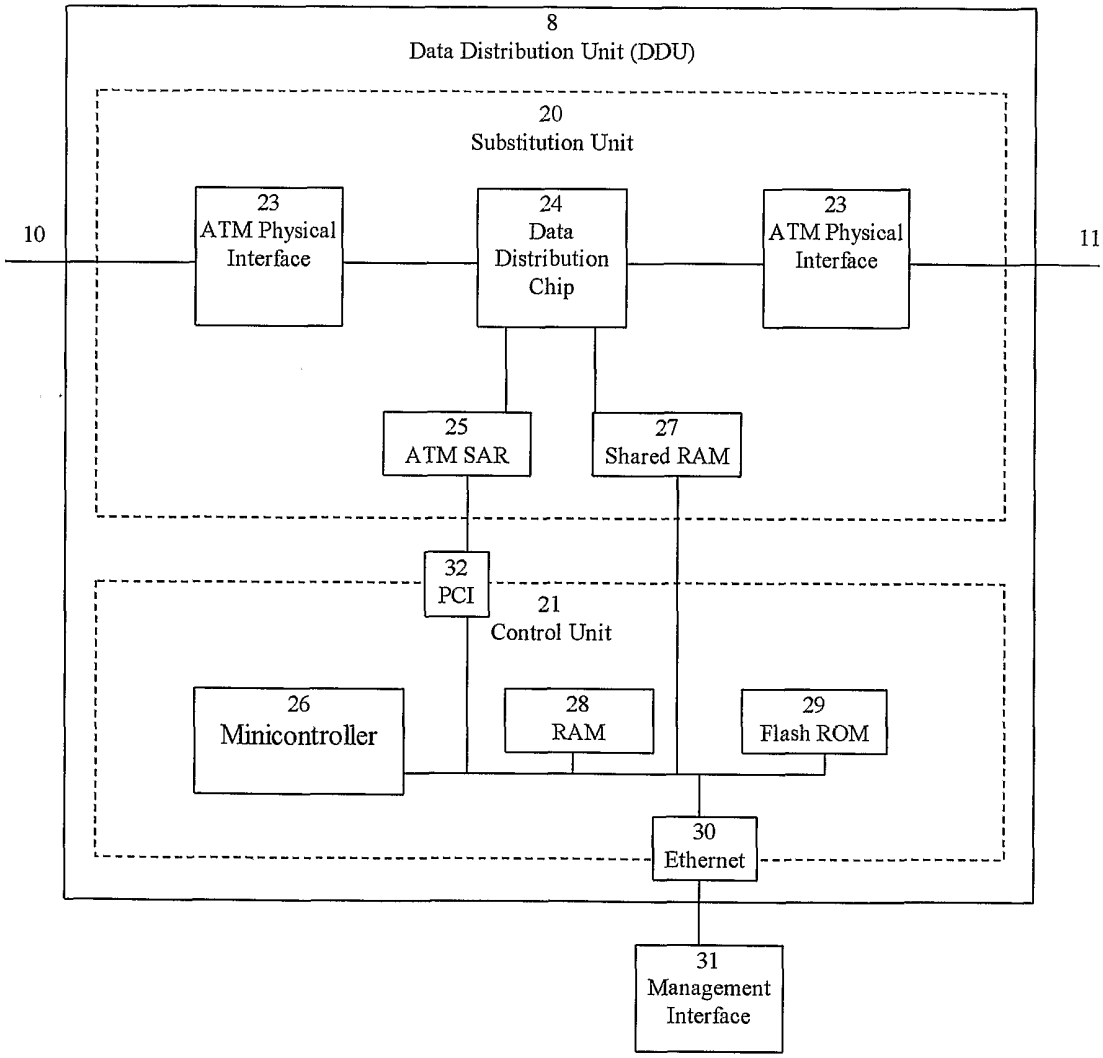
8
Data Distribution Unit (DDU)
20
Substitution Unit
10
23
ATM Physical
Interface
24
Data
Distribution
Chip
23
ATM Physical
Interface
11
25
ATM SAR
27
Shared RAM
32
PCI
21
Control Unit
26
Minicontroller
28
RAM
29
Flash ROM
30
Ethernet
31
Management
Interface

FIG. 5

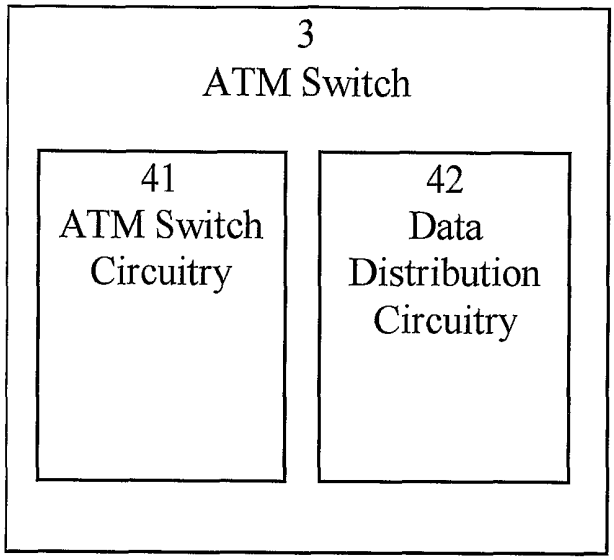
3
ATM Switch
41
ATM Switch
Circuitry
42
Data
Distribution
Circuitry

FIG. 6A

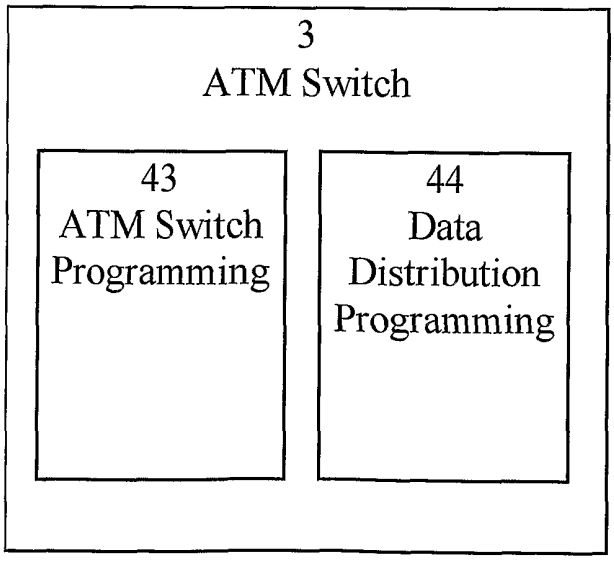
3
ATM Switch
43
ATM Switch
Programming
44
Data
Distribution
Programming

FIG. 6B

102

Incoming address	Outgoing address	Control
1234	7892	None
4323		Leaf
6655	6789	None
8899		Root
5678		Leaf

FIG. 7

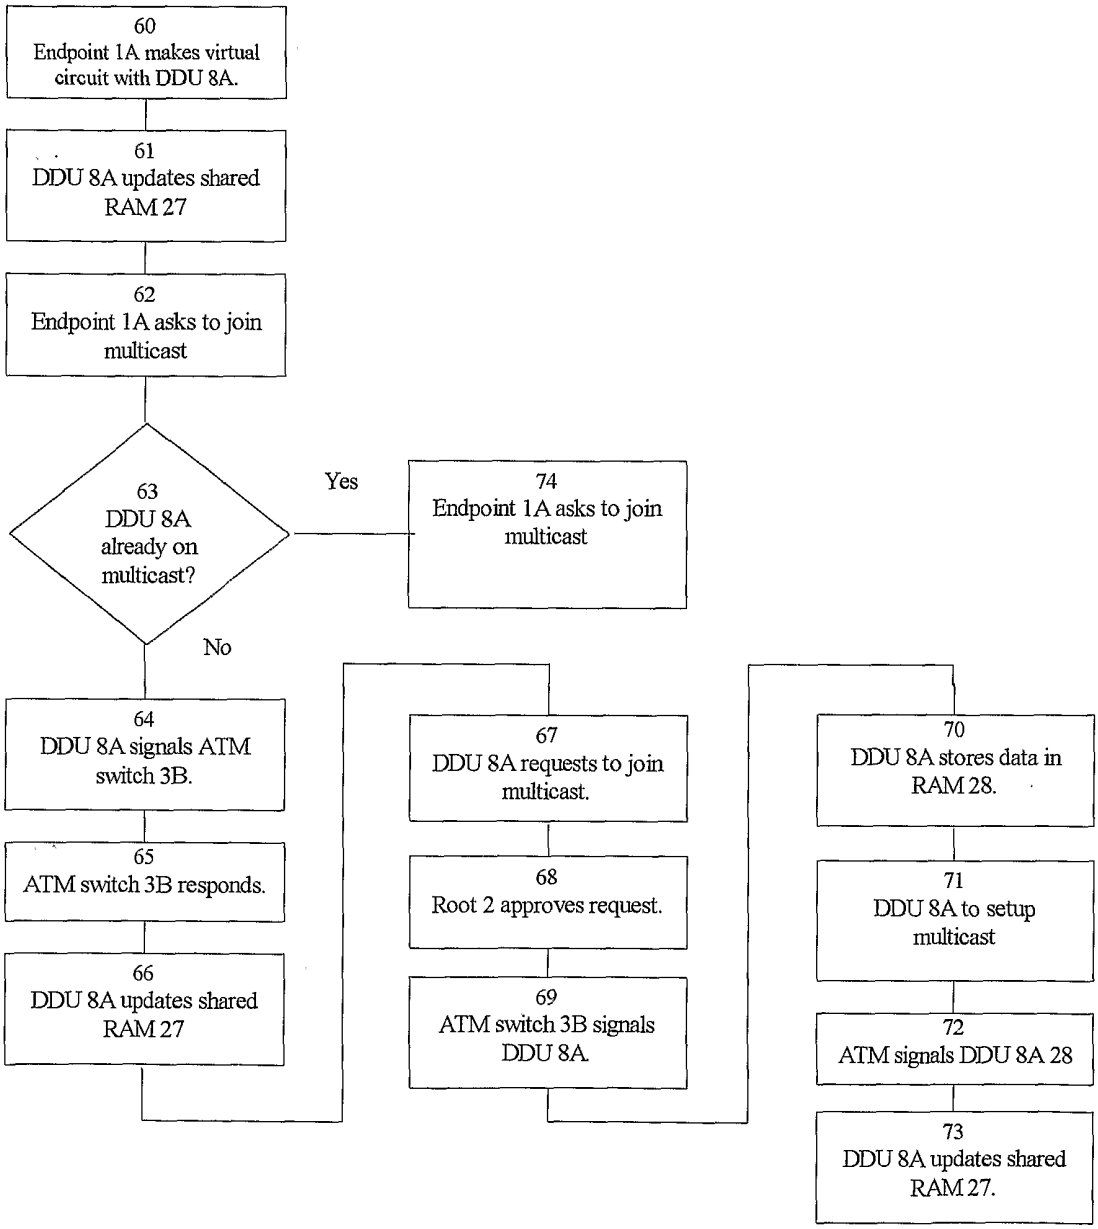
60
Endpoint 1A makes virtual circuit with DDU 8A.
61
DDU 8A updates shared RAM 27
62
Endpoint 1A asks to join multicast
63
DDU 8A already on multicast?
Yes
74
Endpoint 1A asks to join multicast
No
64
DDU 8A signals ATM switch 3B.
65
ATM switch 3B responds.
66
DDU 8A updates shared RAM 27
67
DDU 8A requests to join multicast.
68
Root 2 approves request.
69
ATM switch 3B signals DDU 8A.
70
DDU 8A stores data in RAM 28.
71
DDU 8A to setup multicast
72
ATM signals DDU 8A 28
73
DDU 8A updates shared RAM 27.

FIG. 8

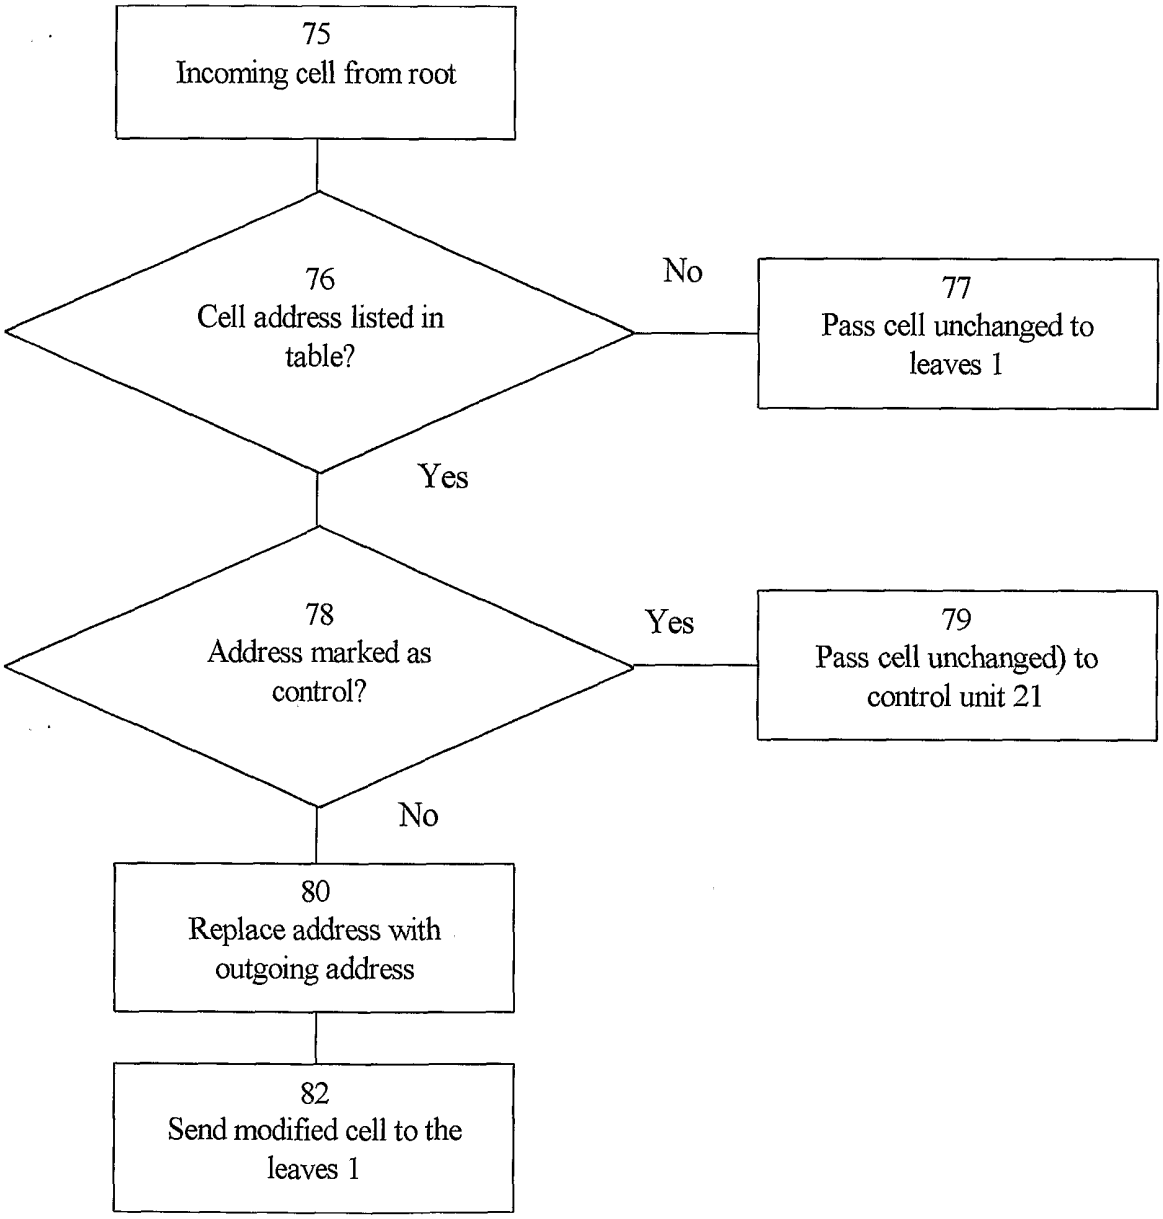
75
Incoming cell from root
76
Cell address listed in table?
No
77
Pass cell unchanged to leaves 1
Yes
78
Address marked as control?
Yes
79
Pass cell unchanged) to control unit 21
No
80
Replace address with outgoing address
82
Send modified cell to the leaves 1

FIG. 9

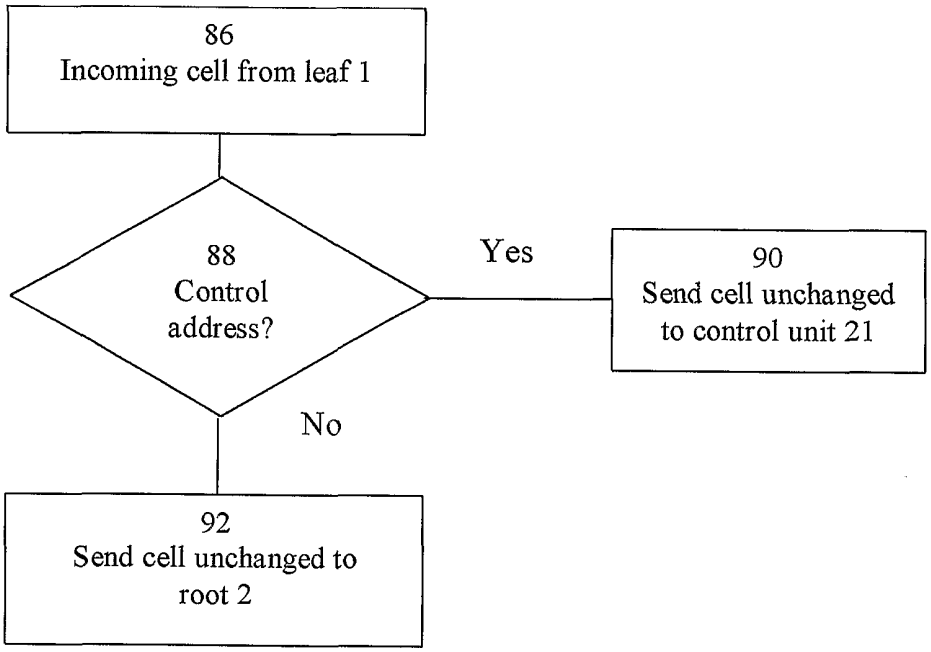
86
Incoming cell from leaf 1
88
Control address?
Yes
90
Send cell unchanged to control unit 21
No
92
Send cell unchanged to root 2

FIG. 10

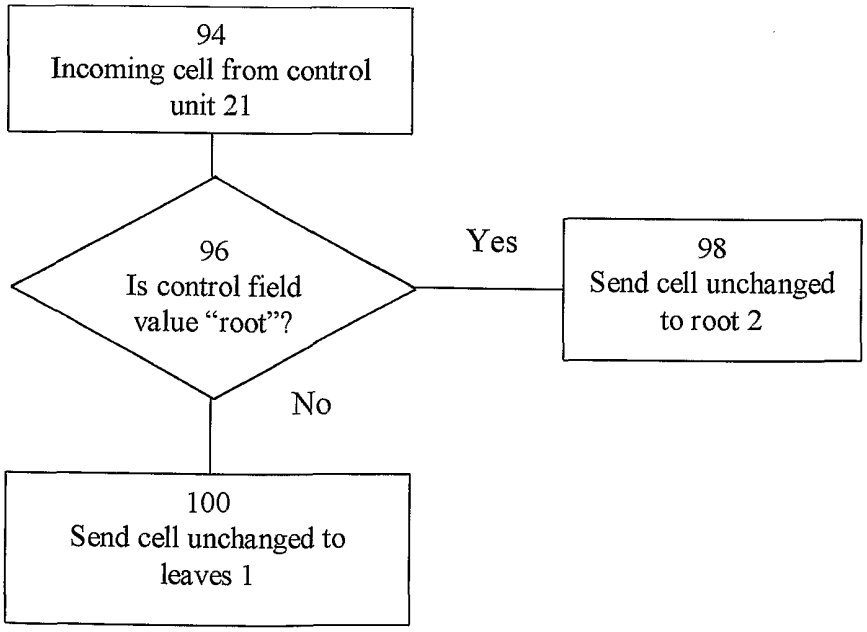
94
Incoming cell from control unit 21
96
Is control field value "root"?
Yes
98
Send cell unchanged to root 2
No
100
Send cell unchanged to leaves 1

FIG. 11