



- (51) International Patent Classification:  
H04L 29/02 (2006.01) H04L 29/06 (2006.01)
- (21) International Application Number:  
PCT/US2015/044815
- (22) International Filing Date:  
12 August 2015 (12.08.2015)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
14/497,954 26 September 2014 (26.09.2014) US
- (71) Applicant: **128 TECHNOLOGY, INC.** [US/US]; One Burlington Woods Drive, Suite 101, Burlington, MA 01803 (US).
- (72) Inventors: **MELAMPY, Patrick, J.**; 1024 Main Street, Dunstable, MA 01827 (US). **BAJ, Michael**; 42 Clyde Street, Apt. 3, Somerville, MA 02145 (US). **KAPLAN, Hadriel**; 9 Kendall Way, Nashua, NH 03062 (US). **KUMAR, Prashant**; 9 Apollo Circle, Andover, MA 01810 (US). **PENFIELD, Robert**; 72 Adin Drive, Concord, MA 01742 (US). **TIMMONS, Patrick**; 1188 Beacon Street, Unit B, Newton, MA 02461 (US).


(74) Agents: **SUNSTEIN, Bruce D.** et al.; Sunstein Kann Murphy & Timbers LLP, 125 Summer Street, Boston, MA 02110 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, 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, 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).

[Continued on next page]

(54) Title: COMPUTER NETWORK PACKET FLOW CONTROLLER



(57) Abstract: An intermediate node obtains a lead packet of a plurality of packets in a session having a unique session identifier, modifies the lead packet to identify at least the intermediate node, and then forwards the lead packet toward the destination node through an intermediate node electronic output interface to the IP network. The intermediate node also receives, through an intermediate node electronic input interface in communication with the IP network, a backward message from a next node having a next node identifier. The backward message includes the next node identifier and the session identifier. The intermediate node forms an association between the next node identifier and the session identifier, stores the association in memory to maintain state information for the session, and obtains (e.g., receives) additional packets of the session. Substantially all of the additional packets in the session are forwarded toward the next node using the stored association.

Fig. 8



**Published:**

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

## Computer Network Packet Flow Controller

### PRIORITY CLAIM

[0001] This application claims the benefit of U.S. Patent Application No. 14/497,954, filed September 26, 2014 and titled “Network Packet Flow Controller,” the entire contents of which are hereby incorporated by reference herein, for all purposes.

### TECHNICAL FIELD

[0002] The present invention relates to data routing and, more particularly, to routing packets in an IP network.

### BACKGROUND ART

[0003] The Internet Protocol (“IP”) serves as the de-facto standard for forwarding data messages (“datagrams”) between network devices connected with the Internet. To that end, IP delivers datagrams across a series of Internet devices, such as routers and switches, in the form of one or more data packets. Each packet has two principal parts: (1) a payload with the information being conveyed (e.g., text, graphic, audio, or video data), and (2) a header, known as an “IP header,” having the address of the network device to receive the packet(s) (the “destination device”), the identity of the network device that sent the packet (the “originating device”), and other data for routing the packet.

[0004] Many people thus analogize packets to a traditional letter using first class mail, where the letter functions as the payload, and the envelope, with its return and mailing addresses, functions as the IP header.

[0005] Current Internet devices forward packets one-by-one based essentially on the address of the destination device in the packet header. Among other benefits, this routing scheme enables network devices to forward different packets of a single datagram along different routes to reduce network congestion, or avoid malfunctioning network devices. Those skilled in the art thus

refer to IP as a “stateless” protocol because, among other reasons, it does not save packet path data, and does not pre-arrange transmission of packets between end points.

[0006] While it has benefits, IP’s statelessness introduces various limitations. For example, without modification, a stateless IP network inhibits or prevents: 1) user mobility in mobile networks, 2) session layer load balancing for packet traffic in the network, and 3) routing between private or overlapping networks. The art has responded to this problem by implementing tunneling protocols, which provide these functions. Specifically, tunneling protocols transport IP datagrams to a destination along a route that normally is different than the route the datagram would have taken if it had not used a tunneling protocol. While nominally accomplishing their goals, tunneling protocols undesirably introduce additional problems into the network. For example, tunneling requires additional overhead that can induce IP packet fragmentation, consequently introducing substantial network inefficiencies into a session. In addition, tunnels generally use more bandwidth than non-tunneled packets, and tunnel origination and termination requires additional CPU cycles per packet. Moreover, tunnel addresses must be provisioned in advance, reducing flexibility.

[0007] Other attempts to overcome problems introduced by statelessness suffer from similar deficiencies.

#### SUMMARY OF VARIOUS EMBODIMENTS

In accordance with one embodiment of the invention, a packet routing method statefully directs packets of a session from an originating node toward a destination node in an IP network. To that end, using an intermediate node, a method obtains a lead packet of a plurality of packets in a session having a unique session identifier. The method then modifies the lead packet to identify at least the intermediate node, and then forwards the lead packet toward the destination node through an intermediate node electronic output interface to the IP network. This electronic output interface is in communication with the IP network. The method also receives, through an intermediate node electronic input interface in communication with the IP network, a backward message (e.g., a packet) from a next node having a next node identifier. The backward message includes the next node identifier and the session identifier. The method thus forms an association between the next node identifier and the session identifier, stores the association in memory to maintain state information for the session, and obtains (e.g., receives) additional packets of the session. The

method then forwards, through the electronic output interface of the intermediate node, substantially all of the additional packets in the session toward the next node using the stored association.

To forward the lead packet toward the destination node, the method can use the next node identifier to address the lead packet toward the next node. To that end, the method may access a routing information base having routing information for the next node. Other embodiments simply forward the lead packet without addressing it toward the next node. The intermediate node thus may forward the lead packet toward the destination node simply using its conventional routing table. In either case, the lead packet may be addressed so that a plurality of network devices receives the lead packet after it is forwarded and before the next node receives the lead packet.

Moreover, in addition to forward routing, some embodiments execute backward routing toward the originating node. For example, the method may determine a prior node, if any, through which the lead packet traversed, and form a return association between the identifier of the prior node (i.e., the “prior node identifier”) and the session identifier. To optionally maintain return state information for the session, the method may store, in the memory, the return association. Accordingly, the method may receive a plurality of packets in a return session addressed from the destination toward the originating node, and forward (through the electronic output interface) substantially all the packets in the return session, toward the prior node using the stored return association.

The intermediate node can include any of a wide variety of network devices, such as a routing device.

In accordance with another embodiment of the invention, an intermediate routing device directs packets of a session from an originating node toward a destination node in an IP network. To that end, the routing device has an electronic input interface for receiving IP packets from the IP network, an electronic output interface for forwarding IP packets on the IP network, and state memory for storing session state information. The routing device also has: 1) a lead packet identifier, operatively coupled with the electronic input interface, configured to identify a lead packet of a plurality of packets in a session having a unique session identifier, 2) a lead packet modifier, operatively coupled with the lead packet identifier, configured to modify the lead packet to identify at least the intermediate routing device, and 3) a router, operatively coupled with the electronic output interface, configured to forward the lead packet toward the destination node

through the electronic output interface. In addition, the routing device further has an associator, operatively coupled with the electronic input interface, configured to: a) receive a backward message having a next node identifier and the session identifier from a next node, b) form an association between the next node identifier and the session identifier, and c) store in the state memory the association to maintain state information for the session. The router correspondingly is configured to use the stored association to forward, through the electronic output interface, additional packets in the session toward the next node.

Illustrative embodiments of the invention are implemented as a computer program product having a tangible computer usable medium with computer readable program code thereon. The computer readable code may be read and utilized by a computer system in accordance with conventional processes.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0008] The invention will be more fully understood by referring to the following Detailed Description of Specific Embodiments in conjunction with the Drawings, of which:

[0009] Fig. 1 is a schematic diagram of a hypothetical network, according to the prior art.

[0010] Fig. 2 is a schematic diagram illustrating fragmentation of a message, according to the prior art.

[0011] Fig. 3 is a schematic diagram of a hypothetical internet, according to the prior art.

[0012] Fig. 4 is a schematic diagram of a hypothetical internet that includes a conventional routers and augmented IP routers (AIPRs), according to an embodiment of the present invention.

[0013] Fig. 5 is a schematic layout of an Ethernet header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

[0014] Fig. 6 is a schematic layout of an IP header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

[0015] Fig. 7 is a schematic layout of a TCP header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

[0016] Fig. 8 is a schematic block diagram of an AIPR of Fig. 4, according to an embodiment of the present invention.

[0017] Fig. 9 is a schematic illustration of information stored in an information base by the AIRP of Figs. 4 and 8, according to an embodiment of the present invention.

[0018] Fig. 10 is a schematic diagram of a modified lead packet produced by the AIRP of Figs. 4 and 8, according to an embodiment of the present invention.

[0019] Figs. 11 and 12 contain flowcharts schematically illustrating operations performed by the AIRP of Figs. 4 and 8, according to an embodiment of the present invention.

[0020]

#### DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

[0021] In accordance with preferred embodiments of the invention, a packet flow controller ensures that, at least for flow in a forward direction, packets of a session follow the same path as the lead/first packet of that session, i.e., from a source device to a destination device. Details of various embodiments are discussed below.

#### Networks

[0022] Illustrative embodiments preferably are implemented on a conventional computer network. Among other things, a network includes at least two nodes and at least one link between the nodes. Nodes can include computing devices (sometimes referred to as hosts or devices) and routers. Computers include personal computers, smart phones, automatic teller machines (ATMs) and many other types of equipment that include processors and network interfaces. Links include wired and wireless connections between pairs of nodes. In addition, nodes and/or links may be implemented completely in software, such as in a virtual machine, a software defined network, and using network function virtualization. Many networks include switches, which are largely transparent for purposes of this discussion. However, some switches also perform routing functions. For the present discussion, such routing switches are considered routers. Routers are described below.

[0023] A node can be directly connected to one or more other nodes, each via a distinct link. For example, Fig. 1 schematically shows a Node A directly connected to Node B via Link 1. In a given network (e.g., within a local area network), each node has a unique network address to facilitate sending and receiving data. A network includes all the nodes addressable within the

network according to the network's addressing scheme and all the links that interconnect the nodes for communication according to the network's addressing scheme. For example, in Fig. 1, Node A, Node B, Node C,... Node F and all the links 1-8 together make up a network 100. For simplicity, a network is depicted as a cloud or as being enclosed within a cloud.

**[0024]** Nodes initiate communications with other nodes via the network, and nodes receive communications initiated by other nodes via the network. For example, a node may transmit/forward/send data (a message) to a directly connected (adjacent) node by sending the message via the link that interconnects the adjacent nodes. The message includes the network address of the sending node (the "source address") and the network address of the intended receiving node (the "destination address"). A sending node can send a message to a non-adjacent node via one or more other nodes. For example, Node D may send a message to Node F via Node B. Using well known networking protocols, the node(s) between the source and the destination forward the message until the message reaches its destination. Accordingly, to operate properly, network protocols enable nodes to learn or discover network addresses of non-adjacent nodes in their network.

**[0025]** Nodes communicate via networks according to protocols, such as the well-known Internet Protocol (IP) and Transmission Control Protocol (TCP). The protocols are typically implemented by layered software and/or hardware components, such as according to the well-known seven-layer Open System Interconnect (OSI) model. As an example, IP operates at OSI Layer 3 (Network Layer), while the TCP operates largely at OSI Layer 4 (Transport Layer). Each layer performs a logical function and abstracts the layer below it, therefore hiding details of the lower layer.

**[0026]** For example, Layer 3 may fragment a large message into smaller packets if Layer 2 (Data Link Layer) cannot handle the message as one transmission. Fig. 2 schematically illustrates a large message 200 divided into several pieces 202, 204, 206, 208, 210 and 212. Each piece 202-212 may then be sent in a separate packet, exemplified by packet 214. Each packet includes a payload (body) portion, exemplified by payload 216, and a header portion, exemplified at 218. The header portion 218 contains information, such as the packet's source address, destination address and packet sequence number, necessary or desirable for: 1) routing the packet to its destination, 2) reassembling the packets of a message, and 3) other functions provided according to the protocol. In

some cases, a trailer portion is also appended to the payload, such as to carry a checksum of the payload or of the entire packet. All packets of a message need not be sent along the same path, i.e., through the same nodes, on their way to their common destination. It should be noted that although IP packets are officially called IP datagrams, they are commonly referred to simply as packets.

**[0027]** Some other protocols also fragment data into packets. For example, the TCP fragments data into segments, officially referred to as TCP protocol data units (PDUs). Nevertheless, in common usage, the term packet is used to refer to PDUs and datagrams, as well as Ethernet frames.

**[0028]** Most protocols encapsulate packets of higher level protocols. For example, IP encapsulates a TCP packet by adding an IP header to the TCP packet to produce an IP packet. Thus, packets sent at a lower layer can be thought of as being made up of packets within packets. Conventionally, a component operating according to a protocol examines or modifies only information within a header and/or trailer that was created by another component, typically within another node, operating according to the same protocol. That is, conventionally, components operating according to a protocol do not examine or modify portions of packets created by other protocols.

**[0029]** In another example of abstraction provided by layered protocols, some layers translate addresses. Some layers include layer-specific addressing schemes. For example, each end of a link is connected to a node via a real (e.g., electronic) or virtual interface, such as an Ethernet interface. At Layer 2 (Data Link Layer), each interface has an address, such as a media access control (MAC) address. On the other hand, at Layer 3 using IP, each interface, or at least each node, has an IP address. Layer 3 converts IP addresses to MAC addresses.

**[0030]** A router typically acts as a node that interconnects two or more distinct networks or two or more sub-networks (subnets) of a single network, thereby creating a “network of networks” (i.e., an internet). Thus, a router has at least two interfaces, where each interface connects the router to a different network, as exemplified in Fig. 3. When a router receives a packet via one interface from one network, it uses information stored in its routing table to direct the packet to another network via another interface. The routing table contains network/next hop associations. These associations tell the router that a particular destination can optimally be reached by sending the packet to a specific router that represents a next hop on the way to the final destination. For

example, if Router 1 300 receives a packet, via its Interface 1 304, from Network 1 302, and the packet is destined to a node in Network 3 306, the Router 1 300 consults its router table and then forwards the packet via its Interface 2 308 to Network 2 310. Network 2 310 will then forward the packet to Network 3 306. The next hop association can also be indicated in the routing table as an outgoing (exit) interface to the final destination.

**[0031]** Large organizations, such as large corporations, commercial data centers and telecommunications providers, often employ sets of routers in hierarchies to carry internal traffic. For example, one or more gateway routers may interconnect each organization's network to one or more Internet service providers (ISPs). ISPs also employ routers in hierarchies to carry traffic between their customers' gateways, to interconnect with other ISPs, and to interconnect with core routers in the Internet backbone.

**[0032]** A router is considered a Layer 3 device because its primary forwarding decision is based on the information in the Layer 3 IP packet—specifically the destination IP address. A conventional router does not look into the actual data contents (i.e., the encapsulated payload) that the packet carries. Instead, the router only looks at the Layer 3 addresses to make a forwarding decision, plus optionally other information in the header for hints, such as quality of service (QoS) requirements. Once a packet is forwarded, a conventional router does not retain any historical information about the packet, although the forwarding action may be collected to generate statistical data if the router is so configured.

**[0033]** As noted, when a router receives a packet via one interface from one network, the router uses its routing table to direct the packet to another network. Table 1 lists information typically found in a basic IP routing table.

Table 1

---

Destination	Partial IP address (Expressed as a bit-mask) or Complete IP address of a packet's final destination
Next hop	IP address to which the packet should be forwarded on its way to the final destination
Interface	Outgoing network interface to use to forward the packet
Cost/Metric	Cost of this path, relative to costs of other possible paths
Routes	Information about subnets, including how to reach subnets that are not directly attached to the router, via one or more hops; default routes to use for certain types of traffic or when information is lacking

---

**[0034]** Routing tables may be filled in manually, such as by a system administrator, or dynamically by the router. The router uses routing protocols to exchange information with other routers and, thereby, dynamically learn about surrounding network or internet topology. For example, routers announce their presence in the network(s), more specifically, the range of IP addresses to which the routers can forward packets. Neighboring routers update their routing tables with this information and broadcast their ability to forward packets to the network(s) of the first router. This information eventually spreads to more distant routers in a network. Dynamic routing allows a router to respond to changes in a network or internet, such as increased network congestion, new routers joining an internet and router or link failures.

**[0035]** A routing table therefore provides a set of rules for routing packets to their respective destinations. When a packet arrives, a router examines the packet's contents, such as its destination address, and finds the best matching rule in the routing table. The rule essentially tells the router which interface to use to forward the packet and the IP address of a node to which the packet is forwarded on its way to its final destination IP address.

**[0036]** With hop-by-hop routing, each routing table lists, for all reachable destinations, the address of the next node along a path to that destination, i.e., the next hop. Assuming that the routing tables are consistent, a simple algorithm of each router relaying packets to their destinations'

respective next hop suffices to deliver packets anywhere in a network. Hop-by-hop is a fundamental characteristic of the IP Internetwork Layer and the OSI Network Layer.

[0037] Thus, each router's routing table typically merely contains information sufficient to forward a packet to another router that is "closer" to the packet's destination, without a guarantee of the packet ever being delivered to its destination. In a sense, a packet finds its way to its destination by visiting a series of routers and, at each router, using then-current rules to decide which router to visit next, with the hope that at least most packets ultimately reach their destinations.

[0038] Note that the rules may change between two successive hops of a packet or between two successive packets of a message, such as if a router becomes congested or a link fails. Two packets of a message may, therefore, follow different paths and even arrive out of order. In other words, when a packet is sent by a source node, there is no predetermined path the packet will take between the source node and the packet's destination. Instead, the path typically is dynamically determined as the packet traverses the various routers. This may be referred to as "natural routing," i.e., a path is determined dynamically as the packet traverses the internet.

[0039] It should be noted that conventionally, packets sent by the destination node back to the source node may follow different paths than the packets from the source node to the destination node.

[0040] In many situations, a client computer node establishes a session with a server computer node, and the client and server exchange packets within the session. For example, a client computer executing a browser may establish a session with a web server. The client may send one or more packets to request a web page, and the web server may respond with one or more packets containing contents of the web page. In some types of sessions, this back-and-forth exchange of packets may continue for several cycles. In some types of sessions, packets may be sent asynchronously between the two nodes.

[0041] A session has its conventional meaning; namely, it is a plurality of packets sent by one node to another node, where all the packets are related, according to a protocol. A session may be thought of as including a lead (or initial) packet that begins the session, and one or more subsequent packets of the session. A session has a definite beginning and a definite end. For example, a TCP session is initiated by a SYN packet. In some cases, the end may be defined by a prescribed packet or series of packets. For example, a TCP session may be ended with a FIN

exchange or an RST. In other cases, the end may be defined by lack of communication between the nodes for at least a predetermined amount of time (a timeout time). For example, a TCP session may be ended after a defined timeout period. Some sessions include only packets sent from one node to the other node. Other sessions include response packets, as in the web client/server interaction example. A session may include any number of cycles of back-and-forth communication, or asynchronous communication, according to the protocol, but all packets of a session are exchanged between the same client/server pair of nodes. A session is also referred to herein as a series of packets.

**[0042]** A computer having a single IP address may provide several services, such as web services, e-mail services and file transfer (FTP) services. Each service is typically assigned a port number in the range 0-65,535 that is unique on the computer. A service is, therefore, defined by a combination of the node's IP address and the service's port number. Note that this combination is unique within the network the computer is connected to, and it is often unique within an internet. Similarly, a single node may execute many clients. Therefore, a client that makes a request to a service is assigned a unique port number on the client's node, so return packets from the service can be uniquely addressed to the client that made the request.

**[0043]** The term socket means an IP address-port number combination. Thus, each service has a network-unique, and often internet-unique, service socket, and a client making a request of a service is assigned a network-unique, and sometimes internet-unique, client socket. In places, the terms source client and destination service are used when referring to a client that sends packets to make requests of a service and the service being requested, respectively.

**[0044]** Although natural routing has performed well for many years, natural routing has shortcomings. For example, because each packet of a session may travel along a different path and traverse a different set of routers, it is difficult to collect metrics for the session. Security functions that may be applicable to packets of the session must be widely distributed or risk not being applied to all the packets. Furthermore, attacks on the session may be mounted from many places.

### Forward and Backward Flow Control

[0045] Illustrative embodiments of the present invention overcome these and other shortcomings by ensuring that subsequent packets of a session follow the same path as the lead packet of the session, at least in the forward direction, i.e., from the source client to the destination service. The subsequent packets traverse at least a subset of the routers the lead packet traverses between the source client and the destination service. Each router in the subset is referred to herein as an intermediate node or waypoint, although the waypoints are not necessarily predetermined before the lead packet is sent by the source client. The lead packet may be naturally routed. Nevertheless, the path taken by the lead packet establishes the waypoints, and the subsequent packets traverse the same waypoints, and in the same order, as the lead packet.

[0046] Of course, some packets may be dropped along the way, as is typical in an IP network or internet, such as by an overloaded router or due to corruption of the packet by a link. Thus, all the packets sent by the source client need not reach the session's destination service and, consequently, all the packets sent by the source client need not traverse all the waypoints. However, subsequent packets that do reach the destination service must traverse all the waypoints. For simplicity of explanation, dropped packets are ignored in the remaining discussion, and the term "all the packets" means all the packets that reach their respective destinations.

[0047] As a result of this forward flow control, metrics collected at one of the waypoints represent all the packets of the session. These metrics are not diluted by packets that bypass the waypoint, because no packet of the session can bypass any waypoint. Security functions, such as inspection for malicious packets, performed at one waypoint are sure to be performed on all packets of the session.

[0048] Some embodiments of the present invention also ensure that return packets from the destination service to the source client also follow the same path, i.e., traverse the waypoints, but in reverse order. This reverse flow control enables use of paths, such as via proprietary networks, that might not otherwise be available by naturally routing the return packets.

[0049] A packet flow controller (also referred to herein as an augmented IP router ("AIPR")) ensures that subsequent packets of a session follow the same path as the lead packet of the session, as discussed above. An AIPR also performs conventional routing functions. Fig. 4 is a schematic diagram illustrating a hypothetical set of interconnected networks 400, 402, 404 and 406,

i.e., an internet. Each network 400-406 includes a number of routers and AIPRs, not all of which are necessarily shown. Network 400 includes AIPR 1 408 and router 410. Network 400 may be, for example, a network of a telecommunications carrier. Network 402 includes a router 412 and AIPR 2 414. Network 402 may be, for example, a network of a first ISP. Network 404 includes a router 416 and AIPR 3 418. Network 404 may be, for example, the Internet backbone or a portion thereof. Network 406 includes a router 420, AIPR 4 422 and another router 424. Network 406 may be, for example, a network of a second ISP.

**[0050]** Assume a source client node 426 initiates a session with a destination service node 428. For example, the source client 426 may request a web page, and the destination service node 428 may include a web server. The source client 426 may, for example, be part of a first local area network (LAN) (not shown) within a first corporation, and the LAN may be connected to the telecommunications carrier network 400 via a gateway router 430 operated by the corporation. Similarly, the destination service node 428 may be operated by a second corporation, and it may be part of a second LAN (not shown) coupled to the network 406 of the second ISP via a gateway router 432 operated by the second corporation. As a lead packet of the session traverses the internet, each AIPR (waypoint) the packet traverses records information that eventually enables the waypoint to be able to identify its immediately previous waypoint and its immediately next waypoint, with respect to the session.

**[0051]** The lead packet of the session is naturally routed. Assume the lead packet reaches AIPR 1 408 before it reaches network 402, 404 or 406. AIPR 1 408 automatically identifies the lead packet as being an initial packet of the session. AIPR 1 408 may use various techniques to identify the beginning of a session, as noted above and as discussed in more detail below. AIPR 1 408 becomes the first waypoint along a path the lead packet eventually follows.

**[0052]** AIPR 1 408 assigns a unique identifier to the session and stores information about the session in the AIPR's database to enable the AIPR 1 408 to identify subsequent packets of the session. In some embodiments, AIPR 1 408 reads the client socket/service socket number pair in the lead packet and stores the client socket/service socket number pair in a database to uniquely identify the session. This enables the AIPR 1 408 to identify the subsequent packets as being part of the session, because all subsequent packets of the session will contain the same client socket/service socket number pair.

[0053] In some embodiments, AIPR 1 408 sets a flag in its database to indicate the lead packet has not traversed any other AIPR before reaching AIPR 1 408. This flag may be used later, for example when the AIPR 1 408 handles return packets. AIPR 1 408 may be able to identify the lead packet as not having traversed any other AIPR by lack of any modification to the packet. Packet modification is described below.

[0054] AIPR 1 408 modifies the lead packet to indicate the lead packet has been handled by an AIPR. In some embodiments, the AIPR 1 408 stores the unique identifier of the session and, if not included in the unique identifier, the AIPR's network address in the packet to produce a modified lead packet. Subsequent AIPRs, if any, that handle the (now modified) lead packet use this modification to identify the lead packet as a lead packet that has been handled by an AIPR, and to indicate that subsequent packets of the session should be routed the same way as the lead packet is routed.

[0055] In some embodiments, AIPR 1 408 assigns a port number on the interface over which AIPR 1 408 will forward the lead packet. The AIPR's network address and this port number, in combination, may be used as a unique identifier of the session, at least from the point of view of the next AIPR along the path. AIPR 1 408 may include the AIPR's network address-port number combination in the modified lead packet. Thus, the next AIPR along the path may assume that subsequent packets sent from this network address-port number combination are part of, or likely to be part of, the session.

[0056] AIPR 1 408 then forwards the lead packet naturally. The lead packet traverses an unspecified number of nodes of network 400 until it reaches router 410, which naturally routes the lead packet to network 402. Assume the router 410 forwards the lead packet to AIPR 2 414 in network 402.

[0057] AIPR 2 414 detects the modification to the lead packet, identifying a need for special treatment. AIPR 2 414 becomes the second waypoint along the path the lead packet will follow. AIPR 2 414 stores in its database the network address of AIPR 1 408 and the port number assigned by AIPR 1 408, in association with a unique identifier of the session, such as the client and server socket number pair, thus identifying the previous waypoint along the path in association with the session. In this way, each waypoint learns the network address and port number of the previous waypoint along this session's path and uses a related association device (an "associator") to

associate this information with a session identifier. This information may be used later to forward return packets, from waypoint to waypoint, back to the source client 426.

**[0058]** In some embodiments, AIPR 2 414 assigns a port number on the interface over which the lead packet was received. The AIPR's network address and this port number, in combination, may be used as a unique identifier of the session, at least from the point of view of AIPR 1 408. Thus, subsequent packets addressed to this network address-port number combination may be assumed to be, or at least are likely to be, part of the session.

**[0059]** In some embodiments, AIPR 2 414 sends a packet back to AIPR 1 408 to inform AIPR 1 408 of the network address-port number combination, in association with the identification of the session. In some embodiments, the network address-port number combination are sent to AIPR 1 408 later, in connection with a return packet, as described below. In either case, AIPR 1 408 learns a network address-port number combination unique to the session, and AIPR 1 408 sends subsequent packets to that address-port combination, rather than naturally forwarding the subsequent packets. In this way, each waypoint learns the network address and port number of the next waypoint along this session's path. This information is used to forward subsequent packets, from waypoint to waypoint, forward to the destination service 428, along the same path as the lead packet. This kind of routing is unlike any routing taught by the prior art known to the inventors.

**[0060]** AIPR 2 214 modifies the lead packet to include the network address of AIPR 2 214, and then forwards the lead packet naturally. As with AIPR 1 408, in some embodiments AIPR 2 214 assigns a port number on the interface over which AIPR 2 214 forwards the packet, and the network address of AIPR 2 214 and the port number are included in the modified lead packet AIPR 2 214 sends.

**[0061]** The lead packet traverses an unspecified number of nodes of network 402, until it reaches router 412, which naturally routes the lead packet to network 404. Assume the router 416 forwards the lead packet to AIPR 3 418.

**[0062]** AIPR 3 418 becomes the third waypoint along the path the lead packet will follow. AIPR 3 418 operates much as AIPR 2 414. The lead packet is then forwarded to network 406, where it traverses AIPR 4 422, which becomes the fourth waypoint.

**[0063]** Three scenarios are possible with respect to the last AIPR 422 (AIPR 4) along the path to the destination service 428.

[0064] In the first scenario, one or more AIPRs relatively close to a destination service are provisioned to handle lead packets for the destination service. The AIPRs may be so provisioned by storing information in their databases to identify the destination service, such as by the service socket number or other unique identifier of the service. These “terminus” AIPRs broadcast their ability to forward packets to the destination service. A terminus AIPR is an AIPR that can forward packets to a destination service, without the packets traversing another AIPR. A terminus AIPR recognizes a lead packet destined to a service that terminates at the AIPR by comparing the destination service socket number to the information provisioned in the AIPR’s database.

[0065] If AIPR 4 422 has been so provisioned, AIPR 4 422 may restore the lead packet to its original form, i.e., the form the lead packet had when the source client 426 sent the lead packet, or as the packet might have been modified by the router 430, such as a result of network address translation (NAT) performed by the router 430. Thus, the lead packet may be restored to a form that does not include any of the modifications made by the waypoints 408, 414 and 418. AIPR 4 422 then forwards the lead packet to the destination service 428. Like AIPR 3 418, AIPR 4 422 stores information in its database identifying AIPR 3 418 as the previous AIPR for this session.

[0066] In the second scenario, AIPR 4 422 is not provisioned with information about the destination service 428. In such embodiments, AIPR 4 422 may operate much as AIPR 2 414 and AIPR 3 418 operate. AIPR 4 422 modifies and naturally forwards the lead packet, and the lead packet is eventually delivered to the destination service 428. The destination service 428 responds to the lead packet. For example, if the lead packet is a SYN packet to initiate a TCP session, the destination service 428 responds with an ACK or SYN/ACK packet. AIPR 4 422 recognizes the return packet as being part of the session, such as based on the source client/destination service network address/port number pairs in the return packet. Furthermore, because the return packet was sent by the destination service 428, and not another AIPR, AIPR 4 422 recognizes that it is the last AIPR along the path for this service.

[0067] AIPR 4 422 stores information in its database indicating AIPR 4 422 is a terminus AIPR. If AIPR 4 422 receives subsequent packets of the session, AIPR 4 422 may restore the subsequent packets to their original forms, i.e., the forms the subsequent packets had when the source client 426 sent the subsequent packets, or as the packets might have been modified by the

router 430, such as a result of network address translation (NAT) performed by the router 430. AIPR 4 422 forwards the subsequent packets to the destination service 428.

**[0068]** AIPR 4 422 modifies the return packet to include a port number on the interface AIPR 4 422 received the lead packet from AIPR 3 418, as well as the network address of AIPR 4 422. AIPR 4 422 then forwards the return packet to AIPR 3 418. Although the return packet may be forwarded by other routers, AIPR 4 422 specifically addresses the return packet to AIPR 3 418. This begins the return packet's journey back along the path the lead packet traveled, through all the waypoints traversed by the lead packet, in reverse order. Thus, the return packet is not naturally routed back to the source client 426. This kind of return packet routing is unlike any routing taught by the prior art known by the inventors.

**[0069]** AIPR 3 418 receives the modified return packet and, because the return packet was addressed to the port number AIPR 3 418 previously assigned and associated with this session, AIPR 3 418 can assume the return packet is part of, or likely part of, the session. AIPR 3 418 copies the network address and port number of AIPR 4 422 from the return packet into the AIPR's database as the next waypoint for this session. If AIPR 3 418 receives subsequent packets of the session, AIPR 3 418 forwards them to the network address and port number of the next waypoint, i.e., AIPR 4 422.

**[0070]** Thus, once an AIPR is notified of a network address and port number of a next AIPR along a session path, the AIPR forwards subsequent packets to the next AIPR, rather than naturally routing the subsequent packets.

**[0071]** AIPR 3 418 forwards the return packet to AIPR 2 414, whose network address and port number were stored in the database of AIPR 3 418 and identified as the previous waypoint of the session. Likewise, each of the waypoints along the path back to the source client 426 forwards the return packet to its respective previous waypoint.

**[0072]** When the first waypoint, i.e., AIPR 1 408, receives the return packet, the waypoint may restore the return packet to its original form, i.e., the form the return packet had when the destination service 428 sent the return packet, or as the packet might have been modified by the router 430, such as a result of network address translation (NAT) performed by the router 430. Recall that the first waypoint set a flag in its database to indicate the lead packet had not traversed any other waypoint before reaching the first waypoint. This flag is used to signal the first waypoint

to restore the return packet and forward the restored return packet to the source client 426. The first waypoint forwards the return packet to the source client 426. Subsequent return packets are similarly handled.

[0073] In the third scenario, not shown in Fig. 4, the last AIPR to receive the lead packet has a network address equal to the network address of the destination service. For example, the destination service network address may be given to a gateway router/AIPR, and the gateway router/AIPR may either process the service request or its router table may cause the packet to be forwarded to another node to perform the service. The last AIPR may restore the lead packet and subsequent packets, as described above.

#### Lead Packet Identification

[0074] As noted, a waypoint should be able to identify a lead packet of a session. Various techniques may be used to identify lead packets. Some of these techniques are protocol-specific. For example, a TCP session is initiated according to a well-known three-part handshake involving a SYN packet, a SYN-ACK packet and an ACK packet. By statefully following packet exchanges between pairs of nodes, a waypoint can identify a beginning of a session and, in many cases, an end of the session. For example, A TCP session may be ended by including a FIN flag in a packet and having the other node send an ACK, or by simply including an RST flag in a packet. Because each waypoint stores information about each session, such as the source client/destination service network address/port number pairs, the waypoint can identify the session with which each received packet is associated. The waypoint can follow the protocol state of each session by monitoring the messages and flags, such as SYN and FIN, sent by the endpoints of the session and storing state information about each session in its database. Such stateful monitoring of packet traffic is not taught by the prior art known to the inventor. Instead, the prior art teaches away from this type of session.

[0075] It should be noted that a SYN packet may be re-transmitted—each SYN packet does not necessarily initiate a separate session. However, the waypoint can differentiate between SYN packets that initiate a session and re-transmitted SYN packets based on, for example, the response packets.

[0076] Where a protocol does not define a packet sequence to end a session, the waypoint may use a timer. After a predetermined amount of time, during which no packet is handled for a session, the waypoint may assume the session is ended. Such a timeout period may also be applied to sessions using protocols that define end sequences.

[0077] Table 2 describes exemplary techniques for identifying the beginning and end of a session, according to various protocols. Similar techniques may be developed for other protocols, based on the definitions of the protocols.

Table 2

Protocol	Destination Port	Technique for Start/End Determination
TCP	Any	Detect start on the first SYN packet from a new address/port unique within the TCP protocol's guard time between address/port reuse. Following the TCP state machine to determine an end (FIN exchange, RST, or guard timeout).
UDP - TFTP	69	Trap on the first RRQ or WRQ message to define a new session, trap on an undersized DAT packet for an end of session.
UDP-SNMP	161, 162	Trap on the message type, including GetRequest, SetRequest, GetNextRequest, GetBulkRequest, InformRequest for a start of session, and monitor the Response for end of session. For SNMP traps, port 162 is used, and the flow of data generally travels in the "reverse" direction.
UDP-SYSLOG	514	A single message protocol, thus each message is a start of session, and end of session.
UDP-RTP	Any	RTP has a unique header structure, which can be reviewed/analyzed to identify a start of a session. This is not always accurate, but if used in combination with a guard timer on the exact same five-tuple address, it should work well enough. The end of session is detected through a guard timer on the five-tuple session, or a major change in the RTP header.
UDP-RTCP	Any	RTCP also has a unique header, which can be reviewed, analyzed, and harvested for analytics. Each RTCP packet is sent periodically and can be considered a "start of session" with the corresponding RTCP response ending the session. This provides a very high quality way of getting analytics for RTCP at a network middle point, without using a Session Border Controller
UDP-DNS		

(Nameserver)	53	Each DNS query is a single UDP message and response. By establishing a forward session (and subsequent backward session) the Augmented router gets the entire transaction. This allows analytics to be gathered and manipulations that are appropriate at the Augmented router.
UDP-NTP	123	Each DNS query/response is a full session. So, each query is a start, and each response is an end.

[0078] Fig. 5 is a schematic layout of an Ethernet header 500, including a Destination MAC Address 502 and an 802.1q VLAN Tag 504. Fig. 6 is a schematic layout of an IP header 600, including a Protocol field 602, a Source IP Address 604 and a Destination IP Address 606. Fig. 7 is a schematic layout of a TCP header 700, including a Source Port 702, a Destination Port 704, a Sequence Number 706, a SYN flag 708 and a FIN flag 710. These packets and the identified fields may be used to identify the beginning of a session, as summarized in Table 3.

Table 3

Data Item	Where From	Description
Physical Interface	Ethernet Header	This is the actual port that the message was received on, which can be associated or discerned by the Destination MAC Address
Tenant	Ethernet Header OR Source MAD Address & Previous Advertisement	Logical association with a group of computers.
Protocol	IP Header	This defines the protocol in use and, for the TCP case, it must be set to a value that corresponds to TCP
Source IP Address	IP Header	Defines the source IP Address of the initial packet of a flow.
Destination IP Address	IP Header	Defines the destination IP Address of the initial packet of a flow.
Source Port	TCP Header	Defines the flow instance

Destination Port	TCP Header	<p>from the source. This may reflect a client, a firewall in front of the client, or a carrier grade NAT. This defines the desired service requested, such as 80 for HTTP.</p>
Sequence Number	TCP Header	<p>This is a random number assigned by the client. It may be updated by a firewall or carrier grade NAT.</p>
SYN Bit On	TCP Header	<p>When the SYN bit is on, and no others, this is an initial packet of a session. It may be retransmitted if there is no response to the first SYN message.</p>

Augmented IP Router (AIPR)

**[0079]** Fig. 8 is a schematic block diagram of an AIPR (waypoint) 800 configured in accordance with illustrative embodiments of the invention. The AIPR 800 includes at least two network interfaces 802 and 804, through which the AIPR 800 may be coupled to two networks. The interfaces 802 and 804 may be, for example, Ethernet interfaces. The AIPR 800 may send and receive packets via the interfaces 802 and 804.

**[0080]** A lead packet identifier 806 automatically identifies lead packets, as discussed herein. In general, the lead packet identifier 806 identifies a lead packet when the lead packet identifier 806 receives a packet related to a session that is not already represented in the AIPR's information base 810, such as a packet that identifies a new source client/destination service network address/port number pair. As noted, each lead packet is an initial, non-dropped, packet of a series of packets (session). Each session includes a lead packet and at least one subsequent packet. The lead packet and all the subsequent packets are sent by the same source client toward the same destination service, for forward flow control. For forward and backward flow control, all the packets of the session are sent by either the source client or the destination service toward the other.

[0081] A session (packet series) manager 808 is coupled to the lead packet identifier 806. For each session, the session manager assigns a unique identifier. The unique identifier may be, for example, a combination of the network address of the AIPR 800 or of the interface 802, in combination with a first port number assigned by the session manager 808 for receiving subsequent packets of this session. The unique identifier may further include the network address of the AIPR 800 or of the other interface 804, in combination with a second port number assigned by the session manager 808 for transmitting the lead packet and subsequent packets. This unique identifier is associated with the session. The session manager 808 stores information about the session in an information base 810. This information may include the unique identifier, in association with the original source client/destination service network address/port number pairs.

[0082] Fig. 9 is a schematic layout of an exemplary waypoint information base 900. Each row represents a session. A session identification column 902 includes sub-columns for the source client 904 and the destination service 906. For each client 904, its network address 908 and port number 910 are stored. For each destination service 906, its network address 912 and port number 914 are stored. This information is extracted from the lead packet.

[0083] State information about the session may be stored in a state column 915. This information may be used to statefully follow a series of packets, such as when a session is being initiated or ended.

[0084] A backward column includes sub-columns for storing information 916 about a portion of the backward path, specifically to the previous AIPR. The backward path information 916 includes information 918 about the previous AIPR and information 920 about the present AIPR 800. The information 918 about the previous AIPR includes the AIPR's network address 922 and port number 924. The session manager 808 extracts this information from the lead packet, assuming the lead packet was forwarded by an AIPR. If, however, the present AIPR 800 is the first AIPR to process the lead packet, the information 918 is left blank as a flag. The information 920 about the present AIPR 800 includes the network address 926 of the interface 802 over which the lead packet was received, as well as the first port number 928 assigned by session manager 808.

[0085] The waypoint information base 900 is also configured to store information 930 about a portion of the forward path, specifically to the next AIPR. This information 930 includes information 932 about the present AIPR 800 and information 934 about the next AIPR along the

path, assuming there is a next AIPR. The information 932 includes the network address 936 of the interface over which the present AIPR will send the lead packet and subsequent packets, as well as the second port number 938 assigned by the session manager 808. The information 934 about the next AIPR along the path may not yet be available, unless the AIPR is provisioned with information about the forward path. The information 934 about the next AIPR includes its network address 940 and port number 942. If the information 934 about the next AIPR is not yet available, the information 934 may be filled in when the AIPR 800 processes a return packet, as described below.

**[0086]** Some embodiments of the waypoint information base 900 may include the forward information 930 without the backward information 916. Other embodiments of the waypoint information base 900 may include the backward information 916 without the forward information 930.

**[0087]** Returning to Fig. 8, a lead packet modifier 812 is coupled to the session manager 808. The lead packet modifier 812 modifies the lead packet to store the unique identifier associated with the session. The original source client network address/port number pair, and the original destination service network address/port number pair, are stored in the modified lead packet, if necessary. The lead packet may be enlarged to accommodate the additional information stored therein, or existing space within the lead packet, such a vendor specific attribute field, may be used. Other techniques for transmitting additional information are protocol specific, for example with TCP, the additional information could be transmitted as a TCP Option field, or added to the SYN packet as data. In either case, the term session data block is used to refer to the information added to the modified lead packet.

**[0088]** Fig. 10 is a schematic diagram of an exemplary modified lead packet 1000 showing the original source and destination IP addresses 1002 and 1004, respectively, and the original source and destination port numbers 1006 and 1008, respectively. Fig. 10 also shows a session data block 1010 in the modified lead packet 1000. Although the session data block 1010 is shown as being contiguous, it may instead have its contents distributed throughout the modified lead packet 1000. The session data block 1010 may store an identification of the sending AIPR, i.e., an intermediate node identifier 1012, such as the network address of the second network interface 804 and the second port number.

[0089] Returning to Fig. 8, the lead packet modifier 812 updates the packet length, if necessary, to reflect any enlargement of the packet. The lead packet modifier 812 updates the checksum of the packet to reflect the modifications made to the packet. The modified lead packet is then transmitted by a packet router 814, via the second network interface 804. The modified lead packet is naturally routed, unless the AIPR 800 has been provisioned with forward path information.

[0090] Eventually, the destination service sends a return packet. The AIPR 800 receives the return packet via the second interface 804. If another AIPR (downstream AIPR) between the present AIPR 800 and the destination service handles the lead packet and the return packet, the downstream AIPR modifies the return packet to include the downstream AIPR's network address and a port number. A downstream controller 816 identifier uses stateful inspection, as described herein, to identify the return packet. The downstream controller 816 stores information 934 (Fig. 9), specifically the network address and port number, about the next AIPR in the waypoint information base 900.

[0091] The present AIPR 800 may use this information to address subsequent packets to the next AIPR. Specifically, a subsequent packet modifier 818 may set the destination address of the subsequent packets to the network address and port number 940 and 942 (Fig. 9) of the next waypoint, instead of directly to the destination service. The packet router 814 sends the subsequent packets, according to their modified destination addresses. Thus, for each series of packets, subsequent packets flow through the same downstream packet flow controllers as the lead packet of the series of packets.

[0092] A last packet identifier 820 statefully follows each session, so as to identify an end of each stream, as discussed above. As noted, in some cases, the end is signified by a final packet, such as a TCP packet with the RST flag set or a TCP ACK packet in return to a TCP packet with the FIN flag set. In other cases, the end may be signified by a timer expiring. When the end of a session is detected, the packet series manager 808 disassociates the unique identifier from the session and deletes information about the session from the waypoint information base 900.

[0093] Where the AIPR 800 is provisioned to be a last AIPR before a destination service, the lead packet modifier 806 restores the lead packet to the state the lead packet was in when the source client sent the lead packet, or as the lead packet was modified, such as a result of network

address translation (NAT). Similarly, the subsequent packet modifier 818 restores subsequent packets.

[0094] Similarly, if the destination address of the lead packet is the same as the network address of the AIPR 800, or its network interface 802 over which it receives the lead packets, the lead packet modifier 806 and the subsequent packet modifier 818 restore the packet and subsequent packets.

[0095] As noted, in some protocols, several packets are required to initiate a session, as with the SYN-SYN/ACK-ACK handshake of the TCP. Thus, the downstream controller identifier 816 may wait until a second return packet is received from the destination service before considering a session as having started.

[0096] As noted, some embodiments of the waypoint 800 also manage return packet paths. The lead packet identifier 806 automatically ascertains whether a lead packet was forwarded to the waypoint 800 by an upstream waypoint. If the lead packet includes a session data block, an upstream waypoint forwarded the lead packet. The packet series manager 808 stores information about the upstream waypoint in the waypoint information base 810. A return packet identifier 822 receives return packets from the second network interface 804 and automatically identifies return packets of the session. These return packets may be identified by destination address and port number being equal to the information 932 (Fig. 9) in the waypoint information base corresponding to the session. A return packet modifier modifies the return packets to address them to the upstream waypoint for the session, as identified by the information 918 in the waypoint information base 900.

[0097] It should be noted that statefully monitoring packets is not done by conventional routers. The prior art known to the inventors teaches away from routers statefully monitoring packets. Statefully monitoring packets is, however, one aspect of the disclosed waypoint. This type of monitoring distinguishes embodiments of the present invention from the prior art.

[0098] Fig. 11 contains a flowchart 1100 schematically illustrating some operations performed by the AIPR 800 (Fig. 8) in accordance with illustrative embodiments of the invention. The flowchart 1100 illustrates a packet routing method for directing packets of a session from an originating node toward a destination node in an IP network. At 1102, an intermediate node obtains a lead packet of a plurality of packets in a session. The intermediate node may include a routing device or a switching device that performs a routing function.

[0099] The packets in the session have a unique session identifier. At 1104, a prior node, through which the lead packet traversed, is determined. The prior node has a prior node identifier. At 1106, a return association is formed between the prior node identifier and the session identifier. At 1108, the return association is stored in memory to maintain state information for the session.

[00100] At 1110, the lead packet is modified to identify at least the intermediate node. At 1112, the lead packet is forwarded toward the destination node through an intermediate node electronic output interface to the IP network. The electronic output interface is in communication with the IP network. At 1114, a backward message (e.g., a packet, referred to as a “backward packet”) is received through an electronic input interface of the intermediate node. The backward message is received from a next node. The next node has a next node identifier. The backward message includes the next node identifier and the session identifier. The electronic input interface is in communication with the IP network.

[00101] At 1116, a forward association is formed between the next node identifier and the session identifier. At 1118, the forward association is stored in memory, to maintain state information for the session. At 1120, additional packets of the session are obtained. At 1122, substantially all of the additional packets in the session are forwarded toward the next node, using the stored forward association. The additional packets are forwarded through the electronic output interface of the intermediate node.

[00102] At 1124, a plurality of packets is received in a return session, or a return portion of the session, from the destination. The return session is addressed toward the originating node. At 1126, substantially all the packets in the return session are forwarded toward the prior node, using the stored return association. The packets are forwarded through the electronic output interface.

[00103] As shown at 1200 in Fig. 12, forwarding the lead packet 1112 toward the destination node may include accessing a routing information base having routing information for the next node. As shown at 1202, the intermediate node may have a routing table, and forwarding the lead packet 1112 toward the destination node may include using the routing table to forward the lead packet toward the destination node. As shown at 1204, forwarding the lead packet 1112 toward the destination node may include using the next node identifier to address the lead packet toward the next node.

[00104] The lead packet may be addressed so that a plurality of network devices receive the lead packet after it is forwarded and before the next node receives the lead packet.

[00105] An AIPR 800 and all or a portion of its components 802-824 may be implemented by a processor executing instructions stored in a memory, hardware (such as combinatorial logic, Application Specific Integrated Circuits (ASICs), Field-Programmable Gate Arrays (FPGAs) or other hardware), firmware or combinations thereof.

[00106] While the invention is described through the above-described exemplary embodiments, modifications to, and variations of, the illustrated embodiments may be made without departing from the inventive concepts disclosed herein. Furthermore, disclosed aspects, or portions thereof, may be combined in ways not listed above and/or not explicitly claimed. Accordingly, the invention should not be viewed as being limited to the disclosed embodiments.

[00107] Although aspects of embodiments may be described with reference to flowcharts and/or block diagrams, functions, operations, decisions, etc. of all or a portion of each block, or a combination of blocks, may be combined, separated into separate operations or performed in other orders. All or a portion of each block, or a combination of blocks, may be implemented as computer program instructions (such as software), hardware (such as combinatorial logic, Application Specific Integrated Circuits (ASICs), Field-Programmable Gate Arrays (FPGAs) or other hardware), firmware or combinations thereof. Embodiments may be implemented by a processor executing, or controlled by, instructions stored in a memory. The memory may be random access memory (RAM), read-only memory (ROM), flash memory or any other memory, or combination thereof, suitable for storing control software or other instructions and data. Instructions defining the functions of the present invention may be delivered to a processor in many forms, including, but not limited to, information permanently stored on tangible non-writable storage media (e.g., read-only memory devices within a computer, such as ROM, or devices readable by a computer I/O attachment, such as CD-ROM or DVD disks), information alterably stored on tangible writable storage media (e.g., floppy disks, removable flash memory and hard drives) or information conveyed to a computer through a communication medium, including wired or wireless computer networks. Moreover, while embodiments may be described in connection with various illustrative data structures, systems may be embodied using a variety of data structures.

What is claimed is:

1. A packet routing method for directing packets of a session from an originating node toward a destination node in an IP network, the method comprising:
  - an intermediate node obtaining a lead packet of a plurality of packets in a session, the packets in the session having a unique session identifier;
  - modifying the lead packet to identify at least the intermediate node;
  - forwarding the lead packet toward the destination node through an intermediate node electronic output interface to the IP network, the electronic output interface being in communication with the IP network;
  - receiving, through an electronic input interface of the intermediate node, a backward message from a next node, the next node having a next node identifier, the backward message including the next node identifier and the session identifier, the electronic input interface being in communication with the IP network;
  - forming an association between the next node identifier and the session identifier;
  - storing, in memory, the association to maintain state information for the session;
  - obtaining additional packets of the session; and
  - forwarding, through the electronic output interface of the intermediate node, substantially all of the additional packets in the session toward the next node using the stored association.
2. The packet routing method as defined by claim 1 wherein forwarding the lead packet toward the destination node comprises using the next node identifier to address the lead packet toward the next node.
3. The packet routing method as defined by claim 2 wherein forwarding the lead packet toward the destination node comprises accessing a routing information base having routing information for the next node.
4. The packet routing method as defined by claim 1 further comprising:
  - determining a prior node, if any, through which the lead packet traversed, the prior node having a prior node identifier;

forming a return association between the prior node identifier and the session identifier; and storing, in the memory, the return association to maintain state information for the session.

5. The packet routing method as defined by claim 4 further comprising:
  - receiving a plurality of packets in a return session from the destination node, the return session being addressed toward the originating node; and
  - forwarding, through the electronic output interface, substantially all the packets in the return session, toward the prior node using the stored return association.
6. The packet routing method as defined by claim 1 wherein the intermediate node has a routing table, forwarding the lead packet toward the destination node comprising using the routing table to forward the lead packet toward the destination node.
7. The packet routing method as defined by claim 1 wherein the intermediate node comprises a routing device or a switching device.
8. The packet routing method as defined by claim 1 wherein the lead packet is addressed so that a plurality of network devices receive the lead packet after it is forwarded and before the next node receives the lead packet.
9. An intermediate routing device for directing packets of a session from an originating node toward a destination node in an IP network, the routing device comprising:
  - an electronic input interface for receiving IP packets from the IP network;
  - an electronic output interface for forwarding IP packets on the IP network;
  - state memory for storing session state information;
  - a lead packet identifier operatively coupled with the electronic input interface, the lead packet identifier configured to identify a lead packet of a plurality of packets in a session, the packets in the session having a unique session identifier;

a lead packet modifier operatively coupled with the lead packet identifier, the lead packet modifier being configured to modify the lead packet to identify at least the intermediate routing device;

a router operatively coupled with the electronic output interface, the router being configured to forward the lead packet toward the destination node through the electronic output interface; and

an associator operatively coupled with the electronic input interface, the associator being configured to: a) receive a backward message from a next node having a next node identifier, the backward message including the next node identifier and the session identifier, b) form an association between the next node identifier and the session identifier, and c) store in the state memory the association to maintain state information for the session,

the router also being configured to forward, through the electronic output interface, additional packets in the session toward the next node using the stored association.

10. The intermediate routing device as defined by claim 9 wherein the router is configured to use the next node identifier to address the lead packet toward the next node.

11. The intermediate routing device as defined by claim 10 wherein the router includes a routing information base having routing information for forwarding the lead packet toward the next node.

12. The intermediate routing device as defined by claim 9 wherein the associator is configured to 1) determine a prior node, if any, through which the lead packet traversed, the prior node having a prior node identifier, 2) form a return association between the prior node identifier and the session identifier, and 3) store the return association in the state memory to maintain state information for the session.

13. The intermediate routing device as defined by claim 12 wherein the router is configured to forward, through the electronic output interface, packets in a return session toward the prior node using the stored return association, the packets in the return session being received from the destination node and addressed toward the originating node.

14. The intermediate routing device as defined by claim 9 further comprising a routing table, the router being configured to forward the lead packet toward the destination node comprising using the routing table.

15. A computer program product for use on a computer system for directing packets of a session from an originating node toward a destination node in an IP network, the computer program product comprising a tangible, non-transient computer usable medium having computer readable program code thereon, the computer readable program code comprising:

program code for causing an intermediate node to obtain a lead packet of a plurality of packets in a session, the packets in the session having a unique session identifier;

program code for modifying the lead packet to identify at least the intermediate node;

program code for forwarding the lead packet toward the destination node through an intermediate node electronic output interface to the IP network, the electronic output interface being in communication with the IP network;

program code for receiving, through an electronic input interface of the intermediate node, a backward message from a next node, the next node having a next node identifier, the backward message including the next node identifier and the session identifier, the electronic input interface being in communication with the IP network;

program code for forming an association between the next node identifier and the session identifier;

program code for storing, in memory, the association to maintain state information for the session;

program code for obtaining additional packets of the session; and

program code for forwarding, through the electronic output interface of the intermediate node, substantially all of the additional packets in the session toward the next node using the stored association.

16. The computer program product as defined by claim 15 wherein the program code for forwarding the lead packet toward the destination node comprises program code for using the next node identifier to address the lead packet toward the next node.

17. The computer program product as defined by claim 16 wherein the program code for forwarding the lead packet toward the destination node comprises program code for accessing a routing information base having routing information for the next node.


18. The computer program product as defined by claim 15 further comprising:  
program code for determining a prior node, if any, through which the lead packet traversed, the prior node having a prior node identifier;  
program code for forming a return association between the prior node identifier and the session identifier; and  
program code for storing, in the memory, the return association to maintain state information for the session.

19. The computer program product as defined by claim 18 further comprising:  
program code for receiving a plurality of packets in a return session from the destination node, the return session being addressed toward the originating node; and  
program code for forwarding, through the electronic output interface, substantially all the packets in the return session, toward the prior node using the stored return association.

20. The computer program product as defined by claim 15 wherein the intermediate node has a routing table, the program code for forwarding the lead packet toward the destination node comprising using the routing table to forward the lead packet toward the destination node.


21. The computer program product as defined by claim 15 wherein the intermediate node comprises a routing device or a switching device.

22. The computer program product as defined by claim 15 wherein the lead packet is addressed so that a plurality of network devices receive the lead packet after it is forwarded toward the next node and before the next node receives the lead packet.




(Prior Art)

**Fig. 1**



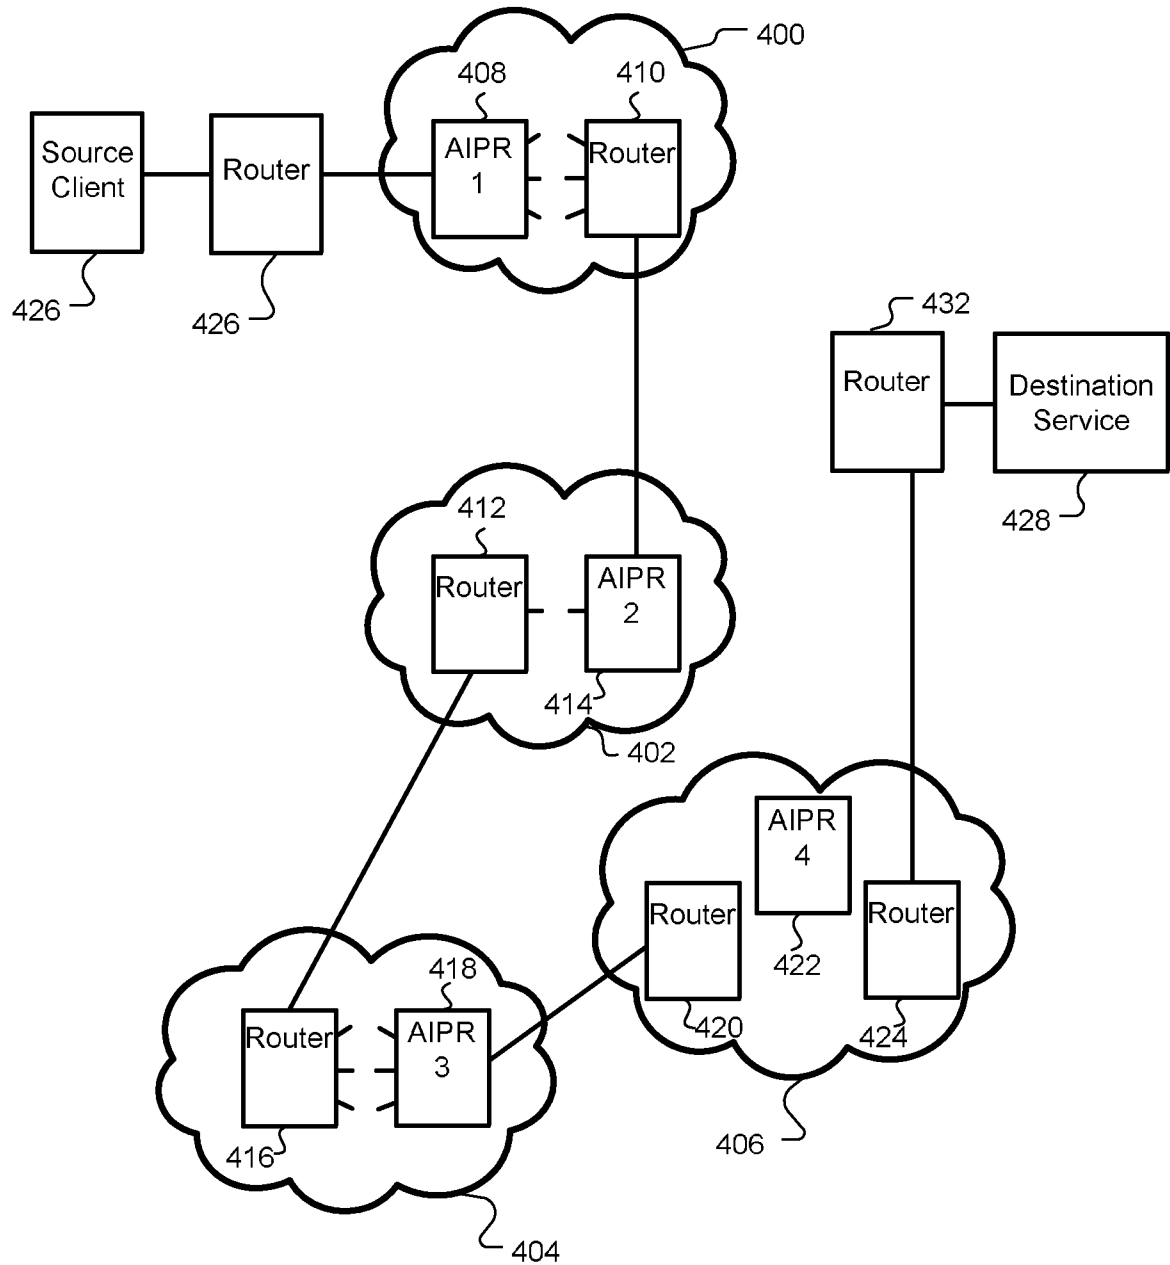
(Prior Art)

**Fig. 2**



(Prior Art)

**Fig. 3**



**Fig. 4**

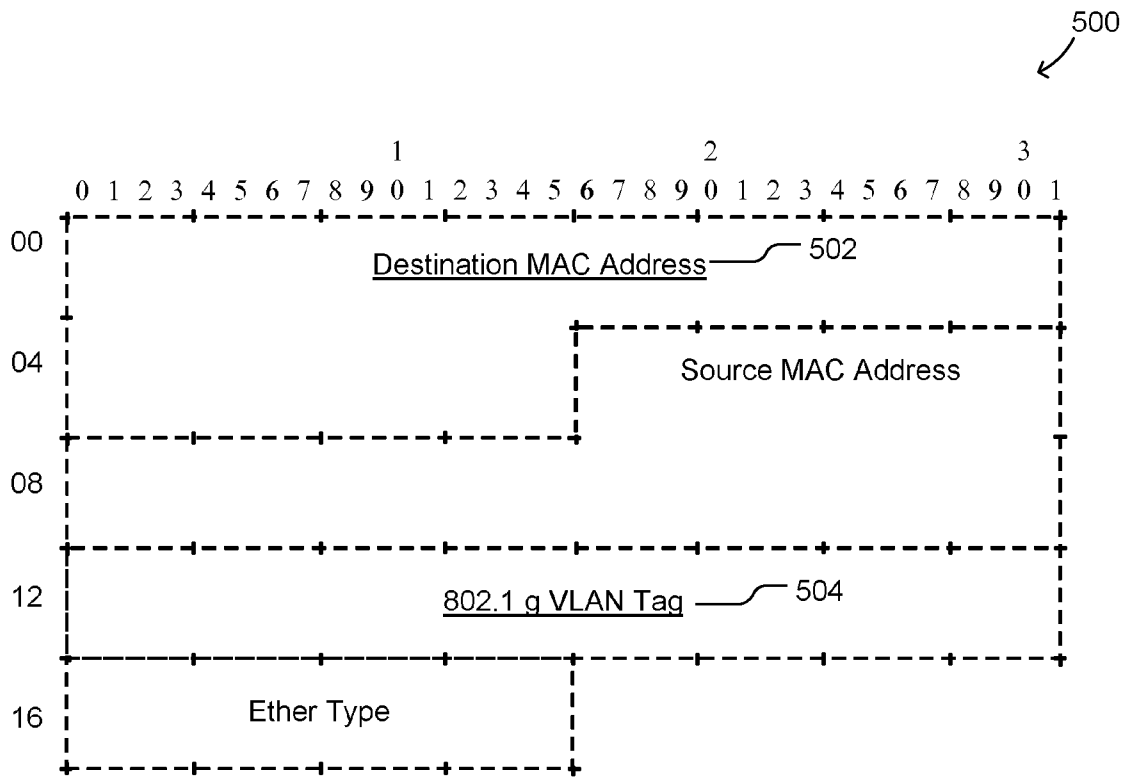


Fig. 5

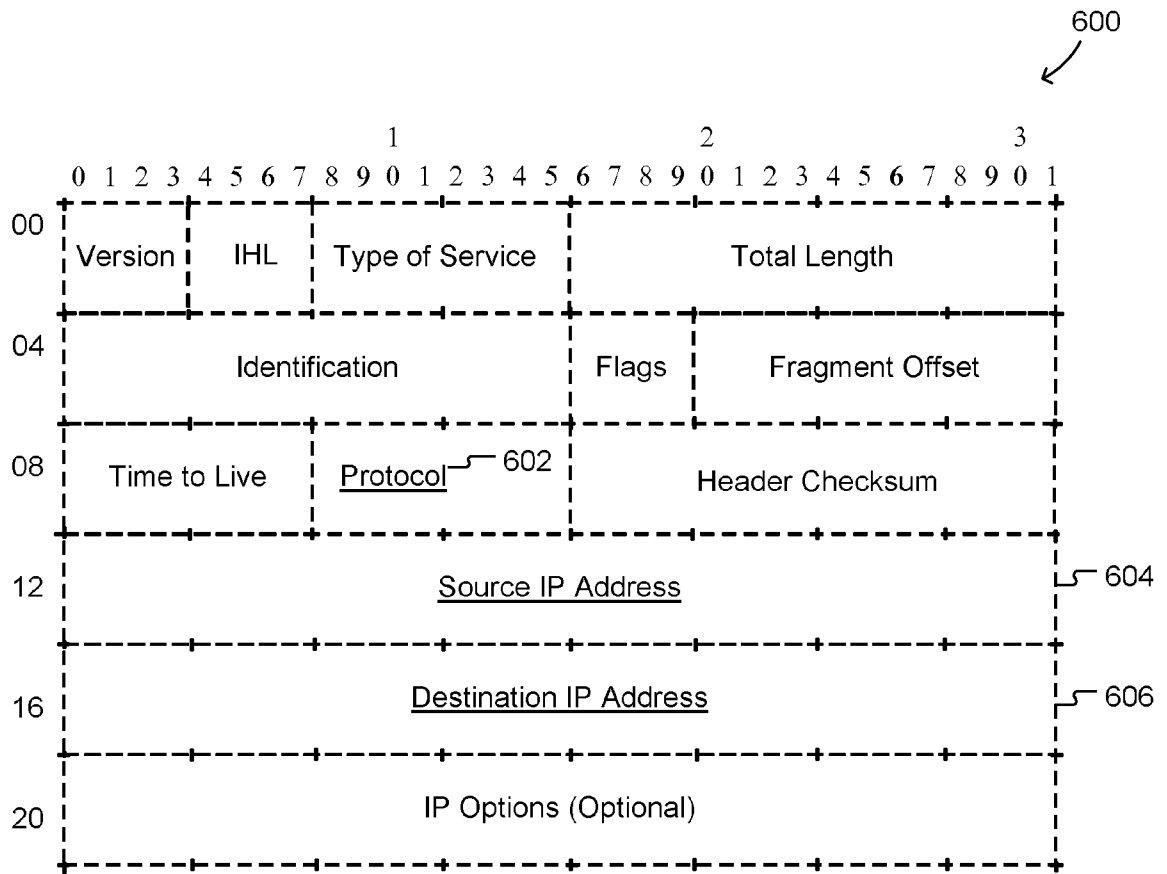


Fig. 6

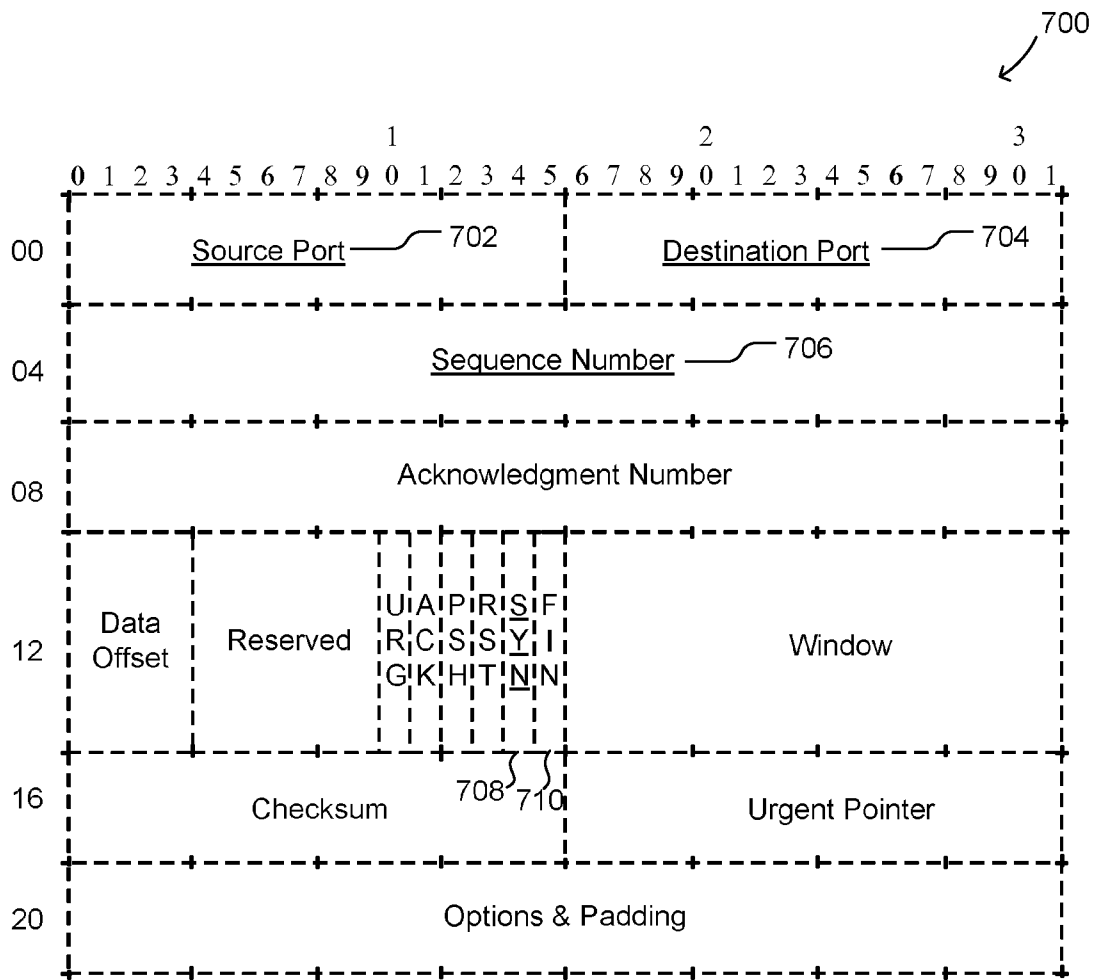


Fig. 7




Fig. 8

Waypoint Information Base


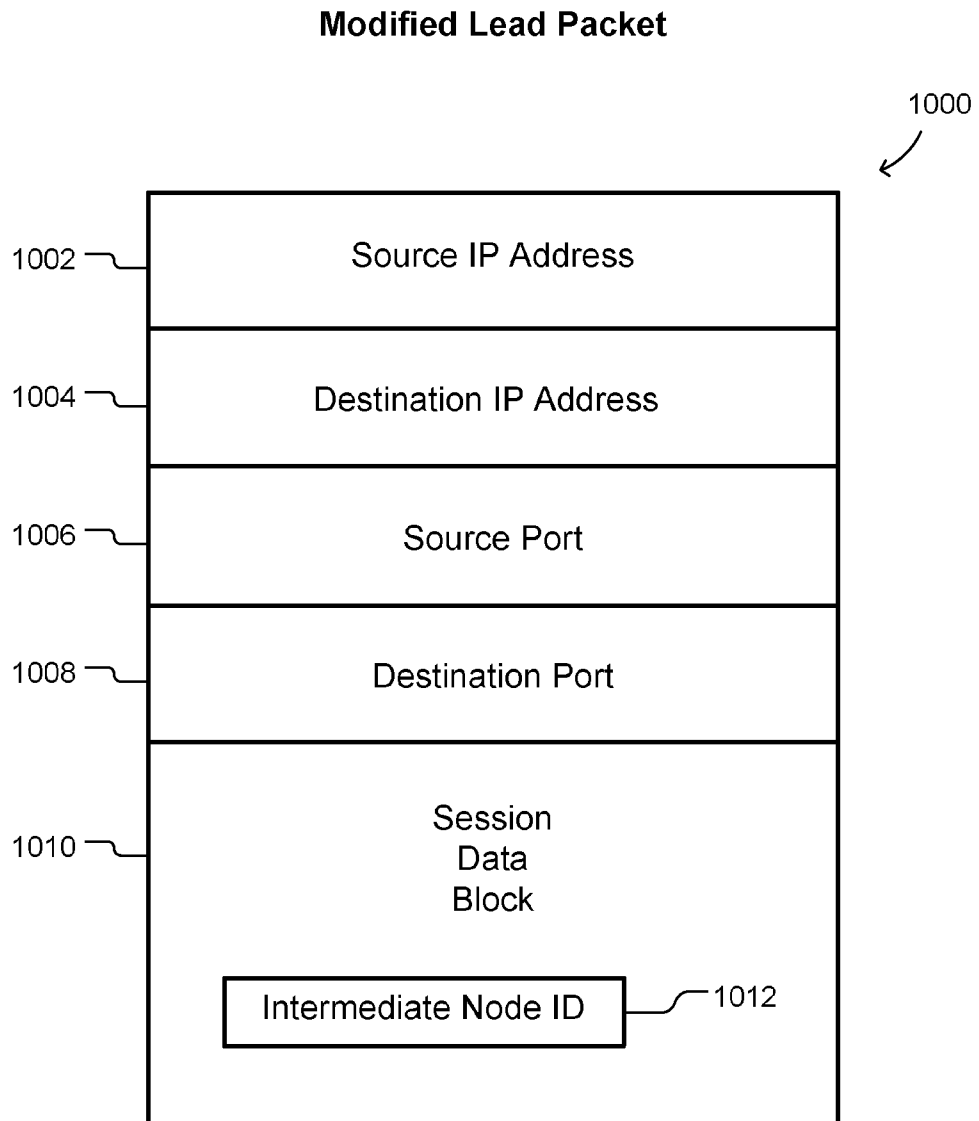
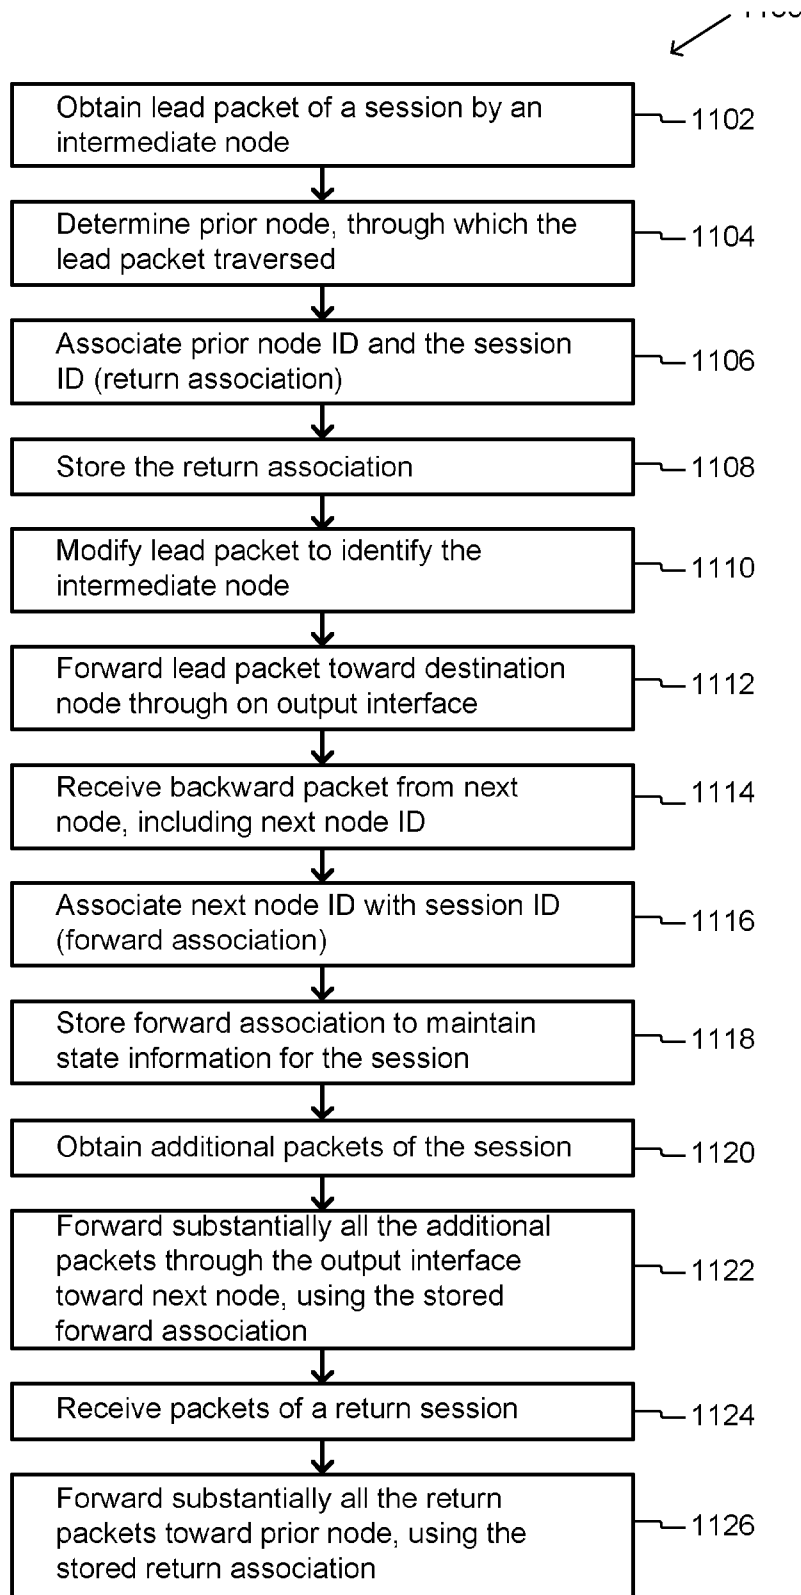


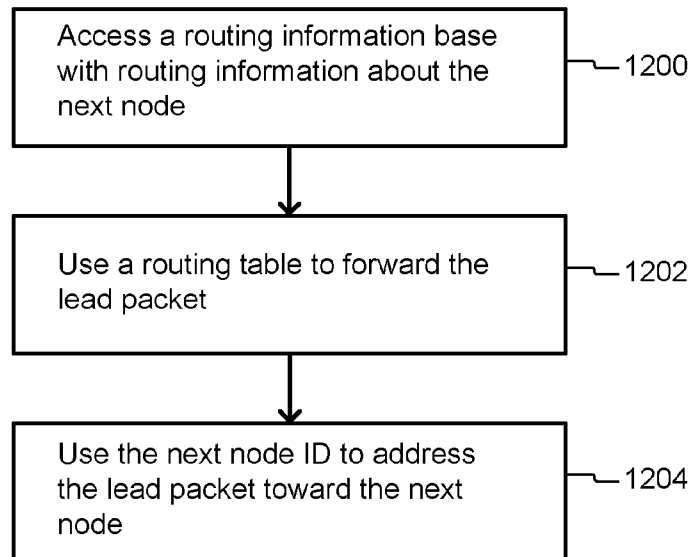
Fig. 9



**Fig. 10**



**Fig. 11**

**Fig. 12**

**INTERNATIONAL SEARCH REPORT**

International application No.

PCT/US2015/044815

<p><b>A. CLASSIFICATION OF SUBJECT MATTER</b>                  IPC (2015.01) H04L 29/02, H04L 29/06</p> <p>According to International Patent Classification (IPC) or to both national classification and IPC</p>								
<p><b>B. FIELDS SEARCHED</b></p> <p>Minimum documentation searched (classification system followed by classification symbols)                  IPC (2015.01) H04L 29/02, H04L 29/06</p> <p>Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched</p> <p>Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)                  Databases consulted: THOMSON INNOVATION, Esp@cenet, Google Patents                  Search terms used: control, intermediate node, first, lead, packet, modify, insert, add, next, identifier</p>								
<p><b>C. DOCUMENTS CONSIDERED TO BE RELEVANT</b></p> <table border="1" style="width:100%; border-collapse: collapse;"> <thead> <tr> <th style="width:10%;">Category*</th> <th style="width:70%;">Citation of document, with indication, where appropriate, of the relevant passages</th> <th style="width:20%;">Relevant to claim No.</th> </tr> </thead> <tbody> <tr> <td style="vertical-align: top;">A</td> <td>                     EP 1313267 B1 MOTOROLA, INC.                      27 Dec 2006 (2006/12/27)                      The whole document                 </td> <td style="vertical-align: top;">1-22</td> </tr> </tbody> </table>			Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.	A	EP 1313267 B1 MOTOROLA, INC. 27 Dec 2006 (2006/12/27) The whole document	1-22
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.						
A	EP 1313267 B1 MOTOROLA, INC. 27 Dec 2006 (2006/12/27) The whole document	1-22						
<p><input type="checkbox"/> Further documents are listed in the continuation of Box C.      <input checked="" type="checkbox"/> See patent family annex.</p>								
<p>* Special categories of cited documents:</p> <table style="width:100%;"> <tr> <td style="width:50%;">                     "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                 </td> <td style="width:50%;">                     "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                      "&amp;" document member of the same patent family                 </td> </tr> </table>			"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				
"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 03 Dec 2015		Date of mailing of the international search report 06 Dec 2015						
Name and mailing address of the ISA: Israel Patent Office Technology Park, Bldg.5, Malcha, Jerusalem, 9695101, Israel Facsimile No. 972-2-5651616		Authorized officer GORBUNOVA Yelena Telephone No. 972-2-5651669						

**INTERNATIONAL SEARCH REPORT**  
Information on patent family members

International application No.  
PCT/US2015/044815

Patent document cited search report	Publication date	Patent family member(s)	Publication Date
EP 1313267 B1	27 Dec 2006	EP 1313267 A1	21 May 2003
		EP 1313267 B1	27 Dec 2006
		AT 349837 T	15 Jan 2007
		AU 2002337180 A1	10 Jun 2003
		DE 60125554 D1	08 Feb 2007
		DE 60125554 T2	04 Oct 2007
		WO 03045018 A1	30 May 2003