




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁷ : H04J 3/26	A1	(11) International Publication Number: WO 00/62454 (43) International Publication Date: 19 October 2000 (19.10.00)
<p>(21) International Application Number: PCT/US00/05390</p> <p>(22) International Filing Date: 1 March 2000 (01.03.00)</p> <p>(30) Priority Data: 09/291,705 14 April 1999 (14.04.99) US</p> <p>(71) Applicant: GTE LABORATORIES INCORPORATED [US/US]; 1209 Orange Street, Wilmington, DE 19801 (US).</p> <p>(72) Inventor: DENG, Shuang; 7 Dolphin Circle, Nashua, NH 03062 (US).</p> <p>(74) Agents: SUCHYTA, Leonard, Charles et al.; GTE Service Corporation, 600 Hidden Ridge Road, MC HQE03G13, Irving, TX 75038 (US).</p>		<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, 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, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, 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), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report.</p>

(54) Title: DATA LINK LAYER SWITCH WITH MULTICAST CAPABILITY

(57) Abstract

A data link layer switch includes a switching mechanism coupled to a plurality of port interface controllers (300). Each of the port interface controllers (300) comprises a multicast address table (320), a multicast matcher (310), and a timer (330). The multicast address table (320) stores multicast addresses for hosts attached to the port interface controller (300). The multicast matcher (310) matches an incoming packet to a target pattern, generates a multicast address from the incoming packet if the incoming packet matches the target pattern, and stores the generated multicast address in the multicast address table. The timer (330) determines an amount of time that the generated multicast address remains in the multicast address table.

300

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	TM	Turkmenistan
BF	Burkina Faso	GR	Greece		Republic of Macedonia	TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's	NZ	New Zealand		
CM	Cameroon		Republic of Korea	PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

DATA LINK LAYER SWITCH WITH MULTICAST CAPABILITY

Technical Field

5 The present invention relates generally to network switches and more particularly to a data link layer switch that supports Internet Protocol (IP) multicasting.

Background Art

10 Multicast communication includes the transmission of identical data packets to selected, multiple destinations. In contrast, broadcast communication includes the indiscriminate transmission of data packets to all destinations, and unicast communication includes the transmission of data packets to a single destination.

15 Every participant in a multicast receives information transmitted by any other participant in the multicast. Users connected to the network who are not participants in a particular multicast do not receive the information transmitted by the participants of the multicast. In this way, the multicast communication uses only the network components (*e.g.*, switches and trunks) actually needed for the multicast transmission.

20 With conventional techniques, a switch transmits multicast packets on all of its ports. Fig. 1 is a block diagram of an Ethernet switch 100. The Ethernet switch 100 includes a switching mechanism 110 connected to several port interface controllers 120. The switching mechanism 110 may include any conventional shared medium, shared memory, or space-division device.

25 The port interface controller 120 transports packets between a port and the switching mechanism 110. Fig. 2 is a block diagram of a conventional port interface controller 120. The port interface controller 120 includes an Ethernet interface 210, a Media Access Control (MAC) address learner 220, a MAC address table 230, a MAC address matcher 240, and a controller 250.

30 The Ethernet interface 210 receives an inbound packet transmitted from a host attached to the switch 100 over an Ethernet channel, and passes the packet to the MAC address learner 220. The MAC address learner 220 reads the Ethernet

address of the originating host and stores the address in the MAC address table 230, if the address is not already stored there. The controller 250 maintains addresses of attached hosts in the MAC address table 230. If a host has not transmitted for a certain period of time, the controller 250 removes the MAC address of the host from
5 the MAC address table 230.

The MAC address learner 220 sends the received packet to the controller 250, which, in turn, forwards the packet to the switching mechanism 110 (Fig. 1). The switching mechanism 110 identifies the appropriate port interface controller 120 for the packet and sends it to this controller.

10 The controller 250 of the identified port interface controller 120 receives the packet and sends it to the MAC address matcher 240. The MAC address matcher 240 compares the packet's destination address to the MAC addresses stored in the MAC address table 230. If the outbound packet's destination address is of a broadcast or multicast type, or is unicast and matches an entry in the MAC address
15 table 230, the MAC address matcher 240 forwards the packet to its port for transmission to the attached host(s).

Conventional Ethernet switches do not support Internet Protocol (IP) multicasting. IP multicasting first found its application in audio and video conferencing. Each IP multicast group has a unique class-D IP address ranging from
20 224.0.0.1 to 239.255.255.255. Multicast data is sent to a group based on this unique address. For an IP multicast packet transmitted on the Ethernet, the multicast Ethernet address includes the least significant 23 bits of the IP address.

To join a particular IP multicast group, a host sends a "request to join" message to the nearest multicast-capable router to request receiving from the
25 multicast group. The router propagates the request up to the multicast source if the data path is not already in place. Upon receiving an IP packet for this group, the router maps the class-D IP multicast group address into an Ethernet multicast address, and sends the resultant Ethernet group to the Ethernet port from which the original request was received.

30 The current Internet Group Management Protocol (IGMP) does not specify an explicit message for withdrawing membership from the multicast group. A host's

membership expires when the router does not receive a periodic membership report from the host.

Some conventional Ethernet switches route IP packets and, therefore, support IP multicasting. However, despite the words "Ethernet switch" in their names, they
5 are, in fact, not Ethernet switches, but routers because the actual switching is carried out at the IP or the network layer, not at the Ethernet or data link layer.

Consequently, their performance and cost are comparable with routers and much higher than that of Ethernet switches due to the overhead of IP decoding and routing.

No multicast mechanism currently exists at the data link layer, corresponding
10 to the IP layer. As a result, a conventional Ethernet switch sends multicast packets to all ports regardless of whether the hosts attached to these ports need them. It is the responsibility of the receiving host to determine whether a packet belongs to the groups that have been requested at the IP layer. The receiving host makes this determination through the use of an address filter either implemented in hardware on
15 a Local Area Network (LAN) controller or in communication control software.

Thus, the current technologies for switching multicast packets at the data link layer are inefficient and will render the network inoperable when several hosts receive different multicast streams concurrently. To illustrate this problem, suppose that a conventional Ethernet switch includes one network port connecting a multicast
20 router at 100 Mb/s and 24 local ports each connecting to a personal computer (PC) at 10 Mb/s. Assume that each of the PCs is tuned to a different video channel via IP multicasting, and each video stream is coded at 1.5 Mb/s. Since the Ethernet switch sends each multicast packet to all ports, the switch incurs a load of $24 \times 1.5 = 36$ Mb/s at each local port, far exceeding the PCs' 10 Mb/s capacity. As a result, the
25 network becomes bogged down by congestion.

Therefore, a need exists to overcome the deficiencies of the conventional data link layer switches.

Disclosure of Invention

Systems and methods consistent with the principles of the present invention
30 address this need by providing true multicasting at an Ethernet switch by sending packets to only those ports having attached hosts that have previously joined the

multicast group. In this way, the need for expensive IP decoding and routing is eliminated.

In accordance with the purpose of the invention as embodied and broadly described herein, a system consistent with the present invention includes a switching
5 mechanism coupled to a plurality of port interface controllers. Each of the port interface controllers comprises a multicast address table, a multicast matcher, and a timer. The multicast address table stores multicast addresses for hosts attached to the port interface controller. The multicast matcher matches an incoming packet to a target pattern, generates a multicast address from the incoming packet if the
10 incoming packet matches the target pattern, and stores the generated multicast address in the multicast address table. The timer determines an amount of time that the generated multicast address remains in the multicast address table.

Brief Description of Drawings

15 In the drawings,
Fig. 1 is a block diagram of an Ethernet switch;
Fig. 2 is a block diagram of a port interface controller in the switch of Fig. 1;
Fig. 3 is an Ethernet switch consistent with the present invention;
Fig. 4 is a flowchart of multicast processing of an incoming packet
20 performed by the port interface controller 300 of Fig. 3; and
Fig. 5 is a flowchart of multicast processing of an outgoing packet performed by the port interface controller 300 of Fig. 3.

Best Mode for Carrying Out the Invention

25 The following detailed description of the invention refers to the accompanying drawings. The same reference numbers in different drawings identify the same or similar elements. Also, the following detailed description does not limit the invention. Instead, the scope of the invention is defined by the appended claims.

Systems and methods consistent with the present invention support true
30 multicasting in a data link layer switch, such as an Ethernet switch, by transmitting packets on only those ports with hosts that have previously joined the multicast

group. This Ethernet switch multicasting reduces network traffic load and allows the Ethernet switch to be used for a large number of multicast streams. In particular, this switch can be used to provide near video-on-demand and broadcast video services.

5 Referring to Fig. 1, an Ethernet switch consistent with the present invention contains a standard configuration. That is, the Ethernet switch includes a switching mechanism 110 coupled to several port interface controllers 120. The elements comprising the port interface controllers 120 differ, however, from those elements shown in Fig. 2.

10 Fig. 3 is a diagram of a port interface controller 300 consistent with the present invention. The port interface controller 300 includes some of the same elements included in the conventional port interface controller 120 shown in Fig. 2. These elements contain the same reference numerals.

The port interface controller 300 includes an Ethernet interface 210, a MAC
15 address learner 220, a MAC address table 230, a controller 250, a multicast request matcher 310, a multicast Ethernet address table 320, a timer 330, and an address matcher 340.

The multicast request matcher 310 includes a standard bit-pattern matcher of fixed length whose target pattern is updatable through a software download to
20 support a variety of protocols and different versions. For the current version of IP and Ethernet protocols, the target pattern operates on the 13th, 25th, and 35th bytes of the Ethernet packet (*i.e.*, the protocol-type field of the Ethernet and IP headers, and the version and type fields of the IGMP header). The matcher 310 finds a match, for example, if the protocol fields are IP and IGMP, and the byte for version
25 and type is 0x21.

The multicast address table 320 contains a list of Ethernet multicast addresses of hosts attached to the Ethernet channel. The table is similar in construction and format to the MAC address table 230. The timer 330 is a counting mechanism implemented in hardware, using a conventional counter, or software,
30 using a conventional counting algorithm. The address matcher 340 includes a

mechanism for comparing a destination address to the addresses stored in the MAC address table 230 and the multicast address table 320.

Fig. 4 is a flowchart of multicast processing of an incoming packet performed by the port interface controller 300 of Fig. 3. The Ethernet interface 210 receives an incoming packet from a host attached to an Ethernet channel [step 410]. The interface 210 sends the packet to the MAC address learner 220 and the multicast request matcher 310 [step 420].

The MAC address learner 220 reads the Ethernet address of the originating host and stores the address in the MAC address table 230, if the address is not already stored there [step 430]. The controller 250 maintains addresses of attached hosts in the MAC address table 230. If a host has not transmitted for a certain period of time, the controller 250 removes the MAC address of the host from the MAC address table 230.

The multicast request matcher 310 determines whether the packet matches a target pattern [step 440]. For example, the matcher 310 may compare the packet to a target pattern that operates on the 13th, 25th, and 35th bytes of the packet. If the packet matches the target pattern [step 450], the matcher 310 performs no decoding, but simply extracts the 37th through 39th bytes of the packet and sets the highest bit to zero. The matcher 310 then appends these bytes to an Ethernet multicast address header, such as the Ethernet multicast header 0x01005e000, and stores the newly formed Ethernet multicast address in the multicast address table 320 [step 460]. The Ethernet multicast address is derived according to an IP to Ethernet address mapping method defined in IETF RFC-1112.

The multicast matcher 310 stores the multicast address in the multicast address table 320 only if the packet is a request to join a multicast group. If the packet is, instead, a request to leave a multicast group, the multicast matcher 310 removes the multicast address from the multicast address table 320.

In addition to storing the multicast address, the matcher 310 starts or restarts the timer 330. The timer 310 is used to determine the length of time that an address remains in the multicast address table 320. An address is removed from the table upon the expiration of the timer.

If the packet does not match the target pattern [step 450] or does match and the multicast address has already been stored in the multicast address table 320 [step 460], the multicast address matcher 310 sends the packet to the controller 250 [step 470]. The controller 250 forwards the packet to the switching mechanism [step 480], ending the processing of the incoming packet by the port interface controller 300.

Fig. 5 is a flowchart of multicast processing of an outgoing packet performed by the port interface controller 300 of Fig. 3. The controller 250 receives an outgoing packet from the switching mechanism [step 510]. The controller 250 forwards the packet to the address matcher 340 [step 520].

The address matcher 340 determines whether the packet is of a unicast, multicast, or broadcast type [step 530]. If the packet is not a broadcast packet [step 540], the address matcher 340 compares the destination address to the addresses stored in the MAC address table 230 and the multicast address table 320 [step 550]. If the destination does not match an address in either of the tables [step 560], the address matcher 340 discards the packet [step 570].

If the packet is a broadcast packet [step 540] or is a unicast or multicast packet and the destination address matches an address in either the MAC address table 230 or the multicast address table 320 [step 560], the address matcher 340 transmits the packet on the Ethernet channel via the Ethernet interface 210 [step 580]. The transmission ends the processing of the outgoing packet by the port interface controller 300.

The systems and methods consistent with the present invention support multicasting in a data link layer switch by sending packets to only those ports having an attached host that has joined the multicast group

The foregoing description of preferred embodiments of the present invention provides illustration and description, but is not intended to be exhaustive or to limit the invention to the precise form disclosed. Modifications and variations are possible in light of the above teachings or may be acquired from practice of the invention. The scope of the invention is defined by the claims and their equivalents.

For example, the foregoing description was directed toward multicasting in an Ethernet switch. However, the present invention is not limited to an Ethernet switch. The foregoing description applies equally well to other data link layer switching mechanisms.

Claims:

1. In a data link layer switch having a switching mechanism coupled to a plurality of port interface controllers, each of the port interface controllers
- 5 comprising:
- a multicast address table that stores multicast addresses for hosts attached to the port interface controller;
 - a multicast matcher that matches an incoming packet to a target pattern, generates a multicast address from the incoming packet if the incoming packet
 - 10 matches the target pattern, and stores the generated multicast address in the multicast address table; and
 - a timer that determines an amount of time that the generated multicast address remains in the multicast address table.
- 15 2. The port interface controller of claim 1, wherein the multicast matcher includes
- means for determining whether the incoming packet is a packet requesting to join or leave a multicast group.
- 20 3. The port interface controller of claim 2, wherein the multicast matcher further includes
- means for storing the generated multicast address in the multicast address table when the incoming packet is a packet requesting to join the multicast group.
- 25 4. The port interface controller of claim 2, wherein the multicast matcher further includes
- means for removing the generated multicast address from the multicast address table when the incoming packet is a packet requesting to
 - 30 leave the multicast group.

5. The port interface controller of claim 1, wherein the multicast matcher includes means for comparing predetermined bytes of the incoming packet to the target pattern to determine if there is a match.
- 5 6. The port interface controller of claim 1, wherein the multicast matcher includes
- means for extracting predetermined bytes from the incoming packet,
- and
- means for appending the predetermined bytes to a multicast header to
- 10 form the generated multicast address.
7. The port interface controller of claim 1, wherein the multicast matcher includes
- means for starting the timer when the multicast address is stored in the
- 15 multicast address table.
8. The port interface controller of claim 1, wherein the timer includes
- means for removing the generated multicast address from the
- multicast address table upon expiration of the determined amount of time.
- 20
9. In a data link layer switch having a switching mechanism coupled to a plurality of port interface controllers, each of the port interface controllers including a multicast address table connected to a multicast matcher, a method for processing an incoming packet by the multicast matcher, the method comprising the steps of:
- 25 matching the incoming packet to a target pattern;
- generating a multicast address from the incoming packet if the incoming packet matches the target pattern; and
- storing the generated multicast address in the multicast address table.
- 30 10. The method of claim 9, wherein the storing step includes the substep of:

determining whether the incoming packet is a packet requesting to join or leave a multicast group.

11. The method of claim 10, wherein the storing step further includes the
5 substep of:

writing the generated multicast address in the multicast address table when the incoming packet is a packet requesting to join the multicast group.

12. The method of claim 10, wherein the storing step further includes the
10 substep of:

removing the generated multicast address from the multicast address table when the incoming packet is a packet requesting to leave the multicast group.

13. The method of claim 9, wherein the matching step includes the
15 substep of:

comparing predetermined bytes of the incoming packet to the target pattern to determine if there is a match.

14. The method of claim 9, wherein the generating step includes the
20 substeps of:

extracting predetermined bytes from the incoming packet, and
appending the predetermined bytes to a multicast header to form the
generated multicast address.

25

15. A data link layer switch comprising:
a switching mechanism; and
a plurality of port interface controllers coupled to the switching mechanism,
each of the port interface controllers including:

30 a multicast address table that stores multicast addresses for hosts
attached to the port interface controller,

a multicast matcher that matches an incoming packet to a target pattern, generates a multicast address from the incoming packet if the incoming packet matches the target pattern, and stores the generated multicast address in the multicast address table, and

5 a timer that determines an amount of time that the generated multicast address remains in the multicast address table.

16. The data link layer switch of claim 15, wherein the multicast matcher includes

10 means for determining whether the incoming packet is a packet requesting to join or leave a multicast group.

17. The data link layer switch of claim 16, wherein the multicast matcher further includes

15 means for storing the generated multicast address in the multicast address table when the incoming packet is a packet requesting to join the multicast group.

18. The data link layer switch of claim 16, wherein the multicast matcher further includes

means for removing the generated multicast address from the multicast address table when the incoming packet is a packet requesting to leave the multicast group.

19. The data link layer switch of claim 15, wherein the multicast matcher includes

means for comparing predetermined bytes of the incoming packet to the target pattern to determine if there is a match.

20. The data link layer switch of claim 15, wherein the multicast matcher includes

means for extracting predetermined bytes from the incoming packet,
and

means for appending the predetermined bytes to a multicast header to
form the generated multicast address.

5

21. The data link layer switch of claim 15, wherein the multicast matcher
includes

means for starting the timer when the multicast address is stored in the
multicast address table.

10

22. The data link layer switch of claim 15, wherein the timer includes
means for removing the generated multicast address from the
multicast address table upon expiration of the determined amount of time.

15

23. The data link layer switch of claim 15, wherein the port interface
controller further includes

a controller for forwarding the incoming packet to the switching
mechanism for routing to a destination.

20

24. In a data layer switching means having a plurality of port controlling
means, each of the port controlling means comprising:

means for storing multicast addresses for hosts attached to the port
controlling means;

means for matching an incoming packet to a target pattern;

25

means for generating a multicast address from the incoming packet if the
incoming packet matches the target pattern;

means for writing the generated multicast address in the storing means; and

means for determining an amount of time that the generated multicast
address remains in the storing means.

25. In a data link layer switch having a switching mechanism coupled to a plurality of port interface controllers, each of the port interface controllers comprising:

5 a multicast address table that stores multicast addresses for hosts attached to the port interface controller;

a MAC address table that stores unicast addresses for hosts attached to the port interface controller; and

10 an address matcher that matches a destination address of an outgoing packet to the multicast and unicast addresses stored in the multicast and MAC address tables, respectively, and transmits the outgoing packet to the destination address if the destination address matches an address stored in one of the multicast address table and the MAC address table.

26. The port interface controller of claim 25, wherein the address matcher
15 includes

means for determining whether the outgoing packet is a broadcast packet.

27. The port interface controller of claim 26, wherein the address matcher
20 further includes

means for transmitting the outgoing packet to the destination address regardless of whether the destination address matches an address in the multicast and MAC address tables if the outgoing packet is determined to be a broadcast packet.

25

28. The port interface controller of claim 25, wherein the address matcher includes

means for discarding the outgoing packet if the destination address fails to match an address in the multicast and MAC address tables.

30

29. In a data link layer switch having a plurality of port interface controllers, each of the port interface controllers including a multicast address table storing multicast addresses for hosts attached to the port interface controller, a MAC address table storing unicast addresses for hosts
5 attached to the port interface controller, and an address matcher, a method for processing an outgoing packet by the address matcher, comprising the steps of:

matching a destination address of the outgoing packet to the multicast and unicast addresses stored in the multicast and MAC address tables, respectively; and
10 transmitting the outgoing packet to the destination address if the destination address matches an address stored in one of the multicast address table and the MAC address table.

30. The method of claim 29, further comprising the step of:
15 determining whether the outgoing packet is a broadcast packet.

31. The method of claim 30, further comprising the step of:
transmitting the outgoing packet to the destination address regardless of whether the destination address matches an address in the multicast and MAC
20 address tables if the outgoing packet is determined to be a broadcast packet.

32. The method of claim 29, wherein the transmitting step includes the substep of:
discarding the outgoing packet if the destination address fails to
25 match an address in the multicast and MAC address tables.

33. A data layer switch comprising:
a switching mechanism; and
a plurality of port interface controllers, each of the port interface controllers
30 including

a multicast address table that stores multicast addresses for hosts attached to the port interface controller,

a MAC address table that stores unicast addresses for hosts attached to the port interface controller, and

5 an address matcher that matches a destination address of an outgoing packet to the multicast and unicast addresses stored in the multicast and MAC address tables, respectively, and transmits the outgoing packet to the destination address if the destination address matches an address stored in one of the multicast address table and the MAC address table.

10

34. The data layer switch of claim 33, wherein the address matcher includes

means for determining whether the outgoing packet is a broadcast packet.

15

35. The data layer switch of claim 34, wherein the address matcher further includes

means for transmitting the outgoing packet to the destination address regardless of whether the destination address matches an address in the multicast and MAC address tables if the outgoing packet is determined to be a broadcast packet.

20

36. The data layer switch of claim 33, wherein the address matcher includes

25 means for discarding the outgoing packet if the destination address fails to match an address in the multicast and MAC address tables.

37. In a data layer switching means having a plurality of port controlling means, each of the port controlling means comprising:

30 means for storing multicast addresses for hosts attached to the port controlling means;

means for storing unicast addresses for hosts attached to the port controlling means;

means for matching a destination address of an outgoing packet to the stored multicast and unicast addresses; and

5 means for transmitting the outgoing packet to the destination address if the destination address matches an address stored in one of the storing means.

38. A data link layer switch comprising:

a switching mechanism; and

10 a plurality of port interface controllers coupled to the switching mechanism, each of the port interface controllers including:

a multicast address table that stores multicast addresses for hosts attached to the port interface controller,

15 a MAC address table that stores unicast addresses for hosts attached to the port interface controller,

a multicast matcher that matches an incoming packet to a target pattern, generates a multicast address from the incoming packet if the incoming packet matches the target pattern, and stores the generated multicast address in the multicast address table, and

20 an address matcher that matches a destination address of an outgoing packet to the multicast and unicast addresses stored in the multicast and MAC address tables, respectively, and transmits the outgoing packet to the destination address if the destination address matches an address stored in one of the multicast address table and the MAC address table.

100

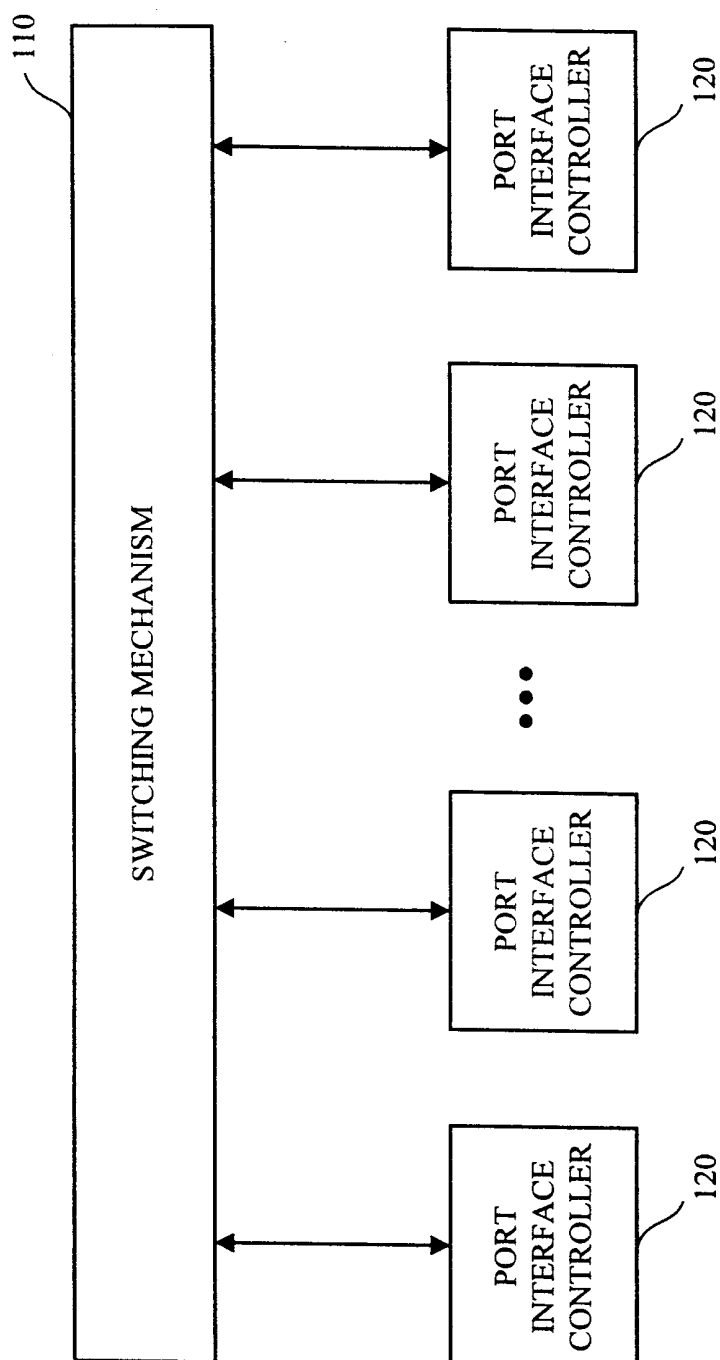


FIG. 1

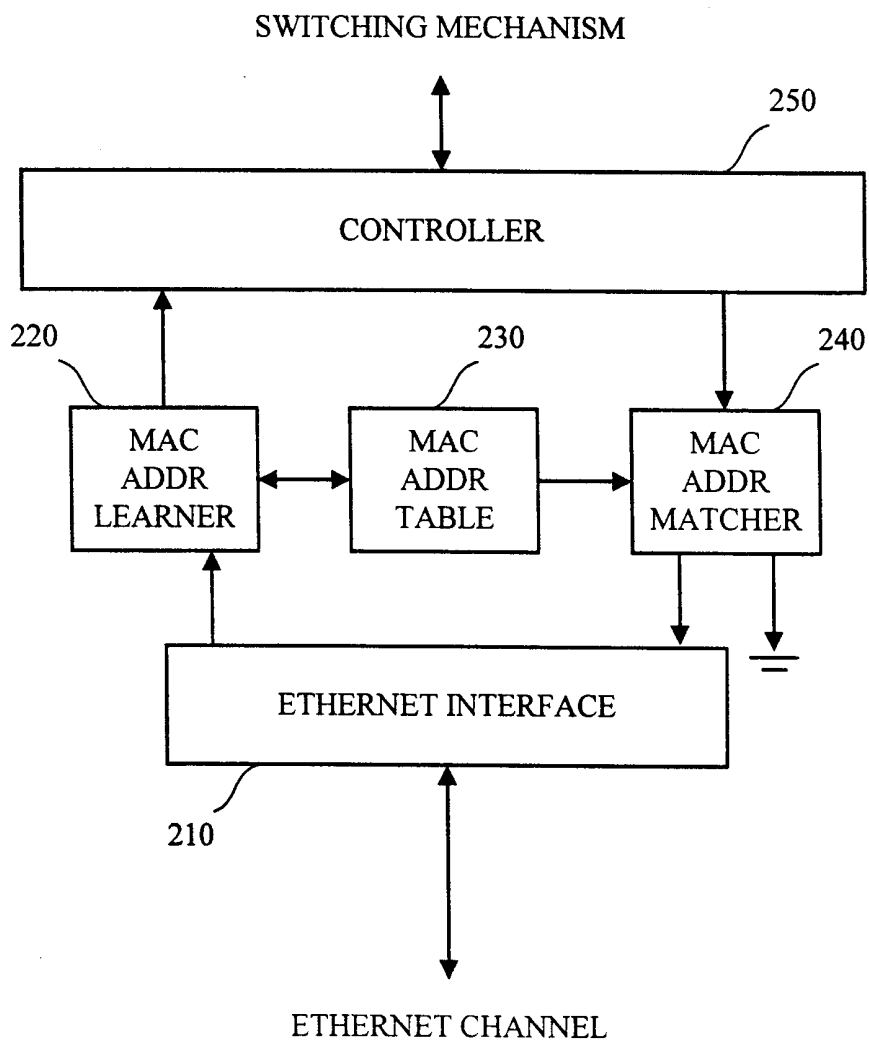

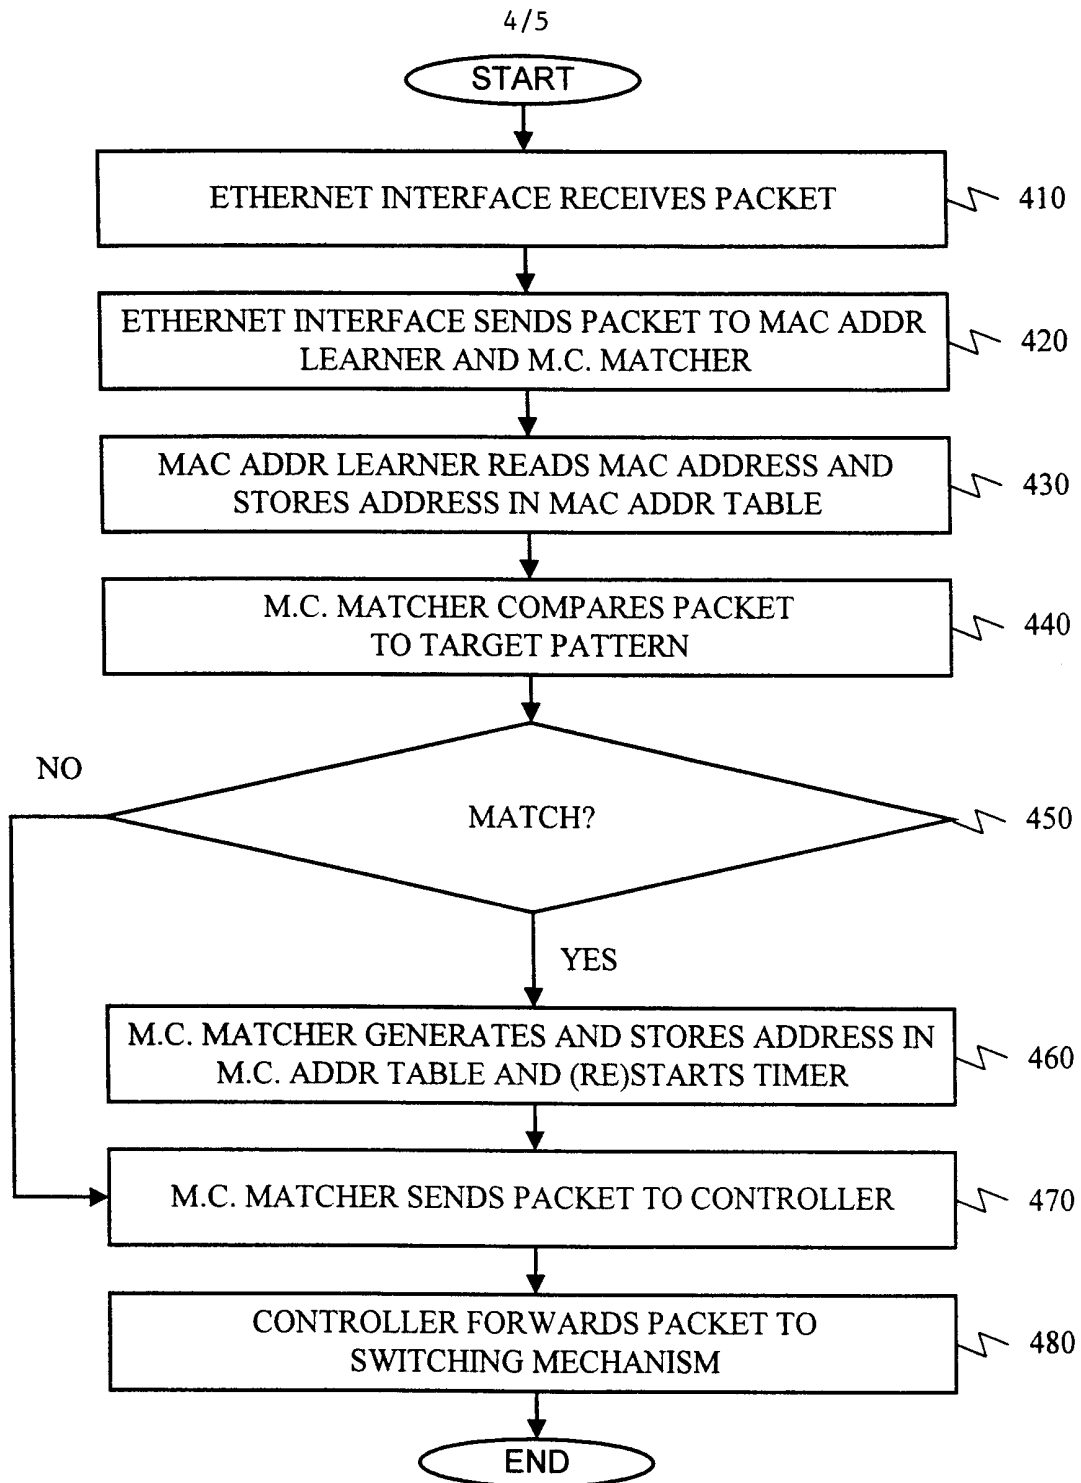
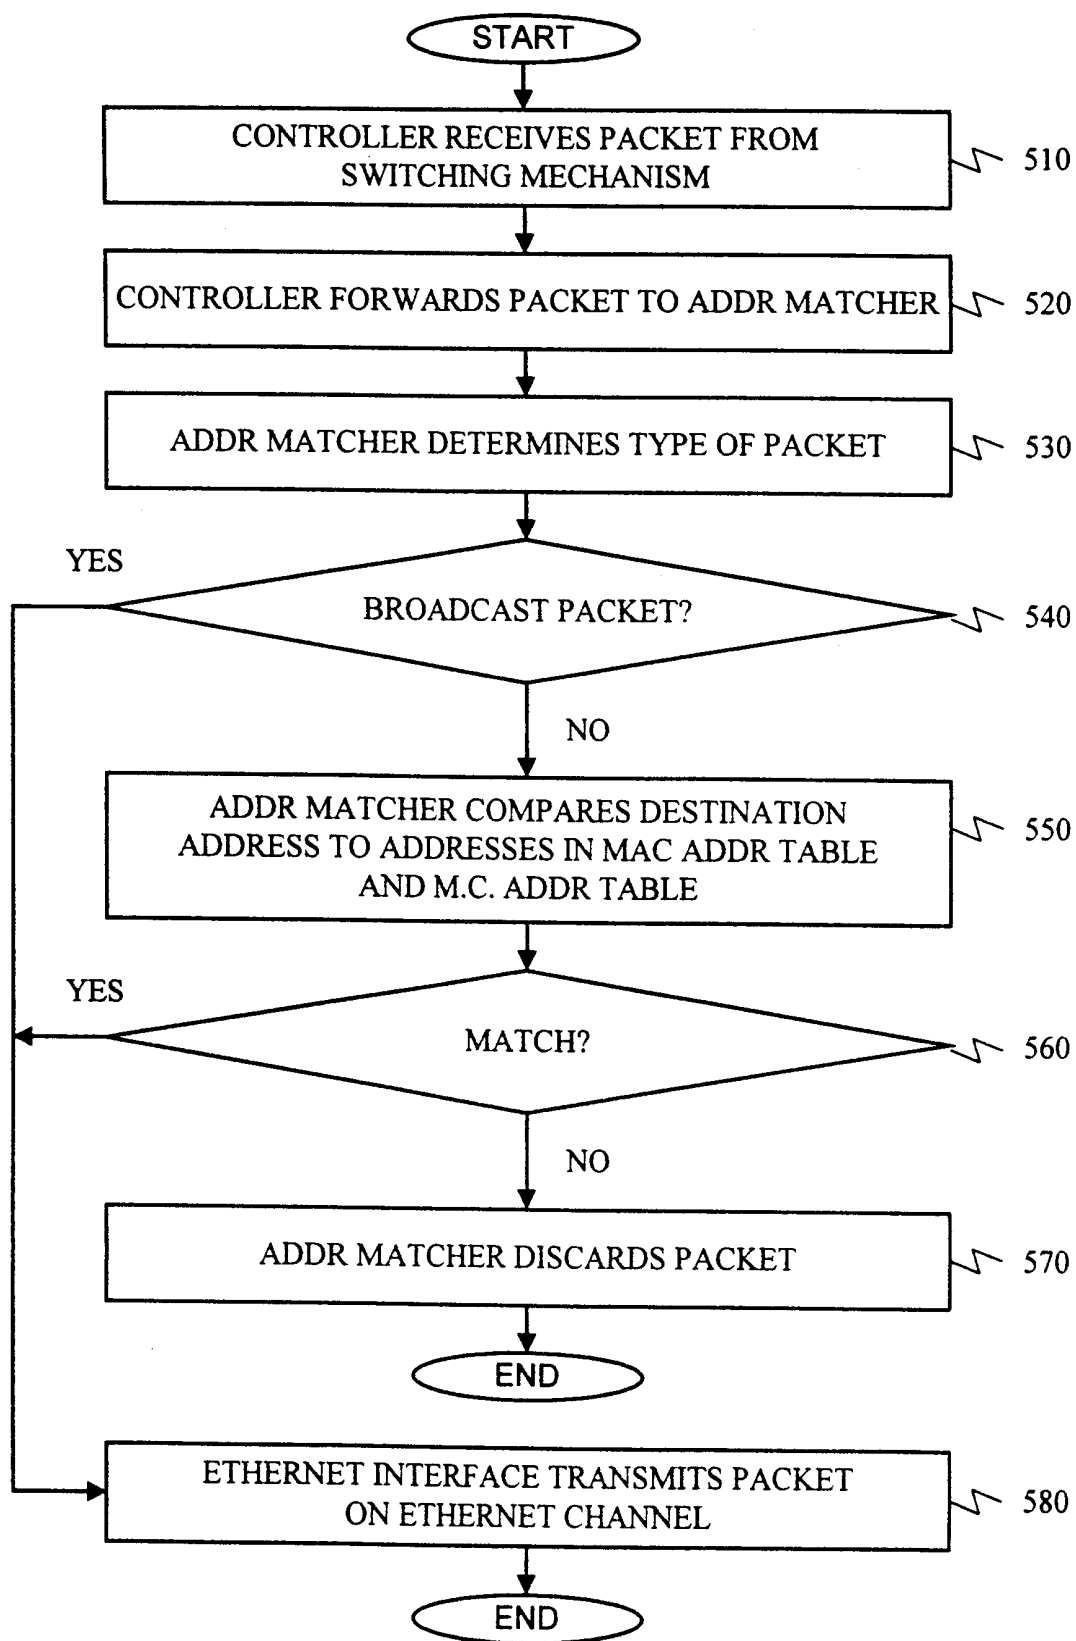
120

FIG. 2
(PRIOR ART)

300**FIG. 3**

**FIG. 4**

**FIG. 5**

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US00/05390

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : H04J 3/26

US CL : 370/432

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. : 370/432, 250,252, 340, 351, 356, 359, 360, 367, 379, 382, 389, 390, 392, 396, 400, 428, 445

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)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P ----- Y,P	US 5,959,989 A (GLEESON et al) 28 September 1999, see Figs. 2A-2C and 3, col. 8, lines 20-29, col. 9, lines 57-64, and col. 14, lines 14-36.	25, 29, 33, and 37-38 ----- 26-28, 30-32, and 34-36
Y	US 5,530,703 A (LIU et al) 25 June 1996, see Figs. 2 and 4A-4C, col. 5, lines 18-67, col. 5, lines 48-53, 55-59, col. 6, lines 1-22, col. 9, lines 11-67, and col. 10, lines 1-56.	1-24, 26-28, 30-32, and 34-36
Y	US 5,835,723 A (ANDREWS et al) 10 November 1998, see Figs. 10-16, col. 23, lines 4-67, col. 24, lines 1-67, and col. 25, lines 1-58.	1-24

☐ 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

06 JUNE 2000

Date of mailing of the international search report

06 JUL 2000

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

Facsimile No. (703) 305-3230

Authorized officer

PHIRIN SAM

Telephone No. (703) 308-9294