



- (51) International Patent Classification:
H04L 12/743 (2013.01)
- (21) International Application Number:
PCT/US2013/036423
- (22) International Filing Date:
12 April 2013 (12.04.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/492,383 8 June 2012 (08.06.2012) US
- (71) Applicant: CISCO TECHNOLOGY, INC. [US/US];
SJC/10/2/1, 170 West Tasman Drive, San Jose, California
95134-1706 (US).
- (72) Inventors: RAJAGOPALAN, Santosh; 6061 Burchell
Ave, San Jose, California 95120 (US). SANE, Sanjay;
3862 Burton Cmn, Fremont, California 94536 (US).
TRACY, Leonard T.; 5250 Keystone Drive, Vancouver,
Washington 98661 (US). BANERJEE, Ayan; 5250 Key-
stone Drive, Fremont, California 94536 (US).
- (74) Agent: FRAME, Thomas J.; Patent Capital Group, 2816
Lago Vista Lane, Rockwall, Texas 75032 (US).

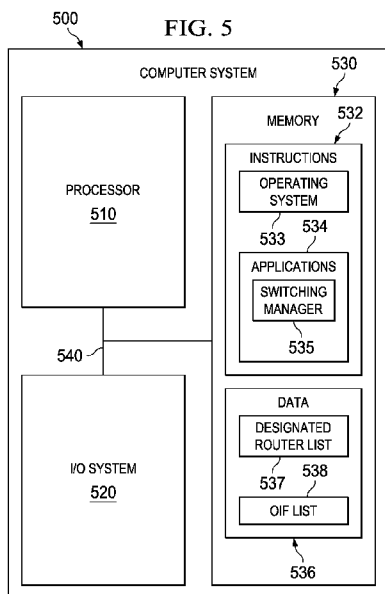
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR LAYER-2 MULTICAST MULTIPATHING



(57) Abstract: An example method is provided and includes a multicast data message from a data source, the message in a first virtual local area network and being associated with a multicast group. The method also includes calculating a hash value based on the virtual local area network, the data source, and the multicast group, determining a port for a designated router in a Layer-2 network based on the hash value, and switching the multicast data message to the port that was determined.

WO 2013/184236 A2



— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

SYSTEM AND METHOD FOR LAYER-2 MULTICAST MULTIPATHING

TECHNICAL FIELD

[0001] This disclosure relates in general to the field of communications and, more particularly, to a system and a method for Layer-2 multicast multipathing.

BACKGROUND

[0002] Ethernet architectures have grown in complexity in recent years. This is due, at least in part, to diverse technologies that have emerged to accommodate a plethora of end users. For example, Data Center Ethernet (DCE) represents an extension to Classical Ethernet (CE), and it can offer a lower cost, lower latency, high-bandwidth configuration. The forwarding methodology adopted by networks is generally scalable and, further, provides forwarding paths with equal-cost multipathing with support for different forwarding topologies.

[0003] Layer-2 link state routing protocols can be used in data centers to overcome the drawbacks of the Spanning Tree Protocol (STP). Transparent Interconnect of Lots of Links (TRILL) from the IETF and Fabric Path (from Cisco Systems, Inc. of San Jose, CA) and are examples of such Layer-2 link state routing techniques. The link state protocols allow for better use of network resources by calculating shortest path forwarding trees to the nodes in the network. Unicast traffic can be forwarded along multiple equal cost paths if available. In certain network scenarios, topology information may not be current, accurate, and/or consistent. Hence, optimally managing network topologies presents a significant challenge to system designers, network operators, and service providers alike.

BRIEF DESCRIPTION OF THE DRAWINGS

[0004] To provide a more complete understanding of the present disclosure and features and advantages thereof, reference is made to the following description, taken in conjunction with the accompanying figures, wherein like reference numerals represent like parts, in which:

[0005] FIGURE 1 is a simplified block diagram of an example system for multicast multipathing in a Layer-2 network in accordance with one embodiment of the present disclosure;

[0006] FIGURE 2 is a simplified block diagram of an access switch for multicast multipathing in a Layer-2 network in accordance with one embodiment of the present disclosure;

[0007] FIGURE 3 is a simplified flow diagram of an example process for multicast multipathing in a Layer-2 network in accordance with one embodiment of the present disclosure;

[0008] FIGURE 4 is a simplified flow diagram illustrating another example process for multicast multipathing in a Layer-2 network in accordance with one embodiment of the present disclosure; and

[0009] FIGURE 5 is a simplified block diagram of a network node multicast multipathing in a Layer-2 network.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

OVERVIEW

[0010] A method is provided in one example embodiment and includes receiving, for example at an access switch of a Layer-2 network, a multicast data message from a data source, the message being in a first virtual local area network and being associated with a multicast group. The method also includes calculating a hash value based on the virtual local area network, the data source, and the multicast group, determining a port for a designated router in the Layer-2 network based on the hash value, and switching the multicast data message to the port that was determined.

EXAMPLE EMBODIMENTS

[0011] FIGURE 1 illustrates an example system 100 for Layer-2 multicast multipathing in accordance with one embodiment of the present disclosure. System 100 includes a number data sources 110, data receivers 120, access switches 130, and switch-routers 140. Data sources 110 may be any logical devices for presenting, storing, and/or processing data. For example, data sources 110 may be personal computers, laptops, servers, mobile devices, and/or tablet computers. In certain implementations, one or more of data sources 110 may be virtual machines. Data sources 110 typically need to communicate data to other logical devices

and, thus, are coupled to one or more communication networks. Data sources 110 may be in the same or different virtual local area networks (VLANs). Data sources 110 may also receive data. Thus, they are only labeled as data sources for the ease of discussion.

[0012] Similarly, data receivers 120 may be any logical devices for presenting, storing, and/or processing data. For example, data receivers 120 may be personal computers, laptops, servers, mobile devices, and/or tablet computers. Data receivers 120 typically need to receive data from other logical devices and, thus, are coupled to one or more communication networks. In certain implementations, one or more of data receivers 120 may be virtual machines. Data receivers 120 may be in the same VLAN as one or more of data sources 110 or in different VLANs. Data receivers 110 may also source data. Thus, they are only labeled as data receivers for the ease of discussion.

[0013] Access switches 130 and switch-routers 140 are communicatively coupled to data sources 110 and data receivers 120 and provide switching and routing between them. Access switches 130 and switch-routers 140 together can form a Layer-2 network 180. In particular implementations, network 180 may be a Layer-2 multipath network (e.g., a Fabricpath network). Access switches 130 and switch-routers 140 may use a forwarding paradigm that provides Layer-2 multipathing capability. Access switches 130 and switch-routers 140 may also provide the ability to scale Layer-2 networks to a large number of switches and/or routers (e.g., Fabricpath).

[0014] A Fabricpath network may, for example, be part of an enterprise network or a data center, which could, for example, also include a number of servers, databases, and/or other devices for storing and/or processing data. A data center could also include more communication networks. A communication network is typically a series of points or nodes of interconnected communication paths for receiving and transmitting messages. As used herein, the term “network node” is meant to encompass switches, routers, proxys, gateways, bridges, load balancers, firewalls, inline service nodes, proxies, servers, processors, modules, or any other suitable device, component, element, or object operable to exchange information in a network environment. A network node may include any suitable hardware, software, components, modules, interfaces, or objects that facilitate the operations thereof. This may be inclusive of appropriate algorithms and communication protocols that allow for the effective exchange (reception and/or transmission) of data or information.

[0015] A data center's communication network(s) may offer a communicative interface between network elements (e.g., switches, bridges, gateways, etc.) and may be any IP network, local area network (LAN), virtual LAN (VLAN), wireless LAN (WLAN), metropolitan area network (MAN), wide area network (WAN), extranet, Intranet, virtual private network (VPN), or any other appropriate architecture or system that facilitates communications in a network environment. The networks can support a transmission control protocol (TCP)/IP, or a user datagram protocol (UDP)/IP in particular embodiments of the present disclosure; however, these networks may alternatively implement any other suitable communication protocol for transmitting and receiving data packets within system 100.

[0016] Access switches 130 and switch-routers 140 may use a link state routing (e.g., intermediate system to intermediate system (IS-IS) or Fabric Shortest Path First (FSPF)) for forwarding purposes, whereas classic Ethernet networks commonly use a spanning tree protocol (STP) as their forwarding protocol. Link state protocols generally work at the Layer-2 broadcast domain. Link state routing is a protocol that allows a node in a network to determine network topology by sharing information about a transmission cost to each of its neighboring nodes. Link state routing messages are transmitted to and received from neighbors. The least expensive path to various destinations can be determined using the link state information.

[0017] Link state information can be used to generate network topology information at various network nodes for creating forwarding tables. The forwarding tables allow network nodes (such as switches, bridges, and routers) to forward the received traffic on an appropriate output interface. In order to generate a network topology map and a forwarding table at a specific network node, link state information is distributed from various network nodes. Each network node is configured to create a link state message having information about the distance, delay, or cost to each of its neighbors. A link state record (LSR) can then be transmitted to neighboring nodes.

[0018] Addressing and forwarding can include the use of a locally assigned aggregatable (or hierarchical) Media Access Control (MAC) address for forwarding activities. A link state protocol can be used to determine the forwarding topology and, further, to support shortest path and equal cost multi-path forwarding for unicast traffic. Multicast frames can also readily use multipathing, albeit using a slightly different scheme. Additionally, a single control protocol can be used to compute unicast paths and multi-destination (e.g., multicast or

broadcast) trees. Broadcast techniques can be viewed as a special case of multicast, in which all nodes are interested in a message.

[0019] In the illustrated example, three access switches 130 are shown, but any number may be used in other implementations. Access switches 130 receive messages (e.g., packets) from data sources 110 and switch them to data receivers 120, assuming they are in the same VLAN, and switch-routers 140. Access switches 130 may also receive messages from switch-routers 140 (e.g., from data sources outside network or from data sources in different VLANs) and switch them to data receivers 120.

[0020] Switch-routers 140 provide a gateway between the Layer-2 network and a Layer-3 network 150. Switch-routers 140 forward (or cooperate with each other in order to forward) messages (e.g., packets) in a network environment by switching and routing messages. Switch-routers 140 may be integrated switch-routers, switches that have routing capability, or routers that have switching capability. Switch-routers 140 may use a protocol such as Protocol-Independent Multicast (PIM) for forwarding multicast messages. PIM is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, a WAN, or the Internet. It is termed protocol-independent because it does not include its own topology discovery mechanism, instead relying on routing information supplied by other routing protocols, such as Routing Information Protocol, Open Shortest Path First, Border Gateway Protocol, and Multicast Source Discovery Protocol. For control plane interactions, a protocol such as Fabricpath IS-IS may be used as the underlying transport mechanism.

[0021] Layer-3 network 150 may be any appropriate type of communication network that uses Layer-3 protocols. Network 150 may, for example, include a number of switches, routers, bridges, repeaters, and/or other equipment for conveying information. Network 150 may, for example, be a wide area network (WAN) or the Internet. In particular implementations, network 150 may interconnect data centers and be a data center interconnect (DCI). Layer-3 network 150 is also communicatively coupled with a router 160, which may, for example, be a switch-router. Router 160 is coupled to a data receiver 170. Data receiver 170 may be similar to data receivers 120, except that it is located across Layer-3 network 150 from data sources 110.

[0022] In certain modes of operation, access switches 130 and switch-routers 140 may use a common hash sequence for forwarding multicast messages. The hash sequence may, for example, be shared between access switches 130 and switch-routers 140 by using a sub-TLV (Type, Length, Value) in the IS-IS router-capability TLV. A hash may, for instance, be performed on a {VLAN, Source, Group} 3-tuple, where the source is the source of the multicast data and the group is the multicast group for which the source is sending data. This should, in most instances, provide a unique identifier for each 3-tuple. This identifier may then be used to retrieve a tree identifier (i.e., to determine which tree to forward on) and a tag for a designated router (i.e., one of switch-routers 140). Thus, the tree and the designated router may be linked. The tree identifier may, for example, be a forwarding tag (Ftag) in Fabricpath. This should result in a disjoint set of {VLAN, Source, Group} (the hash outcomes) being assigned to each tree and available router. An access switch 130 receiving a multicast message may then switch the message to any appropriate receivers 120 and to the appropriate one of switch-routers 140.

[0023] The tree identifier and the designated router associated with a {VLAN, Source, Group} 3-tuple may be determined by an access switch in a number of ways. For example, the tree identifier and designated router may be stored in an outgoing interface (OIF) list and be referenced by the hash value of the 3-tuple. The port that leads to the designated switch-router 140 may be made part of the OIF list. Thus, the access switch can determine which ports to forward messages on. The receiving switch-router 140 may also be aware of the hash and use it to determine the groups for which it will perform Layer-3 multicast forwarding. The receiving switch-router 140 may also determine destination interfaces (e.g., ports) based on the hash. It may, for example, examine an OIF list.

[0024] Note that it is possible to run designated routers in active/standby mode. Two designated routers that are in active/standby mode may have the same index in an ordered list of designated routers, therefore allowing both to receive the same traffic. In addition, this implies that PIM on these routers is aware of the fact that it can be the designated router at a {VLAN, Source, Group} granularity.

[0025] A message may, for example, need to be routed to one of receivers 120 if it is in a different VLAN than the data source 110. Additionally, a message may need to be routed

into Layer-3 network 150 if the message is destined for a data receiver that is not part of the Layer-2 network (e.g., data receiver 170).

[0026] To establish multicast operations, access switches 130 may detect that a data receiver 120 wants to join a multicast group. Access switches 130 may, for example, accomplish this using Internet Group Management Protocol (IGMP) snooping (or by using any other appropriate protocol). When a data receiver 120 desiring to receive data for a new group is first detected by one of access switches 130, the access switch will compute the {VLAN, Source, Group} hash and choose a tree identifier (e.g., an Ftag) and one of switch-routers 140 as the designated router.

[0027] Existing hardware may pick the tree identifier using the {VLAN, Source, Group} 3-tuple. A tree identifier may, for example, be selected based on the root of the tree, for which it is typically beneficial to be near the source. The tree identifier does not need to form part of the key for the forwarding lookup, however. It may be used for Incoming Interface Check (IIC) and Color Blocking Logic (CBL) based tree enforcement. Note that the hardware result for a {VLAN, Source, Group} lookup can be programmed with the ports for the tree identifier resulting from the ({VLAN, Source, Group} hash. This can mean that the entries for the other tree identifiers do not need to be programmed or computed.

[0028] For hardware in which the forwarding to router ports and to the receivers is done using an (effectively) single lookup, software may build an OIF list that is the union of the links to reach the appropriate data receivers and the link to reach the designated router. However, since the access switch is pinning a group to a router, only the port that leads to that switch-router needs to be part of the OIF list (in addition to the ports that lead to data receivers 120).

[0029] For other hardware, a two-step lookup may be used. First, the {VLAN, Source, Group} hash may be implemented, which may provide the OIF list consisting of the data receivers as well as a router_tag. Then, a {VLAN, Router_tag} lookup may be performed that provides the list of router ports. The router_tag may, for example, be: 1) an indirection to the designated switch-router for that {VLAN, Source, Group}; 2) an indirection to the list of available routers, which may happen during router addition or removal; or 3) an indirection to a set of routers that have been setup in active/standby mode.

[0030] Hardware can also be enhanced to provide the tree identifier as a result of the forwarding {VLAN, Source, Group} lookup (in addition to the router_tag and the OIF list). This can be used in situations where the hash results in a poor choice of tree identifiers, which may be configuration driven. The designated router may, for example, be selected by performing a modulo operation on the hash value against the tree identifier. In certain implementations, the distance of the designated router from the root of the tree may also be considered, closer designated routers typically being more preferred. The join request may be sent to the designated router, which may build the forwarding tree for the group. Additionally, the rest of the network may be informed of the join request. For example, in a Fabricpath network, the IS-IS protocol on an access will advertise to the rest of the network that it has received a join request for the {VLAN, Group} set in its link state protocol data unit (PDU) packets.

[0031] When one of access switches 130 learns of removal of a designated switch-router 140, the access switch may briefly change the router_tag for that designated switch-router to point to the ports leading to other available switch-routers 140. Then, the multicast groups that had been assigned to the removed switch-router need to be reassigned to the remaining switch-routers 140. Thus, the mapping of the hash outcome to the designated router should be updated (e.g., redistributed among the rest of the designated routers).

[0032] For example, the modulo operation performed on the hash value against the tree identifier may be adjusted to account for the fewer number of routers, and this may be used to reassign the groups. For instance, the designated routers may be sorted by their MAC address into a list [router 0, ..., router X-1], and the HASH value output may be processed by MOD by X, the number of routers. This may not always provide the most optimal forwarding through the network, but if PIM and Layer-2 agree, it should be functional. In certain implementations, the distance of the designated router from the root of the Layer-2 multicast tree may also be considered, closer designated routers typically being more preferred. Thus, the designated router(s) closest to the root of the Layer-2 multicast tree may be associated with tree (e.g., the same hash values that choose tree X may choose the designated router nearest to the root of that tree), which may provide enhanced forwarding.

[0033] As another example, the access switches may know how many switch-routers 140 are available, and the highest source Media Access Control (MAC) address may take the first tree number. Switches in a Layer-2 network may, for example, use the receipt of a PIM

hello packet to detect the presence of a router in the network. After a settling time, the groups that were previously assigned to the removed designated switch-router can then be reassigned to the remaining router-tags. The redistribution algorithm may be used by the access switches and the switch-routers so that hash values are tied to designated routers.

[0034] When a new designated switch-router is available, the learned groups are redistributed among the new list of designated switch-routers. Thus, a new designated router joining the network would take some subset of values from each of the already present designated routers such that the hash outcomes are evenly distributed among the new set of routers, which should result in the mapping of the hash outcomes to the designated routers being updated (e.g., redistributed among the designated routers).

[0035] For example, the modulo operation performed on the hash value against the tree identifier may be adjusted to account for the larger number of routers, and this may be used to reassign the groups. In certain implementations, the distance of the designated router from the root of the tree may also be considered, closer designated routers typically being more preferred. The Layer-2 forwarding protocol and the PIM may interact to determine the router location with respect to the tree root. Thus, the designated router(s) closest to the root of the Layer-2 multicast tree may be associated with tree (e.g., the same hash values that choose tree X will choose the designated router nearest to the root of that tree), which will provide enhanced forwarding. As another example, the access switches may know how many switch-routers 140 are available, and the highest source Media Access Control (MAC) address may take the first tree number. Switches in a Layer-2 network may, for example, use the receipt of a PIM hello packet to detect the presence of a router in the network. After a settling time, the groups that were previously assigned to the removed designated switch-router can then be reassigned to the remaining router-tags. The redistribution algorithm may be used by the access switches and the switch-routers so that hash values are tied to designated routers.

[0036] To redistribute the groups, there may be a communication between PIM and the Layer-2 forwarding protocol (e.g., FabricPath) to assign the designated router to a tree identifier. The tree identifier that a group is assigned to provides the designated router also. A change in the router_tag may be preceded by a settling period where the traffic is sent to all the available routers. In the network transition situations discussed above, it is possible that more than one switch-router 140 receives messages for a given {VLAN, Source, Group} set. The

PIM's assert mechanism may be used by the switch-routers before they begin forwarding to establish designated switch-router ownership for those {VLAN, Source, Group} sets. This may prevent duplicates.

[0037] The discussed scheme may need modifications from PIM for operation in cases in which in the Layer-2 network there are: a) multiple sources, but no receivers; and b) there are multiple sources, but the receivers only issue (*, g) joins. Basically, the Layer-2 network in these cases would be unaware of the various transmitting sources and/or receivers. For both cases, since the Layer-2 network is unaware of the complete {VLA, Source, Group} information and is therefore unable to compute the complete hash, forwarding would happen towards available switch-routers. In this case, access switches 130 may compute paths for the trees for each {VLAN, Group} combo.

[0038] However, PIM can receive the initial data message when the source sends data. Therefore, this problem can be solved if the PIM-designated router creates an IGMP-like {VLAN, Group, Source} state in the Layer-2 multipath network, possibly by redistribution to IS-IS. When this happens, hardware no longer finds the optimized multicast forwarding (OMF) entry (which is a list of outgoing interfaces going to multiple routers) or the generic (*, Group) entry. Instead, it finds the specific {VLAN, Source, Group} OIF list, which will point towards a specific router.

[0039] Certain embodiments of system 100 can offer any number of significant features. For example, the architecture can improve scalability as the number of multicast trees grows. Currently, an increasing number of trees would lead to greater processing needs in software and greater storage needs in hardware. For example, for N {S,G} groups, $M*N$ entries need to be computed and installed in the hardware, where M is number of trees. Installation of all these entries may be important because the same group's multicast message can travel on any of the M trees, based on what hash/tree identifier was chosen at the ingress. Thus, an OIF list should be used for each tree. However, in system 100, the scalability problem is addressed because it only needs to compute and program entries for a single tree identifier for that group since each {VLAN, Source, Group} should hash to the same entry.

[0040] Additionally, while there may exist more than one router on a VLAN, only a single one gets picked as the designated router for forwarding multicast packets on that VLAN for a group. (There may be multiple routers on a VLAN configured such that they are each a

designated router for a disjoint set of groups.) However, switches in a Layer-2 network use the receipt of a PIM Hello packet to detect the presence of a router in the network. Therefore, at the Layer-2 network level, switches are not aware of the designated router election, and they consider all routers on a VLAN to be active. The Layer-2 network forwards all multicast traffic to the router(s) in addition to receivers attached to the Layer-2 network. Thus, since the designated election is not visible to the switches, the multicast traffic gets channeled to all the routers on that network, regardless of whether they are the designated router or not. This is not only suboptimal, but also does not utilize the additional opportunity for load balancing and fault-tolerance that exists in a Layer-2 multipath (e.g., Fabricpath) network. In system 100, however, the multicast traffic is directed to the appropriate switch-router. Moreover, the traffic is spread out over the available switch-routers, thereby making optimal usage of them.

[0041] Furthermore, the links may be used in a more optimal manner. Existing schemes lead to sub-optimal usage of links in typical access-spine fat-tree type networks. This mechanism of load balancing also leads to better link-utilization in typical fat-tree networks. Thus, system 100 provides a better multi-pathing solution for multi-destination traffic.

[0042] Although FIGURE 1 illustrates one example of a system for Layer-2 multicast multipathing, other systems for Layer-2 multicast multipathing may include fewer, additional, and/or a different arrangement of components. For example, one or more data sources may be located outside of Layer-2 network 180. The data sources may send multicast data through Layer-3 network 150 to one or more of switch-routers 140. The switch-routers 140 may then route the multicast traffic to the appropriate switches 130, which may switch the traffic to the appropriate data receivers 120. In some implementations, however, Layer-3 network 150 may not exist. As another example, access switches 130 may be switch-routers.

[0043] As another example, local area networks may exist between the end nodes (e.g., data sources 110 and data receivers 120) and access switches 130. The local area networks may allow the end nodes to communicate messages with access switches 130. The local area networks may, for example, be conventional Ethernet networks. In the illustrated example, data source 110a and data receiver 110b may both be coupled to a local area network even though they are in different virtual LANs (VLANs), and data receiver 120b may be coupled to another local area network even though data receiver 120b is part of the same VLAN as data source 110a. In certain implementations, local area networks, access switches, and switch-

routers 140 may operate at Layer-2 and form a Layer-2 network. In these situations, each local area network and the network formed by access switches and switch routers-140 may, however, have their own Layer-2 domain.

[0044] As discussed, system 100 is designed for switch-routers 140 to be co-located Layer-3 (e.g., PIM) routers and Layer-2 (e.g., L2MP) switches. In order for the same scheme to work in a network where Layer-3 entities are separate from the Layer-2 entities, additional functionality may be provisioned in the routers (e.g., to establish the designated router based on {VLAN, Group, Source} granularity rather than per-VLAN per group) and also extra signaling to extend the tree identifier and {VLAN, Group, Source} pinning into the routers.

[0045] FIGURE 2 illustrates an example access switch 200, which is a type of network node that could, for example, be used in system 100. Access switch 200 includes a data plane 210 and a control plane 220. In data plane 210, access switch 200 includes switching logic 212 connected between two sets of ports 214a and 214b. Switching logic 212 is configured to route or internally switch traffic received on one port set 214a (ingress ports) to another port set 214b (egress ports). Data plane 210 also includes a processor 216 (e.g., an application specific integrated circuit (ASIC)) to perform enhanced operations. Control plane 220 includes a generic or application-specific processor 228 for implementing the switching functionality and any channel protocols. In particular implementations, processor 228 may be implemented in a state machine, a micro-controller, hardware, firmware, programmable logic, or a combination thereof.

[0046] Also included in access switch 200 are a designated router list 222 and an OIF list 224. Designated router list 222 may store a list of available designated routers and be constructed based on registration messages received from designated routers. Switches in a Layer-2 network may, for example, use the receipt of a PIM Hello packet to detect the presence of a router in the network. OIF list 224 may, for example, be built as access switch 200 receives group join requests from its local data sources.

[0047] In certain modes of operation, access switch 200 may use a hash sequence that is shared in common with routers and other access switches for forwarding multicast messages. In certain implementations, the hash may be computed by processor 216. The hash sequence may, for example, be shared between access switches and routers by using IS-IS TLVs. A hash may, for instance, be performed on a {VLAN, Source, Group} 3-tuple. This identifier may then

be used to retrieve a tree identifier (i.e., to determine which tree to forward on) and a designated router. The tree identifier may, for example, be a forwarding tag (Ftag) in Fabricpath. Thus, the tree and the designated router may be linked. When access switch 200 receives a multicast message, the access switch may then switch the message to any appropriate data receivers and to the appropriate designated router.

[0048] The tree identifier and the designated router associated with a {VLAN, Source, Group} 3-tuple may be determined by an access switch in a number of ways. For example, the tree identifier and designated router may be stored in OIF list 224 and be referenced by the hash value of the 3-tuple. The port that leads to the designated router may be made part of the OIF list. Thus, the access switch can determine which ports to forward messages on.

[0049] To establish multicast operations, access switch 200 may detect that an associated data receiver wants to join a multicast group. Access switch 200 may, for example, accomplish this using IGMP snooping. When a data receiver desiring to receive data for a group is first detected by access switch 200, the access switch may compute the {VLAN, Source, Group} hash and choose a tree identifier (e.g., an Ftag) and a router from designated router list 222 as the designated router. The join request may be sent to the designated router, which may build the forwarding tree for the group.

[0050] Access switch 200 may pick the tree identifier using the {VLAN, Source, Group} 3-tuple. The tree identifier does not need to form part of the key for the forwarding lookup, however. It may be used for IIC and CBL based tree enforcement. Processor 228 may build OIF list 224 and load it into switch fabric 212, where it may be accessed by processor 216. For hardware in which the forwarding to router ports and to the receivers is done using an (effectively) single lookup, software may build an OIF list that is the union of the links to reach the data receivers and the link to reach the designated router.

[0051] Note that the hardware result for a {VLAN, Source, Group} lookup can be programmed with the ports for the tree identifier resulting from the {VLAN, Source, Group} hash. This can mean that the entries for the other tree identifiers do not need to be programmed or computed.

[0052] For hardware in which the forwarding to router port and to the receivers is done using an (effectively) single lookup, software may build an OIF list that is the union of the links to reach the data receivers and the links to reach the designated router. However, since

the access switch is pinning a group to a router, the port that leads to that switch-router should be part of the OIF list (in addition to the ports that lead to the data receivers).

[0053] For other hardware, a two-step lookup may be used. First, the {VLAN, Source, Group} hash may be implemented, which may provide the OIF list consisting of the data receivers as well as the router_tag. Then, a {VLAN, Router_tag} lookup may be performed that provides the list of router ports. The router_tag may, for example, be: 1) an indirection to the designated switch-router for that {VLAN, Source, Group}; 2) an indirection to the list of available routers, which may happen during router insertion or removal; or 3) an indirection to a set of routers that have been setup in active/standby mode.

[0054] Hardware can also be enhanced to provide the tree identifier as a result of the forwarding {VLAN, Source, Group} lookup (in addition to the router_tag and the OIF list). This can be used in situations where the hash results in a poor choice of tree identifiers, which may be configuration driven.

[0055] When access switch 200 learns of removal of a designated router, the access switch may briefly change the router_tag for that designated router to point to ports leading to other available designated routers. Then, the groups that had been assigned to that router may be reassigned to the remaining routers. The access switch may, for example, know how many routers are available, and the highest source Media Access Control (MAC) address may take the first tree number. In other implementations, the router_tags may be redistributed in a balanced manner (e.g., using a Modulo operation) and/or by taking into account the distance of a designated router from a tree root. After a settling time, the groups that were previously assigned to the removed designated router can then be reassigned to the remaining router-tags.

[0056] When an additional designated router is available, the learned groups are redistributed among the updated list of designated switch-routers. In certain implementations, the router_tags may be redistributed in a balanced manner (e.g., using a Modulo operation) and/or by taking into account the distance of a designated router from a tree root. A change in the router_tag may be preceded by a settling period where the traffic is sent to all the available routers.

[0057] The discussed scheme may need modifications from PIM for optimal operation for cases in which in the Layer-2 network there are: a) multiple sources, but no receivers; and b)

there are multiple sources, but the receivers only issue (*, g) joins. Basically, the Layer-2 network in these cases would be unaware of the various transmitting sources and/or receivers. For both cases, since the Layer-2 network is unaware of the complete {VLAN, Source, Group} information and is therefore unable to compute the complete hash, forwarding would happen towards all available routers. However, PIM receives the initial data message when the source sends data. Therefore, this problem can be solved if the PIM-designated router creates an IGMP-like {VLAN, Group, Source} state in the Layer-2 multipath network, possibly by redistribution to IS-IS. When this happens, hardware no longer finds the OMF entry (which goes to all routers) or the generic (*, Group) entry. Instead, it finds the specific {VLAN, Source, Group} OIF list, which will point towards a specific router.

[0058] FIGURE 3 illustrates an example flow diagram 300 for Layer-2 multicast multipathing. The activities illustrated in FIGURE 3 may, for example, be implemented by an access switch 130. If a multicast message has been received, then a hash operation is performed on the VLAN, the source, and the multicast group associated with the message (operation 308). The flow can also include determining a port for a designated router and port(s) for one or more data receivers based on the hash result (operation 312). Determining the router port and one or more data receiver port(s) may, for example, be accomplished by using the hash result as an index into an OIF list. The flow further calls for switching the message within the Layer-2 network (operation 316) using the determined ports.

[0059] If a Layer-2 multicast message has not been received, the flow can include determining whether a data message has been received from a data source (operation 320). A data message may, for example, be received for a data source if the data source is also a data receiver. If a data message has not been received for a data source, then a check is made whether a Layer-2 multicast message has been received from the data source (operation 304). If a data message has been received for a data source, the flow includes switching the message to the data source (operation 324). The flow can also include checking whether a Layer-2 multicast message has been received from the data source (operation 304). The flow activities may be performed many times during the operation of a system for Layer-2 multicast messaging. In certain implementations, the flow may be performed continually.

[0060] Although FIGURE 3 illustrates one process for Layer-2 multicast multipathing, other processes for Layer-2 multicast multipathing may include fewer, additional, and/or a

different arrangement of operations. For example, a process may not include determining whether a data message has been received for the data source. As another example, a process may include handling group join requests from data receivers. As a further example, a process may include handling a change in designated routers (e.g., removal or addition). As an additional example, an access switch may not have to switch a multicast message to a data receiver (e.g., because they are on different VLANs or the data receiver is not in the Layer-2 network).

[0061] FIGURE 4 illustrates an example flow diagram 400 for Layer-2 multicast multipathing. The flow may be, for example, implemented by an access switch 130. The flow can also include determining whether a designated router (e.g., a switch router) has been removed from the Layer-2 network (operation 404). Determining whether a designated router has been removed may, for example, be accomplished by a reachability test. If a designated router has been removed, the flow includes adjusting the designated router port identifier for a multicast group assigned to the designated router so that the port identifier points to ports for other designated routers (operation 408). For example, the port identifier may be adjusted so that it points to all of the other designated router ports.

[0062] The flow can also include selecting a designated router for the multicast group assigned to the removed designated router (operation 412). Selecting a designated router may, for example, be accomplished by redistributing them in a balanced manner (e.g., using a Modulo operation on the hashed tuple) and/or by taking into account the distance of a designated router from a tree root. Then, the group that had been assigned to the removed router may be reassigned to a remaining router. The flow further calls for adjusting the ports identifier for the multicast group assigned to the removed router to point to the port for the selected router.

[0063] Returning to operation 404, if a designated router has not been removed, the flow calls for determining whether a designated router has been added to the Layer-2 network (operation 420). If a designated router has not been added, the flow calls for again checking whether a designated router has been removed (operation 404). If, however, designated router has been added, the flow calls for identifying a multicast group to assign to the added designated router (operation 424). Identifying a multicast group to assign to the added designated router may, for example, be accomplished by load sharing using the network

hashing algorithm. In certain implementations, the router_tag may be redistributed using a Modulo operation and/or by taking into account the distance of a designated router from a tree root.

[0064] The flow can also include adjusting the port identifier for the identified multicast group to point to the port for the added designated router (operation 428). Adjusting the port identifier for the identified multicast group to point to the added designated router may, for example, be accomplished by changing a port identifier in an OIF list. The flow may be performed many times during the operation of a system for Layer-2 multicast messaging. In certain implementations, the flow of FIGURE 4 may be performed continually.

[0065] Although FIGURE 4 illustrates one process for Layer-2 multicast multipathing, other processes for Layer-2 multicast multipathing may include fewer, additional, and/or a different arrangement of operations. For example, a process may not include determining whether a designated router has been removed. As another example, a process may not include determining whether a designated router has been added. As an additional example, a process may include a settling time (e.g., in which router traffic is sent to all routers) before adjusting router tags for groups after a router event (addition or removal).

[0066] Note that in certain example implementations, the switching/routing functions outlined herein may be implemented by logic encoded in one or more non-transitory tangible media (e.g., embedded logic provided in an application specific integrated circuit (ASIC), digital signal processor (DSP) instructions, software (potentially inclusive of object code and source code) to be executed by a processor, or other similar machine, etc.). In some of these instances, a memory element can store data used for the switching/routing operations described herein. This includes the memory element being able to store software, logic, code, or processor instructions that can be executed to carry out the activities described in this Specification. A processor can execute any type of instructions associated with the data to achieve the switching/routing operations detailed herein in this Specification. In one example, the processor could transform an element or an article (e.g., data) from one state or thing to another state or thing.

[0067] In another example, the switching/routing activities outlined herein may be implemented with fixed logic or programmable logic (e.g., software/computer instructions executed by a processor) and the elements identified herein could be some type of a

programmable processor, programmable digital logic (e.g., a field programmable gate array (FPGA), an erasable programmable read only memory (EPROM), an electrically erasable programmable ROM (EEPROM)) or an ASIC that includes digital logic, software, code, electronic instructions, or any suitable combination thereof. These devices may further keep information in any suitable memory element (random access memory (RAM), ROM, EPROM, EEPROM, ASIC, etc.), software, hardware, or in any other suitable component, device, element, or object where appropriate and based on particular needs. Any possible memory items (e.g., database, table, cache, etc.) should be construed as being encompassed within the broad term 'memory element.' Similarly, any of the potential processing elements, modules, and machines described in this Specification should be construed as being encompassed within the broad term 'processor.'

[0068] Note that with the examples provided herein, interaction may be described in terms of two or three elements. However, this has been done for purposes of clarity and example only. In certain cases, it may be easier to describe one or more of the functionalities of a given set of flows by only referencing a limited number of network elements. It should be appreciated that the communication systems are readily scalable and can accommodate a large number of clouds, networks, and/or switches, as well as more complicated/sophisticated arrangements and configurations. Accordingly, the examples provided herein should not limit the scope or inhibit the broad teachings of the specification as potentially applied to a myriad of other architectures.

[0069] FIGURE 5 illustrates an example computer system 500 for routing in a Layer-2 network. Computer system 500 may, for example, illustrate some of the components of the control plane of a switch/router. System 500 includes a processor 510, an input/output system 520, and a memory 530, which can be coupled together by a network 540 (e.g., an internal framework that facilitates single, or bidirectional communications between these components). Alternatively, these elements can be suitable linked (or interact) in any other appropriate fashion. Processor 510 can include a logical processing unit (e.g., an arithmetic logic unit) that processes data under the direction of program instructions (e.g., from software). For example, processor 510 may be a microprocessor, a microcontroller, or an application specific integrated circuit. The processor may operate by reduced instruction set computer (RISC) or complex

instruction set computer (CISC) principles. In general, the processor may be any device that manipulates data in a logical manner.

[0070] Input/output system 520 may include one or more communication interfaces and/or one or more user interfaces. A communication interface may, for instance, be a network interface card (whether wireless or wired), a modem, or a bus interface. A user interface could, for instance, be a user input device (e.g., a keyboard, a keypad, a touchpad, a mouse, a stylus, or a microphone) or a user output device (e.g., a monitor, a display, or a speaker). In general, system 520 may be any combination of devices by which a computer system can receive and output data.

[0071] Memory 530 may, for example, include random access memory (RAM), read-only memory (ROM), flash memory, and/or disc memory. Various items may be stored in different portions of the memory at various times. Memory 530, in general, may be any combination of devices for storing data.

[0072] Memory 530 includes instructions 532 and data 536. Instructions 532 include an operating system 533 (e.g., Windows, Linux, or Unix) and applications 534, which include a switching manager 535. Data 536 includes the data required for and/or produced by applications 534, including a designated router list 537 and an Outgoing Interface (OIF) list 538. Network 540 can be responsible for communicating data between processor 510, input/output system 520, and memory 530. Network 540 may, for example, include a number of different types of busses (e.g., serial and parallel) that are somewhat internal to a given device.

[0073] In certain modes of operation, computer system 500, according to switching manager 535, use a hash sequence shared in common in a Layer-2 network for forwarding multicast messages. A hash may, for instance, be performed on a {VLAN, Source, Group} 3-tuple. This should, in most instances, provide a unique identifier for each 3-tuple. This identifier may then be used to retrieve a tree identifier (i.e., to determine which tree to forward on) and a designated router (i.e., for a switch-router in the Layer-2 network). The tree identifier may, for example, be a forwarding tag (Ftag) in Fabricpath. Thus, the tree and the designated router may be linked.

[0074] The tree identifier and the designated router associated with a {VLAN, Source, Group} 3-tuple may be determined by an access switch in a number of ways. For example, the tree identifier and designated router may be stored in OIF list 538 and be referenced by the

hash value of the 3-tuple. The port that leads to the designated switch-router may be made part of the OIF list. Thus, the access switch can determine which ports to forward messages on.

[0075] To establish multicast operations, computer system 500 may detect that a data receiver wants to join a multicast group (e.g., by using IGMP snooping). When a data receiver desiring to receive data for a new group is first detected, computer system 500 may compute the {VLAN, Source, Group} hash and choose a tree identifier (e.g., an Ftag) and a designated router from designated router list 537. The join request may be sent to the designated router, which may build the forwarding tree for the group. Additionally, the rest of the Layer-2 network may be informed of the join request.

[0076] For hardware in which the forwarding to router ports and to the receivers is done using an (effectively) single lookup, computer system 500 may build an OIF list that is the union of the links to reach the appropriate data receivers and the link to reach the designated router. However, since a group is being pinned to a router, the port that leads to that router should be part of OIF list 538 (in addition to the ports that lead to data receivers 120).

[0077] For other hardware, a two-step lookup may be used. First, computer system 500 may implement the {VLAN, Source, Group} hash, which may provide the OIF list portion consisting of the data receivers as well as a router_tag. Then, a {VLAN, Router_tag} lookup may be performed that provides the list of router ports. The router_tag may, for example, be: 1) an indirection to the designated switch-router for that {VLAN, Source, Group}; 2) an indirection to the list of available routers, which may happen during router addition or removal; or 3) an indirection to a set of routers that have been setup in active/standby mode. Hardware can also be enhanced to provide the tree identifier as a result of the forwarding {VLAN, Source, Group} lookup (in addition to the router_tag and the OIF list). This can be used in situations where the hash results in a poor choice of tree identifiers, which may be configuration driven. The designated router may, for example, be selected by performing a modulo operation on the hash value against the tree identifier. In certain implementations, the distance of the designated router from the root of the tree may also be considered, closer designated routers typically being more preferred.

[0078] When computer system 500 learns of removal of a designated router, the processor may briefly change the router_tag for that designated router to point to ports leading to other available routers. Then, the multicast groups that had been assigned to the removed

router need to be reassigned to the remaining routers. For example, computer system 500 may adjust the modulo operation performed on the hash value against the tree identifier to account for the fewer number of routers, and this may be used to reassign the groups. For instance, the designated routers may be sorted by their MAC address into a list [router 0, ..., router X-1], and the hash value output may be processed by MOD by X, the number of routers. In certain implementations, the distance of the designated router from the root of the Layer-2 multicast tree may also be considered, closer designated routers typically being more preferred. Thus, the designated router(s) closest to the root of the Layer-2 multicast tree may be associated with tree (e.g., the same hash values that choose tree X may choose the designated router nearest to the root of that tree), which may provide enhanced forwarding. As another example, computer system 500 may determine how many designated routers are available, and the highest source MAC address may take the first tree number. Switches in a Layer-2 network may, for example, use the receipt of a PIM hello packet to detect the presence of a router in the network.

[0079] After a settling time, computer system 500 may assign the groups that were previously assigned to the removed designated router to the remaining router-tags. The redistribution algorithm may be used by computer system 500 and the designated routers so that hash values are tied to designated routers. When a new designated router is available, the multicast groups are redistributed among the revised designated router list. For example, computer system 500 may adjust the modulo operation performed on the hash value against the tree identifier to account for the larger number of routers, and this may be used to reassign the groups. In certain implementations, processor may also evaluate the distance of the designated router from the root of the tree, closer designated routers typically being more preferred. A Layer-2 forwarding protocol and the PIM may interact to determine the router location with respect to the tree root. Thus, the designated router(s) closest to the root of the Layer-2 multicast tree may be associated with tree (e.g., the same hash values that choose tree X will choose the designated router nearest to the root of that tree), which will provide enhanced forwarding.

[0080] As another example, computer system 500 may determine how many designated routers are available, and the highest source MAC address may take the first tree number. Switches in a Layer-2 network may, for example, use the receipt of a PIM hello packet

to detect the presence of a router in the network. After a settling time, computer system 500 may reassign the groups that were previously assigned to the removed designated router to the remaining router-tags. The redistribution algorithm may be used by processor 510 and the designated routers so that hash values are tied to designated routers. Computer system 500 may, for example, accomplish these operations by implementing one or more parts of processes 300-400. Computer system 500 may also use any other techniques discussed herein.

[0081] It is also important to note that the operations discussed with reference to FIGURES 1-5 illustrate only some of the possible scenarios that may be executed by, or within, a communication system. Some of these operations may be deleted or removed where appropriate, or these operations may be modified or changed considerably without departing from the scope of the present disclosure. In addition, a number of these operations have been described as being executed concurrently with, or in parallel to, one or more additional operations. However, the timing of these operations may be altered considerably. The preceding operational flows have been offered for purposes of example and discussion. Substantial flexibility is allowable, however, in that any suitable arrangements, chronologies, configurations, and timing mechanisms may be provided without departing from the teachings of the present disclosure.

[0082] Although the present disclosure has been described in detail with reference to particular embodiments, it should be understood that various other changes, substitutions, and alterations may be made hereto without departing from the spirit and scope of the present disclosure. For example, virtually any configuration that seeks to intelligently switch messages could readily adopt the teachings of the present disclosure.

[0083] Numerous other changes, substitutions, variations, alterations, and modifications may be ascertained to one skilled in the art, and it is intended that the present disclosure encompass all such changes, substitutions, variations, alterations, and modifications as falling within the scope of the appended claims. In order to assist the United States Patent and Trademark Office (USPTO) and, additionally, any readers of any patent issued on this application in interpreting the claims appended hereto, Applicant wishes to note that the Applicant: (a) does not intend any of the appended claims to invoke paragraph six (6) of 25 U.S.C. section 112 as it exists on the date of the filing hereof unless the words "means for" or "step for" are specifically used in the particular claims; and (b) does not intend, by any

statement in the specification, to limit this disclosure in any way that is not otherwise reflected in the appended claims.

WHAT IS CLAIMED IS:

1. A method, comprising:
 - receiving a multicast data message from a data source, the message being in a first virtual local area network and being associated with a multicast group;
 - calculating a hash value based on the virtual local area network, the data source, and the multicast group;
 - determining a port for a designated router in a Layer-2 network based on the hash value; and
 - switching the multicast data message to the port that was determined.

2. The method of Claim 1, further comprising:
 - determining a port for a data receiver in the Layer-2 network based on the hash value;and
 - switching the multicast data message to the port that was determined.

3. The method of Claim 1, further comprising:
 - determining whether a designated router has been removed from the Layer-2 network;and
 - adjusting a port identifier for a multicast group assigned to the designated router to identify ports for other designated routers in the Layer-2 network.

4. The method of Claim 3, further comprising:
 - selecting one of the other designated routers to be the designated router for the multicast group assigned to the removed router; and
 - adjusting the port identifier for the multicast group assigned to the removed router to point to a port for the selected one of the other designated routers.

5. The method of Claim 1, further comprising:
identifying an added designated router that has been added to the Layer-2 network;
identifying a particular multicast group to assign to the added designated router; and
adjusting a port identifier for the particular multicast group to point to a port for the added designated router.

6. A network node, comprising:
a memory element configured for storing instructions; and
a processor coupled to the memory element and configured to execute the instructions such that the network node is configured to:
determine whether a multicast data message from a data source has been received, the message being in a first virtual local area network and being associated with a multicast group;
calculate a hash value based on the virtual local area network, the data source, and the multicast group;
determine a port for a designated router in a Layer-2 network based on the hash value; and
switch the multicast data message to the port that was determined.

7. The network node of Claim 6, wherein the network node is further configured to:
determine a port for a data receiver in the Layer-2 network based on the hash value;
and
switch the multicast data message to the port that was determined.

8. The network node of Claim 6, wherein the network node is further configured to:
determine whether a designated router has been removed from the Layer-2 network;
and
adjust a port identifier for a multicast group assigned to the designated router to identify ports for other designated routers in the Layer-2 network.

9. The network node of Claim 8, wherein the network node is further configured to:
select one of the other designated routers to be the designated router for the multicast group assigned to the removed router; and

adjust the port identifier for the multicast group assigned to the removed router to point to a port for the selected one of the other designated routers.

10. The network node of Claim 6, wherein the network node is further configured to:
identifying an added designated router that has been added to the Layer-2 network;
identify a particular multicast group to assign to the added designated router; and
adjust a port identifier for the particular multicast group to point to a port for the added designated router.

11. The network node of Claim 6, wherein the network node is further configured to:
construct an outgoing interface (OIF) list to be referenced by a particular hash value.

12. The network node of Claim 11, wherein the OIF list is constructed based on group join requests from a plurality of data sources.

13. The network node of Claim 6, wherein the network node is further configured to:
use a first identifier to retrieve a second identifier to determine which tree should be used for a forwarding activity and to determine a particular designated router in the Layer-2 network.

14. Logic encoded on one or more non-transitory tangible computer readable media for execution and when executed operable to:

receive a multicast data message from a data source, the message being in a first virtual local area network and being associated with a multicast group;

calculate a hash value based on the virtual local area network, the data source, and the multicast group;

determine a port for a designated router in a Layer-2 network based on the hash value;
and

switch the multicast data message to the port that was determined.

15. The logic of Claim 14, wherein the logic is further operable to:

determine a port for a data receiver in the Layer-2 network based on the hash value;
and

switch the multicast data message to the port that was determined.

16. The logic of Claim 14, wherein the logic is further operable to:

determine whether a designated router has been removed from the Layer-2 network;
and

switch a port identifier for a multicast group assigned to the designated router to identify ports for other designated routers in the Layer-2 network.

17. The logic of Claim 16, wherein the logic is further operable to:

select one of the other designated routers to be the designated router for the multicast group assigned to the removed router; and

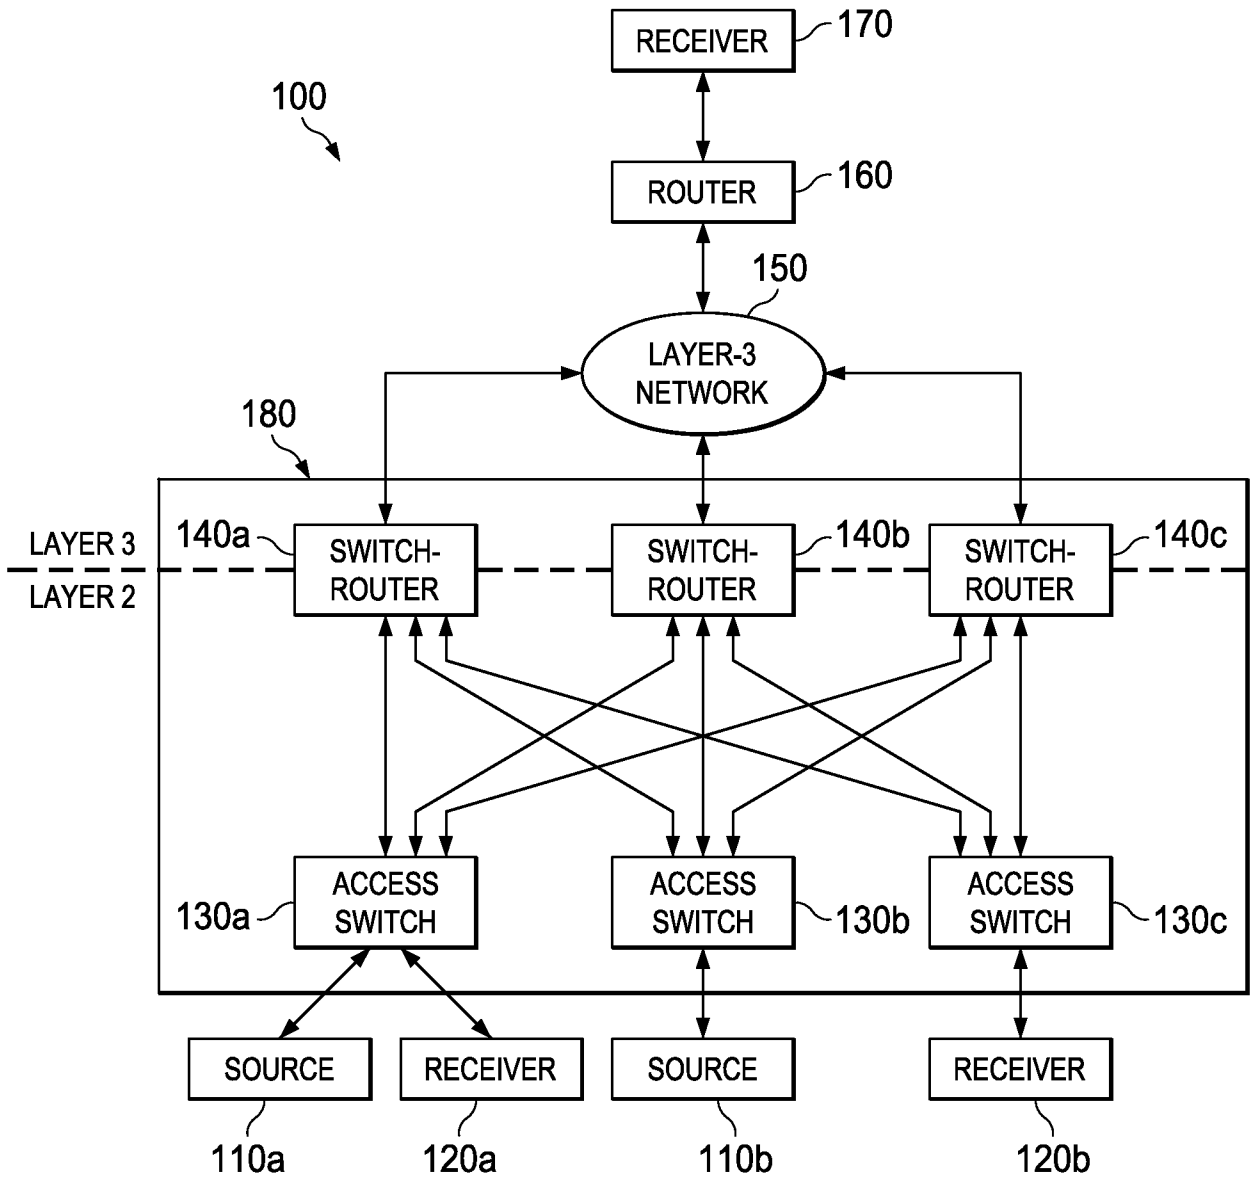
adjust the port identifier for the multicast group assigned to the removed router to point to a port for the selected one of the other designated routers.

18. The logic of Claim 14, wherein the logic is further operable to:
identify an added designated router that has been added to the Layer-2 network;
identify a particular multicast group to assign to the added designated router; and
adjust a port identifier for the particular identified multicast group to point to a port for
the added designated router.

19. The logic of Claim 14, wherein the logic is further operable to:
construct a designated router list, which identifies available designated routers in the
Layer-2 network, wherein the router list is constructed based on registration messages received
from a plurality of network nodes of the Layer-2 network.

20. The logic of Claim 14, wherein the logic is further operable to:
construct an outgoing interface (OIF) list to be referenced by a particular hash value,
wherein the OIF list is constructed based on group join requests from a plurality of data
sources.

FIG. 1






FIG. 2

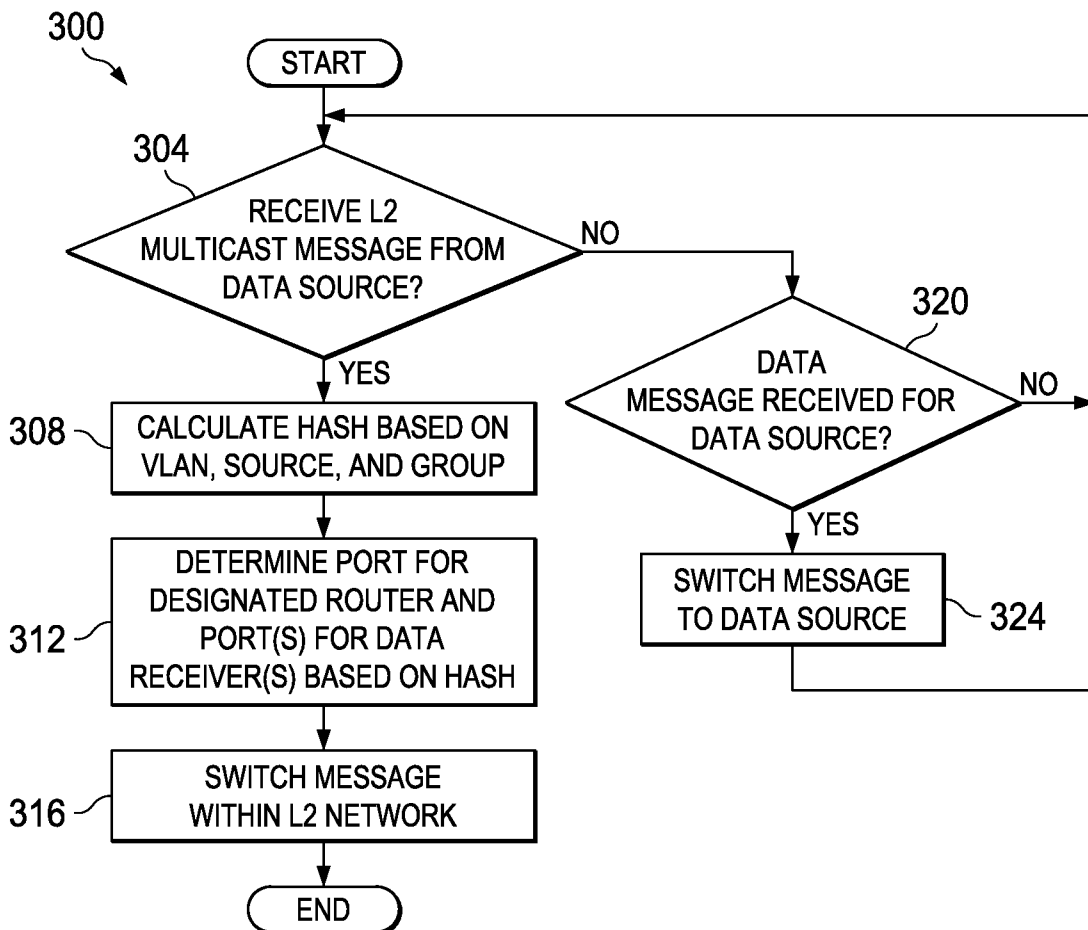


FIG. 3

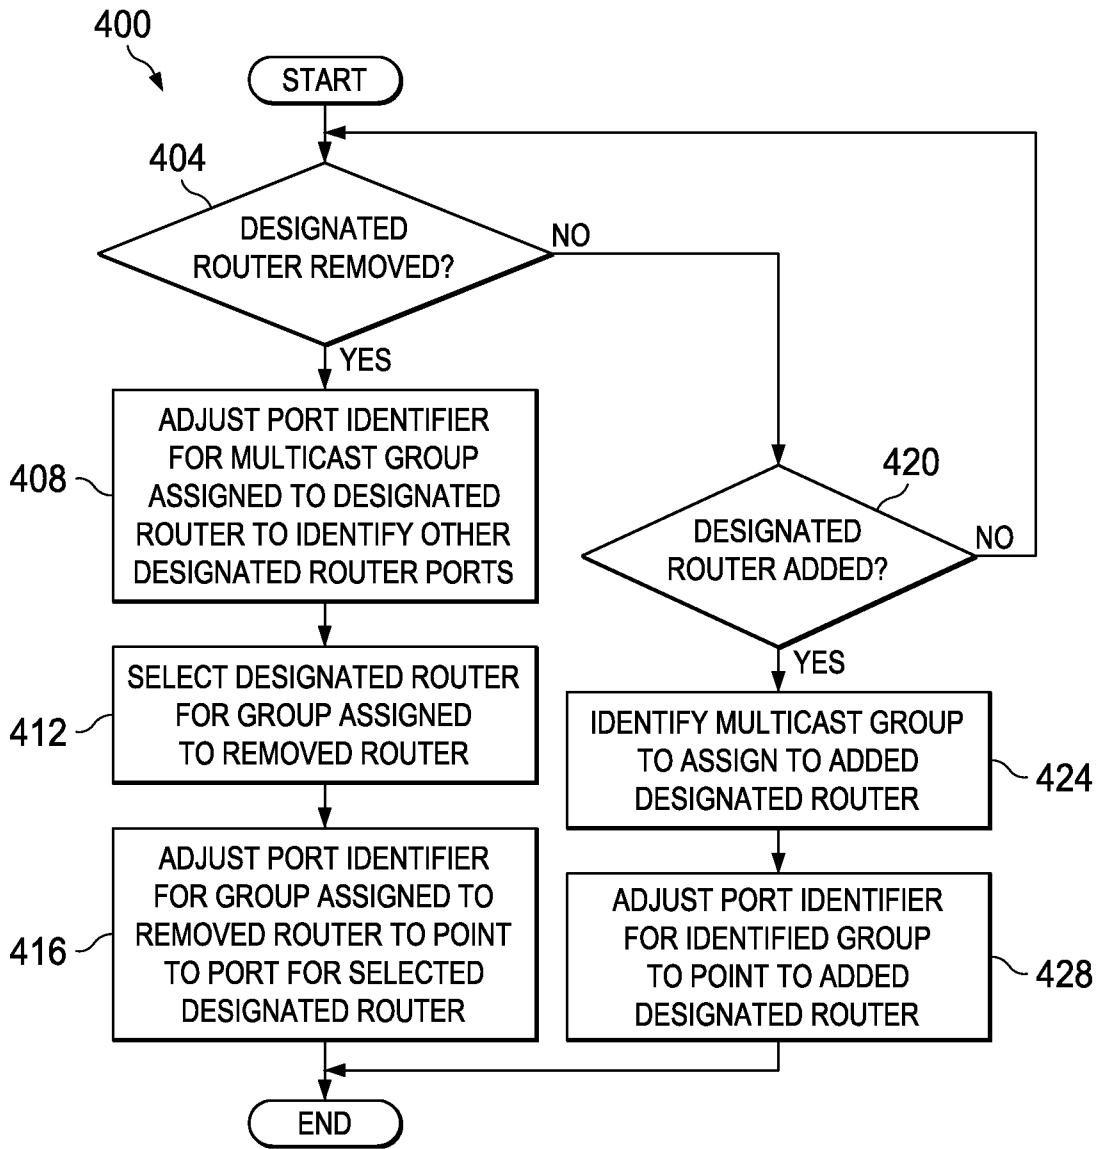


FIG. 4

FIG. 5

