



(43) International Publication Date
12 December 2019 (12.12.2019)

(51) International Patent Classification:

H04L 12/721 (2013.01) H04L 12/707 (2013.01)
H04L 12/751 (2013.01) H04L 12/717 (2013.01)

(21) International Application Number:

PCT/US2019/030440

(22) International Filing Date:

02 May 2019 (02.05.2019)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/680,356 04 June 2018 (04.06.2018) US

(71) Applicant: **FUTUREWEI TECHNOLOGIES, INC.**
[US/US]; 5700 Tennyson Parkway, Suite 600, Plano, Texas
75024 (US).

(72) Inventors: **ECKERT, Toerless**; 777 W Middlefield Road,
Apt 79, Mountain View, California 94043 (US). **CHUN-
DURI, Uma S.**; 41055 Corriea Court, Fremont, Califor-

nia 94539 (US). **QU, Yingzhen**; 2193 Barrett Avenue, San
Jose, California 95124 (US).

(74) Agent: **DIETRICH, William H.** et al.; Conley Rose, P.C.,
5601 Granite Parkway, Suite 500, Plano, Texas 75024 (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, DJ, DK, DM, DO,
DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN,
HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP,
KR, KW, KZ, LA, LC, LK, LR, LS, 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, SA,
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, ST, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,

(54) Title: PREFERRED PATH ROUTE GRAPHS IN A NETWORK

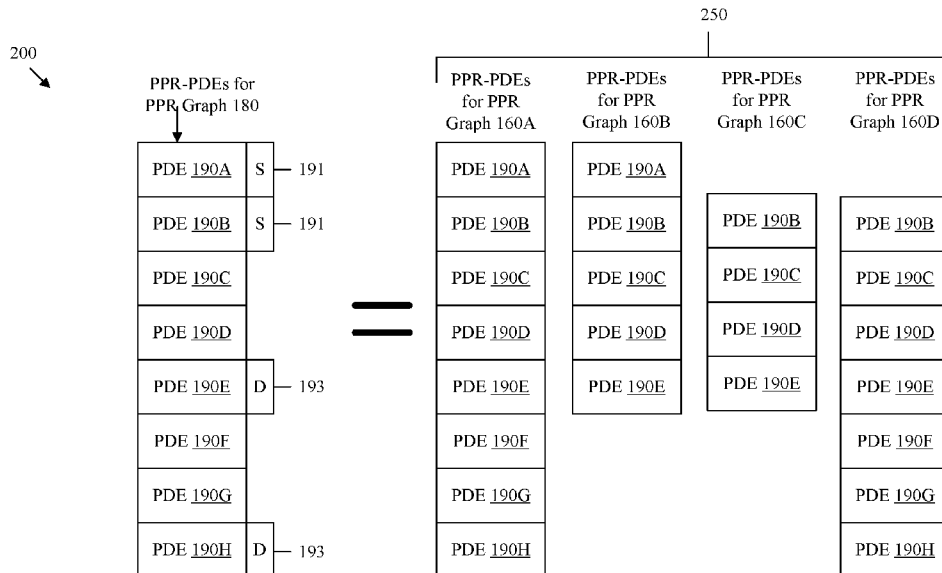


FIG. 2

(57) Abstract: A method implemented by a network element (NE) in a network, comprising receiving, by the NE, preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, and updating, by the NE, a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.



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,
KM, ML, MR, NE, SN, TD, TG).

Published:

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

Preferred Path Route Graphs in a Network

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims the benefit of U.S. Provisional Patent Application No. 62/680,356 filed June 4, 2018 by Toerless Eckert, et. al., and entitled “Preferred Path Routing (PPR) with Graph Structures,” which is incorporated herein by reference as if reproduced in its entirety.

FIELD OF INVENTION

[0002] The present disclosure pertains to the field of preferred path routing. In particular, the present disclosure relates to establishing and provisioning preferred path route (PPR) graphs in a network, in which a PPR graph represents multiple paths from any source to any destination included in the PPR graph.

BACKGROUND

[0003] In a network implementing source routing, a controller that has knowledge of a complete topology of the underlying network can program an ingress node of the network with a custom path that certain traffic has to travel to reach a destination. This custom path may not necessarily be the shortest path between the ingress node and egress node. An ingress node in the network may use a separate flow classification technique (e.g., source and/or destination addresses or transport port numbers) to associate certain traffic flow with a custom path.

[0004] In a network implementing segment routing (SR), packets are steered through the network using segment identifiers (SIDs) that uniquely identify segments in the network. A segment may include one or more nodes, interfaces, and links between two nodes in the network. The SIDs are typically carried in the header of the packet.

[0005] Currently there are two data planes that use segment routing to provision custom paths in a network - Segment Routing (SR) Multi-Protocol Label Switching (MPLS) (SR-MPLS) and SR-Internet Protocol (IP) Version 6 (IPv6) (SR-IPv6). In SR-MPLS, each segment is encoded as a label, and an ordered list of segments are encoded as a stack of labels in the header of the packet. Similarly, in SR-IPv6, each segment is encoded as an IPv6 address within a segment routing header (SRH).

SUMMARY

[0006] According to a first aspect of the present disclosure, there is provided a method implemented by a network element NE in a network. The method comprises receiving, by the NE, preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag, and updating, by the NE, a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

[0007] In a first implementation of the method according to the first aspect, the PPR graph comprises a plurality of branches, wherein each branch comprises a plurality of elements on a path between two NEs included in the PPR graph.

[0008] In a second implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR graph comprises a plurality of ingress NE and only the egress NE, wherein the PPR graph comprises a plurality branches, wherein each of the branches comprises a plurality of elements on a path connecting one of the ingress NEs to an NE along the path to the egress NE, and wherein each of the branches includes at least one common NE.

[0009] In a third implementation of the method according to the first aspect or any preceding implementation of the first aspect, the method further comprises determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information, sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs, determining, by the NE, the egress NE of the first branch based on the destination flag in response to being identified in the one of the plurality of PPR-PDEs, and generating, by the NE, the forwarding entry for the egress NE indicating that a next hop toward the egress NE is a next PPR-PDE of the plurality of PPR-PDEs.

[0010] In a fourth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR graph comprises a plurality of ingress NEs and a

plurality of egress NEs, wherein the PPR graph comprises a plurality branches, wherein each of the branches comprises a plurality of elements on a path connecting one of the ingress NEs to an NE along the path to one of the egress NEs, and wherein each of the branches includes at least one common NE.

[0011] In a fifth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the method further comprises determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information, sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs, determining, by the NE, a plurality of egress NEs of the first branch in response to being identified in the one of the plurality of PPR-PDEs, and generating, by the NE, a plurality of forwarding entries for the plurality of egress NEs.

[0012] In a sixth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information includes a flag indicating that the PPR graph is a bidirectional forest, wherein the PPR graph comprises a plurality of ingress NEs and a plurality egress NEs, wherein the PPR graph comprises a plurality branches, wherein each of the plurality of branches comprises a plurality of elements on a bidirectional path connecting one of the ingress NEs to an NE along the bidirectional path to one of the egress NEs, and wherein each of the branches includes at least one common NE.

[0013] In a seventh implementation of the method according to the first aspect or any preceding implementation of the first aspect, the method further comprises determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information, sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs, determining, by the NE, a first plurality of destinations along a first direction of the first branch in response to being identified in the one of the plurality of PPR-PDEs, generating, by the NE, a plurality of forwarding entries for the first plurality of destinations, determining, by the NE, a second plurality of destinations along a second direction of the first branch, and generating, by the NE, a plurality of forwarding entries for the second plurality of destinations.

[0014] In an eighth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the forwarding database is not updated in response to

identifying the NE in the one of the plurality of PPR-PDEs and determining that a next hop is on a different branch.

[0015] In a ninth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR graph comprises a single source and a single destination.

[0016] In a tenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information describing the PPR graph is received from a central entity of the network or an administrator of the network.

[0017] In an eleventh implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information is transmitted to a plurality of other NEs in the network.

[0018] In a twelfth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the forwarding entry comprises an address prefix indicating the next hop, wherein the address prefix represents a plurality of NEs along a branch of the PPR graph.

[0019] In a thirteenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, a first PPR-PDE on a branch of the PPR graph includes a flag indicating that the branch is cyclical.

[0020] In a fourteenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information includes a type field indicating that the plurality of PPR-PDEs describes a tree, a forest, or a bidirectional forest.

[0021] In a fifteenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR graph comprises a plurality of ingress NEs and a plurality of egress NEs, wherein each of the plurality of egress NEs is represented by a different PPR-ID, and wherein a forwarding entry is created for each of the plurality of egress NEs.

[0022] In a sixteenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information includes a policy group tag associated with one or more PPR graphs, wherein the one or more PPR graphs are assigned the policy group tag based on an attribute associated with the one or more PPR graphs, wherein the attribute is at least one of a traffic engineering (TE) parameter, network characteristic, or service requirement that is associated with a resource to be reserved on each of one or more PPR graphs.

[0023] In a seventeenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the PPR information is received from another NE in the network as a plurality of fragments, wherein each fragment includes a PPR graph identifier (PPG-ID), a fragment identifier, a last fragment flag, and a subset of the plurality of PPR-PDEs, wherein the PPG-ID identifies the PPR graph described by the PPR information, wherein the fragment identifier indicates a sequence number of a respective fragment, and wherein the last fragment flag indicates whether the respective fragment is a last fragment representing the PPR graph of the plurality of fragments.

[0024] In a eighteenth implementation of the method according to the first aspect or any preceding implementation of the first aspect, the forwarding entry is generated based on the plurality of fragments after the respective fragment having the last fragment flag set is received.

[0025] According to a second aspect of the present disclosure, there is provided a method implemented by a central entity in a network. The method comprises determining, by the central entity, preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag, and transmitting, by the central entity, the PPR information describing the PPR graph to a headend NE of the network.

[0026] According to a third aspect of the present disclosure, there is provided a NE implemented in a network. The NE comprises a memory configured to store instructions, and a processor coupled to the memory and configured to execute the instructions to cause the processor to be configured to receive preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag, and update a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of

PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

[0027] According to a fourth aspect of the present disclosure, there is provided a central entity implemented in a network. The NE comprises a memory configured to store instructions, and a processor coupled to the memory and configured to execute the instructions to cause the processor to be configured to determine preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag, and transmit the PPR information describing the PPR graph to a headend NE of the network.

[0028] According to a fifth aspect of the present disclosure, there is provided an apparatus comprising a means for receiving preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag, and a means for updating a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

[0029] According to a sixth aspect of the present disclosure, there is provided an apparatus comprising a means for determining preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag, and a means for transmitting the PPR information describing the PPR graph to a headend NE of the network.

[0030] For the purpose of clarity, any one of the foregoing embodiments may be combined with any one or more of the other foregoing embodiments to create a new embodiment within the scope of the present disclosure.

[0031] These and other features will be more clearly understood from the following detailed description taken in conjunction with the accompanying drawings and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0032] For a more complete understanding of this disclosure, reference is now made to the following brief description, taken in connection with the accompanying drawings and detailed description, wherein like reference numerals represent like parts.

[0033] FIG. 1A is a diagram illustrating a network configured to implement preferred path routing.

[0034] FIG. 1B is a diagram illustrating a network configured to implement preferred path routing and PPR graphs according to various embodiments of the disclosure.

[0035] FIG. 2 is a diagram illustrating the PPR path description elements (PPR-PDEs) describing elements on the PPR graph of FIG. 1B according to various embodiments of the disclosure.

[0036] FIG. 3 is an embodiment of an NE in the network configured to implement preferred path routing and PPR graphs according to various embodiments of the disclosure.

[0037] FIG. 4 is a diagram illustrating a PPR graph implemented similar to a standard PPR according to various embodiments of the disclosure.

[0038] FIG. 5 is a diagram illustrating a PPR graph in the network of FIG. 1B according to various embodiments of the disclosure.

[0039] FIG. 6 is a diagram of a PPR graph implemented as a tree according to various embodiments of the disclosure.

[0040] FIG. 7 is a method for processing PPR information describing the PPR graph of FIG. 6 according to various embodiments of the disclosure.

[0041] FIG. 8 is a diagram illustrating another PPR graph in the network of FIG. 1B according to various embodiments of the disclosure.

[0042] FIG. 9 is a diagram of a PPR graph implemented as a forest according to various embodiments of the disclosure.

[0043] FIG. 10 is a diagram of a PPR graph implemented as a bidirectional forest according to various embodiments of the disclosure.

[0044] FIG. 11 is a diagram of an advertisement including PPR information according to various embodiments of the disclosure.

[0045] FIGS. 12A-C are type-length-value (TLV) diagrams illustrating various fields of an example advertisement according to various embodiments of the disclosure.

[0046] FIG. 13 is a diagram illustrating fragments used to carry PPR information for a PPR graph according to various embodiments of the disclosure.

[0047] FIGS. 14A-B are TLV diagrams showing examples of fields within a fragment according to various embodiments of the disclosure.

[0048] FIGS. 15-16 are flowcharts illustrating methods for advertising and programming PPR graphs according to various embodiments of the disclosure.

[0049] FIGS. 17-18 are apparatuses for advertising and programming PPR graphs according to various embodiments of the disclosure.

DETAILED DESCRIPTION

[0050] It should be understood at the outset that although an illustrative implementation of one or more embodiments are provided below, the disclosed systems and/or methods may be implemented using any number of techniques, whether currently known or in existence. The disclosure should in no way be limited to the illustrative implementations, drawings, and techniques illustrated below, including the exemplary designs and implementations illustrated and described herein, but may be modified within the scope of the appended claims along with their full scope of equivalents.

[0051] FIG. 1A is a diagram illustrating a network 100 (also referred to herein as a “domain or “subnetwork”) configured to implement preferred path routing. The network 100 comprises a central entity 165 (also referred to herein as a “controller”) and multiple network entities (NEs) 101-121. The NEs 101-121 are interconnected via links 122-154. The central entity 165 is connected to one of the NEs 104 via a central entity-to-domain link 166.

[0052] In an embodiment, the central entity 165 may be a network or domain controller that maintains a topology of the network 100 to craft paths (shortest paths, PPRs, and PPR graphs) between edge NEs 101-121 in the network 100, as will be further described below. In an

embodiment, the central entity 165 is substantially similar to a Path Computation Element (PCE), which is further described in Internet Engineering Task Force (IETF) Request for Comments (RFC) 8281, entitled “Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model,” by E. Crabbe, dated December 2017. In an embodiment, the central entity 165 may be substantially similar to a Software Defined Network Controller (SDNC), which is further described in the IETF RFC 8402 entitled “Segment Routing Architecture,” by C. Filsfils, dated July 2018. In an embodiment, the central entity 165 may be substantially similar to an Application Layer Traffic Optimization (ALTO) server, which is further described in the IETF RFC 7285, entitled “Application Layer Traffic Optimization (ALTO) Protocol,” by R. Alimi, dated September 2014.

[0053] In an embodiment, NEs 101-121 (also referred to herein as “nodes”) may be a topological devices (or physical devices) such as a router, a bridge, a network switch, or a logical device configured to perform switching and routing using the preferred path routing mechanisms disclosed herein. In an embodiment, one or more of the NEs 101-121 may be non-topological NEs such as, for example, a function, context, service, or a virtual machine. A non-topological NE may be implemented by the NEs 101-121 or implemented by another device attached to the NEs 101-121.

[0054] In an embodiment, NEs 101-121 may be headend nodes or edge nodes positioned at an edge of the network 100. While NEs 101-121 are shown in FIG. 1A as headend nodes, it should be appreciated that NEs 101-121 may otherwise be an intermediate node or any other type of NE. Although only twenty one NEs 101-121 are shown in FIG. 1A, it should be appreciated that the network 100 shown in FIG. 1A may include any number of NEs. In an embodiment, the central entity 165 and NEs 101-121 are configured to implement various packet forwarding protocols, such as, but not limited to, Multi-Protocol Label Switching (MPLS), Segment Routing-MPLS (SR-MPLS), Internet Protocol (IP) Version 4 (IPv4), IP Version 6 (IPv6), Next Generation Explicit Routing (NGER), or any future packet forwarding protocol.

[0055] The links 122-154 may be wired links, wireless links, or interfaces interconnecting the NEs 101-121 together. Similarly, the central entity-to-domain link 166 is a wired link, wireless link, or interfaces interconnecting at least one of the NEs 101-121 to the central entity 165.

[0056] In operation, the central entity 165 is configured to determine one or more shortest paths between two edge NEs 101-121 in the network 100 and one or more PPRs 160A-D between

different edge NEs 101-121 in the network 100. A shortest path refers to a path between two NEs 101-121 that is determined based on a metric, such as, for example, a cost or weight associated with each link on the path, a number of NEs on the path, a number of links on the path, etc. In an embodiment, a shortest path may be computed for a destination using a Dijkstra's Shortest Path First (SPF) algorithm.

[0057] A PPR 160A-D (also referred to herein as a "Non-Shortest Path" (NSP)) refers to a custom path or any other path that is determined based on an application or server request for a path between an ingress NE 101-121 and an egress NE 101-121 (or between a source and destination). In an embodiment, the PPR 160A-D deviates from the shortest path. However, the PPR 160A-D may also be the same as the shortest path in some circumstances. The PPR 160A-D includes a sequential ordering of elements 121-154 (e.g., NEs 101-121 and/or links 122-154) along a path in the network 100.

[0058] In an embodiment, the central entity 165 determines the PPRs 160A-D based on a network topology of network 100, which is maintained at the central entity 165. In this embodiment, the central entity 165 generates PPR information 170, describing each of the PPRs 160A-D, and sends the PPR information 170 to an NE 104 via the central entity-to-domain link 166.

[0059] As will be further described below, the PPR information 170 may include details regarding each of the PPRs 160A-D, such as, for example, a PPR-identifier (PPR-ID) of each PPR 160A-D, attributes associated with resources to be reserved on each PPR 160A-D, and multiple PPR-path description elements (PPR-PDEs) describing one or more elements on each PPR 160A-D. In this embodiment, NE 104 floods the PPR information 170 to the remaining NEs 101-103 and 105-121 in the network 100 using the underlying Interior Gateway Protocol (IGP) of the network 100. For example, NE 104 transmits the PPR information 170 to neighboring NEs 103 and 105. NE 103 forwards the PPR information 170 to neighboring NE 102, and NE 105 forwards the PPR information 170 to neighboring NE 106. In this way, the remaining NEs 101-103 and 105-121 continue to forward the PPR information 170 to all the remaining NEs 101-103 and 105-121 in the network 100. The IGP implemented by the network 100 may be Open Shortest Path First (OSPF) Version 2 (OSPFv2), OSPF Version 3 (OSPFv3), Intermediate System – Intermediate System (IS-IS), or direct SDN.

[0060] In another embodiment, an operator or administrator of the network 100 may determine the PPRs 160A-D and send the PPR information 170 describing the PPRs 160A-D to one of the NEs 101-121 in the network 100. The PPR information 170 may then be flooded to all the remaining NEs 101-121 in the network 100.

[0061] After receiving the PPR information 170, each of the NEs 101-121 is configured to determine whether the respective NE 101-121 is identified in the PPR information 170 describing one or more of the PPRs 160A-D. If so, then the respective NE 101-121 updates the local forwarding database by adding a forwarding entry defining a next hop for a particular destination or egress NE identified by the PPR information 170. If not, then the respective NE 101-121 ignores the PPR information 170.

[0062] In a network 100 implementing preferred path routing, a separate PPR 160A-D is determined and provisioned for each ingress and egress NE 101-121 pair. As shown by FIG. 1A, PPR 160A is a path provisioned between ingress NE 101 and egress NE 107. PPR 160A includes the following elements: ingress NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and egress NE 107. PPR 160B is a path provisioned between ingress NE 101 and egress NE 118. PPR 160B includes the following elements: ingress NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, and NE 118. PPR 160C is a path provisioned between ingress NE 121 and egress NE 118. PPR 160C includes the following elements: ingress NE 121, link 150, NE 120, link 149, NE 119, link 148, and NE 118. PPR 160D is a path provisioned between ingress NE 121 and egress NE 107. PPR 160D includes the following elements: ingress NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and egress NE 107.

[0063] In this case, the central entity 165 or an administrator of the network 100 generates the PPR information 170 to include details regarding each of PPRs 160A-D and each of the elements on each of PPRs 160A-D. For example, a separate PPR-ID is determined for each of the PPRs 160A-D, and a separate PPR-PDE is created for one or more of the elements on each of the PPRs 160A-D. This PPR information 170 is flooded through the network 100 using the underlying IGP of the network such that one more of NEs 101-121 in network 100 stores the PPR information 170.

[0064] Therefore, to provision PPRs 160A-D in network 100, PPR information 170 for each of the PPRs 160A-D has to be created, flooded through the network 100, and stored at one or more of the NEs 101-121. This creates a scalability issue across the network 100, in that the amount of

PPR information 170 that has to be created, forwarded, and stored is extensive when there is a large number of PPRs 160A-D to be provisioned in the network 100.

[0065] Disclosed herein are embodiments directed to creating PPR graphs that represent a plurality of PPRs 160A-D between one or more ingress NEs 101-121 and one or more egress NEs 101-121 in the network 100. Instead of creating PPR information 170 for each PPR 160A-D in a network 100, the PPR information 170 describes PPR graphs, as disclosed herein. The PPR graphs include flags or bits to indicate whether an NE 101-121 is an ingress NE 101-121 or an egress NE 101-121. In this way, the amount of PPR information 170 that has to be created, forwarded, and stored across the network 100 is dramatically decreased. Therefore, the embodiments disclosed herein enable PPRs 160A-D to be provisioned in a more resource efficient manner that saves computing resources and network resources.

[0066] FIG. 1B is a diagram illustrating a network 175 configured to implement preferred path routing and PPR graphs 180 according to various embodiments of the disclosure. Network 175 is similar to network 100, except that network 175 is configured to provision a single PPR graph 180, instead of multiple PPRs 160A-D. In an embodiment, a PPR graph 180 represents multiple PPRs 160A-D between one or more ingress NEs 101-121 (also referred to herein as “sources”) and one or more egress NEs 101-121 (also referred to herein as “destinations”).

[0067] As shown by FIG. 1B, the PPR graph 180 includes the following elements: NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and NE 107. PPR graph 180 includes two ingress NEs 101 and 121 and two egress NEs 118 and 108. The ingress NEs 101 and 121 are represented in FIG. 1B with rectangles around NEs 101 and 121, and the egress NEs 118 and 107 are represented in FIG. 1B with circles around NEs 118 and 107.

[0068] In an embodiment, the PPR graph 180 represents multiple possible paths between the ingress NEs 101 and 121 and the egress NEs 119 and 108. For example, as shown in FIG. 1B, PPR graph 180 includes all of the elements from PPR 160A (FIG. 1A): ingress NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and egress NE 107. PPR graph 180 also includes all of the elements from PPR 160B (FIG. 1A): ingress NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, and NE 118. Similarly, PPR graph 180 includes all of the elements from PPR 160C (FIG. 1A): ingress NE 121, link 150, NE 120, link 149, NE 119, link 148, and NE 118. Finally, PPR graph

180 includes all of the elements from PPR 160D (FIG. 1A): ingress NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and egress NE 107. In this embodiment, the PPR graph 180 includes a PPR 160A between ingress NE 101 and egress NE 107, a PPR 160B between ingress NE 101 and egress NE 118, a PPR 160C between ingress NE 121 and egress NE 118, and a PPR 160D between ingress NE 121 and egress NE 107. Therefore, in various embodiments, a PPR graph 180 includes a plurality of different PPRs 160A-D, having different ingress NEs 101 and 121 and different egress NEs 118 and 107.

[0069] In an embodiment, the central entity 165 or a network administrator determines the PPR graph 180 based on a network topology of network 175 maintained at the central entity 165 and based on a network resource to be reserved for the PPR graph 180. In an embodiment, the central entity 165 or the network administrator generates PPR information 170 describing the PPR graph 180 and sends the PPR information 170 to a headend NE 104 in network 175 via central entity-to-domain link 166.

[0070] As shown by FIG. 1B, the PPR information 170 includes a PPR type 183, one or more PPR-IDs 186, and one or more PPR-PDEs 190. The PPR type 183 indicates a type of PPR graph 180. The different types of PPR graphs 180 will be further described below with reference to FIGS. 4-10. The one or more PPR-IDs 186 includes information identifying various paths within the PPR graph 180. In an embodiment, each of the PPR-IDs 186 includes an address, label, or identifier of each egress NE 118 and 107 included in the PPR graph 180.

[0071] The PPR-PDEs 190 include information identifying one or more elements (e.g., NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and NE 107) on the PPR graph 180. In an embodiment, each of PPR-PDEs 190 includes a label, address, or identifier of one or more of the elements 101-154 (e.g., NEs 101-121 and links 122-154) on the PPR graph 180. In an embodiment, each of the PPR-PDEs 190 includes several flags, such as a source flag 191 and a destination flag 193. The source flag 191 may be a bit that, when set, indicates that the element 101-154 identified by the PPR-PDE 190 is an ingress NE 101 or 121 on the PPR graph 180. The destination flag 193 is also a bit that, when set, indicates that the element 101-154 identified by the PPR-PDE 190 is an egress NE 118 or 107.

[0072] According to various embodiments, transmitting the PPR information 170 describing a single PPR graph 180, which includes all four PPRs 160A-D, instead of separately transmitting PPR information 170 describing multiple PPRs 160A-D, is a more efficient and effective manner

of communicating data regarding multiple different PPRs 160A-D in a network 175. In addition, the amount of data that each of the NEs 101-121 processes and stores is reduced due to the consolidated nature of the PPR information 170 describing multiple PPRs 160A-D. Therefore, the use of PPR graphs 180 enables NEs 101-121 in a network 175 to be more efficiently programmed to forward traffic according to various the PPR graphs 180.

[0073] FIG. 2 is a diagram 200 illustrating the PPR-PDEs 190A-H describing elements on the PPR graph 180 of FIG. 1B, which is included in the PPR information 170, according to various embodiments of the disclosure. In an embodiment, PPR-PDEs 190A-H describe one or more of the elements (e.g., NE 101, link 141, NE 121, link 150, NE 120, link 149, NE 119, link 148, NE 118, link 145, NE 115, link 138, NE 108, link 128, and NE 107) on the PPR graph 180.

[0074] In the example shown in FIG. 2, PPR-PDE 190A includes details regarding NE 101, such as a label, address, or identifier of NE 101. PPR PDE 190A also includes the source flag 191 set to indicate that NE 101 is an ingress NE 101 of the PPR graph 180. An ingress NE 101 is configured to be a source of traffic, or receive traffic from sources external to the network 175, and forward the traffic through the network 175 using the PPR graph 180. PPR-PDE 190B includes details regarding NE 121, such as a label, address, or identifier of NE 121. PPR-PDE 190C also includes the source flag 191 set to indicate that NE 121 is also an ingress NE 121 of the PPR graph 180.

[0075] PPR-PDE 190C includes details regarding NE 120, and PPR-PDE 190D includes details regarding NE 119. PPR-PDE 190E includes details regarding NE 118, and includes the destination flag 193, which is set to indicate that NE 118 is an egress NE 118 of the PPR graph 180. This means that NE 118 is a destination on the PPR graph 180 and configured to forward traffic outside of the network 175, to another network, or to another entity. PPR-PDE 190F includes details regarding NE 115, and PPR-PDE 190G includes details regarding NE 108. PPR-PDE 190H includes details regarding NE 107, and includes the destination flag 193, which again means that the egress NE 107 is a destination on the PPR graph.

[0076] Information 250 shown to the right of diagram 200 shows the PPR-PDEs 190A-H for the four different PPRs 160A-D shown in FIG. 1A, each of which would have had to be forwarded through the network 100, and some of which would have had to be locally stored at one or more NEs 101-121. However, the PPR-PDEs 190A-H including the source flags 191 and destination flags 193 essentially include all of the information 250 shown to the right of diagram 200, in a

much more compact and efficient data structure. Therefore, data that is transmitted through the network 175 to provision PPRs is reduced and the amount of computing resources needed to process the data is also reduced.

[0077] FIG. 3 is a diagram of an embodiment of an NE 300 in a network such as networks 100 and 175. NE 300 may be implemented as the central entity 165 or the NEs 101-121. The NE 300 may be configured to implement and/or support the routing and PPR graph 180 provisioning mechanisms described herein. The NE 300 may be implemented in a single node or the functionality of NE 300 may be implemented in a plurality of nodes. One skilled in the art will recognize that the term NE encompasses a broad range of devices of which NE 300 is merely an example. While NE 300 is described as a physical device, such as a router or gateway, the NE 300 may also be a virtual device implemented as a router or gateway running on a server or a generic routing hardware (whitebox).

[0078] The NE 300 is included for purposes of clarity of discussion, but is in no way meant to limit the application of the present disclosure to a particular NE embodiment or class of NE embodiments. At least some of the features and/or methods described in the disclosure may be implemented in a network apparatus or module such as a NE 300. For instance, the features and/or methods in the disclosure may be implemented using hardware, firmware, and/or software installed to run on hardware. As shown in FIG. 3, the NE 300 comprises one or more ingress ports 310 and a receiver unit (Rx) 320 for receiving data, at least one processor, logic unit, or central processing unit (CPU) 330 to process the data, transmitter unit (Tx) 340 and one or more egress ports 350 for transmitting the data, and a memory 360 for storing the data.

[0079] The processor 330 may comprise one or more multi-core processors and be coupled to a memory 360, which may function as data stores, buffers, etc. The processor 330 may be implemented as a general processor or may be part of one or more application specific integrated circuits (ASICs) and/or digital signal processors (DSPs). The processor 330 may comprise a network configuration module 335, which may perform processing functions of the central entity 165 or the NEs 101-121. The network configuration module 335 may also be configured to perform the steps of methods 700, 1500, and 1600, and/or any other method discussed herein. As such, the inclusion of the network configuration module 335 and associated methods and systems provide improvements to the functionality of the NE 300. Further, the network configuration module 335 effects a transformation of a particular article (e.g., the network) to a different state. In

an alternative embodiment, network configuration module 335 may be implemented as instructions stored in the memory 360, which may be executed by the processor 330.

[0080] The memory 360 may comprise a cache for temporarily storing content, e.g., a random-access memory (RAM). Additionally, the memory 360 may comprise a long-term storage for storing content relatively longer, e.g., a read-only memory (ROM). For instance, the cache and the long-term storage may include dynamic RAMs (DRAMs), solid-state drives (SSDs), hard disks, or combinations thereof. The memory 360 may be configured to store the PPR information 170, which includes the PPR type 183, PPR-IDs 186, and PPR-PDEs 190A-H (hereinafter referred to as “PPR-PDE 190”), which also include source flags 191 and destination flags 193. The PPR information 170 may also include a PPR graph identifier (PPG-ID) 344, which may be a label, address, or identifier uniquely identifying the PPR graph 180. In addition, the memory 360 is configured to store a forwarding database 365 and a link state database 361. In an embodiment, the forwarding database 365 stores forwarding entries 370 describing forwarding rules for how a particular NE 300 (e.g., NE 101-121 of FIGS. 1-2) should forward a data packet that includes a PPR-ID 186 and/or a destination address.

[0081] It is understood that by programming and/or loading executable instructions onto the NE 300, at least one of the processor 330 and/or memory 360 are changed, transforming the NE 300 in part into a particular machine or apparatus, e.g., a multi-core forwarding architecture, having the novel functionality taught by the present disclosure. It is fundamental to the electrical engineering and software engineering arts that functionality that can be implemented by loading executable software into a computer can be converted to a hardware implementation by well-known design rules. Decisions between implementing a concept in software versus hardware typically hinge on considerations of stability of the design and numbers of units to be produced rather than any issues involved in translating from the software domain to the hardware domain. Generally, a design that is still subject to frequent change may be preferred to be implemented in software, because re-spinning a hardware implementation is more expensive than re-spinning a software design. Generally, a design that is stable that will be produced in large volume may be preferred to be implemented in hardware, for example in an ASIC, because for large production runs the hardware implementation may be less expensive than the software implementation. Often a design may be developed and tested in a software form and later transformed, by well-known design rules, to an equivalent hardware implementation in an ASIC that hardwires the instructions

of the software. In the same manner as a machine controlled by a new ASIC in a particular machine or apparatus, likewise a computer that has been programmed and/or loaded with executable instructions may be viewed as a particular machine or apparatus. In some embodiments, the NE 300 may be configured to implement OSPFv2, OSPFv3, IS-IS, or direct SDN controller based on network implementations.

[0082] FIG. 4 is a diagram illustrating a PPR graph 400 configured to replace the standard PPR 160A-D (collectively referred to as PPR 160) according to various embodiments of the disclosure. In an embodiment, the PPR graph 400 is similar to PPR graph 180, in that the PPR graph 400 includes at least one path between at least one ingress NE 401 and at least one egress NE 405. As shown by FIG. 4, the PPR graph 400 includes NEs 401, 402, 403, 404, and 405, each of which are similar to NEs 101-121 of FIGS. 1A-B. NEs 401-405 are interconnected by links 410-413, which are similar to links 122-154 of FIGS. 1A-B.

[0083] PPR graph 400 only includes one ingress NE 401, which is the source of the PPR graph 400 identified by a rectangle surrounding NE 401. PPR graph 400 also only includes one egress NE 405, which is a destination of the PPR graph 400 identified by a circle surrounding NE 405.

[0084] In an embodiment, the PPR information 170 representing PPR graph 400 includes a PPR type 183, a PPR-ID 186, and five PPR-PDEs 490A-E. In an embodiment, the PPR type 183 includes a value identifying that the PPR graph 400 represents a standard PPR 160 with one ingress NE 401 and one egress NE 405. In an embodiment, the PPR-ID 186 includes an identifier of PPR graph 400, such as, for example, a label, address, or identifier representing the egress NE 405. PPR-PDEs 490A-E are similar to the PPR-PDEs 190, in that PPR-PDEs 490A-E describe one or more elements on the PPR graph 400. Each of the PPR-PDEs 490A-E includes information describing NEs 401-405, respectively. For example, each of the PPR-PDEs 490A-E includes a label, address, or identifier identifying the NEs 401-405, respectively. The PPR-PDE 490A describing NE 401 includes the source flag 191 that is set to indicate that NE 401 is an ingress NE 401. The PPR-PDE 490E describing NE 405 includes the destination flag 193 that is set to indicate that NE 405 is an egress NE 405.

[0085] FIG. 5 is a diagram illustrating a PPR graph 500 in network 175 according to various embodiments of the disclosure. Similar to FIG. 1B, network 175 in FIG. 5 includes NEs 101-121 interconnected by links 122-154.

[0086] PPR graph 500 is similar to PPR graphs 180 and 400, in that the PPR graph 500 includes multiple paths between multiple ingress NEs 101, 121, 120, and 118 (also referred to as source NEs) and at least one egress NE 107 (also referred to as a destination NE). The PPR graph 500 shown in FIG. 5 is implemented as a PPR tree, which includes multiple ingress NEs 101, 121, 120, and 118, but only a single egress NE 107. As shown by FIG. 5, the ingress NEs 101, 121, 120, and 118 are represented with rectangles, and the single egress NE 107 is represented with a circle.

[0087] PPR graph 500 includes one or more branches 580A-D, each of which includes a sequence of elements in the network 175 connecting one of the ingress NEs 101, 121, 120, and 118 to another NE in the PPR graph 500. Each of the paths between the multiple ingress NEs 101, 121, 120, and 118 and the single egress NE 107 is represented by one of branches 580A-E. In an embodiment, each branch 580A-D shares at least one common NE.

[0088] As shown by FIG. 5, PPR graph 500 includes four branches 580A-D. The first branch 580A includes the following elements: ingress NE 101, link 132, NE 111, link 131, NE 110, link 130, NE 108, link 129, NE 108, link 128, and egress NE 107. The second branch 580B includes the following elements: ingress NE 121, link 142, NE 112, link 151, NE 113, link 135, and NE 110. In this case, the second branch 580B interconnects with the first branch 580A at NE 110, and thus, the first branch 580A and the second branch 580B share a common NE 110. In this embodiment, traffic originating from ingress NE 121 can reach egress NE 107 by being forwarded through the second branch 580B until the traffic reaches NE 110. In an embodiment, NE 110 continues to forward the traffic along the first branch 580A until the traffic reaches egress NE 107.

[0089] The third branch 580C includes the following elements: ingress NE 120, link 149, NE 119, link 144, NE 114, link 136, and NE 109. In this case, the third branch 580C interconnects with the first branch 580A at NE 109, and thus, the first branch 580A and the third branch 580C share a common NE 109. In this embodiment, traffic originating from ingress NE 120 can reach egress NE 107 by being forwarded through the third branch 580C until the traffic reaches NE 109. In an embodiment, NE 109 continues to forward the traffic along the first branch 580A until the traffic reaches egress NE 107.

[0090] The fourth branch 580D includes the following elements: ingress NE 118, link 145, NE 115, link 138, and NE 108. In this case, the fourth branch 580D interconnects with the first branch 580A at NE 108, and thus, the first branch 580A and the fourth branch 580D share a

common NE 108. In this embodiment, traffic originating from ingress NE 118 can reach egress NE 107 by being forwarded through the fourth branch 580D until the traffic reaches NE 108. In an embodiment, NE 108 continues to forward the traffic along the first branch 580A until the traffic reaches egress NE 107.

[0091] In an embodiment, the PPR information 170 representing PPR graph 500 includes a PPR type 183, a PPR-ID 186, and PPR-PDEs 590A-N spread across four PDE branches 550A-D. In an embodiment, the PPR type 183 includes a value identifying that the PPR graph 500 represents a PPR tree having multiple ingress NEs 101, 121, 120 and 118 and a single egress NE 107. In an embodiment, the PPR-ID 186 includes an identifier of PPR graph 500, such as, for example, a label, address, or identifier representing the egress NE 107.

[0092] In an embodiment in which the PPR graph 500 includes several branches 580A-D, the PPR information 170 may include PDE branches 550A-D, each PDE branch 550A-D corresponding to a branch 580A-D on the PPR graph 500. Each PDE branch 550A-D includes a branch ID 563A-D identifying the branch 580A-D being described by the PDE branch 550A-D and one or more PPR-PDEs 590A-N. PPR-PDEs 590A-N are similar to PPR-PDEs 190 and 490A-E (hereinafter referred to as “PPR-PDEs 490”) in that PPR-PDEs 590A-N each describe an element on the branch 580A-D.

[0093] In the example shown in FIG. 5, PDE branch 550A includes PPR-PDEs 590A-F. PPR-PDE 590A includes a label, address, or identifier describing NE 101 and a source flag 191 indicating that NE 101 is an ingress NE 101 or source NE 101 of traffic to be forwarded along the PPR graph 500. PPR-PDE 590B includes a label, address, or identifier describing NE 111, PPR-PDE 590C includes a label, address, or identifier describing NE 110, PPR-PDE 590D includes a label, address, or identifier describing NE 109, and PPR-PDE 590E includes a label, address, or identifier describing NE 108. PPR-PDE 590F includes a label, address, or identifier describing NE 107 and a destination flag 193 indicating that NE 107 is an egress NE 107 or destination NE 107 of traffic to be forwarded along the PPR graph 500.

[0094] PDE branch 550B includes PPR-PDEs 590G-I and PPR-PDE 590C. PPR-PDE 590G includes a label, address, or identifier describing NE 121 and a source flag 191. PPR-PDE 590H includes a label, address, or identifier describing NE 112, and PPR-PDE 590I includes a label, address, or identifier describing NE 113. PPR-PDE 590C may include the label, address, or

identifier describing the common NE 110 or a pointer to the PPR-PDE 590C included in the PDE branch 550A.

[0095] PDE branch 550C includes PPR-PDEs 590J-L and PPR-PDE 590D. PPR-PDE 590J includes a label, address, or identifier describing NE 120 and a source flag 191. PPR-PDE 590K includes a label, address, or identifier describing NE 119, and PPR-PDE 590L includes a label, address, or identifier describing NE 114. PPR-PDE 590D may include the label, address, or identifier describing the common NE 109 or a pointer to the PPR-PDE 590D included in the PDE branch 550A.

[0096] PDE branch 550D includes PPR-PDEs 590M-N and PPR-PDE 590E. PPR-PDE 590M includes a label, address, or identifier describing NE 118 and a source flag 191. PPR-PDE 590N includes a label, address, or identifier describing NE 115. PPR-PDE 590D may include the label, address, or identifier describing the common NE 108 or a pointer to the PPR-PDE 590E included in the PDE branch 550A.

[0097] In an embodiment, the central entity 165 generates the PPR information 170 and then sends the PPR information 170 to one of the NEs 104 via the central entity-to-domain link 166. The headend NE 104 processes the PPR information 170 and floods the PPR information 170 to the other NEs 101-103 and 105-121 in the network 175. In another embodiment, an administrator of network 175 sends the PPR information 170 to one of the NEs 101-121, which then floods the network 175 with the PPR information 170.

[0098] FIG. 6 is a diagram of a PPR graph 600 implemented as a PPR tree according to various embodiments of the disclosure. As shown by FIG. 6, PPR graph 600 includes NEs 601-612, which are each similar to NEs 101-121 of FIGS. 1A-B. The PPR graph 600 is similar to PPR graph 500, in that PPR graph 600 includes multiple ingress NEs 601, 603, 606, 608, 610, and 612 (also referred to as source NEs) and at least one egress NE 605 (also referred to as a destination NE). The ingress NEs 601, 603, 606, 608, 610, and 612 are represented in FIG. 6 with rectangles, and the single egress NE 605 is represented with a circle.

[0099] The PPR graph 600 also includes three branches 680A-C. Branch 680A includes the following sequence of NEs: NEs 601, 602, 603, 604, and 605. Branch 680B includes the following sequence of NEs: NEs 612, 611, 610, and 609. Branch 680C includes the following sequence of NEs: NEs 606, 607, 608, 609, and 605.

[00100] FIG. 7 is a method 700 for processing PPR information 170 describing the PPR graph 600 of FIG. 6 according to various embodiments of the disclosure. Method 700 may be implemented by any of NEs 601-612 of PPR graph 600, which again is implemented as a PPR tree. Method 700 is implemented when an NE 601-612 receives the PPR information 170 describing PPR graph 600 from at least one of a central entity 165, another NE 601-612, or an administrator of the network 175.

[00101] At step 703, the PPR-PDEs 590A-N (hereinafter referred to as “PPR-PDE 590”) of a first PDE branch 550A-D (hereinafter referred to as “PDE branch 550”) included in the PPR information 170 representing the PPR graph 600 are searched to determine whether the NE 601-612 is identified in one of the PPR-PDEs 590. For example, the network configuration module 335 executed by the processor 330 is configured to search the plurality of PPR-PDEs 590 of a first PDE branch 550 included in the PPR information 170.

[00102] For example, suppose that NE 609 is implementing method 700. NE 609 is configured to search through the PPR-PDEs 590 of a first PDE branch 550, which may describe branch 680A. In this case, NE 609 is not identified in any of the PPR-PDEs 590 describing branch 680A and, therefore, ignores the PPR-PDEs 590 describing branch 680A.

[00103] Continuing, NE 609 is configured to search through the PPR-PDEs 590 of a second PDE branch 550, which describes branch 680B. In this case, NE 609 is identified in the last PPR-PDE 590 describing branch 680B. However, NE 609 ignores the last PPR-PDE 590 describing branch 680B because NE 609 is not an egress NE, is the last hop of branch 680B, and is the common NE 609 between branch 680B and 680C. Thereafter, NE 609 searches through PPR-PDEs 590 of a third PDE branch 550, which describes branch 680C. In this case, NE 609 determines that a PPR-PDEs 590 describing branch 680C identifies NE 609 as being on a path towards an egress NE 605.

[00104] At step 706, a determination of a next hop on the PPR graph 600 is made in response to the NE 601-612 being identified in one of the PPR-PDEs 590. For example, the network configuration module 335 executed by the processor 330 is configured to determine a next hop on the PPR graph 600.

[00105] Continuing with the example, after NE 609 determines that a PPR-PDE 590 describing branch 680C identifies NE 609 as being on a path towards an egress NE 605, the NE 609 determines that the next PPR-PDE 590 on this path describes egress NE 605.

[00106] At step 709, NE 609 updates its forwarding database 365 to include a forwarding entry 370 for the egress NE 605 describing the next hop on the path to the egress NE 605. For example, the network configuration module 335 executed by the processor 330 is configured to update the forwarding database 365. In an embodiment, the forwarding entry 370 may include the next PPR-PDE 590 describing egress NE 605. In an embodiment, the forwarding entry 370 also includes a source flag 191 and destination flag 193, when applicable.

[00107] As another example, supposing NE 612 is implementing method 700, NE 612 may first determine that NE 612 is identified in the PPR-PDEs 590 describing branch 680B. NE 612 then determines that the egress NE 605 is the egress NE for branch 680B, and that NE 611 is the next hop (e.g., defined as the next PPR-PDE 590) on the path to the egress NE 605. NE 612 updates the forwarding database 365 to include a forwarding entry 370 corresponding the next PPR-PDE 590 describing NE 611.

[00108] Method 700 is one example of how an NE 601-612 processes the PPR information 170 describing the PPR graph 600. However, it should be appreciated that NEs 601-612 may process the PPR information 170 describing the PPR graph 600 using various other methods and implementations.

[00109] FIGS. 5-7 describe a PPR tree having multiple ingress NEs but only a single egress NE or destination NE. FIGS. 8-9 describe a PPR forest with multiple ingress NEs and multiple egress NEs.

[00110] FIG. 8 is a diagram illustrating a PPR graph 800 in network 175 according to various embodiments of the disclosure. Similar to FIG. 1B, network 175 in FIG. 8 includes NEs 101-121 interconnected by links 122-154. PPR graph 800 is similar to PPR graph 500 of FIG. 8, except that PPR graph 800 includes multiple egress NEs 109 and 107, instead of just a single egress NE 107.

[00111] The PPR graph 800 shown in FIG. 8 is implemented as PPR forest, which includes multiple ingress NEs 101, 121, and 119, and multiple egress NEs 109 and 107. As shown by FIG. 8, the ingress NEs 101, 121, and 119 are represented with rectangles, and the multiple egress NEs 109 and 107 are represented with circles.

[00112] The PPR graph 800 includes one or more branches 880A-C. Each branch 880A-C includes a sequence of elements in the network 175 connecting one of the ingress NEs 101, 121, and 119 to another NE in the PPR graph 800. In an embodiment, each branch 880A-C shares at least one common NE.

[00113] As shown by FIG. 8, PPR graph 800 includes three branches 880A-C. The first branch 880A includes the following elements: ingress NE 101, link 132, NE 111, link 131, NE 110, link 130, NE 109, link 129, NE 108, link 128, and egress NE 107. The second branch 880C includes the following elements: ingress NE 121, link 142, NE 112, link 151, NE 113, link 135, and NE 110. In this case, the second branch 880B interconnects with the first branch 880A at NE 110, and thus, the first branch 880A and the second branch 880B share a common NE 110. In this embodiment, traffic originating from ingress NE 121 can reach either egress NE 109 or egress NE 107 by being forwarded through the second branch 880B until the traffic reaches either egress NE 109 or egress NE 107. In an embodiment, NE 110 continues to forward the traffic along the first branch 880A until the traffic reaches either egress NE 109 or egress NE 107.

[00114] The third branch 880C includes the following elements: ingress NE 119, link 148, NE 118, link 145, NE 115, link 137, and NE 109. In this case, the third branch 880C interconnects with the first branch 880A at NE 109, and thus, the first branch 880A and the third branch 880C share a common NE 109. In this embodiment, traffic originating from ingress NE 119 can reach either egress NE 109 or egress NE 107 by being forwarded through the third branch 880C until the traffic reaches NE 109. In an embodiment, NE 109 can act as the egress NE and forward the traffic external to network 175. In another embodiment, NE 109 continues to forward the traffic along the first branch 880A until the traffic reaches egress NE 107.

[00115] In an embodiment, the PPR information 170 representing PPR graph 800 includes a PPR type 183, a PPR-ID 186, and PPR-PDEs 890A-L spread across three PDE branches 850A-C. In an embodiment, the PPR type 183 includes a value identifying that the PPR graph 800 represents a PPR forest having multiple ingress NEs 101, 121, and 119 and multiple egress NEs 109 and 107.

[00116] In an embodiment in which the PPR graph 800 includes several branches 880A-C, the PPR information 170 may include PDE branches 850A-C, each PDE branch 850A-C corresponding to a branch 880A-C on the PPR graph 800. Each PDE branch 850A-C includes a branch ID 863A-C identifying the branch 880A-C being described by the PDE branch 850A-C and one or more PPR-PDEs 890A-L. PPR-PDEs 890A-L are similar to PPR-PDEs 190, 490, 590 in that PPR-PDEs 890A-L each describe an element on the PPR graph 800.

[00117] In the example shown in FIG. 8, PDE branch 850A includes PPR-PDEs 890A-F. PPR-PDE 890A includes a label, address, or identifier describing NE 101 and a source flag 191

indicating that NE 101 is an ingress NE 101 or source NE 101 of traffic to be forwarded along the PPR graph 800. PPR-PDE 890B includes a label, address, or identifier describing NE 111, and PPR-PDE 890C includes a label, address, or identifier describing NE 110. PPR-PDE 890D includes a label, address, or identifier describing NE 109 and a destination flag 193 indicating that NE 109 is an egress NE 109 or destination NE 109 of traffic to be forwarded along the PPR graph 800. In an embodiment, PPR-PDE 890D also includes the PPR-ID 186, which may be a label, address, or identifier identifying egress NE 109.

[00118] PPR-PDE 890E includes a label, address, or identifier describing NE 108. PPR-PDE 890F includes a label, address, or identifier describing NE 107 and a destination flag 193 indicating that NE 107 is an egress NE 107 or destination NE 107 of traffic to be forwarded along the PPR graph 800. In an embodiment, PPR-PDE 890F also includes the PPR-ID 186, which may be a label, address, or identifier identifying egress NE 107.

[00119] PDE branch 850B includes PPR-PDEs 890G-I and PPR-PDE 890C. PPR-PDE 890G includes a label, address, or identifier describing NE 121 and a source flag 191. PPR-PDE 890H includes a label, address, or identifier describing NE 112, and PPR-PDE 890I includes a label, address, or identifier describing NE 113. PPR-PDE 890C may include the label, address, or identifier describing the common NE 110 or a pointer to the PPR-PDE 890C included in the PDE branch 850A.

[00120] PDE branch 850C includes PPR-PDEs 890J-L and PPR-PDE 890D. PPR-PDE 890J includes a label, address, or identifier describing NE 119 and a source flag 191. PPR-PDE 890K includes a label, address, or identifier describing NE 118, and PPR-PDE 890L includes a label, address, or identifier describing NE 115. PPR-PDE 890D may include the label, address, or identifier describing the common NE 109 or a pointer to the PPR-PDE 890D included in the PDE branch 850A.

[00121] FIG. 9 is a diagram of a PPR graph 900 implemented as a PPR forest according to various embodiments of the disclosure. As shown by FIG. 9, PPR graph 900 includes NEs 901-912, which are each similar to NEs 101-121 of FIGS. 1A-B. The PPR graph 900 is similar to PPR graph 800, in that PPR graph 900 includes multiple ingress NEs 901, 903, 906, 908, 910, and 912 (also referred to as source NEs) and multiple egress NEs 902, 905, and 907 (also referred to as a destination NEs). The ingress NEs 901, 903, 906, 908, 910, and 912 are represented in FIG. 9 with rectangles, and the egress NEs 902, 905, and 907 are represented with circles.

[00122] The PPR graph 900 also includes three branches 980A-C. Branch 980A includes the following sequence of NEs: NEs 901, 902, 903, 904, and 905. Branch 980B includes the following sequence of NEs: NEs 912, 911, 910, and 905. Branch 980C includes the following sequence of NEs: NEs 906, 907, 908, 909, and 910. As shown by the dotted arrow in FIG. 9, branch 980C is cyclical, in that NE 905 and NE 906 are connected within the PPR graph 900 such that traffic can be transmitted from NE 905 to NE 906 in a circular fashion.

[00123] Programming each of the NEs 901-912 to forward traffic along the PPR graph 900 may be implemented in a manner similar to that described above with reference to method 700, except that each NE 901-912 adds a forwarding entry 370 for each egress NE 905 and 907 in the PPR graph 900. For example, after NE 901 receives PPR information 170 describing PPR graph 900, NE 901 determines whether NE 901 is identified in any of the PPR-PDEs 890 describing the elements of branch 980A. As NE 901 is included in branch 980A, NE 901 is also identified in one of the PPR-PDEs 890 describing the elements of branch 980A. At this NE 901 identifies that there are two egress NEs 902 and 907 included in branch 980A. In this case, NE 901 determines a next hop for each of the egress NEs 902 and 907 included in the branch 980A and creates a forwarding entry 370 for each of the egress NEs 902 and 907. The forwarding entry 370 for egress NE 902 may include the next PPR-PDE 890 describing NE 902, and the forwarding entry 370 for egress NE 905 may also include the next PPR-PDE 890 describing NE 902.

[00124] As another illustrative example, assume that NE 903 receives PPR information 170 describing PPR graph 900. In this case, NE 903 only creates a forwarding entry 370 for egress NE 905 on branch 980A. This is because NE 903 may not be configured to forward traffic backwards to egress NE 902 and because branch 980A is not a cyclical branch.

[00125] As yet another example, with regard to the cyclical branch 980C, assume that NE 909 receives PPR information 170 describing PPR graph 900. In this case, NE 909 creates a forwarding entry 370 for both egress NE 905 and egress NE 907 because branch 980C is a cyclical branch. In this example, after NE 909 receives the PPR information 170 describing PPR graph 900, NE 909 may determine that NE 909 is not identified in the PPR-PDEs 890 for branches 980A-B. However, NE 909 determines that NE 909 is identified in the PPR-PDEs 890 for branch 980C. NE 909 first determines that NE 905 is an egress NE 905 of branch 980C. Since branch 980C is cyclical, NE 909 may continue to the beginning of the branch NE 980C and then

determine that NE 907 is another egress NE 907 of branch 980C. NE 909 then creates forwarding entries 370 for each of the egress NEs 905 and 907.

[00126] In an embodiment, a PPR graph 900 implemented as PPR forest is only permitted to include a single cyclical loop branch, such as branch 980C. In this embodiment, the central entity 165 is configured to determine the PPR graph 900 with a maximum of one cyclical loop branch, such as branch 980C.

[00127] FIG. 10 is a diagram of a PPR graph 1000 implemented as a bidirectional forest according to various embodiments of the disclosure. As shown by FIG. 10, PPR graph 1000 includes NEs 1000-1012, which are each similar to NEs 101-121 of FIGS. 1A-B. The PPR graph 1000 is similar to PPR graph 900, in that PPR graph 1000 includes three branches 1080A-C with multiple ingress NEs 1001, 1003, 1006, 1008, 1010, and 1012 (also referred to as source NEs) and multiple egress NEs 1002, 1005, and 1007 (also referred to as a destination NEs). The ingress NEs 1001, 1003, 1006, 1008, 1010, and 1012 are represented in FIG. 10 with rectangles, and the egress NEs 1002, 1005, and 1007 are represented with circles. However, unlike the NEs 901-912 of PPR graph 900 that were only configured to transmit traffic in one direction, the NEs 1001-1012 of PPR graph 1000 are configured to transmit traffic in both directions (forwards and backwards).

[00128] Programming each of the NEs 1001-1012 to forward traffic along the PPR graph 1000 may be implemented in a manner similar to that described above with reference to method 700, except that each NE 1001-1112 adds a forwarding entry 370 for every single egress NE 1002, 1005, and 1007 in the PPR graph 1000. This is because each NE 1001-1012 in the PPR graph 1000 has access to every other NE 1001-1012, due to the bidirectional nature of the PPR graph 1000.

[00129] For example, after NE 1001 receives PPR information 170 describing PPR graph 1000, NE 1001 determines whether NE 1001 is identified in any of the PPR-PDEs 890 describing the PPR graph 1000. At this point, NE 1001 identifies that there are three egress NEs 1002, 1005, and 1007 included in PPR graph 1000. In this case, NE 1001 determines a next hop for each of the egress NEs 1002, 1005, and 1007 and creates a forwarding entry 370 for each of the egress NEs 1002, 1005, and 1007. The forwarding entry 370 for egress NE 1002 may include the next PPR-PDE 890 describing NE 1002, forwarding entry 370 for egress NE 1005 may include the next PPR-PDE 890 describing NE 1002, and the forwarding entry 370 for egress NE 1007 may also include the next PPR-PDE 890 describing NE 1002. In this embodiment, each of the NEs 1001-

1012 program a forwarding entry 370 for each egress NE 1002, 1005, and 1007 of PPR graph 1000.

[00130] In an embodiment, a PPR graph 1000 implemented as bidirectional forest is not permitted to include a cyclical branch or loop. In this embodiment, the central entity 165 is configured to determine a PPR graph 1000 implemented as bidirectional forest without the inclusion of a cyclical branch or loop.

[00131] FIG. 11 is a diagram of an advertisement 1100 including the PPR information 170 according to various embodiments of the disclosure. In an embodiment, the advertisement 1100 may be transmitted by the central entity 165 to one of the headend NEs 101-121 of network 175. In an embodiment, the advertisement 1100 may be transmitted between one or more of the NEs 101-121 of network 175.

[00132] The advertisement 1100 includes various fields that are used to carry information describing the PPR graphs 180, 400, 500, 600, 800, 900, or 1000, such as PPR information 170. As shown by FIG. 11, advertisement 1100 includes a graph type field 1103, graph flags 1106, a destination/PPR-ID field 1109, a policy group tag field 1111, one or more PPR-PDE field 1113A-C, and an attributes field 1133.

[00133] The graph type field 1103 may carry a PPR type 183, which is a value indicating whether the PPR graph being described is a standard PPR, a tree, a forest, or a bidirectional forest. The graph flags field 1106 may include one or more flags, or bits, that are set to instruct a receiving NE how to program the information carried by the advertisement 1100. The destination/PPR-ID field 1109 may carry a label, address, or identifier of one or more egress NEs on a PPR graph.

[00134] The policy group tag field 1111 may carry a policy group tag assigned to the PPR graph being described by the advertisement. In an embodiment, a policy group tag is an identifier, or a value, assigned to one or more PPR graphs. In an embodiment, the central entity 165 determines the policy group tag for a particular PPR graph based on an attribute of a resource to be reserved along the PPR graph. An attribute refers to network characteristic or traffic engineering (TE) parameter of the resource, such as bandwidth required to transmit a data packet along the PPR graph, a burst size permitted to be transmitted along the PPR graph, a bounded latency permitted to occur on an NE along the PPR graph, or a lifetime indicating a time period during which the resource is to be reserved on each NE and/or link of the PPR graph. In an embodiment, PPR graphs having one or more common attributes may be assigned the same policy group tag.

[00135] In an embodiment, when a customer requests a path to a destination or an egress NE, the central entity 165 or an ingress NE may determine that the customer is associated with a particular policy group tag. For example, the customer is associated with a particular policy group tag based on attribute that the customer requested. In this embodiment, the ingress NE or the central entity 165 selects one of the PPR graphs having the particular policy group tag for the customer and forwards traffic for the customer along the selected PPR graph.

[00136] Advertisement 1100 further includes PPR-PDE fields 1113A-C, which are each configured to carry PPR-PDEs 190, 490, 590, or 890 that describe an element along the PPR graph. Each of the PPR-PDE fields 1113A-C may also include a branch identifier field 1114, a cyclic flag field 1115, a source flag field 1118, PPR-PDE flags 1121, a destination flag field 1124, and PPR-ID field 1127. The branch identifier field 1114 may carry a branch identifier 563 or 863. The source flag field 1118 may carry the source flag 191, and the destination flag field 1124 may carry the destination flag 193. The cyclic flag field 1115 is a flag, or a bit, included in the PPR-PDE for a first NE on a PDE branch. When the cyclic flag field 1115 is set to 1, the cyclic flag field 115 indicates that the branch is a cyclic branch logically the first NE and the last NE on the branch. The PPR-PDE flags field 1121 includes various flags or bits that describe a condition or state of an element being described by the PPR-PDE. The PPR-ID field 1127 may include a label, address, or identifier of an egress NE being described by the PPR-PDE having the destination flag 193 set.

[00137] The attributes field 1133 includes one or more attributes regarding a resource that should be reserved for one or more elements along the PPR graph. As described above, an attribute may be bandwidth required to transmit a data packet along the PPR graph, a burst size permitted to be transmitted along the PPR graph, a bounded latency permitted to occur on an NE along the PPR graph, or a lifetime indicating a time period during which the resource is to be reserved on each NE and/or link of the PPR graph.

[00138] Advertisement 1100 shown in FIG. 11 only shows a few fields within a message that is sent to the NEs 101-121 in network 175 with information describing a PPR graph. Advertisement 1100 may include other fields that are not explicitly shown and described with reference to FIG. 11.

[00139] FIGS. 12A-C are type-length-value (TLV) diagrams illustrating various fields of an example advertisement, such as advertisement 1100 of FIG. 11, according to various embodiments

of the disclosure. FIG. 12A shows a PPR graph TLV 1200, which includes fields similar to advertisement 1100 that carries the PPR information 170. The PPR graph TLV 1200 includes four logical sections: the encoding of the PPR prefix, the PPG-ID 344, the PPR-PDE sub-TLVs belonging to one or more branch IDs, and a set of optional attributes.

[00140] As shown by FIG. 12A, the PPR graph TLV 1200 includes a type field 1203, a length field 1206, a graph type field 1209, graph flags 1211, a PPR-prefix sub-TLV 1214, a last fragment flag 1217, a fragment identifier field 1220, a PPG-ID field 1223, a branch-ID sub-TLV and PPR-PDE sub-TLV 1226, and an attributes sub-TLV 1229. The length field 1206 may carry a total length value of the PPR graph TLV 1200 in bytes. The graph type field 1209 is a one octet value that defines a PPR graph structure. In an embodiment, the graph type field 1209 carries the PPR type 183, which is a value indicating whether the PPR graph is a standard PPR, a tree, a forest, or a bidirectional forest. The graph flags 1211 will be further described below with reference to FIG. 12C.

[00141] The PPR-prefix sub-TLV 1214 includes a label, address, or identifier of an egress NE or the destination of the PPR graph. The branch-ID sub-TLV and PPR-PDE sub-TLV 1226 is further described below with reference to FIG. 12B.

[00142] In an embodiment, the PPR graph TLV 1200 includes various fields used when the PPR graph TLV 1200 is divided into multiple fragments, as will be further described below with reference to FIG. 13. In this embodiment, each of the fragments is similar to the PPR graph TLV 1200, and each fragment carries different PPR-PDEs. In this embodiment, the PPR graph TLV 1200 includes the last fragment flag 1217, fragment identifier field 1220, and a PPG-ID field 1223. The PPG-ID field 1223 carries a PPG-ID 344, which is a label, address, or identifier identifying the PPR graph being described by the PPR information 170 included in the PPR graph TLV 1200. The PPG-ID 344 remains the same for each of the fragments being sent for a particular PPR graph. The fragment identifier field 1220 indicates an identifier or count of the particular fragment being sent relative to the total number of fragments. The last fragment flag 1217 is a flag, or a bit, that is set to indicate whether the current fragment is a last fragment for a particular PPR graph. The attributes sub-TLV 1229 includes one or more attributes defining a resource to be reserved for a PPR graph.

[00143] Referring now to FIG. 12B, shown is the branch-ID sub-TLV and PPR-PDE sub-TLV 1226 of the PPR graph TLV 1200. The branch-ID sub-TLV and PPR-PDE sub-TLV 1226

includes fields 1250A-D, each of which may include a PDE branch, carrying a branch identifier 563 or 863 and one more PPR-PDEs 190, 490, 590, or 890 describing elements on the branch. As shown by FIG. 12B, one of the fields 1250C describes a branch that is cyclic, denoted using the a cyclic flag 1115. The branch-ID sub-TLV and PPR-PDE sub-TLV 1226 also includes the attributes sub-TLV 1229.

[00144] Although only four fields 1250A-D are shown in FIG. 12B, it should be appreciated that any number of fields 1250A-D are included in the branch-ID sub-TLV and PPR-PDE sub-TLV 1226 describing any number of branches on PPR graph.

[00145] Referring now to FIG. 12C, shown is the graph flags 1211 of the PPR graph TLV 1200. The graph flags 1211 include an S flag 1261 and a D flag 1262. If the S flag 1261 is set, the PPR graph TLV 1200 should be flooded across the entire domain. If the S flag 1261 is not set, the PPR graph TLV 1200 should not be leaked between IS-IS levels. If the D flag 1262 is set, the PPR graph TLV 1200 is leaked from IS-IS level-2 to IS-IS level 1. Otherwise, the bit should be cleared.

[00146] FIG. 13 is a diagram illustrating fragments 1300A-N used to carry PPR information for a PPR graph according to various embodiments of the disclosure. In some cases, networks implementing certain types of IGPs, such as IS-IS, can only send signals having a maximum of 255 bytes. However, some PPR graphs having several branches and numerous elements, which may need much more than 255 bytes to send advertisements 1100 including PPR information 170 describing the PPR graph.

[00147] In this case, the advertisement 1100 describing a single PPR graph can be divided into fragments 1300A-N, in which the collection of fragments 1300A-N represent the entire PPR graph. In an embodiment, each fragment 1300A-N carries a subset of the PPR-PDEs for a particular branch of a PPR graph. In an embodiment, each fragment 1300A-N carries a subset of the PPR-PDEs for the entire PPR graph. In an embodiment, each fragment 1300A-N has a size that is less than the size that is capable of being transported in a network implementing a specific IGP.

[00148] As shown by FIG. 13, fragment 1300A includes various fields, such as a graph type field 1303, a length field 1306, a header 1309, a PPG-ID field 1311, a fragment ID field 1313, a last fragment flag 1316, a subtype field 1319, and a length field 1321. The graph type field 1303 is similar to the graph type field 1103 of FIG. 11. The length field 1306 may carry a length of the fragment 1300A. The header 1309 includes subtypes describing the PPR graph overall, and may include details such as Quality of Service parameters, and may only be included in the last

fragment N. The PPG-ID field 1311 carries a PPG-ID 344, which is a label, address, or identifier identifying the PPR graph. The PPG-ID 344 remains the same for each of the fragments being sent for a particular PPR graph. The fragment ID field 1313 indicates an identifier or count of the particular fragment being sent relative to the total number of fragments. The last fragment flag 1316 is a flag, or a bit, that is set to indicate whether the current fragment is a last fragment for a particular PPR graph. For the fragment 1300A, which may be the first fragment for a PPR graph, the last fragment flag 1316 may be cleared or set to 0. The subtype field 1319 indicates a type of the branch being described by the fragment 1300A. Each branch is one subtype, and the branches may be spread out over the fragments 1300A-N. In an embodiment, a single branch that is too large to fit into a fragment is split up into multiple branches.

[00149] Each of the fragments 1300A-N includes one or more PPR-PDEs describing one or more elements on the PPR graph, or a branch of the PPR graph. The one or more PPR-PDEs may be carried PPR-PDE fields, such as PPR-PDE fields 1113A-C.

[00150] The last fragment 1300N for a set of fragments 1300A-N describing a PPR graph is substantially similar to the first fragment 1300A, except that the last fragment 1300N includes an attributes field 1350. The attributes field 1350 includes one or more attributes for a resource to be reserved on an element of the PPR graph. In addition, since the last fragment 1300N is the last fragment for a PPR graph, the last fragment flag 1316 for the last fragment 1300N may be set to 1. It should be appreciated that fragments 1300A-N may include other fields that are not necessarily shown by FIG. 13.

[00151] FIGS. 14A-B are TLV diagrams showing examples of fields within a fragment 1300A-N according to various embodiments of the disclosure. FIG. 14A shows fragment TLV 1400, which is a new top level IS-IS TLV diagram, with fields similar to those described above with reference to the fragments 1300A-N of FIG. 13. In particular, the fragment TLV 1400 includes a type field 1403, a length field 1406, a flags field 1409, a PPG-ID length field 1411, a PPG-ID field 1414, and a fragment sub-TLV 1417. The type field 1403 is similar to the type field 1303 of FIG. 13, the length field 1406 is similar to the length field 1306 of FIG. 13, the flags field 1409 is similar to the flags field 1121 of advertisement 1100, the PPG-ID field 1414 is similar to the PPG-ID field 1311 of FIG. 13, and the PPG-length field 1411 carries a length of the PPG-ID field 1414. The fragment sub-TLV 1417 is further described below with reference to FIG. 14B. In an

embodiment, each fragment 1300A-N includes the fragment TLV 1400, which remains the same for each fragment 1300A-N.

[00152] Referring now to FIG. 14B, shown is the fragment sub-TLV 1417, which includes a type field 1420, a length field 1423, a fragment identifier field 1426, fragment flags 1429, a PPR-PDEs and PPR attributes fields 1430. In an embodiment, each fragment 1300A-N includes the fragment sub-TLV 1417, and each fragment sub-TLV 1417 is different for each fragment 1300A-N.

[00153] FIG. 15 is a method 1500 for advertising and programming PPR graphs in a network according various embodiments of the disclosure. The PPR graphs may be any one of 180, 400, 500, 600, 800, 900, or 1000. The network may be similar to network 175. In an embodiment, method 1500 may be performed by the central entity 165, which may be implemented as an NE 300. In an embodiment, method 1500 may be performed after a central entity 165 receives instructions or a request to construct a PPR graph.

[00154] At step 1503, PPR information 170 describing a PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network is determined. For example, the network configuration module 335, executed by the processor 330, determines the PPR information 170. In an embodiment, the PPR information 170 includes a PPR-ID 186 associated with the egress NE and a plurality of PPR-PDEs 190, 590, or 890 describing a plurality of elements included in the PPR graph. In an embodiment, a PPR-PDE 190, 590, or 890 describing the ingress NE includes a source flag 191, and a PPR-PDE 190, 590, or 890 describing the egress NE including a destination flag 193.

[00155] In an embodiment, any NE/PPR-PDE can only be a non-last NE/PPR-PDE in one branch but a last-NE/PPR-PDE in more than one branch. In an embodiment, for a PPR graph implemented as a tree, every PDE-ID can be a non-last NE/PPR-PDE in at most one branch and can be a last-PDE in one or more branches. Branches form a tree by joining nodes with the same PDE identifier (PDE-ID). Leaves of the tree may be indicated as ingress NEs or sources, and the root of the tree may be the only egress NE or destination of the tree.

[00156] For a PPR graph implemented as a forest, every NE/PPR-PDE can be a non-last-NE/PPR-PDE in at most one branch, but a last-NE/PPR-PDE in zero or more branches. In an embodiment, branches form a tree with at most one branch indicating a “loop.” Every NE can be marked as source and/or destination, except leaves of the graph, which are a source and not a

destination. There must be at least one destination or egress NE in the looped branch. If there is no looped branch, then the root of the tree is be a destination and cannot be a source.

[00157] For a PPR graph implemented as a bidirectional forest, every NE/PPR-PDE can be a non-last PDE in at most one branch. A NE/PPR-PDE can be a last-PDE in one or more branches. Branches form a tree by joining NEs/PPR-PDEs with the same PDE-ID. Leaves of the tree may be ingress NEs. In this case, all connections between NEs are assumed to be bidirectional.

[00158] At step 1506, the PPR information 170 is transmitted to a headend NE 104 of the network. For example, a Tx 340 transmits the PPR information 170 to a headend NE 104 of the network. In an embodiment, the headend NE 104 floods the PPR information 170 to a remainder of the NEs 101-103 and 105-121 of the network 175.

[00159] FIG. 16 is another method 1600 for advertising and programming PPR graphs in a network according various embodiments of the disclosure. The PPR graphs may be any one of 180, 400, 500, 600, 800, 900, or 1000. The network may be similar to network 175. In an embodiment, method 1600 may be performed by one of the NEs 101-121 in the network 175.

[00160] At step 1603, PPR information 170 describing PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network is received. For example, Rx 320 receives the PPR information 170. In an embodiment, the PPR information 170 includes a PPR-ID 186 associated with the egress NE and a plurality of PPR-PDEs 190, 590, or 890 describing a plurality of elements included in the PPR graph. In an embodiment, a PPR-PDE 190, 590, or 890 describing the ingress NE includes a source flag 191, and a PPR-PDE 190, 590, or 890 describing the egress NE including a destination flag 193.

[00161] At step 1606, a forwarding database 365 is updated to include a forwarding entry 370 for the egress NE in response to identifying the NE in the plurality of PPR-PDEs. For example, the network configuration modules 335, executed by the processor 330, updates to include a forwarding entry 370 to include a forwarding entry 370 for the egress NE in response to identifying the NE in the plurality of PPR-PDEs. In an embodiment, the NE floods the PPR information 170 to a remainder of the NEs 101-121 of the network 175 via an underlying IGP of the network. In an embodiment, the forwarding entry 370 indicates a next hop by which to forward a data packet comprising the PPR-ID.

[00162] In an embodiment, the forwarding entry 370 includes the next PPR-PDE of the plurality of PPR-PDEs. In an embodiment, the forwarding entry 370 includes an address prefix indicating the next hop. In an embodiment, the address prefix represents a plurality of NEs along a branch of the PPR graph.

[00163] FIG. 17 shows an apparatus 1700 for advertising and programming a PPR graph according various embodiments of the disclosure. The apparatus 1700 comprises a means for determining 1703 and a means for transmitting 1706. The means for determining 1703 includes a means for determining PPR information describing a PPR graph, which represents a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-ID associated with the egress NE and a plurality of PPR-PDEs describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag. The means for transmitting 1706 includes a means for transmitting the PPR information describing the PPR graph to a headend NE of the network.

[00164] FIG. 18 shows another apparatus 1800 for advertising and programming a PPR graph according various embodiments of the disclosure. The apparatus 1800 comprises a means for receiving 1803 and a means for updating 1806. The means for receiving 1803 includes a means for receiving PPR information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-ID associated with the egress NE and a plurality of PPR-PDEs describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination. The means for updating 1806 includes a means for a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

[00165] While several embodiments have been provided in the present disclosure, it should be understood that the disclosed systems and methods might be embodied in many other specific forms without departing from the spirit or scope of the present disclosure. The present examples are to be considered as illustrative and not restrictive, and the intention is not to be limited to the details given herein. For example, the various elements or components may be combined or integrated in another system or certain features may be omitted, or not implemented.

[00166] In addition, techniques, systems, subsystems, and methods described and illustrated in the various embodiments as discrete or separate may be combined or integrated with other systems, modules, techniques, or methods without departing from the scope of the present disclosure. Other items shown or discussed as coupled or directly coupled or communicating with each other may be indirectly coupled or communicating through some interface, device, or intermediate component whether electrically, mechanically, or otherwise. Other examples of changes, substitutions, and alterations are ascertainable by one skilled in the art and could be made without departing from the spirit and scope disclosed herein.

CLAIMS

What is claimed is:

1. A method implemented by a network element (NE) in a network, comprising:
 - receiving, by the NE, preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag; and
 - updating, by the NE, a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.
2. The method according to claim 1, wherein the PPR graph comprises a plurality of branches, wherein each branch comprises a plurality of elements on a path between two NEs included in the PPR graph.
3. The method according to any one of claims 1 to 2, wherein the PPR graph comprises a plurality of ingress NE and only the egress NE, wherein the PPR graph comprises a plurality of branches, wherein each of the branches comprises a plurality of elements on a path connecting one of the ingress NEs to an NE along the path to the egress NE, and wherein each of the branches includes at least one common NE.
4. The method according to claim 3, further comprising:
 - determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information;
 - sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs;
 - determining, by the NE, the egress NE of the first branch based on the destination flag in response to being identified in the one of the plurality of PPR-PDEs; and

generating, by the NE, the forwarding entry for the egress NE indicating that a next hop toward the egress NE is a next PPR-PDE of the plurality of PPR-PDEs.

5. The method according to any one of claims 1 to 2, wherein the PPR graph comprises a plurality of ingress NEs and a plurality of egress NEs, wherein the PPR graph comprises a plurality branches, wherein each of the branches comprises a plurality of elements on a path connecting one of the ingress NEs to an NE along the path to one of the egress NEs, and wherein each of the branches includes at least one common NE.

6. The method according to any one of claims 1 to 5, further comprising:

determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information;

sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs;

determining, by the NE, a plurality of egress NEs of the first branch in response to being identified in the one of the plurality of PPR-PDEs; and

generating, by the NE, a plurality of forwarding entries for the plurality of egress NEs.

7. The method according to any one of claims 1 to 2, wherein the PPR information includes a flag indicating that the PPR graph is a bidirectional forest, wherein the PPR graph comprises a plurality of ingress NEs and a plurality egress NEs, wherein the PPR graph comprises a plurality branches, wherein each of the plurality of branches comprises a plurality of elements on a bidirectional path connecting one of the ingress NEs to an NE along the bidirectional path to one of the egress NEs, and wherein each of the branches includes at least one common NE.

8. The method according to 7, further comprising:

determining, by the NE, a plurality of PPR-PDEs corresponding to each of the plurality of branches based on the PPR information;

sequentially searching, by the NE, the plurality of PPR-PDEs on a first branch of the plurality of branches to determine whether the NE is identified in one of the plurality of PPR-PDEs;

determining, by the NE, a first plurality of destinations along a first direction of the first branch in response to being identified in the one of the plurality of PPR-PDEs;

generating, by the NE, a plurality of forwarding entries for the first plurality of destinations;

determining, by the NE, a second plurality of destinations along a second direction of the first branch; and

generating, by the NE, a plurality of forwarding entries for the second plurality of destinations.

9. The method according to any one of claims 1 to 8, wherein the forwarding database is not updated in response to identifying the NE in the one of the plurality of PPR-PDEs and determining that a next hop is on a different branch.

10. The method according to any one of claims 1 to 9, wherein the PPR graph comprises a single source and a single destination.

11. The method according to any one of claims 1 to 10, wherein the PPR information describing the PPR graph is received from a central entity of the network or an administrator of the network.

12. The method according to any one of claims 1 to 11, wherein the PPR information is transmitted to a plurality of other NEs in the network.

13. The method according any one of claims 1 to 11, wherein the forwarding entry comprises an address prefix indicating the next hop, wherein the address prefix represents a plurality of NEs along a branch of the PPR graph.

14. The method according to any one of claims 1 to 13, wherein a first PPR-PDE on a branch of the PPR graph includes a flag indicating that the branch is cyclical.

15. The method according to any one of claims 1 to 14, wherein the PPR information includes a type field indicating that the plurality of PPR-PDEs describes a tree, a forest, or a bidirectional forest.

16. The method according to any one of claims 1 to 15, wherein the PPR graph comprises a plurality of ingress NEs and a plurality of egress NEs, wherein each of the plurality of egress NEs is represented by a different PPR-ID, and wherein a forwarding entry is created for each of the plurality of egress NEs.

17. The method according to any one of claims 1 to 16, wherein the PPR information includes a policy group tag associated with one or more PPR graphs, wherein the one or more PPR graphs are assigned the policy group tag based on an attribute associated with the one or more PPR graphs, wherein the attribute is at least one of a traffic engineering (TE) parameter, network characteristic, or service requirement that is associated with a resource to be reserved on each of one or more PPR graphs.

18. The method according to any one of claims 1 to 17, wherein the PPR information is received from another NE in the network as a plurality of fragments, wherein each fragment includes a PPR graph identifier (PPG-ID), a fragment identifier, a last fragment flag, and a subset of the plurality of PPR-PDEs, wherein the PPG-ID identifies the PPR graph described by the PPR information, wherein the fragment identifier indicates a sequence number of a respective fragment, and wherein the last fragment flag indicates whether the respective fragment is a last fragment representing the PPR graph of the plurality of fragments.

19. The method according to claim 18, wherein the forwarding entry is generated based on the plurality of fragments after the respective fragment having the last fragment flag set is received.

20. A method implemented by a central entity in a network, comprising:

determining, by the central entity, preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag; and

transmitting, by the central entity, the PPR information describing the PPR graph to a headend NE of the network.

21. A network element (NE) implemented in a network, comprising:

a memory configured to store instructions; and

a processor coupled to the memory and configured to execute the instructions to cause the processor to be configured to:

receive preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag; and

update a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

22. A central entity implemented in a network, comprising:

a memory configured to store instructions; and

a processor coupled to the memory and configured to execute the instructions to cause the processor to be configured to:

determine preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an

egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag; and

transmit the PPR information describing the PPR graph to a headend NE of the network.

23. An apparatus comprising:

a means for receiving preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress NE and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag and a PPR-PDE describing the egress NE including a destination flag; and

a means for updating a forwarding database to include a forwarding entry for the egress NE in response to identifying the NE in the plurality of PPR-PDEs, the forwarding entry indicating a next hop by which to forward a data packet comprising the PPR-ID.

24. An apparatus comprising:

a means for determining preferred path route (PPR) information describing a PPR graph, the PPR graph representing a plurality of PPRs between an ingress network entity (NE) and an egress NE in the network, the PPR information including a PPR-identifier (PPR-ID) associated with the egress NE and a plurality of PPR-Path Description Elements (PPR-PDEs) describing a plurality of elements included in the PPR graph, a PPR-PDE describing the ingress NE including a source flag, and a PPR-PDE describing the egress NE including a destination flag; and

a means for transmitting the PPR information describing the PPR graph to a headend NE of the network.




FIG. 1A




FIG. 1B

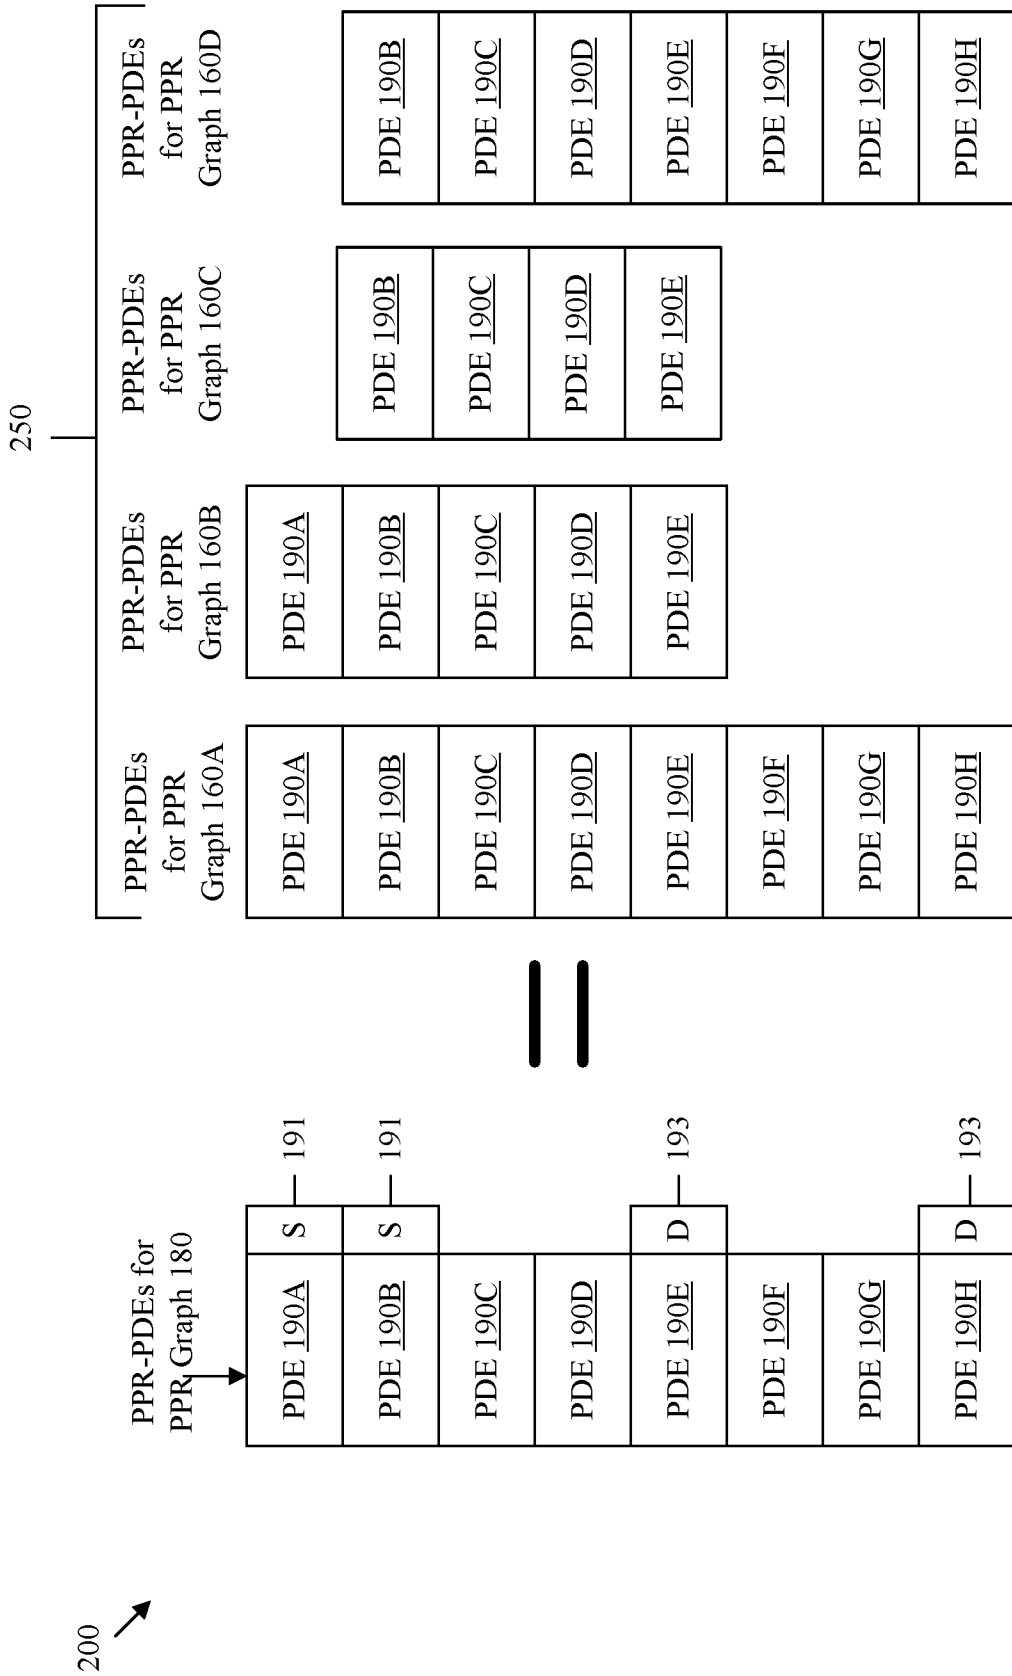


FIG. 2





FIG. 3


400
↙



170

PPR Type <u>183</u>	PPR-ID <u>186</u>	PPR-PDE <u>490A</u>	S <u>191</u>	PPR-PDE <u>490B</u>	PPR-PDE <u>490C</u>	PPR-PDE <u>490D</u>	PPR-PDE <u>490E</u>	D <u>193</u>
---------------------------	----------------------	------------------------	-----------------	------------------------	------------------------	------------------------	------------------------	-----------------

FIG. 4



170

PPR Type 183	PPR-ID 186	B ID 563A	PPR-PDE 590A	S 191	PPR-PDE 590B	PPR-PDE 590C	PPR-PDE 590D	PPR-PDE 590E	PPR-PDE 590F	D 193
550B		B ID 563B	PPR-PDE 590G	S 191	PPR-PDE 590H	PPR-PDE 590I	PPR-PDE 590C			
550C		B ID 563C	PPR-PDE 590J	S 191	PPR-PDE 590K	PPR-PDE 590L	PPR-PDE 590D			
550D		B ID 563D	PPR-PDE 590M	S 191	PPR-PDE 590N	PPR-PDE 590E				

550A

FIG. 5




FIG. 6

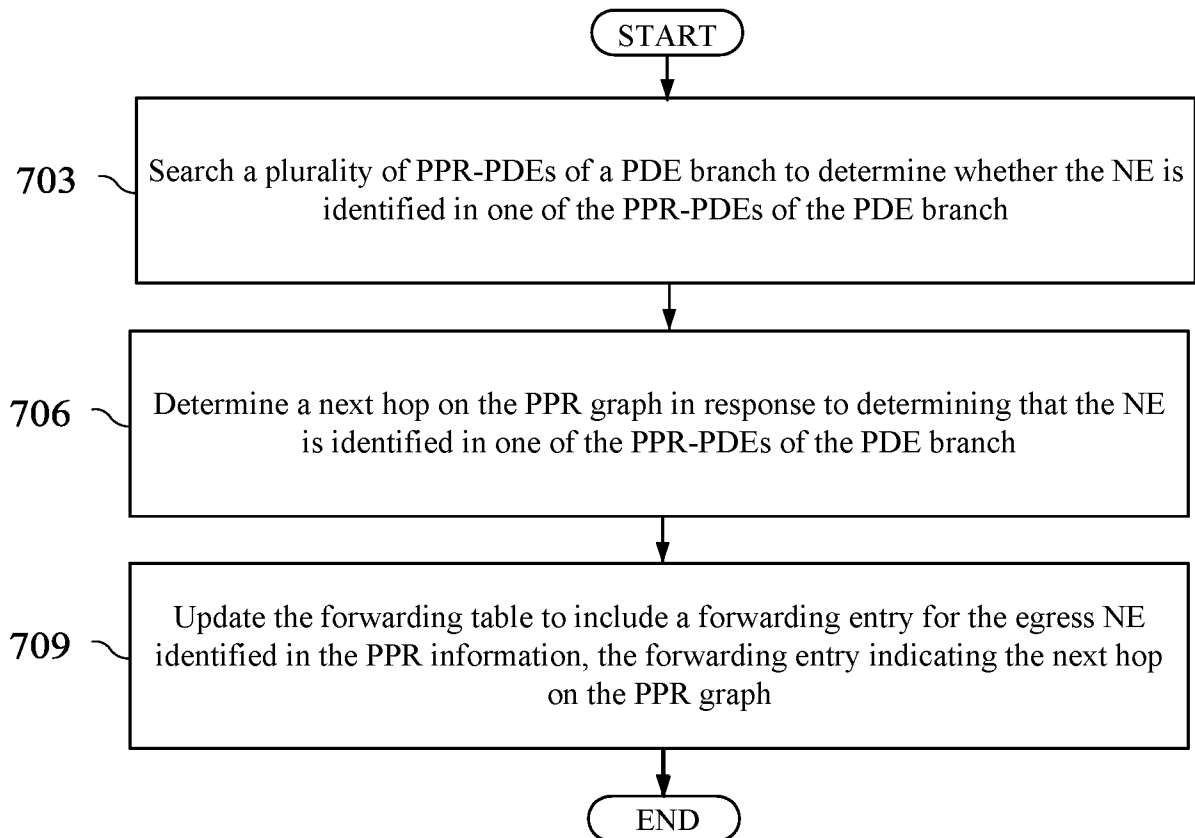

700
↘

FIG. 7

175



PPR Type 183	B ID 863A	PPR-PDE 890A	S 191	PPR-PDE 890B	PPR-PDE 890C	PPR-PDE 890D	D 193	PPR-ID 186	PPR-PDE 890E	PPR-PDE 890F	D 193
850B	B ID 863B	PPR-PDE 890G	S 191	PPR-PDE 890H	PPR-PDE 890I	PPR-PDE 890C					
850C	B ID 863C	PPR-PDE 890J	S 191	PPR-PDE 890K	PPR-PDE 890L	PPR-PDE 890D	D 193				

FIG. 8

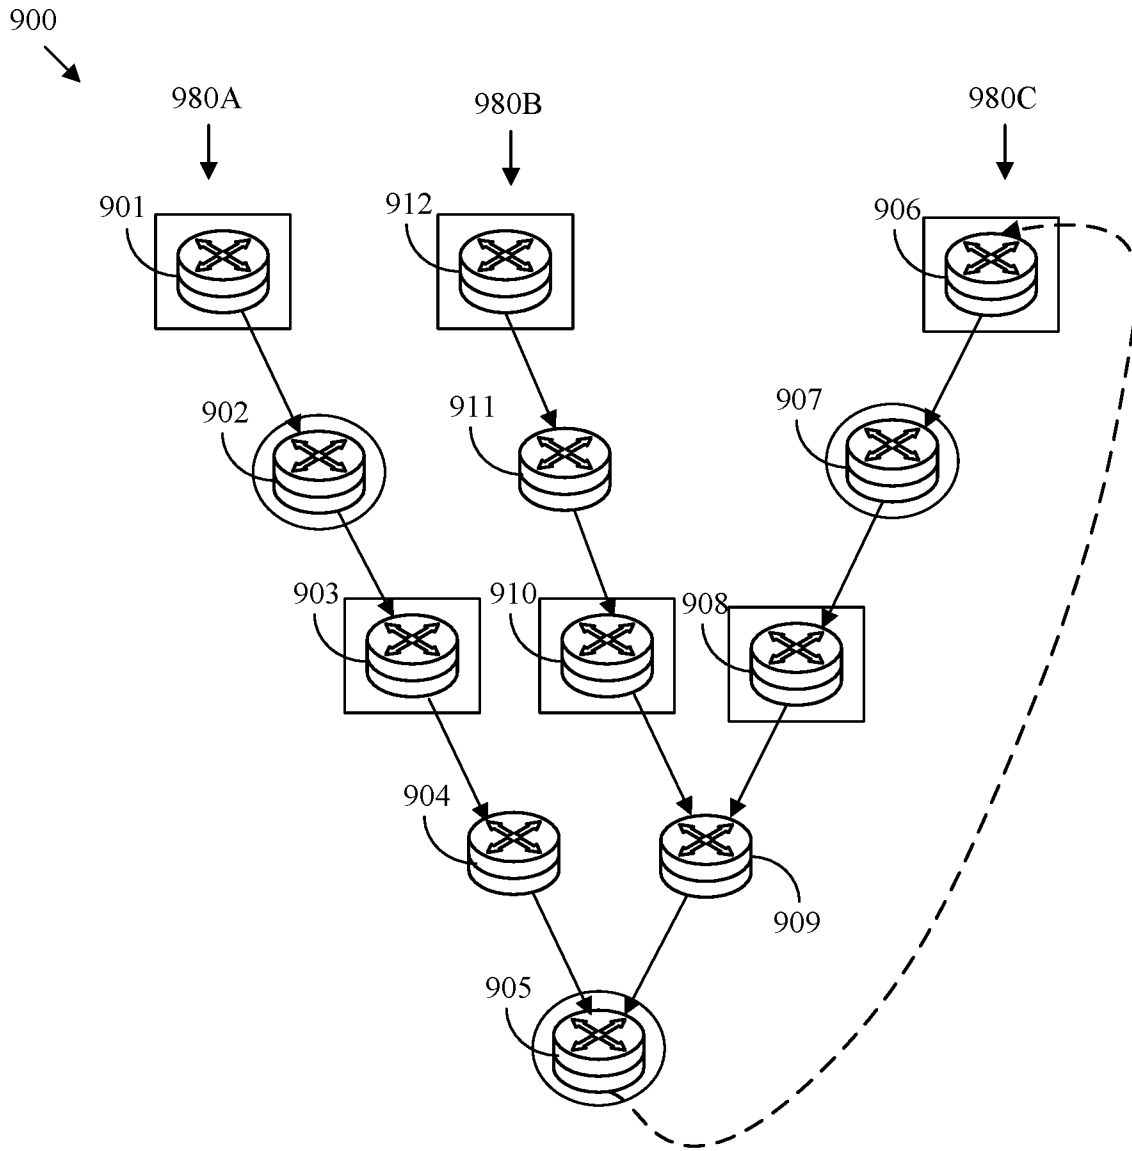


FIG. 9




FIG. 10

1100
↘

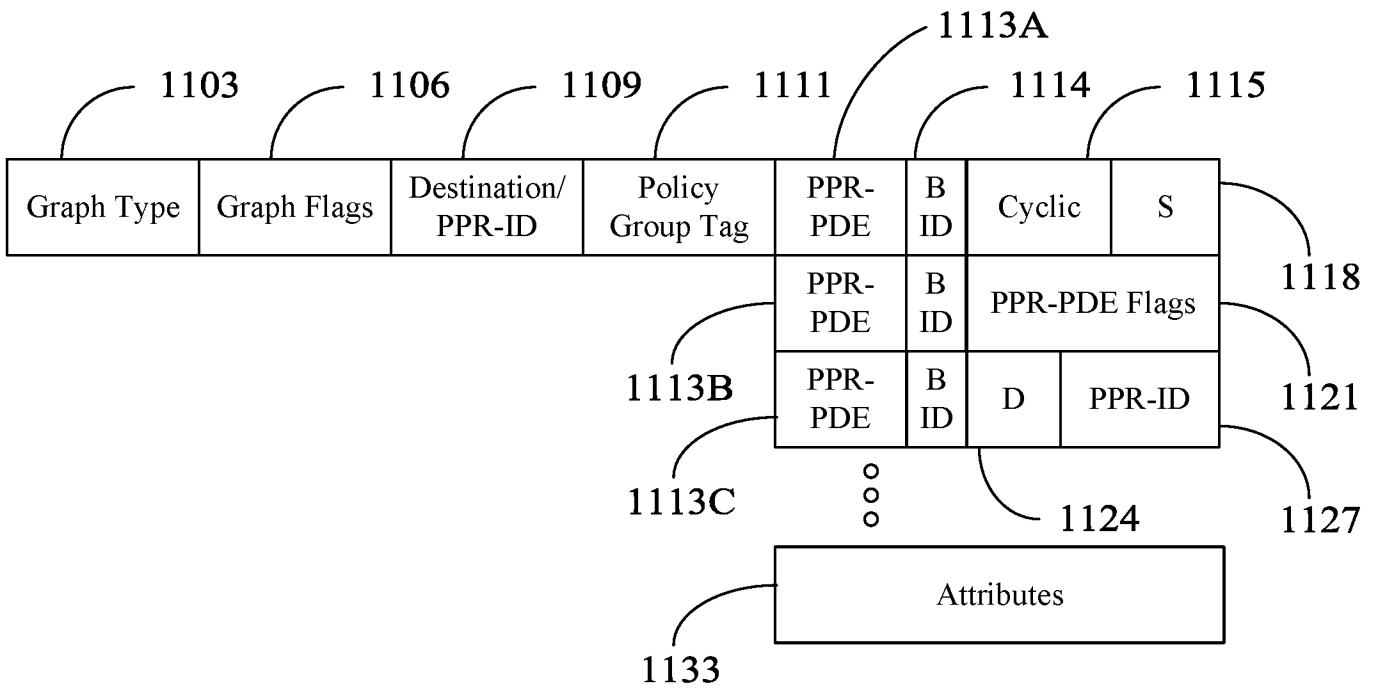


FIG. 11

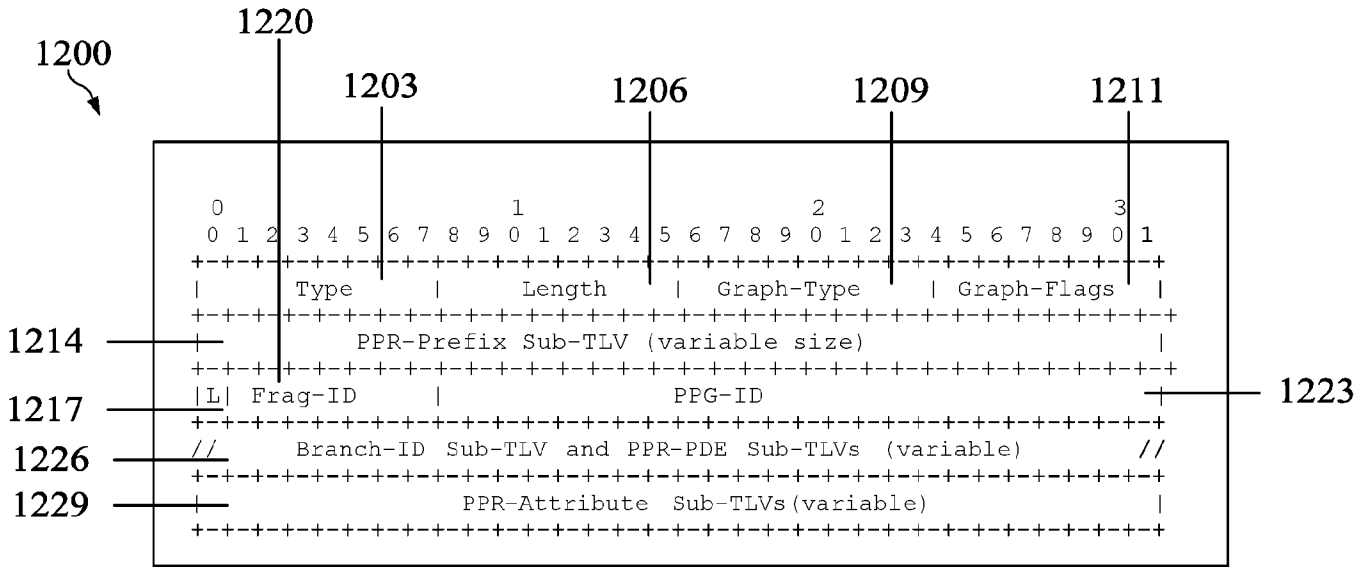


FIG. 12A

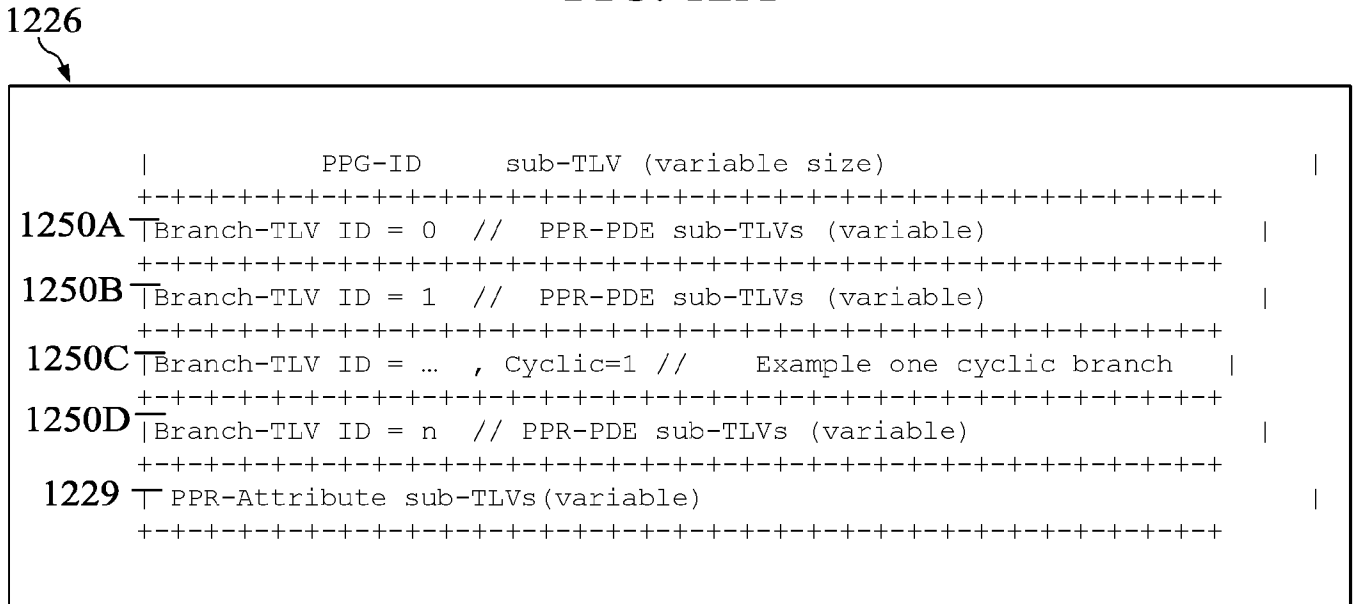



FIG. 12B



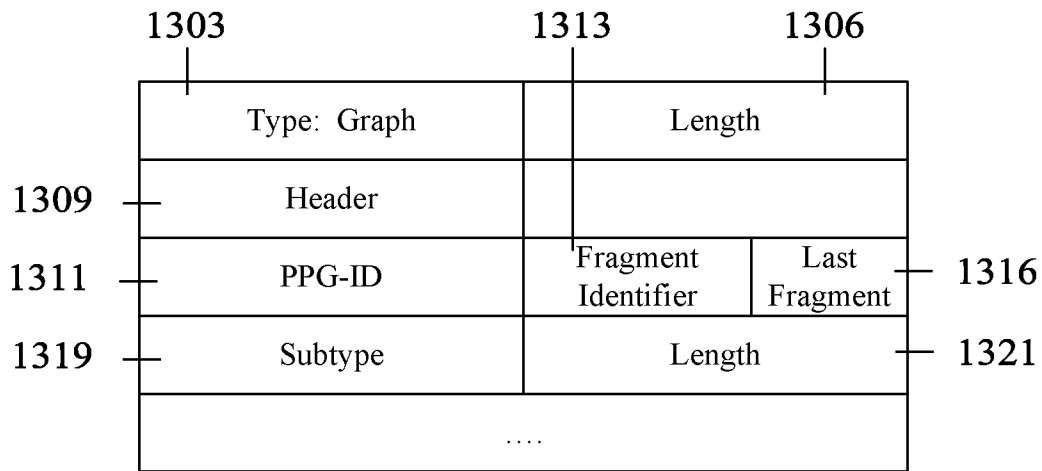
1261

1262

FIG. 12C

1300A
↘

Fragment 0



○
○
○

1300N
↘

Fragment N

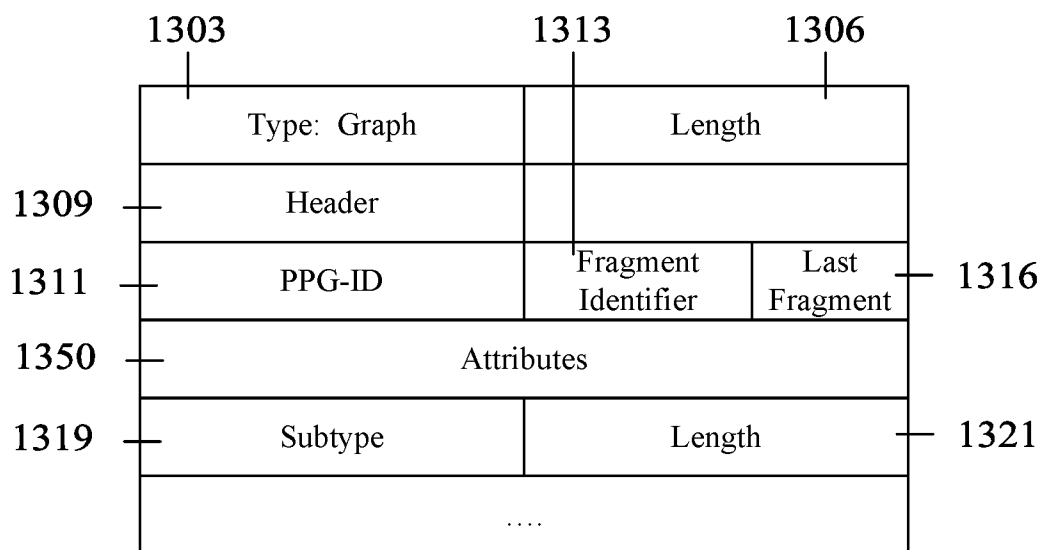


FIG. 13

1400

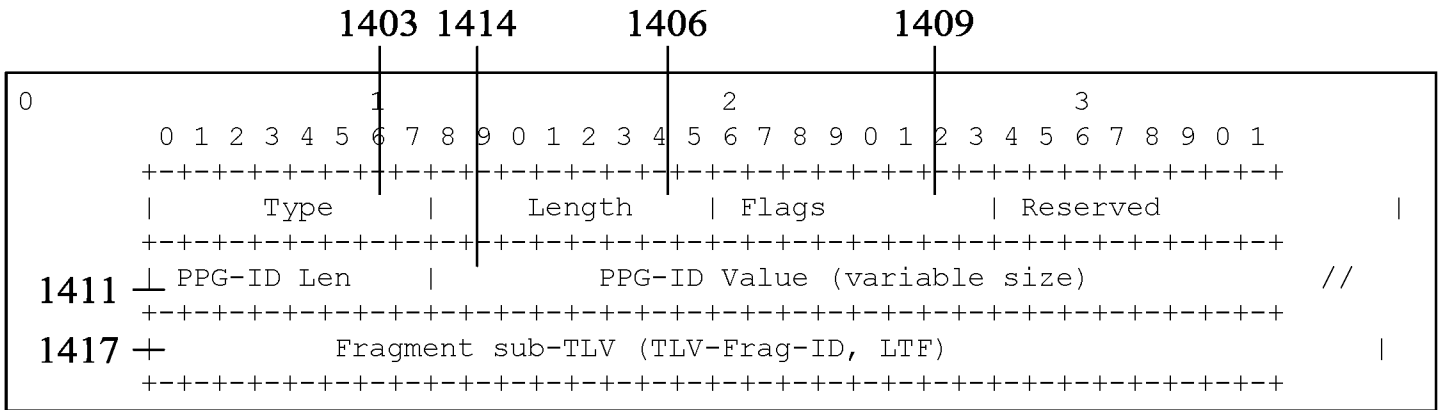


FIG. 14A

1417

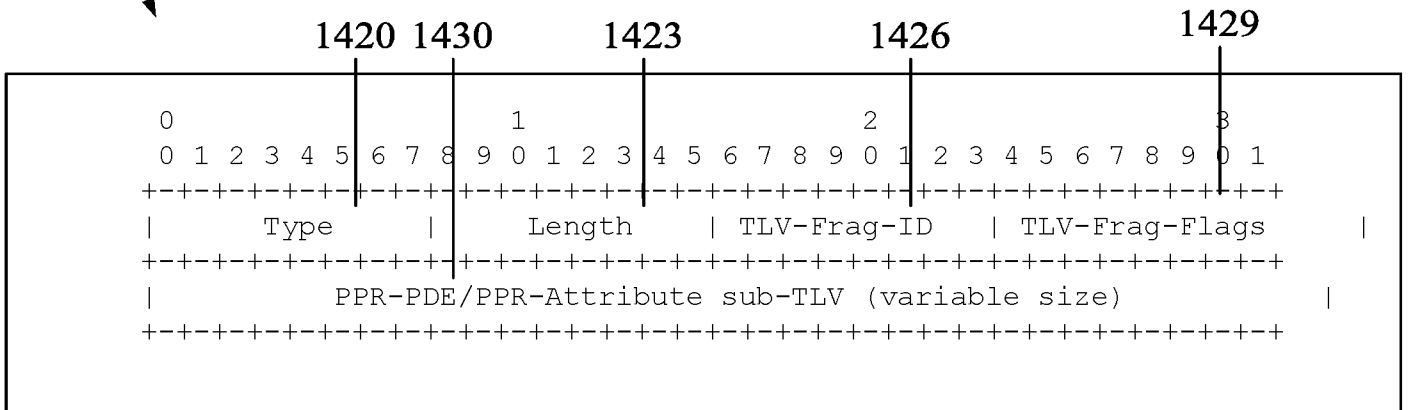


FIG. 14A

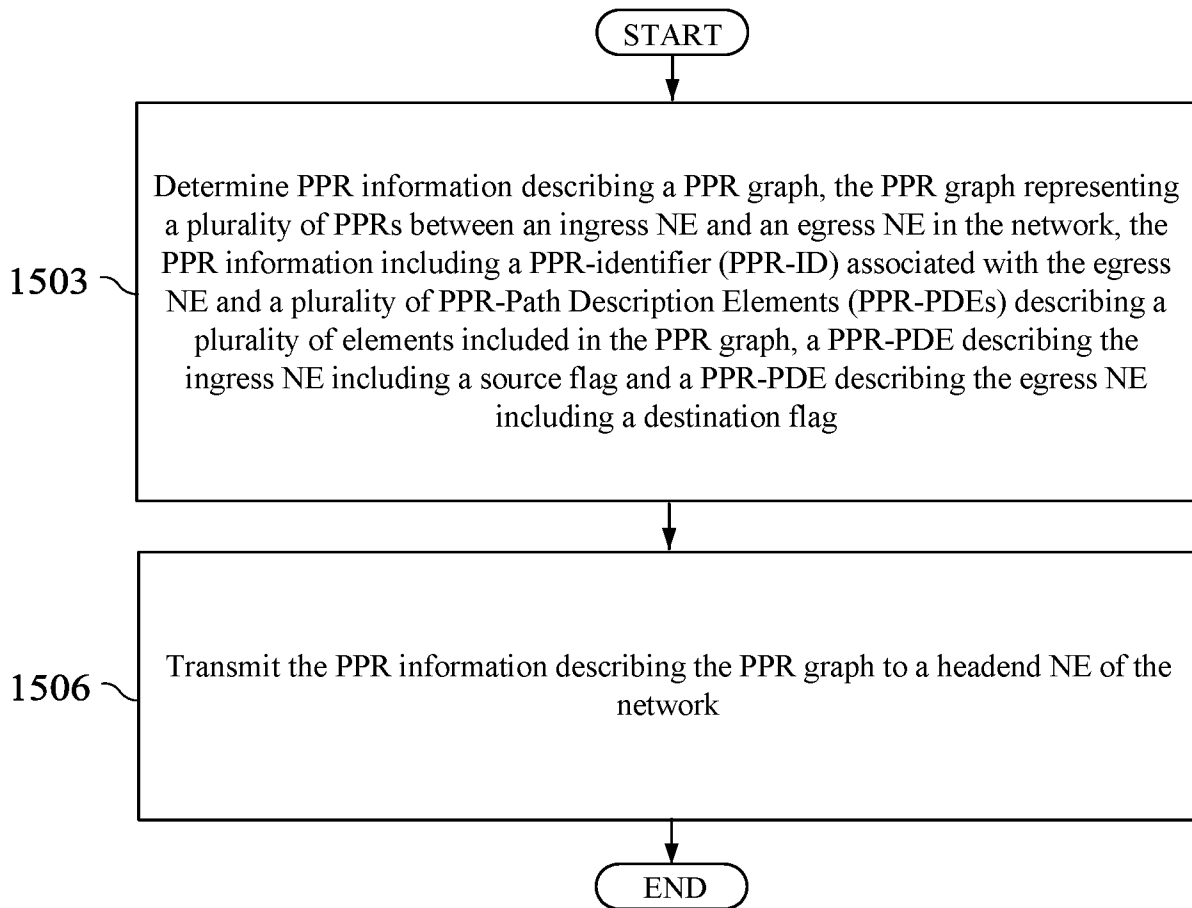
1500
↘

FIG. 15

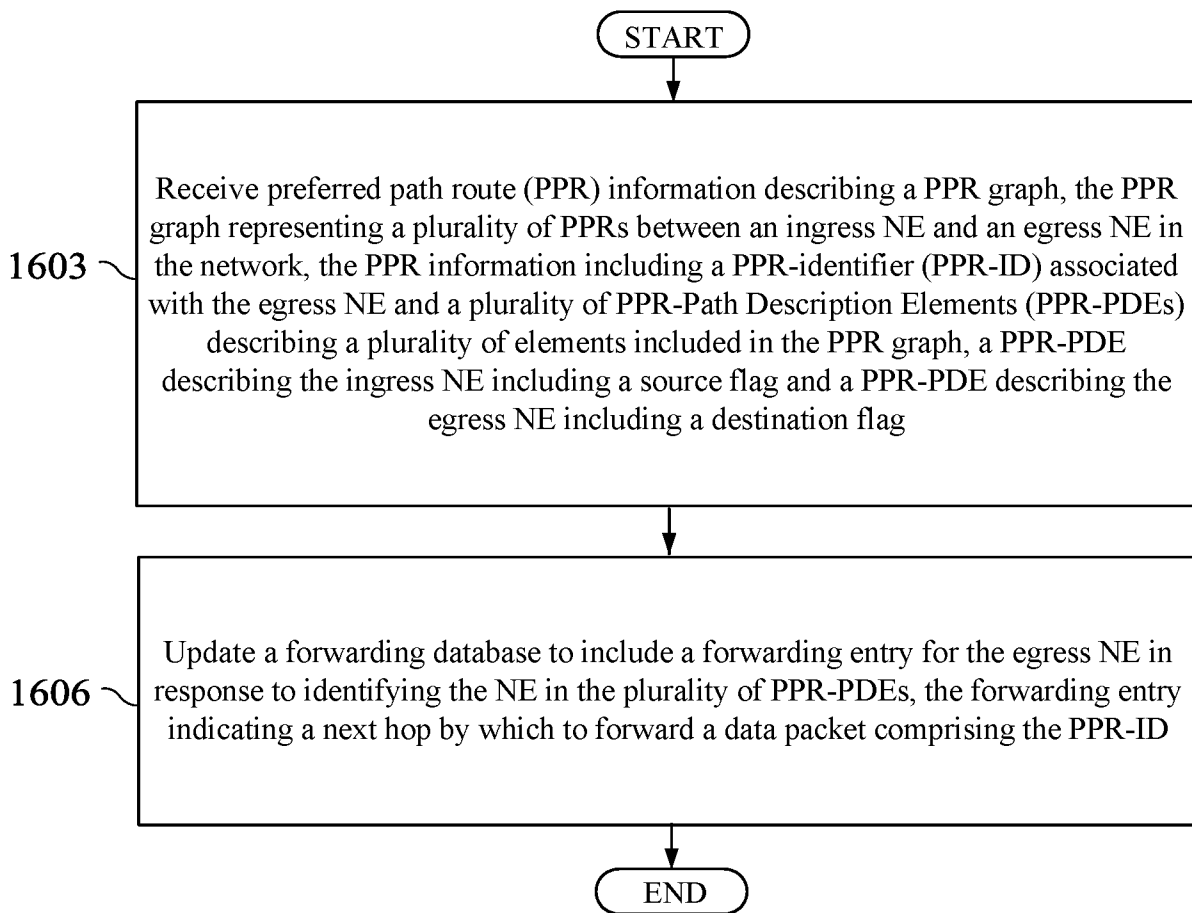
1600
↓

FIG. 16

1700
↘




FIG. 17

1800
↘




FIG. 18

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2019/030440

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/721 H04L12/751 H04L12/707 H04L12/717
ADD.
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)
H04L
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)
EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2016/308758 A1 (LI XU [CA] ET AL) 20 October 2016 (2016-10-20) figures 1-5,7 paragraphs [0019] - [0037] -----	1-24
A	US 2017/104668 A1 (ZHANG QIONG [US] ET AL) 13 April 2017 (2017-04-13) figures 1,2,4 paragraphs [0005], [0023] - [0032], [0036], [0038] -----	1-24

Further documents are listed in the continuation of Box C.

See patent family annex.

* 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
29 May 2019

Date of mailing of the international search report
13/06/2019

Name and mailing address of the ISA/
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer
Tyszka, Krzysztof

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2019/030440

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016308758 A1	20-10-2016	US 2016308758 A1	20-10-2016
		US 2017359256 A1	14-12-2017

US 2017104668 A1	13-04-2017	JP 2017076971 A	20-04-2017
		US 2017104668 A1	13-04-2017
