



(51) International Patent Classification:
H04J 3/06 (2006.01)

(21) International Application Number:
PCT/US2010/023947

(22) International Filing Date:
11 February 2010 (11.02.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/157,501 4 March 2009 (04.03.2009) US
12/421,443 9 April 2009 (09.04.2009) US

(71) Applicant (for all designated States except US): TELEFONAKTIEBOLAGET LM ERICSSON [SE/SE]; Telefonplan, S-164 83 Stockholm (SE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): SRIGANESH, Kini [IN/US]; 45855 Cheyenne Place, Fremont, CA 94539 (US). LU, Wenhui [US/US]; 1028 Belvedere Lane, San Jose, CA 95129 (US).

(74) Agents: VINCENT, Lester, J. et al.; BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN LLP, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (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, 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, PE, PG, PH, PL, PT, RO, RS, RU, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (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, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: LDP IGP SYNCHRONIZATION FOR BROADCAST NETWORKS




FIG. 1

(57) Abstract: A network element that has a broadcast interface to a broadcast network becoming operational determines whether at least one alternate path exists to the broadcast network. The broadcast interface is to carry traffic on a label switched path. If an alternate path exists, the broadcast interface will not be advertised to the broadcast network until label distribution protocol (LDP) is operational with all neighbors on the broadcast interface.

WO 2010/101705 A1

Published:

— *with international search report (Art. 21(3))*

LDP IGP Synchronization for Broadcast Networks

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 61/157,501, filed March 4, 2009, which is hereby incorporated by reference.

BACKGROUND

Field

[0002] Embodiments of the invention relate to the field of networking; and more specifically to LDP (Label Distribution Protocol) IGP (Interior Gateway Protocol) synchronization for broadcast networks.

Background

[0003] LDP (described in “LDP Specification”, RFC 5036, October 2007), which is used to establish LSPs (label switched paths) to destinations, typically relies on IGP (e.g., Open Shortest Path First (OSPF) (defined in “OSPF Version 2”, STD 54, RFC 2328, April 1998), Intermediate system to Intermediate system (IS-IS) (defined in “Intermediate system to intermediate system intra-domain-routing routine information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)” ISO standard 10589, 1992), etc.) to provide the underlying routing information (e.g., the cost metric between the hops in the network). Even though LDP typically relies on the IGP protocol, they are independent from each other. IGP may be operational on a link prior to LDP becoming operational on that link which can lead to packet loss.

[0004] The request for comments (RFC) 5443, “LDP IGP Synchronization”, March 2009 (hereinafter “LDP IGP Synchronization”) describes a mechanism to discourage links from being used for IP forwarding if LDP is not fully operational on that link. RFC 5443 describes that when LDP is not fully operational on a given link (e.g., all the label bindings have not yet been exchanged), IGP will advertise the link with maximum cost to discourage traffic from being sent over the link. When LDP on the link becomes operational (e.g., all label bindings have been exchanged), IGP advertises the link with the correct cost.

[0005] On broadcast links (more than one LDP/IGP peer on the same link), IGP advertises a common cost to the broadcast link, rather than a separate cost to each peer. A broadcast pseudo-node may be implemented in a network to reduce the number of links in the shortest path first (SPF) (e.g., open shortest path first (OSPF), constrained shortest path first (CSPF)) database of each network element. Each network element forms adjacencies with the broadcast pseudo-node and advertises a link and cost to the broadcast pseudo-node and the broadcast pseudo-node advertises a link and cost to each network element.

[0006] For broadcast links, the RFC 5443 “LDP IGP Synchronization” describes that the mechanism can be applied to the link as a whole and not an individual peer. Advertising maximum cost to the link as a whole may cause sub-optimal traffic diversion and/or black-holing of traffic (e.g., VPN traffic).

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

[0008] Figure 1 illustrates an exemplary broadcast network according to one embodiment of the invention;

[0009] Figure 2 is a block diagram of an exemplary network element configured for LDP-IGP synchronization for broadcast networks according to one embodiment of the invention;

[0010] Figure 3 is a flow diagram illustrating exemplary operations for LDP-IGP synchronization for broadcast networks according to one embodiment of the invention;

[0011] Figure 4 is a flow diagram illustrating exemplary operations when LDP-IGP synchronization is complete according to one embodiment of the invention;

[0012] Figure 5A illustrates an exemplary local IGP interface structure according to one embodiment of the invention;

[0013] Figure 5B illustrates an exemplary link state database according to one embodiment of the invention;

[0014] Figure 6 is a flow diagram illustrating exemplary operations when a network element is ready to update its link-state advertisement with an interface to a broadcast network according to one embodiment of the invention; and

[0015] Figure 7 illustrates an exemplary network element used in some embodiments of the invention.

DETAILED DESCRIPTION

[0016] In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description. Those of ordinary skill in the art, with the included descriptions, will be able to implement appropriate functionality without undue experimentation.

[0017] References in the specification to “one embodiment,” “an embodiment,” “an example embodiment,” etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0018] In the following description and claims, the terms “coupled” and “connected,” along with their derivatives, may be used. It should be understood that these terms are not intended as synonyms for each other. “Coupled” is used to indicate that two or more elements, which may or may not be in direct physical or electrical contact with each other, co-operate or interact with each other. “Connected” is used to indicate the establishment of communication between two or more elements that are coupled with each other.

[0019] The techniques shown in the figures can be implemented using code and data stored and executed on one or more electronic devices (e.g., a network element,

etc.). Such electronic devices store and communicate (internally and/or with other electronic devices over a network) code and data using machine-readable media, such as machine-readable storage media (e.g., magnetic disks; optical disks; random access memory; read only memory; flash memory devices; phase-change memory) and machine-readable communication media (e.g., electrical, optical, acoustical or other form of propagated signals – such as carrier waves, infrared signals, digital signals, etc.). In addition, such electronic devices typically include a set of one or more processors coupled to one or more other components, such as a storage device, one or more user input/output devices (e.g., a keyboard, a touchscreen, and/or a display), and a network connection. The coupling of the set of processors and other components is typically through one or more busses and bridges (also termed as bus controllers). The storage device and signals carrying the network traffic respectively represent one or more machine-readable storage media and machine-readable communication media. Thus, the storage device of a given electronic device typically stores code and/or data for execution on the set of one or more processors of that electronic device. Of course, one or more parts of an embodiment of the invention may be implemented using different combinations of software, firmware, and/or hardware.

[0020] As used herein, a network element (e.g., a router, switch, bridge, etc.) is a piece of networking equipment, including hardware and software, that communicatively interconnects other equipment on the network (e.g., other network elements, computer end stations, etc.). Some network elements are “multiple services network elements” that provide support for multiple networking functions (e.g., routing, bridging, switching, Layer 2 aggregation, and/or subscriber management), and/or provide support for multiple application services (e.g., data, voice, and video).

[0021] Within certain network elements, multiple “interfaces” may be configured. As used herein an interface is a connection between a network element and one of its attached networks. An interface has information associated with it that is typically obtained from underlying lower level protocols and the routing protocol (e.g., an IP address and mask). An interface is sometimes referred to as a link. A broadcast interface is an interface that is connected with a broadcast network.

[0022] A method and apparatus for LDP-IGP synchronization for broadcast networks is described. In one embodiment of the invention, when a link to a broadcast network comes up, the network element whose link is coming up will postpone advertising the link to the broadcast network until LDP is operational with all neighbors on that broadcast link (all broadcast neighbors). For example, that network element determines whether that link is the only path to the broadcast network. If that link is the only path to the broadcast network, the interface corresponding to that link is marked as a cut-edge interface. When the network element is ready to update its link-state advertisement (LSA) with the link to the broadcast network, it will postpone adding that link to its LSA if the corresponding interface is not a cut-edge interface and LDP is not operational with all neighbors on that interface. After LDP becomes operational with all neighbors on that interface (e.g., LDP becomes synchronized with IGP), the link will be added to the LSA and will be transmitted to the broadcast network.

[0023] As used herein, the term “link state advertisements” and “LSA” is protocol agnostic. For example, if the IGP protocol is Open Shortest Path First (OSPF), then the link state advertisement may be a Router-LSA. If the IGP protocol is intermediate system to intermediate system (IS-IS), then the link state advertisement may be a Link State PDU.

[0024] In one embodiment, if the IGP protocol is Open Shortest Path First (OSPF), the Router-LSA is not updated with a Link Type 2 (link to transit network) for the subnet until LDP is operational with all network elements on that subnet.

[0025] In one embodiment, if the IGP protocol is intermediate system to intermediate system (IS-IS), the Link State PDU is not updated with an IS Reachability TLV (or an Extended IS Reachability TLV) to the broadcast network until LDP is operational with all network elements.

[0026] Figure 1 illustrates an exemplary broadcast network according to one embodiment of the invention. The network 100 includes the network elements 105, 110, 130, 140, 145, 150, 155, 160, 165, and 170, and the broadcast pseudo-node 120. The network elements 110, 130, 140, 145, and 150 are each directly coupled with the broadcast pseudo-node 120. It should be understood that the topology of

the network 100 is exemplary, and other topologies may be used in embodiments of the invention. For example, in some embodiments, a broadcast pseudo-node is not used.

[0027] In one embodiment, the network elements 105, 160, 165, and 170 are provider edge network elements and the network 100 is a Multiprotocol Label Switching (MPLS) network with Virtual Private Network (VPN) applications. As one example, the network element 105 may have a LDP LSP path to the network elements 160, 155, 170, and/or 165. It should be understood that other LDP LSP paths may be established in the network 100.

[0028] The network elements 110, 130, 140, 145, and 150 each advertise their corresponding interface's cost to the broadcast pseudo-node 120. For purposes of explanation, the network element 110 will advertise a cost of 1 on the broadcast interface 114, the network element 130 will advertise a cost of 1 on the broadcast interface 132, the network element 140 will advertise a cost of 1 on the broadcast interface 142, the network element 145 will advertise a cost of 1 on the broadcast interface 146, and the network element will advertise 152 will advertise a cost of 1 on the broadcast interface 152. In addition, exemplary purposes, each other link in the network 100 has a cost of 1 with the exception of the link between the network element 110 and the network element 165 which has a cost of 10. The broadcast pseudo-node 120 also advertises its interfaces to its neighbors (e.g., through the interfaces 122, 124, 126, and 127) with a cost of zero. It should be understood that the interfaces illustrated in Figure 1 are exemplary and other interfaces are not illustrated in order not to confuse understanding of the invention.

[0029] In one interpretation of the RFC 5443 "LDP IGP Synchronization" as applied to broadcast networks, when a new network element is discovered on the broadcast network, all the network elements with direct links to the broadcast network advertise maximum cost to the broadcast network. For example, with reference to Figure 1, if the broadcast interface 132 of the network element 130 becomes operational (e.g., the network element 130 was added to the network while the network elements 110, 140, 145, and 150 are each already connected with the broadcast pseudo-node 120) and detected by the network elements 110, 140, 145, 150, those network elements each start advertising maximum cost to the broadcast

network (the network elements 110, 132, 140, 145, and 150 advertise maximum cost on the broadcast interfaces 114, 132, 142, 146, and 152 respectively). Since maximum cost is greater than the cost of 10, traffic between the network element 105 and the network element 165 will be diverted to the sub-optimal path of network element 105 to the network element 110 to the network element 165 (instead of the optimal path of network element 105 to network element 110 to the network element 150 to the network element 165). Thus, using the interpretation of the RFC “LDP IGP Synchronization” that has each network element advertising a maximum cost to the broadcast network may result in sub-optimal traffic diversion.

[0030] In addition, as a result of applying the mechanism described in the RFC 5443 “LDP IGP Synchronization” for broadcast networks, traffic may be black-holed (continually dropped) at a network element due to each broadcast network element advertising a maximum cost. For example, using the network topology of Figure 1 and the same costs of the links in the previous example, the network element 105 transmits VPN traffic to the network element 160 prior to the interface 132 becoming operational on the network element 130 and the LDP LSP carrying that VPN traffic will take the following path: network element 105 to the network element 110 to the network element 140 to the network element 155 to the network element 160 (this is the only path possible in the topology). After the interface 132 becomes operational, it should be understood that an additional possible LDP LSP from the network element 105 to the network element 160 will exist once LDP peering has been established along that path (network element 105 to the network element 110 to the network element 130 to the network element 160). Once established, this LDP LSP will be favored over the previous LSP because it will have a lower cost (3 versus 4 in the example). However, since each link of the broadcast network will be advertised as a maximum cost, the path through the network element 130 will have a lower cost than the path through the network elements 140 and 155. As a result, the network element 110 will switch its next-hop to the network element 160 from the network element 140 to the network element 130 before an LDP LSP has been established through the network element 130. Consequently, the network element 110 will try, and fail, to transmit VPN traffic to the network element 160 through the network element 130 until that LDP LSP has

been established (e.g., the network element 110 will not have the proper labels to process the VPN traffic). Thus, the VPN traffic will be black-holed at the network element 110. The amount of traffic loss in this case is at least the order of the time it takes for the LDP LSP to become operational. Although the topology in Figure 1 is rather simple, it should be understood that in some topologies this can be of the order of several minutes (which can violate carrier class reliability metrics).

[0031] In another interpretation of the RFC 5443 “LDP IGP Synchronization” as applied to broadcast networks, when a new network element is discovered on the broadcast network, only that network element transmits a maximum cost to the broadcast network and the other network elements transmit their normal cost. However, this interpretation also suffers from the possibility of black-holing traffic. It should be understood that although the network element 130 advertises a maximum cost to the broadcast network, this only affects traffic being transmitted by the network element 130 and it will not affect traffic being received by the network element 130. Thus, from the point of view of the network element 110, the link between the network element 110 and the network element 130 will have a cost of 1. Using the same example as above and since IGP will converge faster than LDP, the network element 110 will try, and fail, to transmit VPN traffic to the network element through the network element 130 until that LDP LSP has been established. Thus, similar to the above example, the VPN traffic will be black-holed at the network element 110.

[0032] In contrast, as will be described later herein, embodiments of the invention allow for LDP-IGP synchronization in broadcast networks without sub-optimally diverting traffic or black-holing traffic.

[0033] Figure 2 is a block diagram of an exemplary network element configured for LDP-IGP synchronization for broadcast networks according to one embodiment of the invention. While Figure 2 illustrates the network element 130, it should be understood that one or more of the network elements in the network 100 may include similar features. As illustrated in Figure 2, the network element 130 includes the control plane 210 coupled with the data plane 260. The control plane 210 includes the command line interface 230, the IGP module 220, the interface state manager 235, the LDP module 240, and the label manager 250. The IGP

module 200 includes the SPF and cut-edge detection module 280 and the LDP-IGP synchronization LSA module 285. The IGP module 220 manages the neighbor adjacency table 222, the link state database 224, the local IGP RIB (routing information base) 226, and the local IGP interface structure 228.

[0034] The IGP module 220 may receive LDP-IGP synchronization for broadcast networks configuration parameters from the command line interface 230. For example, a network administrator may use the command line interface 230 to configure the LDP-IGP synchronization for broadcast networks on the network element 130 (e.g., whether the LDP-IGP synchronization for broadcast networks is enabled, one or more interfaces to monitor for LDP-IGP synchronization for broadcast networks, etc.). In another embodiment, the LDP-IGP synchronization for broadcast networks mechanism is installed for each broadcast interface on the network element 130.

[0035] The interface state manager 235 manages the interfaces of the network element 130, including the interfaces 132 and 134. For example, the interface state manager 235 detects when an interface is operational. The interface state manager 235 is coupled with the IGP module 220. The IGP module 220 registers those broadcast interfaces (e.g., the interfaces as specified during configuration) with the interface state manager 235. The interface state manager 235 notifies the IGP module 220 upon a state change of one of those registered interfaces (e.g., a broadcast interface becoming operational, a broadcast interface going down, etc.). The IGP module 220 may then update the local IGP interface structure 228 with those interfaces.

[0036] The IGP module 220 establishes and maintains neighbor adjacencies with other network elements in the network 100. For example, the IGP module 220 transmits and receives hello packets from other network elements in the network 100. From these hello packets, the IGP module creates and maintains the neighbor adjacency table 222.

[0037] The IGP module 220 also transmits and receives link-state information (typically in the form of link-state advertisements (LSAs)) to construct a topology of the network 100. From the LSAs it receives and transmits, the IGP module 220

creates and maintains the link state database 224 (thus the link state database 224 is a representation of the network topology of the network 100). Figure 5B illustrates an exemplary link state database 224 according to one embodiment of the invention. As illustrated in Figure 5B, the link state database 224 includes a From field 530, a To field 540, and a Link field 550.

[0038] The IGP module 220 also includes the SPF and cut-edge detection module 280 to determine the optimum paths to the destinations of the link state database 224 (thus, the SPF algorithm is applied to the information of the link state database 224). The resulting routes are added to the local IGP RIB 226, which are then programmed to one or more FIBs (forwarding information bases) on the data plane 260. For example, the routes in the IGP RIB 226 are programmed to one or more packet processing units of one or more line cards in the network element 130.

[0039] In addition to determining the optimum paths to the destinations in the link state database 224, the SPF and cut-edge detection module 280 determines whether interfaces are broadcast cut-edge interfaces. A broadcast cut-edge interface is an interface to the broadcast network that represents the only path to that broadcast network. In other words, besides connecting to the broadcast network through the broadcast interface, if an alternate path to the broadcast network does not exist, that broadcast interface is a cut-edge interface. For example, with reference to Figure 1, the interface 132 is not a cut-edge interface because an alternate IGP path exists to the broadcast pseudo-node 120 (network element 130 to the network element 160 to the network element 155 to the network element 140 to the broadcast pseudo-node 120). The interface 146 is a cut-edge interface because an alternate IGP path does not exist to the broadcast pseudo-node 120. If an interface is a cut-edge interface, a cut-edge property is set for that interface in the local IGP interface structure 228.

[0040] In one embodiment of the invention, the SPF and cut-edge detection module 280 makes the determination of a broadcast cut-edge interface by augmenting the step of SPF calculation that adds a network prefix to the local IGP RIB 226 to include an additional check whether that prefix belongs to an interface that is down (e.g., if the prefix matches a down interface, the prefix will be viable from somewhere else and the interface will not be a cut-edge) or is a stub network

through which traffic cannot be further transited beyond. Of course, other methods may be used to determine whether alternate paths to the broadcast network exist.

[0041] As previously described, the IGP module 220 transmits and receives LSAs to construct the topology of the network 100. The IGP module 220 includes the LDP-IGP synchronization LSA module 285 which excludes links from the LSAs whose corresponding broadcast interfaces are not cut-edge interfaces until LDP becomes operational to all neighbors on that interface. In other words, for a broadcast interface, if there is an alternate path to the broadcast network, the LDP-IGP synchronization LSA module 285 will exclude that link corresponding to the broadcast interface from its LSAs until the LDP and IGP have synchronized (e.g., when LDP is operational with all neighbors on that broadcast interface). In one embodiment, if LDP is not fully operational for a broadcast interface, a suppress link from LSA flag is set in the local IGP interface structure 228.

[0042] Figure 5A illustrates an exemplary local IGP interface structure 228 according to one embodiment of the invention. As illustrated in Figure 5A, the local IGP interface structure 228 includes the interface field 510 which identifies a local IGP interface, the cut-edge interface property field 515 which indicates whether an interface is a cut-edge interface, and the suppress link from LSA flag 520 which indicates whether the interface should be suppressed from the LSA.

[0043] The LDP module 240 negotiates labels with the other network elements in the network 100. In one embodiment, the LDP module 240 determines when LDP is fully operational for a particular broadcast interface and notifies the IGP module 220 that LDP is fully operational for that interface. The IGP module 220 may then clear the suppress link from LSA flag 520 if it has been set (if the interface was not a cut-edge interface) and advertise that link to the other network elements. The LDP module 240 is also coupled with the label manager 250, which creates and maintains the LSP structure(s) 255 which, among other things, create and manage the labels for the LSPs. The labels are programmed to one or more label forwarding information bases (LFIBs) in the data plane 260. For example, the labels stored in the LSP structure(s) 255 are programmed to one or more packet processing units of one or more line cards in the network element 130.

[0044] It should be understood that in some embodiments cut-edge interfaces are not excluded from the LSA since they are the only route to the broadcast network. Thus, since there is no other path for the traffic to flow, the traffic for any LSP will only start after LDP is operational.

[0045] Figure 3 is a flow diagram illustrating exemplary operations for LDP-IGP synchronization for broadcast networks according to one embodiment of the invention. The operations of Figure 3 will be described with reference to the exemplary embodiments of Figures 1 and 2. However, it should be understood that the operations of Figure 3 can be performed by embodiments of the invention other than those discussed with reference to Figures 1 and 2, and the embodiments discussed with reference to Figures 1 and 2 can perform operations different than those discussed with reference to Figure 3.

[0046] At operation 310, a neighbor up event is detected on a broadcast interface. For example, the IGP module 220 may detect neighbor up event by receiving a hello packet message from a network element on the interface 132 which is not represented in the neighbor adjacency table 222, where the hello packet includes an identifier of the network element 132, and their LSA information is exchanged and synchronized (it should be understood that this link state advertisement information does not include the interface to the broadcast network, i.e., the broadcast interface 132). Flow moves from block 310 to block 320, where the IGP module 220 determines if the new neighbor is the first neighbor on the interface. For example, if the neighbor adjacency table 222 does not include an entry for the interface in the neighbor adjacency table 222, then the new neighbor is the first neighbor on the interface.

[0047] If the discovered neighbor is the first neighbor on the interface, then flow moves to block 335, otherwise flow moves to block 350. At block 315, an interface to the broadcast network is detected as up. For example, the interface state manager 235 detects that the interface 132 is operational. Flow moves from block 315 to block 335, where the IGP module 220 determines whether the detected interface 132 is a cut-edge interface to the broadcast pseudo-node 120. For example, the IGP module 220 accesses the local IGP interface structure 228 to check the cut-edge property of the interface 132. With reference to Figure 1, the interface 132 is not a

cut-edge interface (and the cut-edge property will indicate as such) since an alternate path exists to the broadcast pseudo-node 120. If the interface is not a cut edge, flow moves to block 340, otherwise flow moves to block 350.

[0048] At block 340, the IGP module 220 determines if LDP-IGP synchronization is complete for the interface 132. For example, the IGP module 220 accesses the local IGP interface structure 228 to check the status of the suppress link to LSA flag for the interface 132. The flag will indicate whether IGP module 220 will treat LDP and IGP as being synchronized (e.g., if the flag is set, then the IGP module 220 treats LDP and IGP as not being synchronized and if the flag is not set, then the IGP module 220 treats LDP and IGP as being synchronized). If LDP-IGP synchronization is complete, flow moves to block 325; otherwise flow moves to block 345. At block 345, since LDP-IGP synchronization is not complete and the interface is not a cut-edge interface, the suppress link to LSA flag is set. Flow moves from block 345 to block 350 where processing is continued and only those links which are not suppressed (as indicated by the suppress link LSA flag) are advertised to the broadcast network..

[0049] For example, the IGP module 220 postpones transmitting an LSA that includes the broadcast interface 132 to the broadcast pseudo-node 120 until LDP is operational on all neighbors on the interface 132. Thus, unlike the RFC “LDP IGP Synchronization” which would advertise the link with maximum cost until LDP is operational, in embodiments of the invention the link is not advertised until LDP is operational. By not advertising the link until LDP and IGP are synchronized, the other network elements in the network 100 cannot use that link because the bidirectional connectivity check would fail until LDP and IGP are synchronized. Thus, unlike the examples previously discussed with reference to the RFC “LDP IGP Synchronization,” traffic will not be redirected to sub-optimal links and traffic will not be black-holed using embodiments of the invention.

[0050] Figure 4 is a flow diagram illustrating exemplary operations when LDP-IGP synchronization is complete according to one embodiment of the invention. The operations of Figure 4 will be described with reference to the exemplary embodiments of Figures 1 and 2. However, it should be understood that the operations of Figure 4 can be performed by embodiments of the invention other than

those discussed with reference to Figures 1 and 2, and the embodiments discussed with reference to Figures 1 and 2 can perform operations different than those discussed with reference to Figure 4.

[0051] At block 410, responsive to the LDP module 240 exiting from the LDP-IGP synchronization process for an interface, the LDP module 240 sends a LDP-IGP sync notification message to the IGP module 220. In one embodiment, the LDP-IGP synchronization process will end upon a configurable timer that allows LDP session establishment expiring. For example, based on empirical data, the network administrator may configure a timer which provides a worst case (or best guess) of the time it should take for LDP to become operational on all neighbors of the interface. The configuration value for the timer may be included in the LDP-IGP synchronization for broadcast networks configuration parameters received by the IGP module 220.

[0052] In another embodiment the network element 130 may implement the LDP End-of-LIB mechanism as specified in IETF draft "LDP End-of-LIB: draft-ietf-mpls-end-of-lib-03.txt", January 2009, to determine when the LDP-IGP synchronization process exits. For example, in LDP End-of-LIB, each LDP peer (neighbor) may signal completion of its label advertisements following session establishment. After receiving all completion signals from each LDP peer, the LDP will be operational with all neighbors and the LDP-IGP synchronization process will exit.

[0053] Flow moves from block 410 to block 415, where the IGP module 220 determines if the suppress link to the LSA flag 520 is set for the interface. If the flag is not set, then flow moves to block 420 where alternative action is taken (e.g., nothing needs to be performed). If the flag is set, then flow moves to block 425 where the IGP module 220 clears the suppress link to the LSA flag 520 for the interface. Flow moves from block 425 to block 430.

[0054] In one embodiment of the invention, an interface will be excluded from the link state database 224 if the suppress link to LSA flag is set. Thus, at block 430, the IGP module 220 updates the link state database 224 with the interface. Flow moves from block 430 to block 435, where the IGP module 220 runs a SPF

algorithm on the link state database to possibly update the local IGP RIB 226 with a better route. Flow next moves to block 440, where the data plane (e.g., the FIB(s) 270) is programmed with the updated route information. Flow next moves to block 445, where the IGP module 220 includes the interface in its LSA and then flow moves to block 450 where the IGP module 220 transmits the LSA including the interface to its neighbors.

[0055] Figure 6 is a flow diagram illustrating exemplary operations when a network element is ready to update its link-state advertisement with an interface to a broadcast network according to one embodiment of the invention. The operations of Figure 6 will be described with reference to the exemplary embodiments of Figure 2. However, it should be understood that the operations of Figure 6 can be performed by embodiments of the invention other than those discussed with reference to Figure 2, and the embodiments discussed with reference to Figure 2 can perform operations different than those discussed with reference to Figure 6.

[0056] At block 610, the LDP-IGP synchronization LSA module 285 begins an update link state advertisement process. The update link state advertisement process will include an attempt to add a broadcast interface to the broadcast network (e.g. the broadcast pseudo-node 120). Flow moves from block 610 to block 620, where the LDP-IGP synchronization LSA module 285 determines whether the suppress link to LSA flag is set for the interface. If the flag is not set, then flow moves to block 630 where the interface is included in the link state advertisement. If the IGP is OSPF, the Router-LSA is updated with a Link Type 2 (link to transit network) for the subnet. If the IGP is ISIS, the Link State PDU is updated with an IS Reachability TLV (or an Extended IS Reachability TLV). If the flag is set, then flow moves to block 640 where that interface will be excluded from the link state advertisement.

[0057] Unlike the RFC 5443 “LDP IGP Synchronization,” the LDP-IGP synchronization for broadcast networks mechanism described herein is not required to be implemented on each of the network elements 110 130, 140, 145, and 150 in the network 100. In some embodiments, at a minimum, only a network element that is being added to the broadcast network implements the LDP-IGP synchronization for broadcast networks mechanism described herein. Thus, the LDP-IGP

synchronization for broadcast networks mechanism described herein may be introduced into a network in a gradual manner without any backward compatibility issues.

[0058] Furthermore, unlike the RFC 5443 “LDP IGP Synchronization” which requires that each network element advertising a maximum cost to the broadcast network revert back to advertising their normal cost after LDP peering between all the network elements on the broadcast network is operational (which is a non-trivial process that requires coordination between all the network elements), embodiments of the invention do not require any action to be taken by the other network elements other than the network element whose link to the broadcast network is coming up.

[0059] Figure 7 illustrates an exemplary network element used in some embodiments of the invention. As illustrated in Figure 7, the network element 130 includes the control cards 715 and 720 (e.g., one control card is active the other is a backup), the resource cards 725A-725N, and the line cards 730A-730N. It should be understood that the architecture of the network element 130 illustrated in Figure 7 is exemplary, and different combinations of cards may be used in other embodiments of the invention.

[0060] Each of the cards illustrated in Figure 7 include one or more processors and one or more memories. For example, the line cards 730A-730B typically include one or more packet processing units to process packets including forwarding and/or switching packets at high speed, and include one or more memories to store a forwarding information base (sometimes referred to as a routing table) and a label forwarding information base. The control cards 715 and 720 also include one or more processors to perform signaling, routing (including creation of and/or management of routing tables), connection setup, session setup, etc. For example, among other things, the control card 715 execute instructions stored in memory to execute the IGP module 220 and the LDP module 240.

[0061] While the flow diagrams in the figures show a particular order of operations performed by certain embodiments of the invention, it should be understood that such order is exemplary (e.g., alternative embodiments may perform

the operations in a different order, combine certain operations, overlap certain operations, etc.)

[0062] While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described, can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting.

CLAIMS

What is claimed is:

1. A method for LDP-IGP (label distribution protocol-interior gateway protocol) synchronization on a broadcast network performed on a network element, comprising:
 - detecting that a broadcast interface of the network element has become operational, wherein the broadcast interface will be used to carry traffic through a label switched path (LSP) between the network element and the broadcast network;
 - determining that there is at least one alternate path to reach the broadcast network besides through the broadcast interface; and
 - postponing including the broadcast interface in a link state advertisement (LSA) until label distribution protocol (LDP) is operational with all neighbors on that broadcast interface.
2. The method of claim 1, further comprising:
 - responsive to determining that LDP is operational with all neighbors on that broadcast interface, including that broadcast interface in a LSA and transmitting that LSA to the broadcast network.
3. The method of claim 1, further comprising:
 - setting a suppress link to LSA flag for the broadcast interface until LDP is operational with all neighbors on that broadcast interface; and
 - wherein postponing of including the broadcast interface in the LSA includes checking the suppress link to LSA flag when updating a LSA to determine whether to add the broadcast interface to the updated LSA.
4. The method of claim 1, further comprising:
 - postponing updating a link state database (LSDB) of the network element with the broadcast interface until LDP is operational with all neighbors on the broadcast interface.

5. The method of claim 1, wherein determining that there is at least one alternate path to reach the broadcast network includes the following:
 - during execution of a shortest path first (SPF) algorithm, determining that a network prefix associated with the broadcast interface belongs to a down interface.
6. The method of claim 1, wherein LDP is determined to be operational with all neighbors after a LDP-IGP synchronization timer has expired.
7. The method of claim 1, wherein the network element is implementing Open Shortest Path First (OSPF), and wherein postponing including the broadcast interface in an LSA includes postponing updating a router-LSA with a Link Type 2 for the broadcast network until LDP is operational with all neighbors on that broadcast interface.
8. The method of claim 1 wherein the network element is implementing intermediate system to intermediate system (IS-IS), and wherein postponing including the broadcast interface in a LSA includes postponing updating a Link State PDU with an IS Reachability TLV for the broadcast network until LDP is operational with all neighbors on that broadcast interface.
9. The method of claim 1, wherein the traffic is Virtual Private Network (VPN) traffic.
10. A network element to perform LDP-IGP (label distribution protocol-interior gateway protocol) synchronization on a broadcast network, comprising:
 - one or more processors to execute instructions; and
 - one or more memories coupled with the processor, the one or more memories having instructions that, when executed cause the one or more processors to perform operations including,
 - detecting that a broadcast interface of the network element has become operational, wherein the broadcast interface will be used to carry traffic through a label switched path (LSP) between the network element and the broadcast network,

determining that there is at least one alternate path to reach the broadcast network besides through the broadcast interface, and postponing including the broadcast interface in a link state advertisement (LSA) until label distribution protocol (LDP) is operational with all neighbors on that broadcast interface. .

11. The network element of claim 10, wherein the one or more memories further has instructions which when executed cause the one or more processors to perform the following operations:

responsive to determining that LDP is operational with all neighbors on that broadcast interface, including that broadcast interface in a LSA and transmitting that LSA to the broadcast network.

12. The network element of claim 10, wherein the one or more memories further has instructions which when executed cause the one or more processors to perform the following operations:

setting a suppress link to LSA flag for the broadcast interface until LDP is operational with all neighbors on that broadcast interface; and wherein postponing of including the broadcast interface in the LSA includes checking the suppress link to LSA flag when updating a LSA to determine whether to add the broadcast interface to the updated LSA.

13. The network element of claim 10, wherein the one or more memories further has instructions which when executed cause the one or more processors to perform the following operations:

postponing updating a link state database (LSDB) of the network element with the broadcast interface until LDP is operational with all neighbors on the broadcast interface.

14. The network element of claim 10, wherein determining that there is at least one alternate path to reach the broadcast network includes the following:

during execution of a shortest path first (SPF) algorithm, determining that a network prefix associated with the broadcast interface belongs to a stub network.

15. The network element of claim 10, wherein LDP is determined to be operational with all neighbors after a LDP-IGP synchronization timer has expired.
16. The network element of claim 10, wherein the network element is implementing Open Shortest Path First (OSPF), and wherein postponing including the broadcast interface in an LSA includes postponing updating a router-LSA with a Link Type 2 for the broadcast network until LDP is operational with all neighbors on that broadcast interface.
17. The network element of claim 10 wherein the network element is implementing intermediate system to intermediate system (IS-IS), and wherein postponing including the broadcast interface in a LSA includes postponing updating a Link State PDU with an IS Reachability TLV for the broadcast network until LDP is operational with all neighbors on that broadcast interface.
18. A network element to perform label distribution protocol (LDP)-interior gateway protocol (IGP) synchronization for broadcast networks, comprising:
 - an IGP module to perform the following:
 - determine whether broadcast interfaces of the network element are cut-edge interfaces to broadcast networks, wherein a cut-edge interface is an interface to a broadcast network and no other paths to that broadcast network exist,
 - postpone including those broadcast interfaces that not cut-edge interfaces in link state advertisements until LDP is operational with all neighbors on those broadcast interfaces;
 - and
 - an LDP module coupled with the IGP module, the LDP module to exchange label bindings with neighbors of the network element for the broadcast interfaces.

19. The network element of claim 18, wherein the IGP module further is to postpone updating a link state database with those broadcast interfaces that are not cut-edge interfaces until LDP is operational with all neighbors on those broadcast interfaces.

20. The network element of claim 18, wherein the IGP module is further to perform a shortest path first (SPF) algorithm that includes determining whether broadcast interfaces are cut-edge interfaces.

21. The network element of claim 18, wherein the IGP module is further to include those broadcast interfaces that are cut-edge interfaces in link state advertisements regardless whether LDP is operational with all neighbors on those broadcast interfaces.




FIG. 1




FIG. 2

3/7

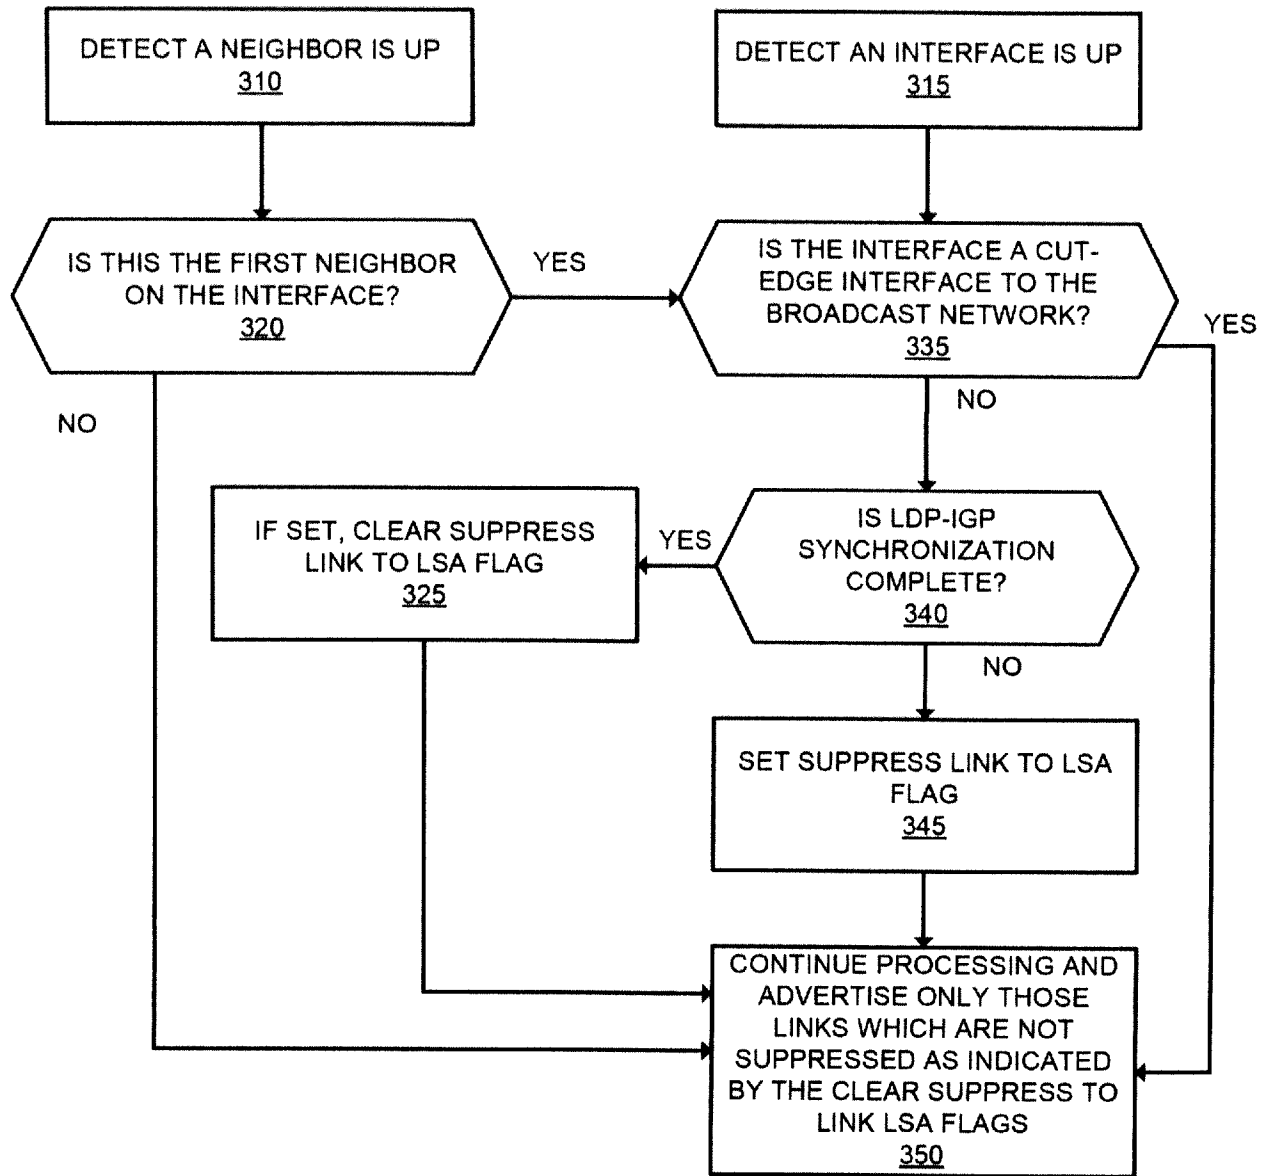


FIG. 3

4/7

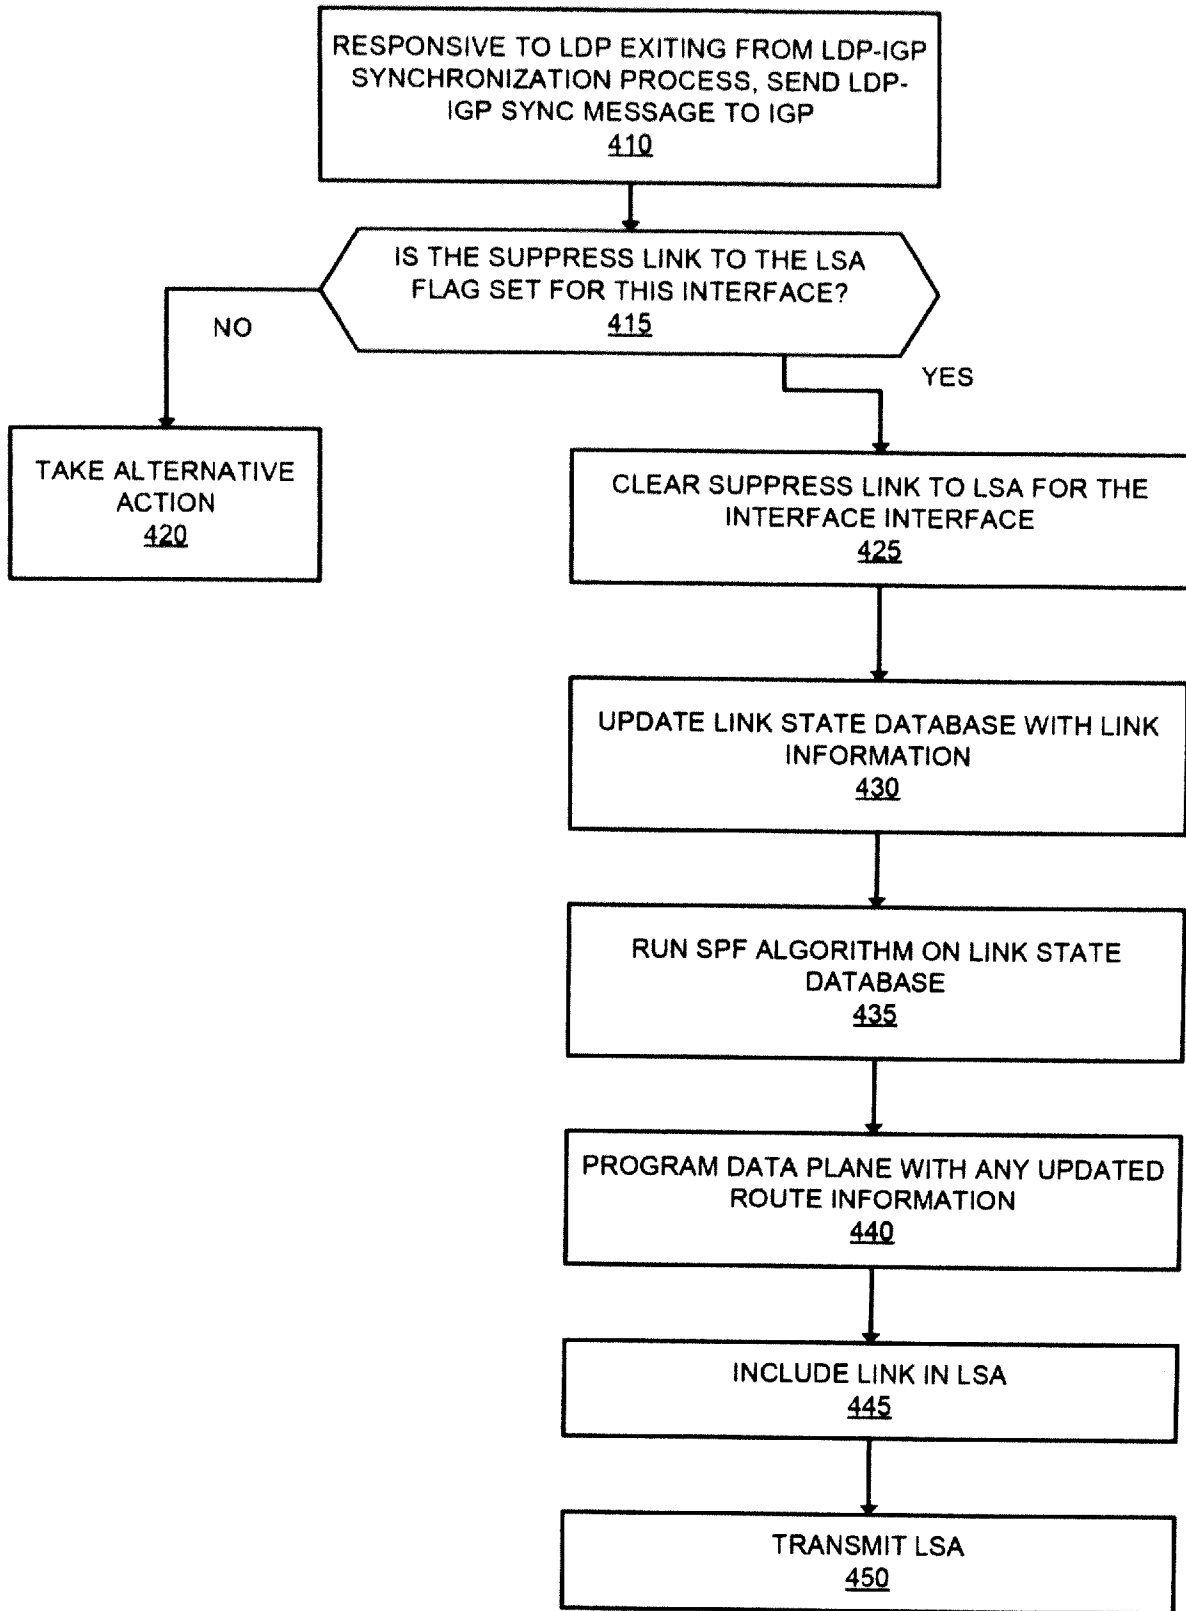


FIG. 4

LOCAL IGP INTERFACE STRUCTURE 228

INTERFACE 510	CUT-EDGE 515	SUPPRESS LINK FLAG 520
132	NO	SET

FIG. 5A

LINK STATE DATABASE 224

FROM 530	TO 540	LINK 550
130	120	132
130	160	134
140	120	142
140	155	144
145	120	146
145	170	148
150	120	152
150	165	154
110	120	114
110	105	112
110	165	116
120	110	122
120	130	124
120	140	126
120	145	129
120	150	128

FIG. 5B

6/7

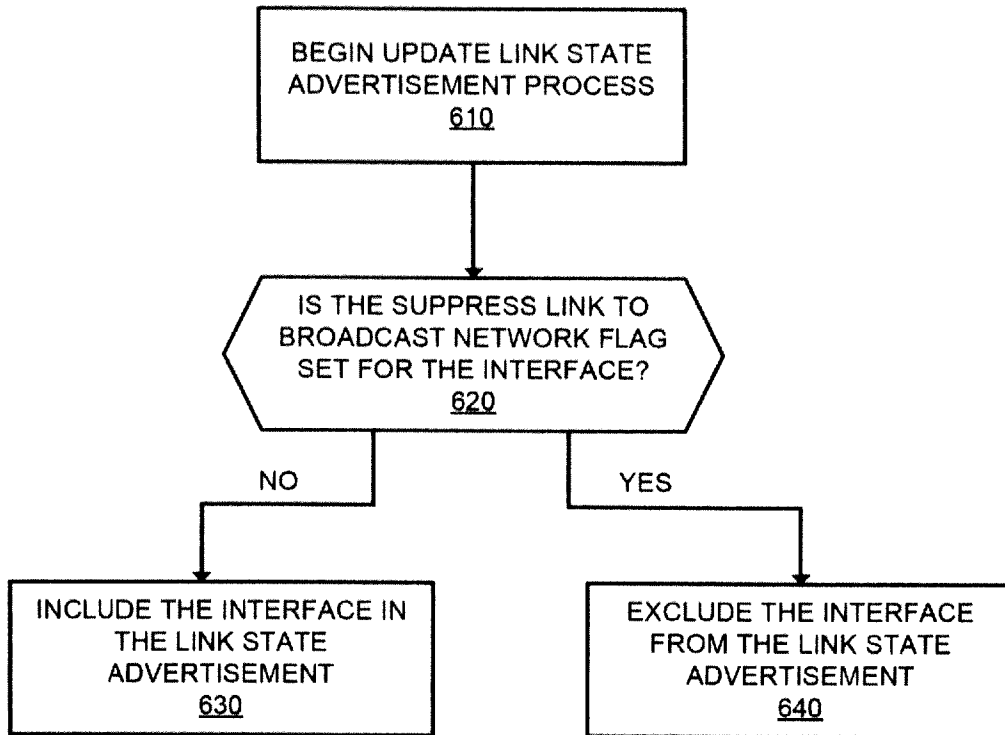


FIG. 6

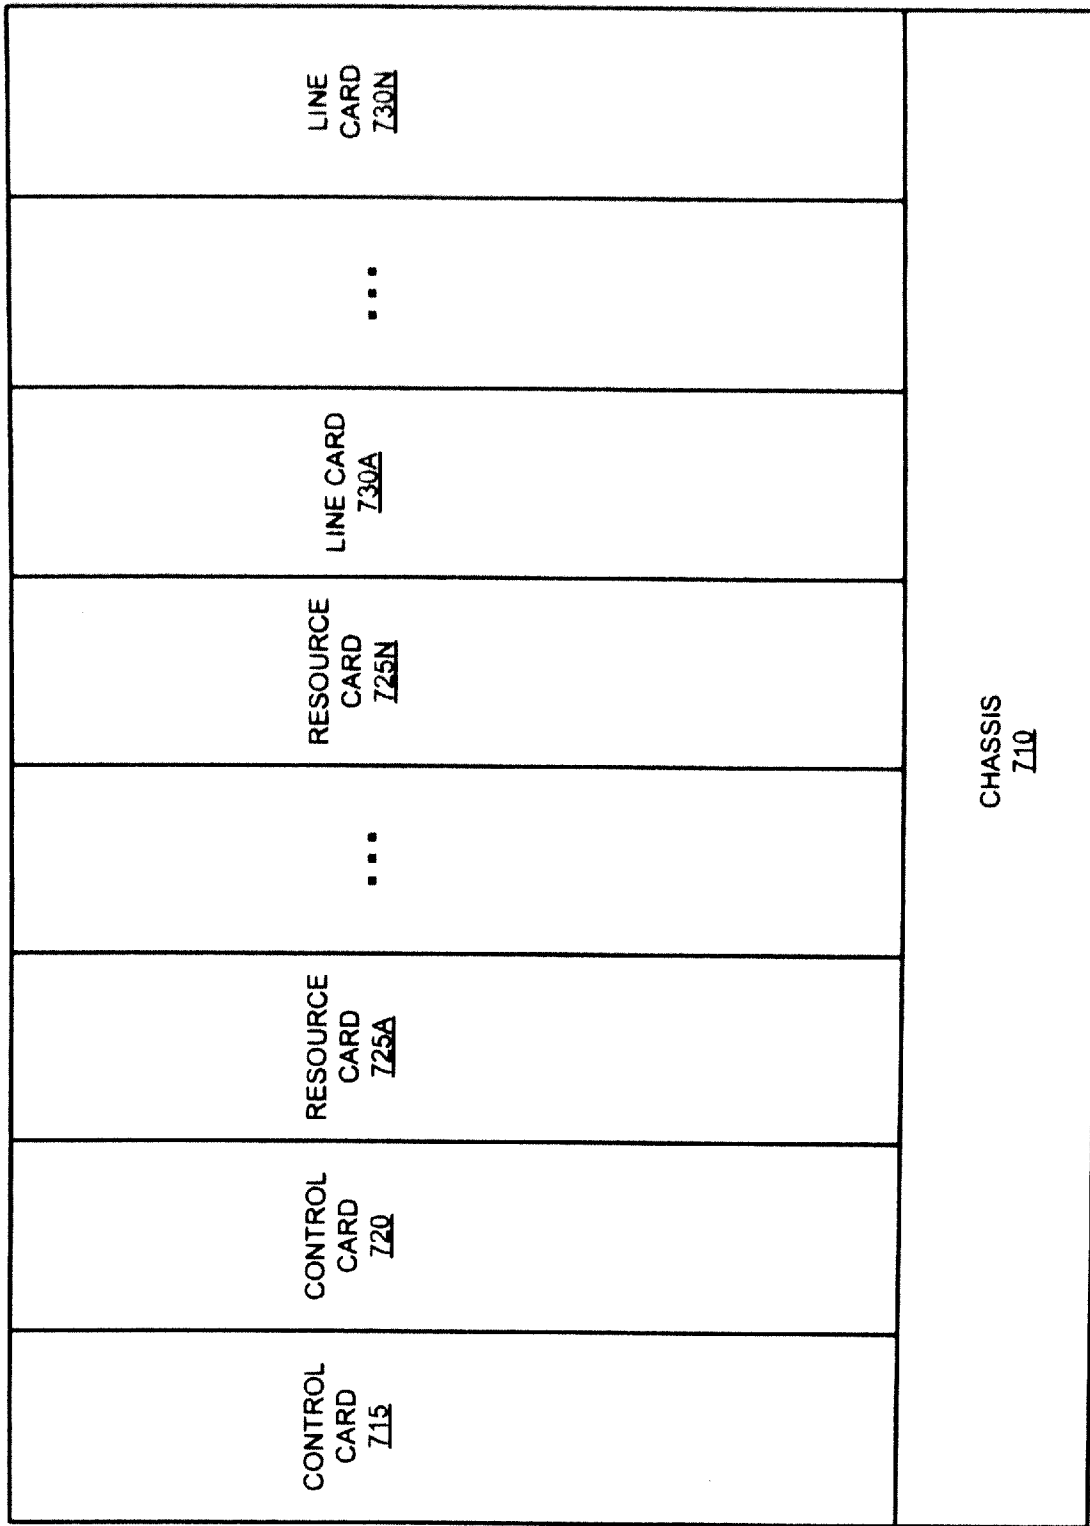


FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2010/023947

A. CLASSIFICATION OF SUBJECT MATTER IPC(8) - H04J 3/06 (2010.01) USPC - 370/503 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) IPC(8) - H04J 3/06 (2010.01) USPC - 370/503, 392 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) Patbase		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2008/0101418 A1 (VASSEUR) 01 March 2008 (01.03.2008) entire document	1-21
Y	US 2007/0019647 A1 (ROY et al) 25 January 2007 (25.01.2007) entire document	1-21
Y	US 7,042,850 B2 (STEWART) 09 May 2006 (09.05.2006) entire document	1-21
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/>		
* Special categories of cited documents: "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "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) "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 "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 "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 "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 "&" document member of the same patent family		
Date of the actual completion of the international search 01 April 2010		Date of mailing of the international search report 16 APR 2010
Name and mailing address of the ISA/US Mail Stop PCT, Attn: ISA/US, Commissioner for Patents P.O. Box 1450, Alexandria, Virginia 22313-1450 Facsimile No. 571-273-3201		Authorized officer: Blaine R. Copenheaver PCT Helpdesk: 571-272-4300 PCT OSP: 571-272-7774