

PATENT ASSIGNMENT COVER SHEET

Electronic Version v1.1
 Stylesheet Version v1.2

EPAS ID: PAT7351428

SUBMISSION TYPE:	CORRECTIVE ASSIGNMENT
NATURE OF CONVEYANCE:	Corrective Assignment to correct the EXECUTION DATE OF ROBERT PENFIELD previously recorded on Reel 059824 Frame 0854. Assignor(s) hereby confirms the EXECUTION DATE FOR ROBERT PENFIELD IS 04/26/2022.

CONVEYING PARTY DATA

Name	Execution Date
PATRICK TIMMONS	04/26/2022
MICHAEL BAJ	04/27/2022
HADRIEL S. KAPLAN	04/27/2022
ROBERT PENFIELD	04/26/2022
PATRICK J. MELAMPY	04/27/2022

RECEIVING PARTY DATA

Name:	128 TECHNOLOGY, INC.
Street Address:	200 SUMMIT DRIVE
Internal Address:	SUITE 600
City:	BURLINGTON
State/Country:	MASSACHUSETTS
Postal Code:	01803

PROPERTY NUMBERS Total: 1

Property Type	Number
Patent Number:	11165863

CORRESPONDENCE DATA

Fax Number: (651)735-1102
Correspondence will be sent to the e-mail address first; if that is unsuccessful, it will be sent using a fax number, if provided; if that is unsuccessful, it will be sent via US Mail.

Phone: 6517351100
Email: pairdocketing@ssiplaw.com
Correspondent Name: SHUMAKER & SIEFFERT P.A.
Address Line 1: 1625 RADIO DRIVE , SUITE 100
Address Line 4: WOODBURY, MINNESOTA 55125

ATTORNEY DOCKET NUMBER:	2014-331US01
NAME OF SUBMITTER:	JULIA A. EHRREICH
SIGNATURE:	/Julia A. Ehrreich/

DATE SIGNED:

05/26/2022

Total Attachments: 41

source=Corrected_Assignment#page1.tif
source=Corrected_Assignment#page2.tif
source=Corrected_Assignment#page3.tif
source=Corrected_Assignment#page4.tif
source=Corrected_Assignment#page5.tif
source=Corrected_Assignment#page6.tif
source=Corrected_Assignment#page7.tif
source=Corrected_Assignment#page8.tif
source=Corrected_Assignment#page9.tif
source=Corrected_Assignment#page10.tif
source=Corrected_Assignment#page11.tif
source=Corrected_Assignment#page12.tif
source=Corrected_Assignment#page13.tif
source=Corrected_Assignment#page14.tif
source=Corrected_Assignment#page15.tif
source=Corrected_Assignment#page16.tif
source=Corrected_Assignment#page17.tif
source=Corrected_Assignment#page18.tif
source=Corrected_Assignment#page19.tif
source=Corrected_Assignment#page20.tif
source=Corrected_Assignment#page21.tif
source=Corrected_Assignment#page22.tif
source=Corrected_Assignment#page23.tif
source=Corrected_Assignment#page24.tif
source=Corrected_Assignment#page25.tif
source=Corrected_Assignment#page26.tif
source=Corrected_Assignment#page27.tif
source=Corrected_Assignment#page28.tif
source=Corrected_Assignment#page29.tif
source=Corrected_Assignment#page30.tif
source=Corrected_Assignment#page31.tif
source=Corrected_Assignment#page32.tif
source=Corrected_Assignment#page33.tif
source=Corrected_Assignment#page34.tif
source=Corrected_Assignment#page35.tif
source=Corrected_Assignment#page36.tif
source=Corrected_Assignment#page37.tif
source=Corrected_Assignment#page38.tif
source=Corrected_Assignment#page39.tif
source=Corrected_Assignment#page40.tif
source=Corrected_Assignment#page41.tif

PATENT ASSIGNMENT COVER SHEET

Electronic Version v1.1
Stylesheet Version v1.2

EPAS ID: PAT7306746

SUBMISSION TYPE:	NEW ASSIGNMENT
NATURE OF CONVEYANCE:	ASSIGNMENT
CONVEYING PARTY DATA	
Name	Execution Date
PATRICK TIMMONS	04/26/2022
MICHAEL BAJ	04/27/2022
HADRIEL S. KAPLAN	04/27/2022
ROBERT PENFIELD	04/26/2022 — 04/27/2022 —
PATRICK J. MELAMPY	04/27/2022
RECEIVING PARTY DATA	
Name:	128 TECHNOLOGY, INC.
Street Address:	200 SUMMIT DRIVE SUITE 600
City:	BURLINGTON
State/Country:	MASSACHUSETTS
Postal Code:	01803
PROPERTY NUMBERS Total: 1	
Property Type	Number
Patent Number:	11165863
CORRESPONDENCE DATA	
Fax Number:	(651)735-1102
<i>Correspondence will be sent to the e-mail address first; if that is unsuccessful, it will be sent using a fax number, if provided; if that is unsuccessful, it will be sent via US Mail.</i>	
Phone:	6517351100
Email:	pairdocketing@ssiplaw.com
Correspondent Name:	SHUMAKER & SIEFFERT P.A.
Address Line 1:	1625 RADIO DRIVE SUITE 100
Address Line 4:	WOODBURY, MINNESOTA 55125
ATTORNEY DOCKET NUMBER:	2014-331US01/JNA0161-US
NAME OF SUBMITTER:	ADAM J. REVOIR
SIGNATURE:	/Adam J. Revoir/
DATE SIGNED:	04/29/2022
Total Attachments: 38	
source=Assignment-Confirmatory_Signed#page1.tif	

CONFIRMATORY ASSIGNMENT

For good and valuable consideration, the receipt of which is hereby acknowledged, the person(s) named below (referred to as "INVENTOR" whether singular or plural) has sold, assigned, and transferred and does hereby confirm the sale, assignment, and transfer to **128 Technology, Inc.**, having a place of business at **200 Summit Drive Suite 600, Burlington, MA 01803, United States of America** ("ASSIGNEE"), for itself and its successors, transferees, and assignees, the following:

1. The entire worldwide right, title, and interest in all inventions and improvements ("SUBJECT MATTER") that are disclosed in the following provisional application filed under 35 U.S.C. § 111(b), non-provisional application filed under 35 U.S.C. § 111(a), international application filed according to the Patent Cooperation Treaty (PCT), or U.S. national phase application filed under 35 U.S.C. § 371 ("APPLICATION"):

Application No. 16/050,722, entitled "NETWORK NEIGHBORHOODS FOR ESTABLISHING COMMUNICATION RELATIONSHIPS BETWEEN COMMUNICATION INTERFACES IN AN ADMINISTRATIVE DOMAIN" filed on July 31, 2018 which claims priority to Application No. 62/541,357 filed on August 04, 2017. (I hereby authorize the Assignee and its representative to hereafter add herein such application number(s) and/or filing date(s) when known.)

2. The entire worldwide right, title, and interest in and to:
(a) the APPLICATION; (b) all applications claiming priority from the APPLICATION; (c) all provisional, utility, divisional, continuation, substitute, renewal, reissue, and other applications related thereto which have been or may be filed in the United States or elsewhere in the world; (d) all patents (including reissues and re-examinations) which may be granted on the applications set forth in (a), (b), and (c) above; and (e) all right of priority in the APPLICATION and in any underlying provisional or foreign application, together with all rights to recover damages for infringement of provisional rights.

3. The entire worldwide right, title, and interest in and to (including all claims of):

U.S. Patent No. 11,165,863 issued November 02, 2021, which is included as an Appendix to this Assignment.

INVENTOR agrees that ASSIGNEE may apply for and receive patents for SUBJECT MATTER in ASSIGNEE's own name.

INVENTOR agrees to do the following, when requested, and without further consideration, in order to carry out the intent of this Assignment: (1) execute all oaths, assignments, powers of attorney, applications, and other papers necessary or desirable to fully secure to ASSIGNEE the rights, titles and interests herein conveyed; (2) communicate to ASSIGNEE all known facts relating to the SUBJECT MATTER; and (3) generally do all lawful acts that ASSIGNEE shall consider desirable for securing, maintaining, and enforcing worldwide patent protection relating to the SUBJECT MATTER and for vesting in ASSIGNEE the rights, titles, and interests herein conveyed. INVENTOR further agrees to provide any successor,

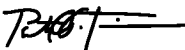
assign, or legal representative of ASSIGNEE with the benefits and assistance provided to ASSIGNEE hereunder.

INVENTOR represents that INVENTOR has the rights, titles, and interests to convey as set forth herein, and covenants with ASSIGNEE that the INVENTOR has not made and will not hereafter make any assignment, grant, mortgage, license, or other agreement affecting the rights, titles, and interests herein conveyed.

INVENTOR grants the attorney of record the power to insert on this Assignment any further identification that may be necessary or desirable in order to comply with the rules of the United States Patent and Trademark Office for recordation of this document.

This Assignment may be executed in one or more counterparts, each of which shall be deemed an original and all of which may be taken together as one and the same Assignment.

Name and Signature



Patrick Timmons

Date of Signature

Apr 26, 2022

Name and Signature



Michael Baj

Date of Signature

Apr 27, 2022

Name and Signature




Hadriel S. Kaplan

Date of Signature

Apr 27, 2022

Name and Signature



Robert Penfield

Date of Signature

Apr 26, 2022

Name and Signature



Patrick J. McLampy

Date of Signature

Apr 27, 2022

Title: NETWORK NEIGHBORHOODS FOR ESTABLISHING COMMUNICATION RELATIONSHIPS
BETWEEN COMMUNICATION INTERFACES IN AN ADMINISTRATIVE DOMAIN

Date Filed: July 31, 2018

Application No.: 16/050,722

APPENDIX



US011165863B1

(12) **United States Patent**
Timmons et al.

(10) **Patent No.:** **US 11,165,863 B1**
(45) **Date of Patent:** **Nov. 2, 2021**

(54) **NETWORK NEIGHBORHOODS FOR ESTABLISHING COMMUNICATION RELATIONSHIPS BETWEEN COMMUNICATION INTERFACES IN AN ADMINISTRATIVE DOMAIN**

(71) Applicant: **128 Technology, Inc.**, Burlington, MA (US)

(72) Inventors: **Patrick Timmons**, Newton, MA (US); **Michael Baj**, Somerville, MA (US); **Hadriel S. Kaplan**, Nashua, NH (US); **Robert Penfield**, Concord, MA (US); **Patrick J. McLampy**, Dunstable, MA (US)

(73) Assignee: **128 Technology, Inc.**, Burlington, MA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 67 days.

(21) Appl. No.: **16/050,722**

(22) Filed: **Jul. 31, 2018**

Related U.S. Application Data

(60) Provisional application No. 62/541,357, filed on Aug. 4, 2017.

(51) **Int. Cl.**
H04L 29/08 (2006.01)
H04L 12/715 (2013.01)
(Continued)

(52) **U.S. Cl.**
CPC **H04L 67/1093** (2013.01); **H04L 41/08** (2013.01); **H04L 41/0803** (2013.01);
(Continued)

(58) **Field of Classification Search**
CPC H04L 67/1093; H04L 41/0803; H04L 67/1042; H04L 45/02; H04L 41/08;
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,699,350 A * 12/1997 Kraslavsky H04L 29/06 370/254
5,917,820 A * 6/1999 Rekhter H04L 45/7457 370/392

(Continued)

FOREIGN PATENT DOCUMENTS

CN 101552703 A 10/2009
CN 101646220 A 2/2010

(Continued)

OTHER PUBLICATIONS

Berners-Lee et al., "Uniform Resource Identifier (URI): Generic Syntax", Network Working Group, Request for Comments 3986, The Internet Society, 61 pages, Jan. 2005.

(Continued)


Primary Examiner — Ramy M Osman

(74) *Attorney, Agent, or Firm* — Shumaker & Sieffert, P.A.

(57) **ABSTRACT**

A method and apparatus form and/or define a network topology in a Layer 3 network with a plurality of nodes, where each node has at least one interface. To that end, the method defines a plurality of neighborhoods, and assigns at least one interface of each node to at least one of the neighborhoods. The method also assigns a communication role to each interface so that each communication role is effective relative to one of the plurality of neighborhoods. The method then enables communication between the interfaces of the plurality of nodes as a function of the neighborhoods and the communication roles.

20 Claims, 17 Drawing Sheets



(51)	Int. Cl.			8,989,020 B2	3/2015	So	
	H04L 12/775	(2013.01)		9,059,920 B2	6/2015	Ravindran et al.	
	H04L 12/24	(2006.01)		9,160,652 B2	10/2015	Taillon et al.	
	H04L 12/751	(2013.01)		9,185,186 B2	11/2015	Hong et al.	
(52)	U.S. Cl.			9,240,953 B2	1/2016	Carlstrom	
	CPC	H04L 41/0893 (2013.01); H04L 41/12		9,276,864 B1	3/2016	Vincent	
		(2013.01); H04L 45/02 (2013.01); H04L 45/04		9,614,756 B2	4/2017	Joshi	
		(2013.01); H04L 45/58 (2013.01); H04L 45/64		9,699,069 B1	7/2017	Kielhofner et al.	
		(2013.01); H04L 67/1042 (2013.01); H04L		9,729,439 B2	8/2017	MeLampy et al.	
		67/1044 (2013.01); H04L 67/1059 (2013.01)		9,871,748 B2	1/2018	Gosselin et al.	
(58)	Field of Classification Search			10,362,121 B1	7/2019	Crausaz	
	CPC	H04L 45/04; H04L 45/58; H04L 41/12;		2001/0030649 A1	10/2001	Mamiya et al.	
		H04L 45/64; H04L 41/0893; H04L		2002/0021689 A1	2/2002	Robbins et al.	
		67/1059; H04L 67/1044		2002/0044553 A1	4/2002	Chakravorty	
	See application file for complete search history.			2002/0075883 A1	6/2002	Dell et al.	
				2002/0176363 A1	11/2002	Durinovic-Johri et al.	
				2003/0041170 A1*	2/2003	Suzuki	H04L 69/16
							709/238
(56)	References Cited			2003/0051048 A1*	3/2003	Watson	H04L 45/00
	U.S. PATENT DOCUMENTS						709/238
				2003/0198189 A1	10/2003	Roberts et al.	
				2003/0214938 A1	11/2003	Jindal et al.	
				2003/0236870 A1*	12/2003	Boivie	H04L 45/64
							709/223
	6,310,883 B1	10/2001	Mann et al.	2004/0088542 A1	5/2004	Daude et al.	
	6,515,963 B1	2/2003	Bechtolsheim et al.	2004/0190456 A1*	9/2004	Haas	H04L 45/04
	6,563,824 B1	5/2003	Bhatia et al.				370/238
	6,584,071 B1	6/2003	Kodialam et al.	2004/0221051 A1*	11/2004	Liong	H04L 41/0893
	6,721,334 B1	4/2004	Ketcham				709/230
	6,738,387 B1	5/2004	Lin et al.	2004/0264481 A1	12/2004	Darling et al.	
	6,778,531 B1	8/2004	Kodialam et al.	2005/0013300 A1	1/2005	Akahane et al.	
	6,785,273 B1*	8/2004	Verma	2005/0036616 A1	2/2005	Huang et al.	
			H04L 45/00	2005/0063307 A1	3/2005	Samuels et al.	
			370/351	2005/0182932 A1	8/2005	Wheeler	
	6,798,743 B1	9/2004	Ma et al.	2005/0238022 A1	10/2005	Panigrahy	
	7,020,087 B2	3/2006	Steinberg et al.	2005/0249206 A1	11/2005	Wybenga et al.	
	7,020,143 B2	3/2006	Zdan	2006/0036762 A1*	2/2006	Vadlakonda	H04L 45/02
	7,035,214 B1	4/2006	Seddigh et al.				709/238
	7,068,600 B2	6/2006	Cain	2006/0176894 A1	8/2006	Oh et al.	
	7,106,739 B2	9/2006	Beier	2007/0171825 A1	7/2007	Roberts et al.	
	7,120,118 B2	10/2006	Rajagopal et al.	2007/0171826 A1	7/2007	Roberts et al.	
	7,154,902 B1	12/2006	Sikdar	2008/0049764 A1*	2/2008	Solomon	H04L 63/162
	7,200,114 B1*	4/2007	Tse-Au				370/401
			H04L 41/0816	2008/0214175 A1	9/2008	Papadoglou et al.	
			370/230.1	2008/0259938 A1	10/2008	Keene et al.	
	7,218,632 B1	5/2007	Bechtolsheim et al.	2009/0007021 A1	1/2009	Hayton	
	7,248,579 B1*	7/2007	Friedman	2009/0059958 A1	3/2009	Nakata	
			H04L 45/02	2009/0086651 A1	4/2009	Luft et al.	
			370/389	2010/0125898 A1	5/2010	Dubuc et al.	
	7,315,541 B1	1/2008	Housel et al.	2010/0191968 A1	7/2010	Patil et al.	
	7,373,660 B1	5/2008	Guichard et al.	2010/0287405 A1*	11/2010	Soon	H04L 43/0817
	7,466,703 B1	12/2008	Arunachalam et al.				714/4.1
	7,525,949 B1*	4/2009	Rampal	2011/0173324 A1	7/2011	Wang et al.	
			H04L 45/02	2012/0144061 A1	6/2012	Song	
			370/352	2012/0236860 A1	9/2012	Kompella et al.	
	7,533,183 B1*	5/2009	Crump	2012/0287932 A1*	11/2012	Haddad	H04L 45/04
			H04L 45/04				370/392
			370/395.31	2013/0227166 A1	8/2013	Ravindran et al.	
	7,536,720 B2	5/2009	Burdett et al.	2013/0297824 A1	11/2013	Lan et al.	
	7,634,805 B2	12/2009	Aroya	2014/0040488 A1	2/2014	Small et al.	
	7,706,411 B2	4/2010	Wakumoto et al.	2014/0108584 A1*	4/2014	Lu	H04L 47/76
	7,730,301 B2	6/2010	Correll et al.				709/208
	7,756,924 B2*	7/2010	Kaler	2014/0161137 A1*	6/2014	Dasari	H04L 41/0803
			H04L 67/104				370/437
			709/204	2014/0269422 A1	9/2014	Filsfils et al.	
	7,773,611 B2	8/2010	Booth, III et al.	2015/0188814 A1	7/2015	Jain et al.	
	7,818,450 B2	10/2010	Chen et al.	2015/0229618 A1	8/2015	Wan et al.	
	7,872,973 B2	1/2011	Sterne et al.	2015/0381324 A1	12/2015	Mirsky et al.	
	8,068,417 B1	11/2011	Roberts	2016/0021032 A1*	1/2016	Maier	H04L 45/42
	8,094,560 B2	1/2012	Bagepalli et al.				370/401
	8,108,882 B1*	1/2012	Haynes, Jr.	2016/0094444 A1	3/2016	MeLampy et al.	
			G06F 13/102	2016/0164780 A1	6/2016	Timmons et al.	
			709/230	2017/0063625 A1*	3/2017	Philip	H04L 41/0803
	8,139,479 B1	3/2012	Raszuk	2018/0026885 A1	1/2018	Jeuk et al.	
	RE44,119 E	4/2013	Wang et al.				
	8,437,248 B2	5/2013	Li et al.				
	8,527,641 B2	9/2013	Degaonkar et al.				
	8,570,893 B2	10/2013	Guo et al.				
	8,584,199 B1	11/2013	Chen et al.				
	8,619,775 B2	12/2013	Amir et al.				
	8,634,428 B2	1/2014	Le Pennec et al.				
	8,660,126 B2	2/2014	Crambert et al.				
	8,804,489 B2	8/2014	Lu et al.				
	8,942,085 B1	1/2015	Pani et al.				
	8,958,298 B2*	2/2015	Zhang				
			H04L 45/74				
			370/235				

(56)

References Cited

U.S. PATENT DOCUMENTS

2018/0063608 A1 3/2018 Prakash et al.
 2018/0254947 A1* 9/2018 Baj H04L 41/0893

FOREIGN PATENT DOCUMENTS

CN	101068242	B	4/2010
CN	102158371	A	8/2011
CN	101207604	B	3/2012
CN	101640629	B	8/2012
CN	102739507	A	10/2012
CN	102769679	B	6/2015
CN	103179192	B	11/2015
CN	105245469	A	1/2016
CN	103188260	B	3/2017
EP	1313267	B1	12/2006
KR	10-2011-0062994	A	6/2011
WO	WO-03058868	A2	7/2003
WO	WO-2007084707	A2	7/2007
WO	WO-2007084755	A2	7/2007
WO	WO-2008043230	A1	4/2008
WO	WO-2015131537	A1	9/2015
WO	WO-2016007052	A1	1/2016

OTHER PUBLICATIONS

Bjorklund, "YANG—A Data Modeling Language for the Network Configuration Protocol (NETCONF)", Internet Engineering Task Force (IETF), Request for Comments 6020, ISSN: 2070-1721, 173 pages, Oct. 2010.

Caida, "Observing routing asymmetry in Internet traffic", (www.caida.org/research/traffic-analysis/asymmetry/1/), 7 pages, dated Jul. 17, 2013.

Chiosi et al., "Network Functions Virtualisation—Introductory White Paper", Issue 1, at the "SDN and OpenFlow World Congress", Darmstadt—Germany, (http://portal.etsi.org/nfv/nfv_white_paper), 16 pages, dated Oct. 22, 2012.

Cisco Systems, "Parallel Express Forwarding on the Cisco 10000 Series", (White Paper) Cisco Systems, printed Jun. 17, 2015, 4 pages.

Data Plane Development Kit, "Programmer's Guide, Release 16.04.0", 216 pages, Apr. 12, 2016.

Davis, "Layer 3 Switches Explained", Happy Router, 6 pages, dated Aug. 30, 2007.

Filsfils et al., "Segment Routing Architecture", Network Working Group, Draft, 28 pages, Oct. 21, 2013.

Hansson et al., "A Unified Approach to Constrained Mapping and Routing on Network-on-Chip Architectures", CODES+ISSS '05 Proceedings of the 3rd IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis, 6 pages, Sep. 19-21, 2005.

Herbert, "xps: Transmit Packet Steering", Eklektix, Inc., 11 pages, Oct. 26, 2010.

Iana, "Transmission Control Protocol (TCP) Parameters", (www.iana.org/assignments/tcp-parameters/tcp-parameters.xhtml), 5 pages, dated Sep. 22, 2014.

International Searching Authority, International Search Report—International Application No. PCT/2015/044815, dated Dec. 6, 2015, together with the Written Opinion of the International Searching Authority, 8 pages.

International Searching Authority, International Search Report—International Application No. PCT/2015/060840, dated Mar. 8, 2016, together with the Written Opinion of the International Searching Authority, 13 pages.

International Searching Authority, International Search Report—International Application No. PCT/US2016/013416, dated Jun. 8, 2016, together with the Written Opinion of the International Searching Authority, 12 pages.

International Searching Authority, International Search Report—International Application No. PCT/US2016/026938, dated Jul. 28, 2016, together with the Written Opinion of the International Searching Authority, 9 pages.

International Searching Authority, International Search Report—International Application No. PCT/US2016/062033, dated Mar. 7, 2017, together with the Written Opinion of the International Searching Authority, 10 pages.

International Searching Authority, International Search Report—International Application No. PCT/US2019/015404, dated Apr. 15, 2019, together with the Written Opinion of the International Searching Authority, 17 pages.

Iyer, "Load Balancing and Parallelism for the Internet: A Dissertation submitted to the Dept. of Computer Science and the Committee on Graduate Studies of Stanford University in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy, retrieved on the internet at http://yuba.stanford.edu/~sundaes/Dissertation/sundar_thesis.pdf, 436 pages, Jul. 2008.

Katz et al., "Bidirectional Forwarding Detection (BFD)", Internet Engineering Task Force (IETF), Request for Comments 5880, ISSN: 2070-1721, Juniper Networks, 49 pages, Jun. 2010.

Klement, "1.2 Overview of a TCP communications session", RPG IV Socket Tutorial (<http://www.scottklement.com/rpg/socketut/overview>), 2 pages, 2001.

Microsoft, "Introduction to Receive Side Scaling", Developer Resources, [https://msdn.microsoft.com/en-us/library/windows/hardware/ff556942\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/hardware/ff556942(v=vs.85).aspx), 3 pages, Apr. 2014.

Microsoft, "Non-RSS Receive Processing", Developer Resources, [https://msdn.microsoft.com/en-us/library/windows/hardware/ff568798\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/hardware/ff568798(v=vs.85).aspx), 2 pages, Jan. 2015.

Microsoft, "RSS with a Single Hardware Receive Queue", Developer Resources, [https://msdn.microsoft.com/en-us/library/windows/hardware/ff570727\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/hardware/ff570727(v=vs.85).aspx), 2 pages, Jan. 2015.

Microsoft, "RSS with Hardware Queuing", Developer Resources, [https://msdn.microsoft.com/en-us/library/windows/hardware/ff570728\(v=vs.85\).aspx](https://msdn.microsoft.com/en-us/library/windows/hardware/ff570728(v=vs.85).aspx), 2 pages, Jan. 2015.

PC Magazine Encyclopedia, "Definition of: TCP/IP abc's", PC Magazine Encyclopedia (www.pcmag.com/encyclopedia/term/52615), 5 pages, 2005.

Previdi, et al., "IPv6 Segment Routing Header (SRH)", Network working Group, Draft, 24 pages, Jul. 3, 2014.

Roberts, "The Next Generation of IP—Flow Routing", SSGRR 2003S International Conference, L'Aquila Italy, 11 pages, Jul. 29, 2003.

Rouse, "What is routing table?", Posted by Margaret Rouse (<http://searchnetworking.techtarget.com/definition/routing-table>), 5 pages, Apr. 2007.

Shang et al., "Making Better Use of All Those TCP ACK Packets", Computer Science Department, Worcester Polytechnic Institute, 10 pages, 2005.

Shaw, "Multi-queue network interfaces with SMP on Linux", Greenhost, <https://greenhost.net/2013/04/10/multi-queue-network-interfaces-with-smp-on-linux/>, 5 pages, Apr. 10, 2013.

Sollins et al., "Functional Requirements for Uniform Resource Names", Network Working Group, Request for Comments 1737, 7 pages, Dec. 1994.

Srinivasan et al., "A Technique for Low Energy Mapping and Routing in Network-on-Chip Architectures", ISLPED '05 Proceedings of the 2005 International Symposium on Low Power Electronics and Design, 6 pages, Aug. 8-10, 2005.

Wikipedia, "LAN switching", 5 pages, dated Jun. 12, 2013.

Wikipedia, "Management information base", 6 pages, dated Jul. 15, 2013.

Wikipedia, "Reverse path forwarding", 3 pages, dated Jul. 31, 2013.

Wikipedia, "Equal-cost multi-path routing", 1 page, dated Sep. 12, 2013.

Wikipedia, "Software-defined networking", 6 pages, dated Sep. 16, 2013.

Wikipedia, "Transmission Control Protocol", 18 pages, dated Sep. 16, 2013.

Wikipedia, "Network socket", 4 pages, dated Sep. 19, 2013.

Wikipedia, "Router (computing)", 8 pages, dated Sep. 23, 2013.

Wikipedia, "Network address translation", 11 pages, dated Sep. 24, 2013.

Wikipedia, "Open vSwitch", 2 pages, dated Nov. 24, 2013.

Wikipedia, "Active queue management", https://en.wikipedia.org/wiki/Active_queue_management, 2 pages, Apr. 22, 2015.

PATENT**REEL: 060202 FRAME: 0009**

(56)

References Cited

OTHER PUBLICATIONS

Wikipedia, "Network interface controller", https://en.wikipedia.org/wiki/Network_interface_controller, 5 pages, May 19, 2015.

"STEP," Solution Note, White paper, Retrieved from: https://www.128technology.com/wp-content/uploads/2017/09/whitepaper_step_sept2017.pdf. Aug. 2017, 9 pp.

* cited by examiner




FIG. 1 PRIOR ART




FIG. 2 PRIOR ART




FIG. 3 PRIOR ART

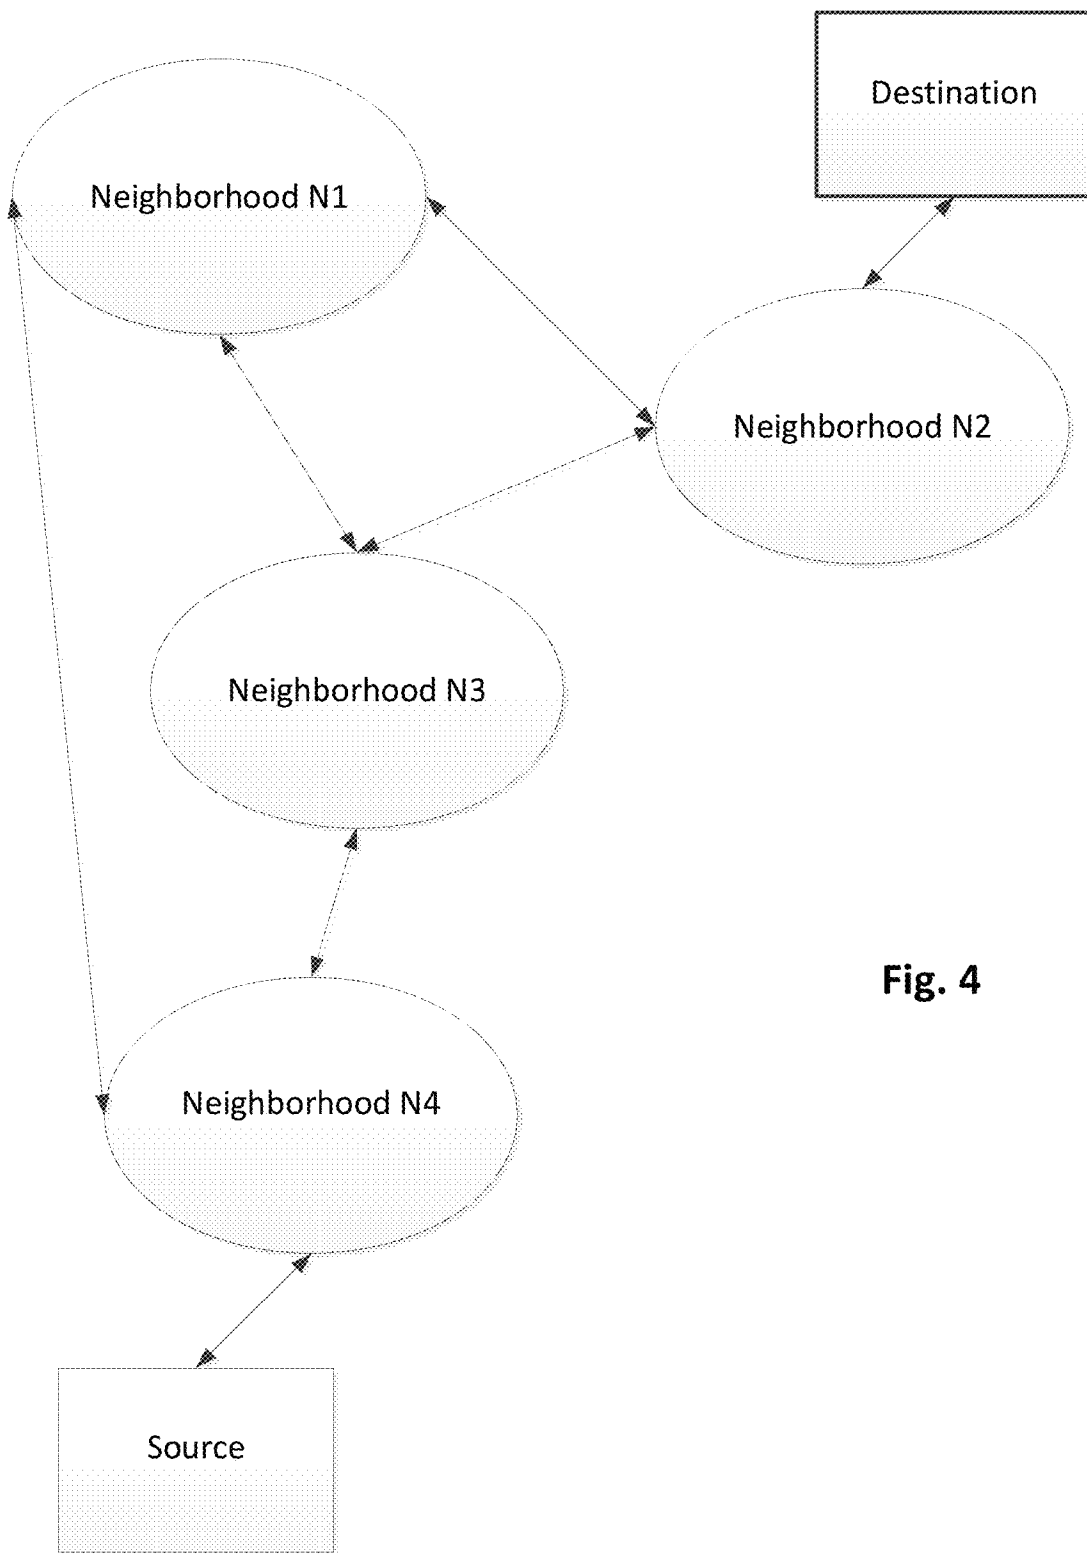


Fig. 4

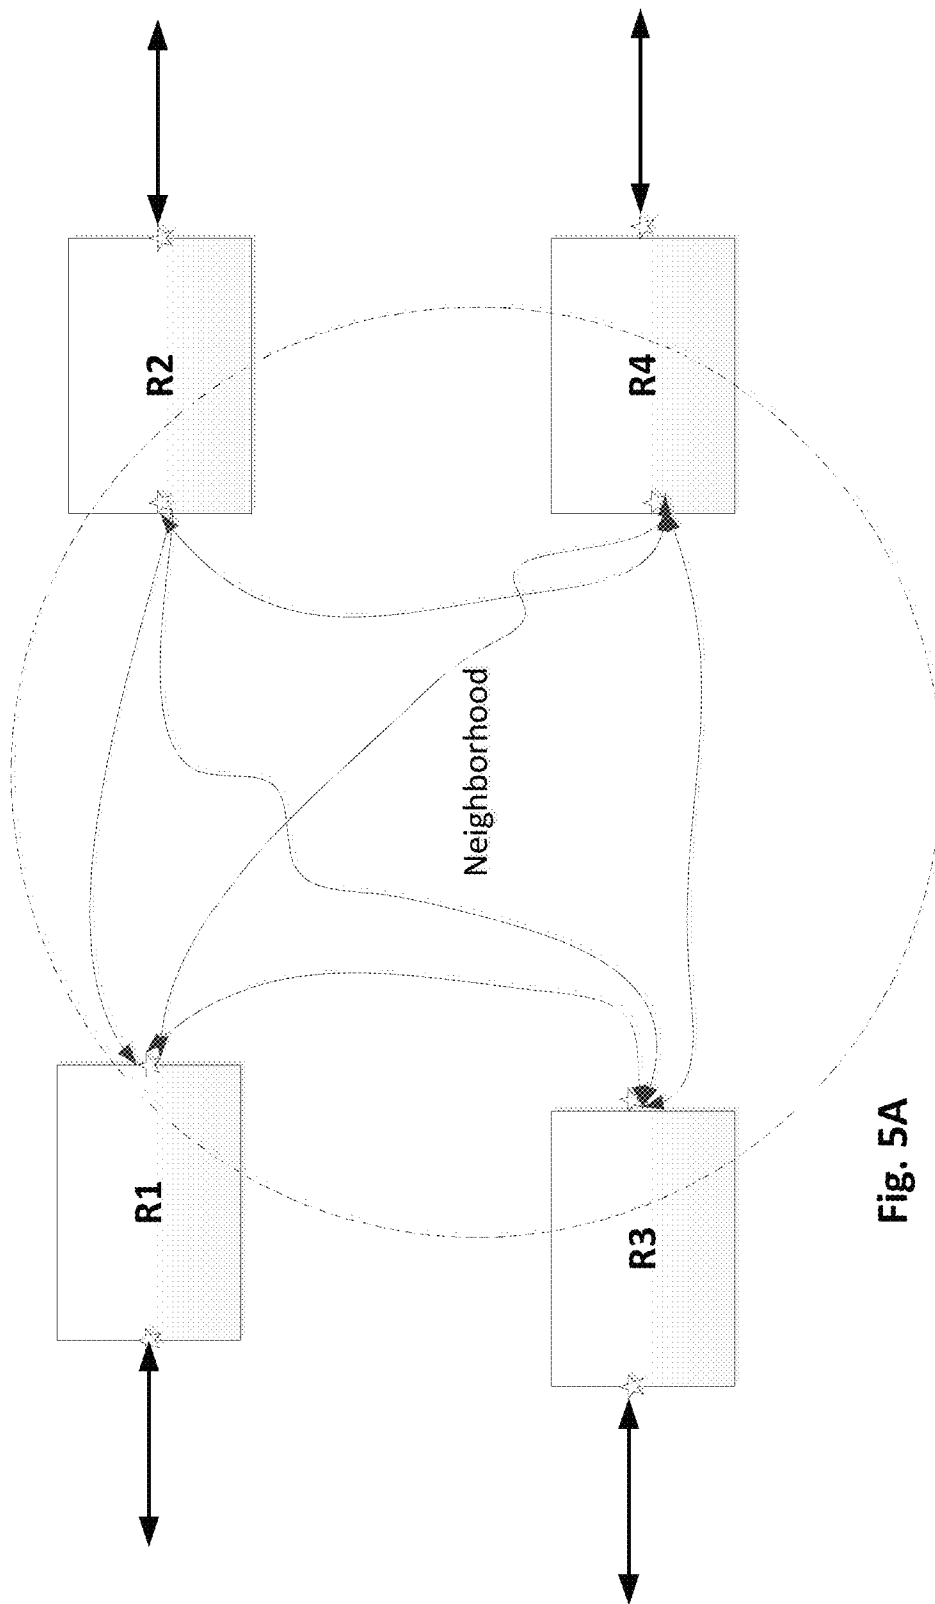


Fig. 5A




Fig. 5B

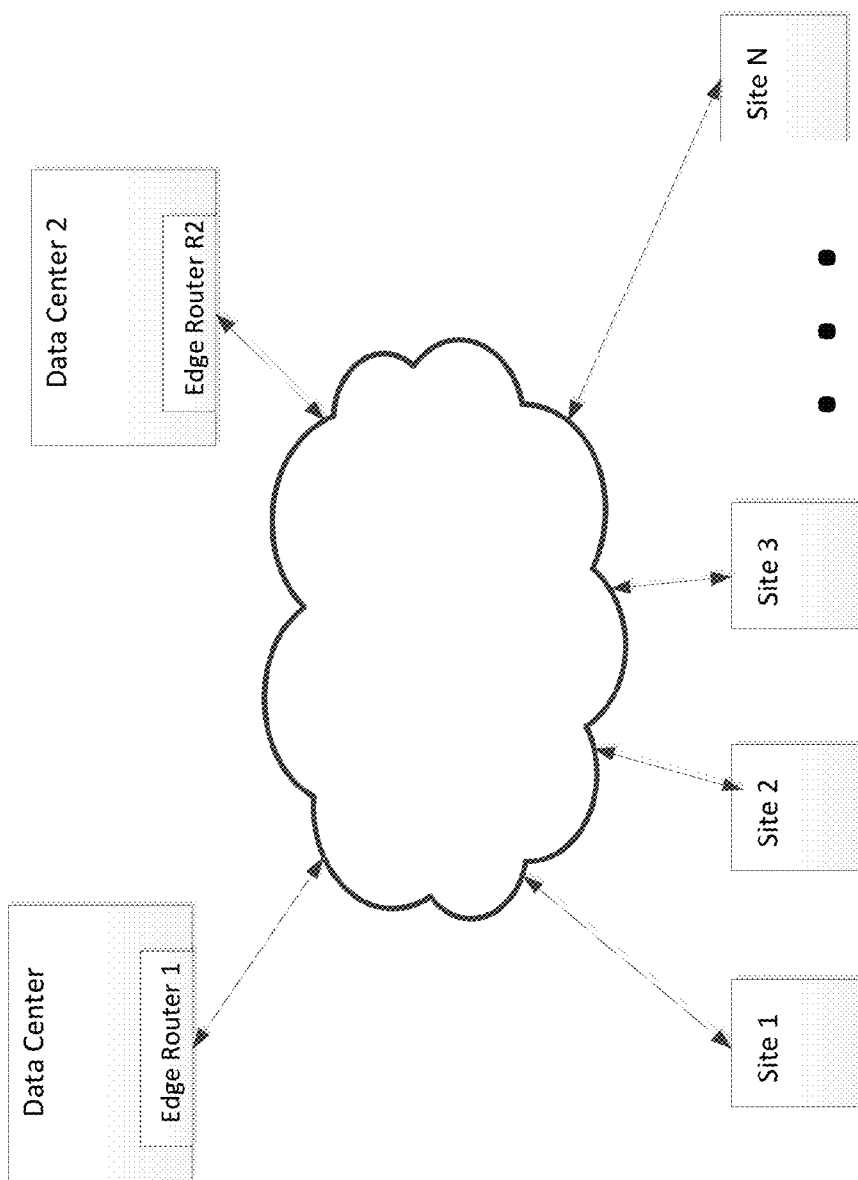


Fig. 5C

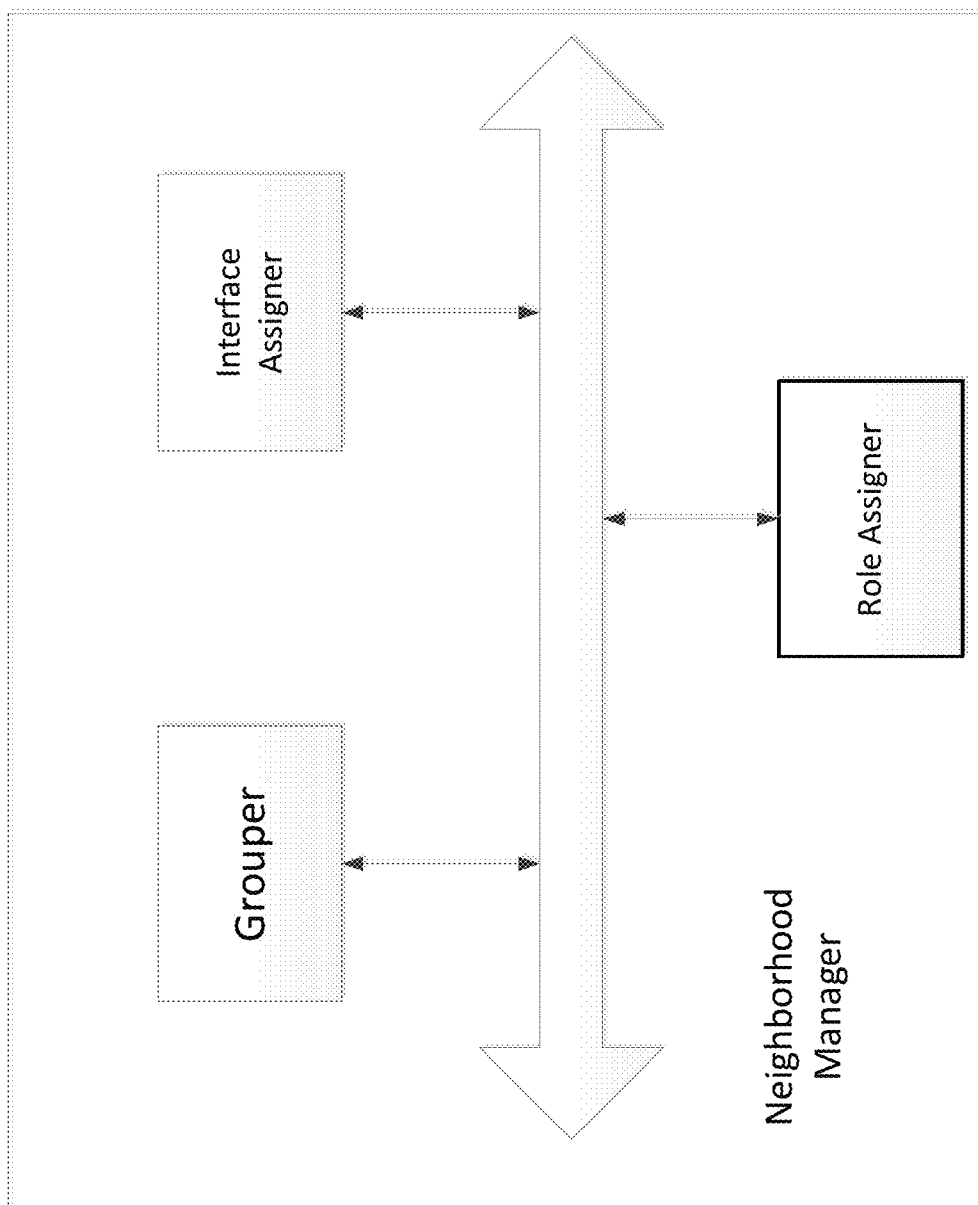


Fig. 6

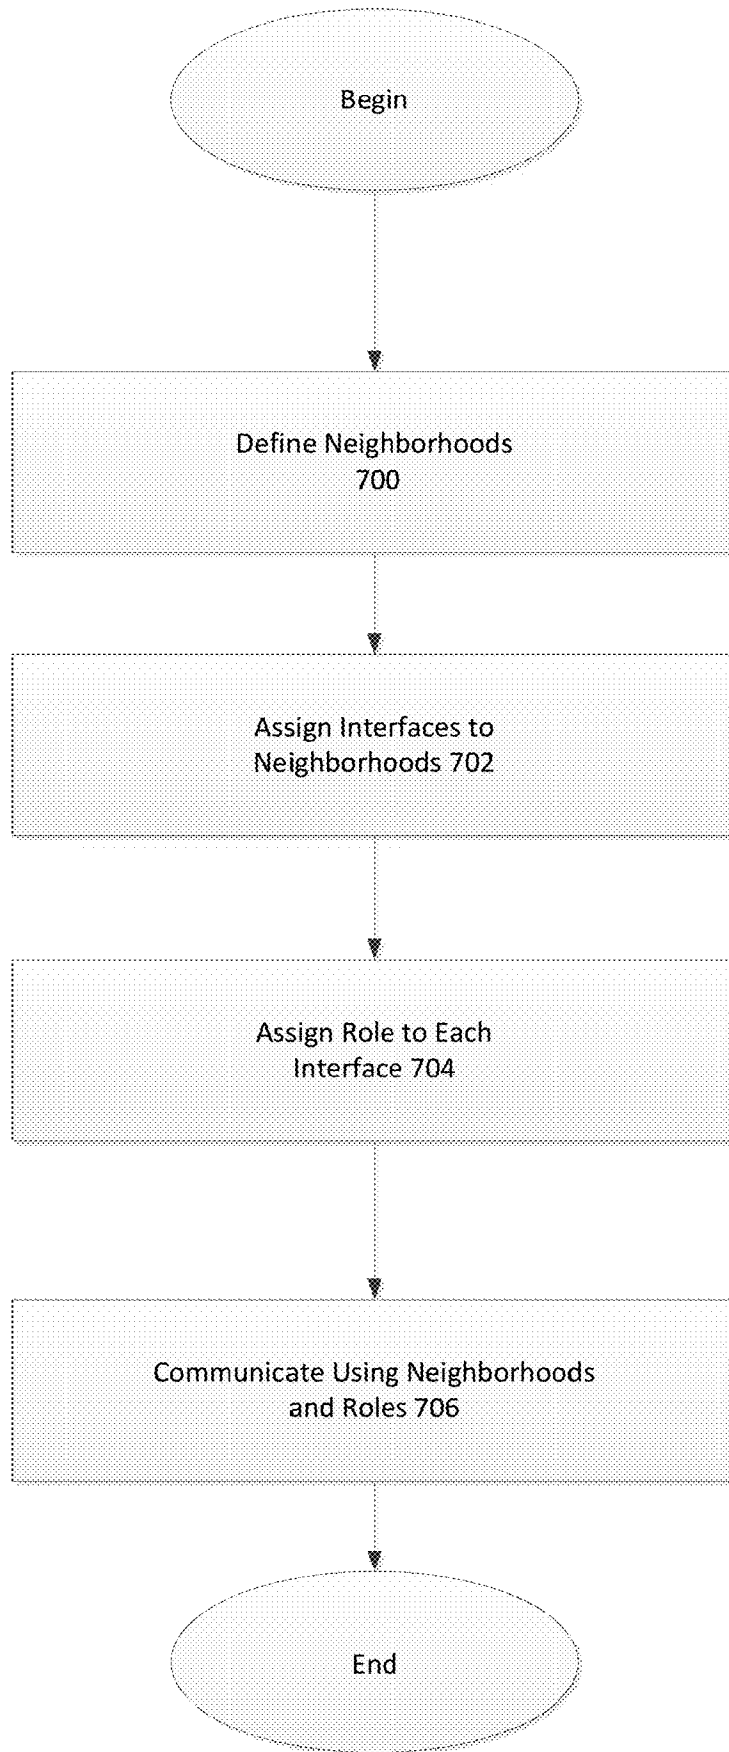


Fig. 7

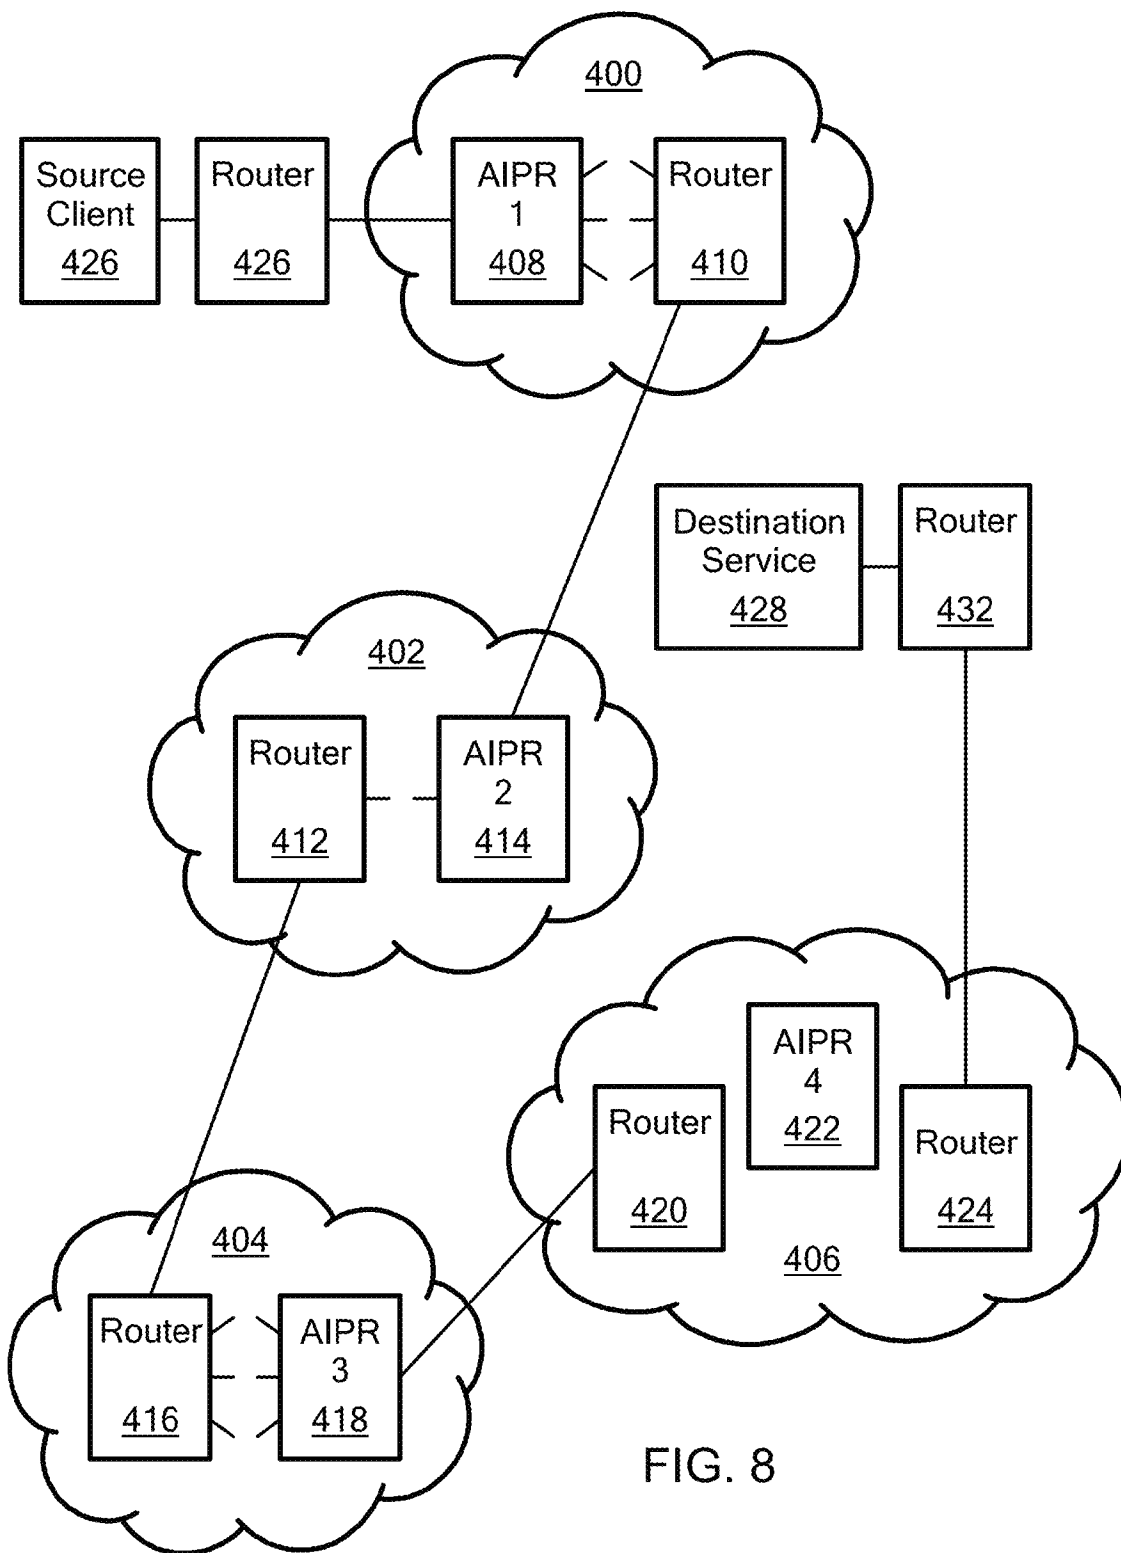


FIG. 8




FIG. 9




FIG. 10

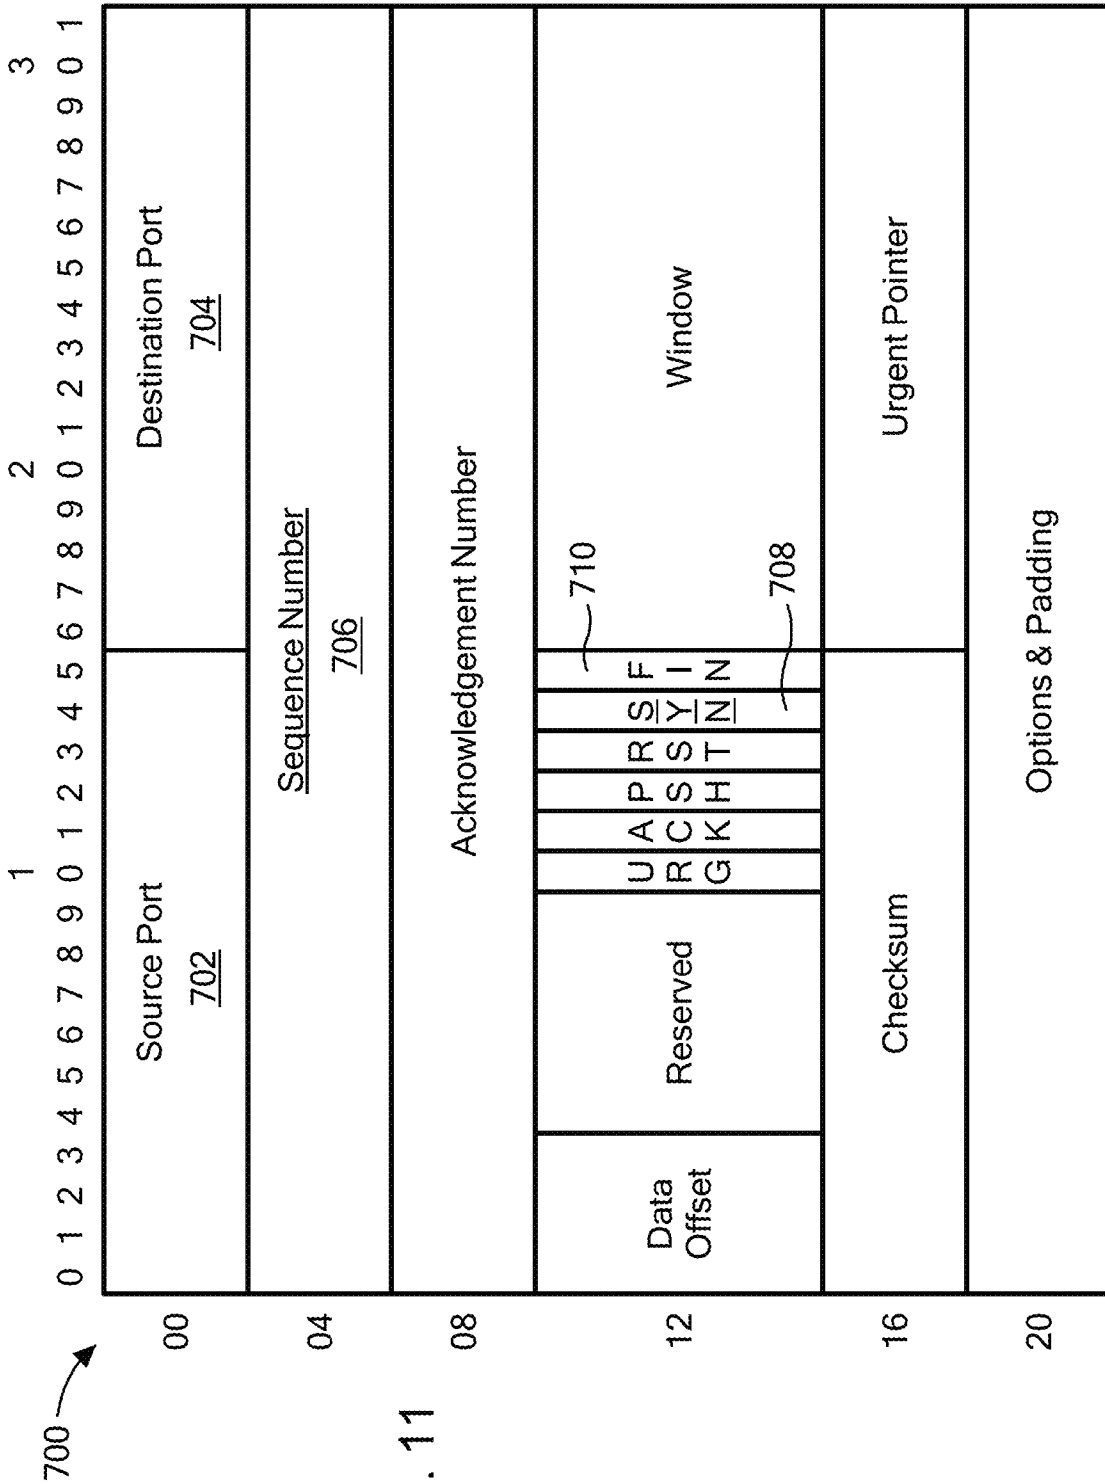


FIG. 11

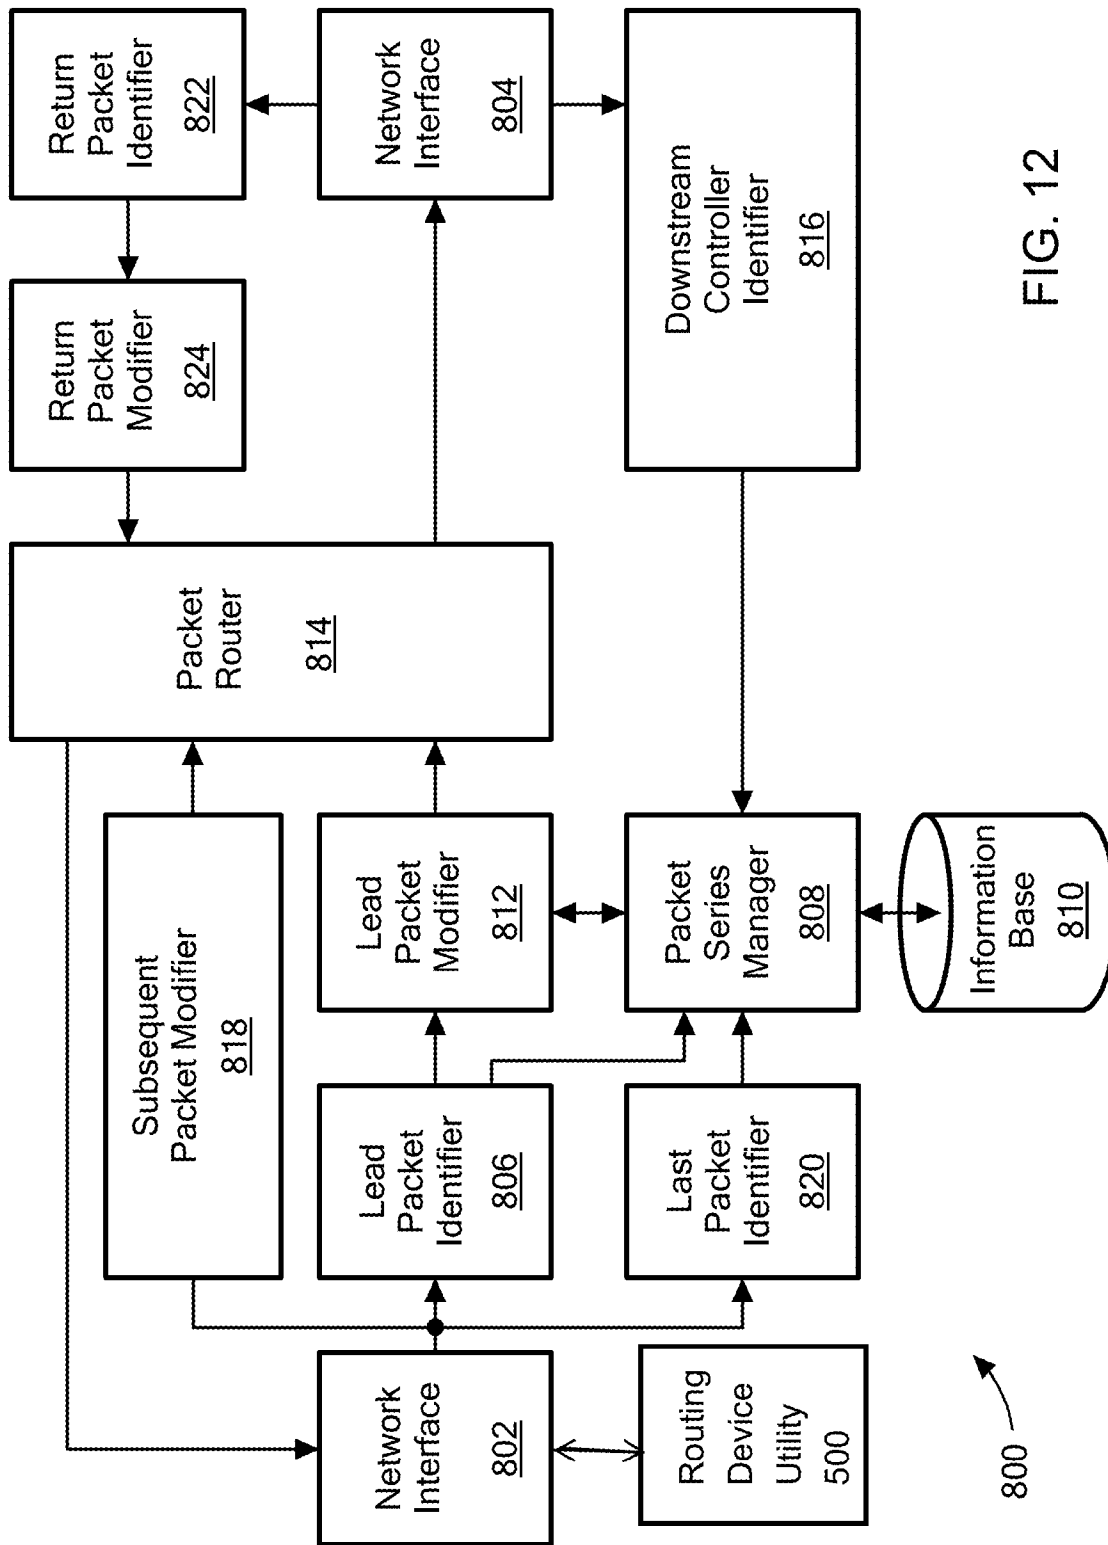


FIG. 12

800




FIG. 13

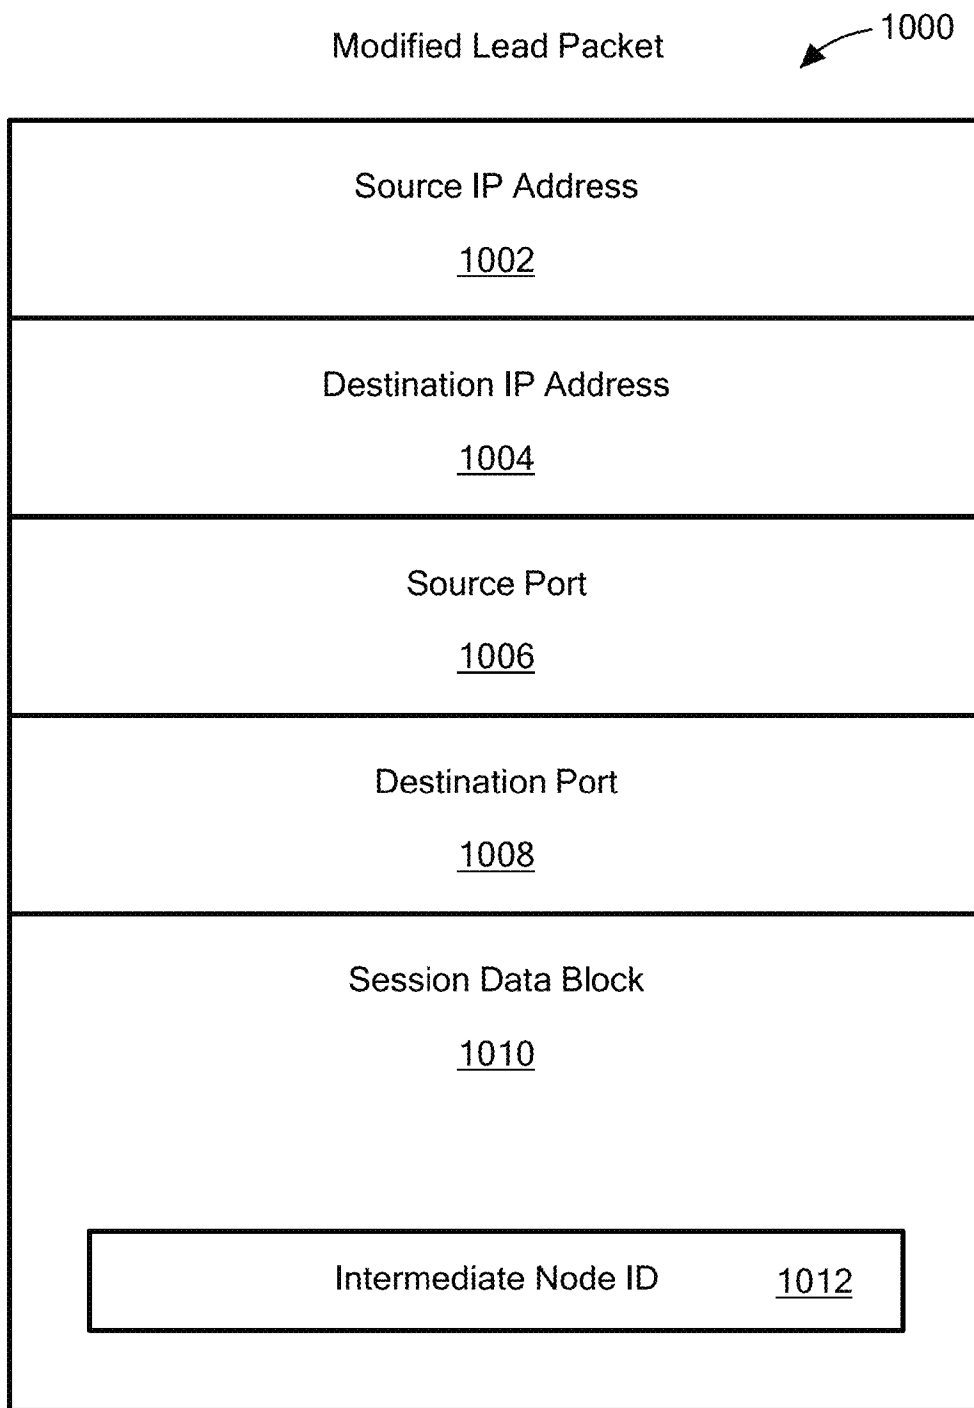


FIG. 14

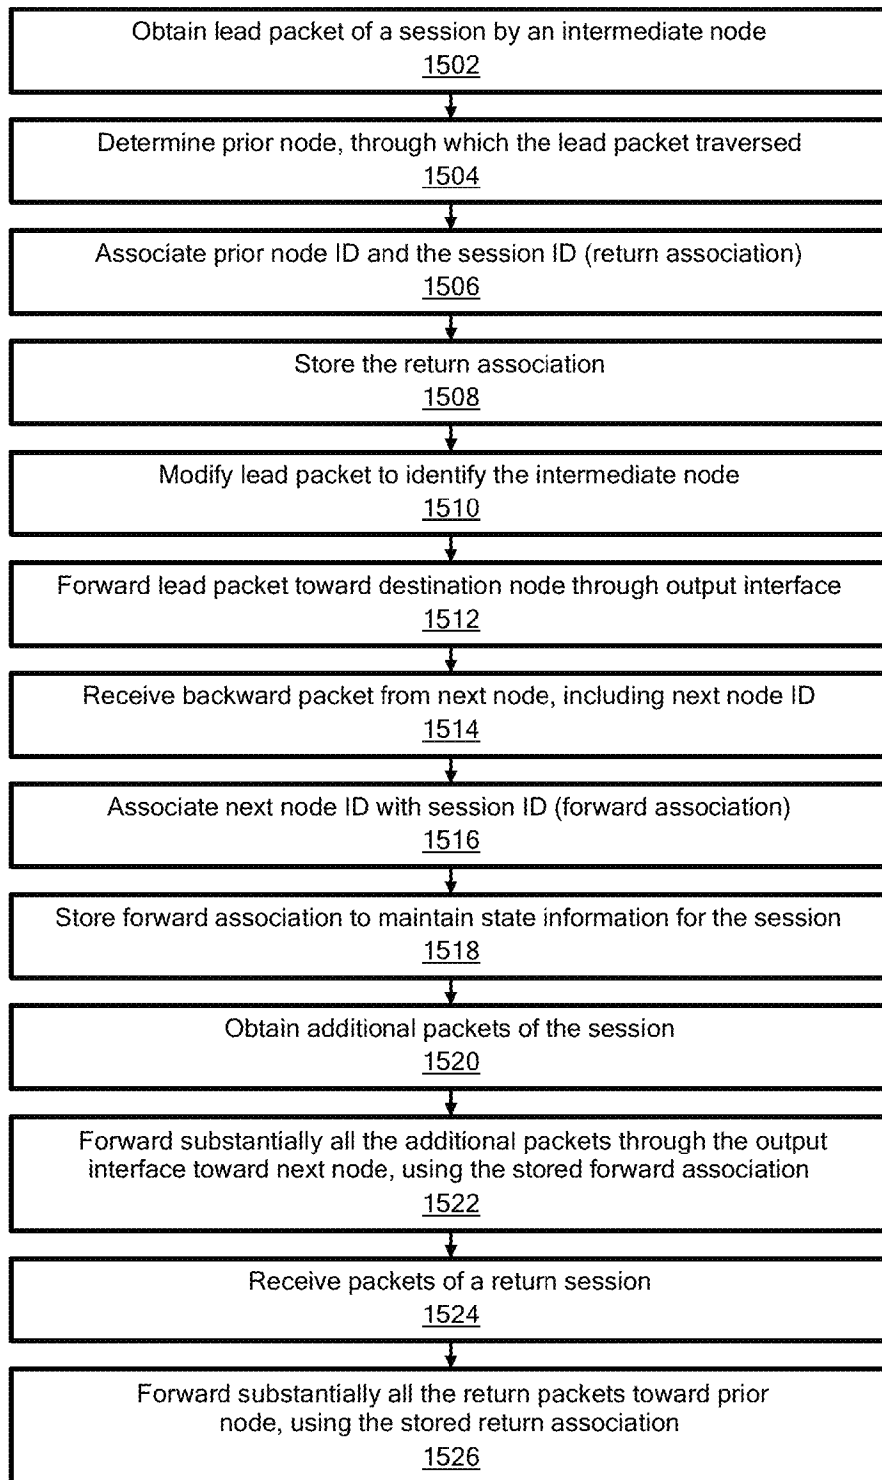


FIG 15

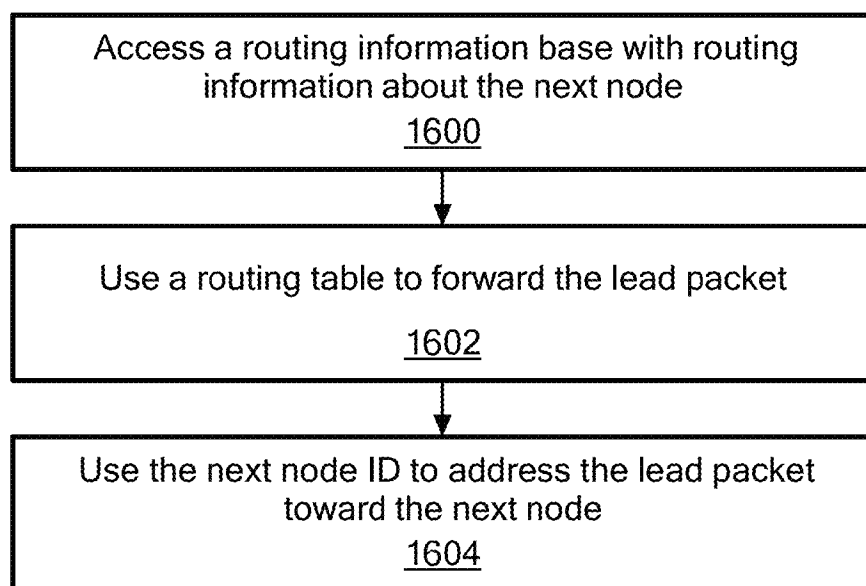


FIG. 16

1

**NETWORK NEIGHBORHOODS FOR
ESTABLISHING COMMUNICATION
RELATIONSHIPS BETWEEN
COMMUNICATION INTERFACES IN AN
ADMINISTRATIVE DOMAIN**

PRIORITY

This patent application claims priority from provisional U.S. patent application No. 62/541,357, filed Aug. 4, 2017, entitled, "NETWORK NEIGHBORHOODS," and naming Patrick Timmons, Michael Baj, Hadriel Kaplan, Robert Penfield, and Patrick McLampy as inventors, the disclosure of which is incorporated herein, in its entirety, by reference.

FIELD OF THE INVENTION

The invention generally relates to computer networks and, more particularly, the invention relates to routing data between computer networks.

BACKGROUND OF THE INVENTION

Routers transmit data across the Internet using the Internet Protocol ("IP"). As known by those in the art, 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, 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.

Routers typically communicate with other routers via their interfaces. Configuring peer relationships between routers and their interfaces can be a cumbersome and time consuming process—particularly when hundreds or thousands of routers are peered with a single router within a single administrative authority. The problem becomes more pronounced, however, upon failure or replacement of single router (a hub router) peering with a plurality of other routers (spoke routers). In that case, the new hub router typically requires extensive configuration to peer with the spoke routers.

SUMMARY OF VARIOUS EMBODIMENTS

In accordance with one embodiment of the invention, a method forms and/or defines a network topology in a Layer 3 network with a plurality of nodes, where each node has at least one interface. To that end, the method defines a plurality of neighborhoods, and assigns at least one interface of each node to at least one of the neighborhoods. The method also assigns a communication role to each interface so that each communication role is effective relative to one of the plurality of neighborhoods. The method then enables communication between the interfaces of the plurality of nodes as a function of the neighborhoods and the communication roles.

A single interface may be assigned to one or more neighborhoods. For example, the act of assigning at least one interface may include assigning a given interface (of a given node) to both a first neighborhood and a second neighborhood of the plurality of neighborhoods. In addition,

2

the act of assigning a communication role may include assigning a first communication role to the given interface relative to the first neighborhood, and assigning a second communication role to the given interface relative to the second neighborhood. The first and second communication roles may be the same or different

The communication roles may vary. For example, for each assigned neighborhood of each interface, the method may assign the communication role as one of a mesh, hub, or spoke role. Thus, a given interface may communicate with 1) all interfaces in its neighborhood (mesh), to hub interface(s) and mesh interface(s), and/or to spoke interface(s) and mesh interface(s). Some embodiments may not affirmatively assign a communication role to a given interface. In that case, the method may simply assign a default communication role (e.g., a mesh) in the event no communication role is affirmatively made.

The plurality of nodes preferably includes a plurality of routers. In a similar manner, each neighborhood preferably is a non-unique layer 3 network and defined as an attribute of interfaces of nodes. The plurality of interfaces in a given neighborhood may form peer relationships with each other. Specifically, where the nodes having interfaces in a given neighborhood are considered peers, the method may cause a first node (interface) to attempt to peer with a second node (interface) in the given neighborhood. Among other things, the Layer 3 network may include a data center and a plurality of remote sites.

To implement various embodiments, the method may form a neighborhood database having ID data identifying a plurality of the interfaces. Each ID datum in the database may be associated with at least one neighborhood datum identifying at least one neighborhood, and at least one role datum identifying at least one communication role.

In accordance with another embodiment, a communication apparatus for a network topology in a Layer 3 network with a plurality of nodes forms or defines the topology. As in the embodiment above, each node has at least one interface. To that end, the apparatus has a grouper configured to define a plurality of neighborhoods, and a pair of assigners. Specifically, the apparatus has an interface assigner, operatively coupled with the grouper, configured to assign at least one interface of each of the nodes to at least one of the neighborhoods, and a role assigner operatively coupled with the grouper. The role assigner is configured to assign a communication role to each interface, where each communication role is effective relative to one of the plurality of neighborhoods. The plurality of neighborhoods is configured to enable communication between the interfaces of the plurality of nodes as a function of the neighborhoods and the communication roles.

Illustrative embodiments of the invention are implemented as a computer program product having a 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

Those skilled in the art should more fully appreciate advantages of various embodiments of the invention from the following "Description of Illustrative Embodiments," discussed with reference to the drawings summarized immediately below.

PATENT

REEL: 060202 FRAME: 0028

3

FIG. 1 schematically shows a hypothetical prior art network that may implement illustrative embodiments of the invention.

FIG. 2 schematically illustrates a prior art technique for fragmenting a message.

FIG. 3 schematically shows a hypothetical internet that may implement illustrative embodiments of the invention.

FIG. 4 schematically shows a plurality of neighborhoods in an administrative domain in accordance with illustrative embodiments.

FIG. 5A schematically shows an example of a single neighborhood using a mesh format in accordance with illustrative embodiments.

FIG. 5B schematically shows an example of three neighborhoods using a mesh format in accordance with illustrative embodiments.

FIG. 5C schematically shows an example of a hub/spoke neighborhood in accordance with illustrative embodiments of the invention.

FIG. 6 schematically shows a neighborhood manager configured to form and use neighborhoods in accordance with illustrative embodiments.

FIG. 7 shows a process of forming and using neighborhoods in accordance with illustrative embodiments.

FIG. 8 schematically shows a hypothetical internet that includes a conventional routers and augmented IP routers (AIPRs), according to an embodiment of the present invention.

FIG. 9 schematically shows a layout of an Ethernet header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

FIG. 10 schematically shows a layout of an IP header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

FIG. 11 schematically shows a layout of a TCP header, identifying fields used for identifying a beginning of a session, according to an embodiment of the present invention.

FIG. 12 schematically shows a block diagram of an AIPR of FIG. 8, according to an embodiment of the present invention.

FIG. 13 shows a schematic illustration of information stored in an information base by the AIPR of FIGS. 8 and 12, according to an embodiment of the present invention.

FIG. 14 schematically shows a modified lead packet produced by the AIPR of FIG. 9, according to an embodiment of the present invention.

FIGS. 15 and 16 show flowcharts illustrating operations performed by the AIPR of FIGS. 8 and 12, according to an embodiment of the present invention.

DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

In illustrative embodiments, an administrative domain more efficiently configures its nodes to improve communication and re-configuration. To that end, a manager for an administrative domain assigns some or all of the interfaces of its nodes (e.g., Layer 3 device, such as routers) to a collection of other interfaces in its domain. This collection of interfaces within a single administrative domain is referred to as a “neighborhood.”

The interfaces in a neighborhood may be considered to be peered with other interfaces in its neighborhood subject to other possible limitations. For example, some interfaces may

4

be assigned a specific role to communicate directly with a subset of interfaces in its neighborhood. Accordingly, routing can be neighborhood based. Addition or removal of an interface to or from a neighborhood also requires less effort than that required by know prior art techniques. Details of various embodiments are discussed below.

Networks

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) 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.

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.

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.

Nodes communicate via networks according to protocols, such as the well-known Internet Protocol (IP) and above noted Transmission Control Protocol (TCP). The protocols are typically implemented by layered software and/or hardware components 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.

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.

Some other protocols also fragment data into packets. For example, the well-known TCP protocol 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.

Most protocols encapsulate packets of higher layer 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.

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 is used to find gateways to get an IP packet from the source to the destination.

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; i.e., one 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 thus 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 asso-

ciation can also be indicated in the routing table as an outgoing (exit) interface to the final destination.

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.

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. Illustrative embodiments discussed below relate to an improved apparatus and method for optimizing statistical data generation and collection.

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

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 learns 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.

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.

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 Inter-

network Layer and the OSI Network Layer. 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.

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.

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.

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.

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.

A computer having a single IP address may provide several services, such as web services, e-mail services and file transfer (FTP) services. The number of addresses a

computer has is orthogonal to the number of services it may provide, although it has at least one. 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.

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.

Neighborhoods

As noted above, illustrative embodiments assign some or all of the interfaces of the nodes in an administrative domain to a collection interfaces in the same domain. As such, these interfaces are considered to be "peered" interfaces as the term is known in the art. FIG. 4 schematically shows a plurality of neighborhoods in an administrative domain in accordance with illustrative embodiments. Specifically, FIG. 4 shows a Source and a Destination that communicate across an administrative domain via a plurality of neighborhoods identified as neighborhoods N1, N2, N3, and N4. The neighborhoods are connected by links, which are identified simply as double sided arrows that extend extending between interfaces to the neighborhoods N1-N4. The links can be Layer 2, Layer 3, or Layer 4 connections. Messages forwarded from the Source toward the Destination can take any of the following routes:

Source-N4-N3-N1-N2-Destination
Source-N4-N2-N2-Destination
Source-N4-N1-N2-Destination

Reverse messages from the Destination toward the Source can take the same routes back to the Source, but in reverse order—thus the double-sided arrows identifying links or connections. The routing protocol therefore can abstractly direct (e.g., route) messages, such as datagrams or IP packets, in terms of neighborhoods rather than routers or router interfaces. Indeed, when hopping from node to node within and between neighborhoods, individual nodes will use their more specific routing processes (e.g., TCP/IP).

To simplify this discussion, illustrative embodiments are discussed in terms of using routers. Thus, when forwarding packets in a session (e.g., a stateful session, discussed below), each router will access its local routing table to determine the next hop router. The illustrative routing protocol, however, directs the hops using the neighborhood data.

FIG. 5A schematically shows an example of a single neighborhood (encompassed by a dashed circle to show the neighborhood) configured in accordance with illustrative embodiments. As shown, this network has four routers R1, R2, R3, and R4, and their interfaces are identified with a five-sided star. The interfaces in this neighborhood are assigned a "mesh" role so that they can peer with all of the interfaces in the neighborhood (roles discussed below). To clearly show these roles, FIG. 5A has double sided inter-

connects extending from each interface to all of the other interfaces in the neighborhood.

For simplicity, this figure only shows two interfaces per router. Indeed, those skilled in the art can extend various embodiments to more interfaces. For each router in FIG. 5A, one of its two shown interfaces is part of the neighborhood, while the other interface is not part of the neighborhood. Those other interfaces may be part of other neighborhoods or part of no neighborhood.

In some embodiments discussed below, those other interfaces also may be part of the neighborhood. For simplicity, those other interfaces are shown with a double sided arrow to indicate two directional data transmission.

More specifically, in illustrative embodiments, a neighborhood may be considered to be a non-unique, Layer 3 network within a single administrative domain. For example, multiple independent Layer 3 networks (e.g., local area networks) may communicate through some intermediary (e.g., the Internet) and effectively form a larger Layer 3 network. Other Layer 3 networks also apply.

As discussed in greater detail below, a neighborhood manifests itself as an attribute of the interface of a router (e.g., a label, such as a string). Interfaces within the same neighborhood may be presumed to be mutually peerable and preferably are configured to attempt to peer together. Although they each extend through a single administrative domain, a neighborhood preferably does not extend to other administrative domains. To communicate between neighborhoods in different administrative domains, illustrative embodiments may provide an inter-authority translation function that maps foreign neighborhood names to local neighborhood names.

In preferred embodiments, every interface (i.e., the “network interfaces”) of a router is assigned to at least one neighborhood. In that and related cases, certain interfaces may default to a specific neighborhood in the absence of an affirmative neighborhood assignment. Alternative embodiments, however, may assign fewer than all of the interfaces to at least one neighborhood, and/or have no default neighborhood assignment.

As noted above, a neighborhood may be considered to represent a non-unique Layer 3 network. Specifically, while each interface in a neighborhood may peer with other interfaces in the neighborhood, those interfaces may be capable of peering with interfaces of other neighborhoods. However, in illustrative embodiments, interfaces are configured not to attempt to peer with interfaces in other neighborhoods (unless an interface is in both neighborhoods). In some embodiments, an administrator or logic may mistakenly assign, to the same neighborhood, two interfaces that cannot communicate with each other. In that case, the two interfaces will not be able to peer even though they likely will attempt to peer.

As noted, a single router interface may be assigned to two or more neighborhoods. For example, FIG. 5B schematically a neighborhood network arrangement in a single administrative domain having three neighborhoods N5, N6, and N7. The neighborhoods include interfaces from the following routers:

Neighborhood N5: routers R1, R2, and R3

Neighborhood N6: routers R3, R4, and R5

Neighborhood N7: routers R4 and R5.

One interface on Router R3 is a member of both neighborhoods N5 and N6 (a “multi-neighborhood interface”), while another interface on Router R3 is a member of neighborhood of N6 only. The multi-neighborhood interface on Router R3 therefore can peer with interfaces in both

neighborhoods N5 and N6, while the other single-neighborhood interface on Router R3 can only peer with interfaces in neighborhood N6.

The system can be configured to cause the multi-neighborhood interface to use different neighborhoods for different purposes, such as for different types of traffic, different destinations, varying costs, and/or load balancing. For multi-neighborhood interfaces, some embodiments may include an ordered list specifying the priority or otherwise ordering the neighborhoods within a list.

The neighborhoods in FIGS. 5A and 5B enable each interface to peer with each other interface in its neighborhood. Interfaces that can do this are referred to as having a “mesh” role. Such a neighborhood also is considered to have a “mesh configuration.” Sometimes, however, it is not practical, economical or desirable to use a mesh configuration. For example, a large retail corporation, such as a retail company, with thousands of branches (e.g., individual, geographically distributed stores) may have a policy that the branches should not peer. In that case, a mesh configuration is impractical and undesirable. Instead, such a neighborhood may use an alternative configuration referred to as a “hub and spoke” configuration.

To illustrate this point, FIG. 5C schematically shows a neighborhood using a hub and spoke configuration. As shown, the neighborhood has two data centers that each has an edge router to communicate with the plurality of stores (identified by their routers R1, R2, R3 . . . RN). The interfaces of the edge routers thus are assigned the role as “hub,” while the interfaces of the site/store routers are assigned the roles of “spokes.” Accordingly, the following designations apply to the three different roles that may be assigned to an interface:

Mesh: interface attempts to peer with all interfaces in its neighborhood

Hub: interface attempts to peer with interfaces having the roles of spokes and mesh only

Spoke: interface attempts to peer with interfaces having the roles of hub and mesh only.

The edge router R1 therefore cannot peer with the edge router R2. Conversely, the routers at each of the sites can peer with the edge routers R1 and R2 at the data centers. Accordingly, communication across the network within an administrative domain may proceed as a function of the network interfaces, their assigned neighborhoods, and their assigned roles.

Logic and/or an administrator therefore implements the neighborhoods using a specialized administrative manager. To that end, FIG. 6 schematically shows a neighborhood manager configured to form and use neighborhoods in accordance with illustrative embodiments. As shown, the network manager has a plurality of components operatively connected to each other by a conventional interconnect mechanism. FIG. 6 simply shows a bus communicating each of the components. Those skilled in the art should understand that this generalized representation can be modified to include other conventional direct or indirect connections. Accordingly, discussion of the bus is not intended to limit various embodiments.

Indeed, it should be noted that FIG. 6 only schematically shows each of the components (i.e., a grouper, a role assigner and an interface assigner) as a single box. Those skilled in the art should understand that each of these components can be implemented in a variety of conventional manners, such as by using hardware, software, or a combination of hardware and software, across one or more other functional components. For example, the grouper may be

11

implemented using a plurality of microprocessors executing firmware. As another example, the grouper may be implemented using one or more application specific integrated circuits (i.e., “ASICs”) and related software, or a combination of ASICs, discrete electronic components (e.g., transistors), and microprocessors. Accordingly, the representation of the grouper and other components in a single box of FIG. 6 is for simplicity purposes only. In fact, in some embodiments, the grouper and other components of FIG. 6 are distributed across a plurality of different machines—not necessarily within the same housing or chassis, and/or logically distinct.

It should be reiterated that the representation of FIG. 6 is a significantly simplified representation of the network manager. Those skilled in the art should understand that such a device may have other physical and functional components, such as central processing units, graphical user interfaces, and short-term memory. Accordingly, this discussion is not necessarily intended to suggest that FIG. 6 represents all of the elements of the network manager.

As shown, the network manager has the above noted grouper configured to define a plurality of neighborhoods, and the noted interface assigner configured to assign interfaces to specific groups. In some embodiments, the functions of the grouper and interface assigner are combined into a single entity that inherently defines the groups merely by assigning interfaces specific neighborhood labels/identification information.

The network manager also has a role assigner configured to assign a role to each of the interfaces in the networks. In illustrative embodiments, each multi-neighborhood interface may have the same role across all of its neighborhoods, or have different roles in different neighborhoods. For example, a given multi-neighborhood interface may have a mesh role in a first neighborhood, and a spoke role in a second neighborhood.

FIG. 7 shows a process of forming and using neighborhoods in accordance with illustrative embodiments. It should be noted that this process is simplified from a longer process that normally would be used to form and use neighborhoods. Accordingly, the process may have many steps, such as policy analysis and conflict resolution steps, which those skilled in the art likely would use. In addition, some of the steps may be performed in a different order than that shown, or at the same time. Those skilled in the art therefore can modify the process as appropriate.

The process begins at step 700, in which the grouper defines one or more neighborhoods in the administrative domain. Each neighborhood, however, may be considered to be essentially defined by its interfaces. Accordingly, at step 702, the interface assigner assigns interfaces of routers to one or more of the defined neighborhoods in its administrative domain. To that end, the administrator, a database manager and/or similar logic may form a neighborhood database (e.g., a relational database). In that database, each interface has one or more associated neighborhoods. Those neighborhoods may be identified by strings defined (or declared) at step 700. As noted above, a priority or ordered list may be assigned to the multi-neighborhood interfaces. Next, at step 704, the role assigner assigns a role to each interface in their various neighborhoods. Again, a specific string may be assigned, such as “mesh,” “hub,” or “string.” For each multi-neighborhood interface, a single role may be assigned to each neighborhood in which it is a member.

The administrator or logic may update the database as new routers are added and removed. For example, if the data center edge router R1 of FIG. 5C is replaced with a new edge

12

router, the administrator simply replaces the old data center edge router information with the new data center edge router information. Accordingly, this eliminates the time error-prone and time consuming process of forming associations between interfaces; i.e., having to manually re-peer the sometimes thousands of spoke routers with the single new data center edge router. This significantly simplifies a complex network technical problem with the simple change of a few database entries.

The process concludes at step 706 by communicating through the administrative domain using the neighborhoods, as discussed above with regard to FIG. 4.

Illustrative embodiments also may group neighborhoods into districts. As such, routing may be district by district in a manner similar to neighborhood by neighborhood routing. Districts can similarly be grouped to super-districts, and so on to even larger groupings.

Stateful Routing

In some embodiments, the packets of a session follow the same path as the lead packet of that session, at least in the forward direction, i.e., from a source client to a 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. Among other benefits, this stateful routing scheme more effectively enhances the benefits of collecting routing statistics. Those skilled in the art may implement such a stateful routing scheme with the configuration functionality of the administrative domain 312.

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, or routed by a pre-specified path. It should be noted that although the discussion below describes use of natural routing, it should not limit various embodiments to such a scheme. If the path is not pre-specified, then the path taken by the lead packet establishes the waypoints. In either case, the subsequent packets traverse the same waypoints, and in the same order, as the lead packet.

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.

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. Importantly, this enables statistics to be calculated for specific sessions.

Some embodiments 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, which might not otherwise be available by naturally routing the return packets.

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. In fact, any of the routers mentioned when discussing FIGS. 4-7 may be implemented as AIPRs/waypoints. An AIPR also performs conventional routing functions, and also those described above with regard to statistics collection. FIG. 8 is a schematic diagram illustrating a hypothetical set of interconnected networks 400, 402, 404 and 406, i.e., an internet. Each network 401-406 includes a number of routers and AIPRs, not all of which are necessarily shown. Network 401 includes AIPR 408 and router 410. Network 401 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.

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 401 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.

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.

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.

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.

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.

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.

AIPR 1 408 then forwards the lead packet naturally. The lead packet traverses an unspecified number of nodes of network 401 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.

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.

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.

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.

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.

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.

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.

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

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.

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.

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.

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.

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.

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.

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.

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.

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.

In the third scenario, not shown in FIG. 8, 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.

It should be noted that although preferred embodiments use stateful routing as noted above, other embodiments do not use stateful routing.

Lead Packet Identification

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

17

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.

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.

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.

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.

18

FIG. 9 is a schematic layout of an Ethernet header 500, including a Destination MAC Address 502 and an 802.1q VLAN Tag 504. FIG. 10 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. 11 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
15 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	Logical association with a group of computers.
20 Protocol	Address & Previous Advertisement IP Header	This defines the protocol in use and, for the TCP case, it must be set to a value that corresponds to TCP
25 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 from the source. This may reflect a client, a

TABLE 3-continued

Data Item	Where From	Description
		firewall in front of the client, or a carrier grade NAT.
Destination Port	TCP Header	This defines the desired service requested, such as 80 for HTTP.
Sequence Number	TCP Header	This is a random number assigned by the client. It may be updated by a firewall or carrier grade NAT.
SYN Bit On	TCP Header	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.

Augmented IP Router (AIPR)

FIG. 12 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**.

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.

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.

FIG. 13 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.

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.

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**.

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.

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**. Statistical information may be gathered and/or calculated using either or both forward and backward information **916**.

Returning to FIG. 12, 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.

FIG. 14 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. 14 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.

Returning to FIG. 12, 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.

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. 13), specifically the network address and port number, about the next AIPR in the waypoint information base **900**.

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. 13) 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.

A last packet identifier **820** statefully follows each session 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**.

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.

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.

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.

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. 13) 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**.

FIG. 15 shows a flowchart schematically illustrating some operations performed by the AIPR **800** (FIG. 12) in accordance with illustrative embodiments of the invention. The flowchart illustrates a packet routing method for directing packets of a session from an originating node toward a destination node in an IP network. At step **1502**, 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.

The packets in the session have a unique session identifier. At step **1504**, a prior node, through which the lead packet traversed, is determined. The prior node has a prior node identifier. At step **1506**, a return association is formed between the prior node identifier and the session identifier. At step **1508**, the return association is stored in memory to maintain state information for the session.

At step **1510**, the lead packet is modified to identify at least the intermediate node. At step **1512**, 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 step **1514**, 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.

At step **1516**, a forward association is formed between the next node identifier and the session identifier. At step **1518**, the forward association is stored in memory, to maintain state information for the session. At step **1520**, additional packets of the session are obtained. At step **1522**, 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.

At step **1524**, 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 step **1526**, 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.

As shown at step **1600** in FIG. 16, forwarding the lead packet **1412** toward the destination node may include accessing a routing information base having routing information for the next node. As shown at step **1602**, the intermediate node may have a routing table, and forwarding the lead packet **1412** toward the destination node may include using the routing table to forward the lead packet toward the destination node. As shown at step **1604**, for-

warding the lead packet 1412 toward the destination node may include using the next node identifier to address the lead packet toward the next node.

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.

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. In a similar manner, the routing device utility 510 also 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.

Various embodiments of the invention may be implemented at least in part in any conventional computer programming language. For example, some embodiments may be implemented in a procedural programming language (e.g., "C"), or in an object oriented programming language (e.g., "C++"). Other embodiments of the invention may be implemented as a pre-configured, stand-alone hardware element and/or as preprogrammed hardware elements (e.g., application specific integrated circuits, FPGAs, and digital signal processors), or other related components.

In an alternative embodiment, the disclosed apparatus and methods (e.g., see the various flow charts described above) may be implemented as a computer program product for use with a computer system. Such implementation may include a series of computer instructions fixed either on a tangible, non-transitory medium, such as a computer readable medium (e.g., a diskette, CD-ROM, ROM, or fixed disk). The series of computer instructions can embody all or part of the functionality previously described herein with respect to the system.

Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Furthermore, such instructions may be stored in any memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology, such as optical, infrared, microwave, or other transmission technologies.

Among other ways, such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (e.g., shrink wrapped software), preloaded with a computer system (e.g., on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (e.g., the Internet or World Wide Web). In fact, some embodiments may be implemented in a software-as-a-service model ("SAAS") or cloud computing model. Of course, some embodiments of the invention may be implemented as a combination of both software (e.g., a computer program product) and hardware. Still other embodiments of the invention are implemented as entirely hardware, or entirely software.

Although the above discussion discloses various exemplary embodiments of the invention, it should be apparent that those skilled in the art can make various modifications that will achieve some of the advantages of the invention without departing from the true scope of the invention.

What is claimed is:

1. A method comprising:

configuring a first interface of a first node of a plurality of nodes with an attribute specifying a neighborhood of a plurality of neighborhoods within an administrative domain, the plurality of nodes forming a Layer 3 network;

configuring the first interface with a first communication role within a network topology formed by nodes of the plurality of nodes configured with the attribute specifying the neighborhood;

forming, by the first node and based on the first communication role of the first interface being compatible with a second communication role of a second interface of a second node of the plurality of nodes configured with the attribute specifying the neighborhood, a peering relationship between the first interface and the second interface, wherein the second communication role is within the network topology; and

in response to forming the peering relationship between the first interface and the second interface, forwarding, by the first node via the first interface, network traffic to the second interface of the second node.

2. The method of claim 1, further comprising:

configuring, by the first node, the first interface of the first node with a second attribute specifying a second neighborhood of the plurality of neighborhoods different from the neighborhood;

configuring, by the first node, the first interface with a third communication role within a second network topology formed by nodes of the plurality of nodes configured with the second attribute specifying the second neighborhood;

forming, by the first node and based on the third communication role of the first interface being compatible with a fourth communication role of a third interface of a third node of the plurality of nodes configured with the second attribute specifying the second neighborhood of the plurality of neighborhoods, a peering relationship between the first interface and the third interface; and

in response to forming the peering relationship between the first interface and the third interface, forwarding, by the first node via the first interface, network traffic to the third interface of the third node.

3. The method of claim 2, wherein the third communication role is different than the first communication role.

4. The method of claim 2, wherein the third communication role is the same as the first communication role.

5. The method of claim 1, wherein the first communication role within the network topology formed by nodes of the plurality of nodes configured with the attribute specifying the neighborhood comprises one of a mesh role, a hub role, or a spoke role.

6. The method of claim 1,

wherein the first node comprises a plurality of interfaces, the plurality of interfaces including the first interface and a third interface, and

wherein the method further comprises:

configuring the third interface of the first node with a second attribute specifying a second neighborhood of the plurality of neighborhoods different from the neighborhood;

configuring the third interface with a third communication role within a second network topology formed by nodes of the plurality of nodes configured with

25

the second attribute specifying the second neighborhood different than the first communication role; forming, by the first node and based on the third communication role for the second neighborhood of the third interface being compatible with a fourth communication role of a fourth interface of a third node of the plurality of nodes configured with the second attribute specifying the second neighborhood, a peering relationship between the third interface and the fourth interface; and

in response to forming the peering relationship between the third interface and the fourth interface, forwarding, by the first node via the third interface, network traffic to the fourth interface of the third node.

7. The method of claim 1, further comprising forming, by the first node, a peering relationship between the first interface and each interface of a plurality of interfaces of the plurality of nodes, the plurality of interfaces configured with the attribute specifying the neighborhood and configured with communication roles that are compatible with the first communication role of the first interface.

8. The method of claim 1,

wherein the network traffic comprises packets of a session comprising a forward packet flow and a reverse packet flow, and

wherein forwarding the network traffic further comprises modifying a first packet of at least one of the forward packet flow and the reverse packet flow of the session to include a session identifier for the session.

9. The method of claim 1, wherein the first node comprises a router.

10. The method of claim 1, wherein the Layer 3 network comprises:

a data center; and

a plurality of remote sites.

11. A first node comprising:

a first interface,

wherein a plurality of nodes comprising the first node are configured to form a Layer 3 network,

wherein the first interface is configured with an attribute specifying a neighborhood of a plurality of neighborhoods within an administrative domain,

wherein the first interface is further configured with a first communication role within a network topology formed by nodes of the plurality of nodes configured with the attribute specifying the neighborhood, and

wherein the first node is configured to:

form, based on the first communication role of the first interface being compatible with a second communication role of a second interface of a second node of the plurality of nodes configured with the attribute specifying the neighborhood, a peering relationship between the first interface and the second interface, wherein the second communication role is within the network topology; and

in response to forming the peering relationship between the first interface and the second interface, forward, via the first interface, network traffic to the second interface of the second node.

12. The first node of claim 11,

wherein the first interface of the first node is further configured with a second attribute specifying a second neighborhood of the plurality of neighborhoods different from the neighborhood,

wherein the first interface is further configured with a third communication role for the second neighborhood, and

26

wherein the first node is further configured to:

form, based on the third communication role of the first interface being compatible with a fourth communication role of a third interface of a third node of the plurality of nodes configured with the second attribute specifying the second neighborhood, a peering relationship between the first interface and the third interface; and

in response to forming the peering relationship between the first interface and the third interface, forward, via the first interface, network traffic to the third interface of the third node.

13. The first node of claim 12, wherein the third communication role is different than the first communication role.

14. The first node of claim 12, wherein the third communication role is the same as the first communication role.

15. The first node of claim 11, wherein the first communication role within the network topology formed by nodes of the plurality of nodes configured with the attribute specifying the neighborhood comprises one of a mesh role, a hub role, or a spoke role.

16. The first node of claim 11,

wherein the first node comprises a plurality of interfaces, the plurality of interfaces including the first interface and a third interface, and

wherein the third interface of the first node is configured with a second attribute specifying a second neighborhood of the plurality of neighborhoods different from the neighborhood,

wherein the third interface is further configured with a third communication role within a second network topology formed by nodes of the plurality of nodes configured with the second attribute specifying the second neighborhood different than the first communication role, and

wherein the first node is further configured to

form, based on the third communication role of the third interface being compatible with a fourth communication role of a fourth interface of a third node of the plurality of nodes configured with the second attribute specifying the second neighborhood, a peering relationship between the third interface and the fourth interface; and

in response to forming the peering relationship between the third interface and the fourth interface, forward, via the third interface, network traffic to the fourth interface of the third node.

17. The first node of claim 11, further configured to form a peering relationship between the first interface and each interface of a plurality of interfaces of the plurality of nodes, the plurality of interfaces configured with the attribute specifying the neighborhood and configured with communication roles that are compatible with the first communication role of the first interface.

18. The first node of claim 11,

wherein the network traffic comprises packets of a session comprising a forward packet flow and a reverse packet flow, and

wherein to forward the network traffic, the first node is configured to modify a first packet of at least one of the forward packet flow and the reverse packet flow of the session to include a session identifier for the session.

19. The first node of claim 11, wherein the Layer 3 network comprises:

a data center; and

a plurality of remote sites.

20. A non-transitory, computer-readable medium comprising instructions for causing one or more processors of a first node of a plurality of nodes to:

form, based on a first communication role with which the first interface is configured being compatible with a second communication role of a second interface of a second node of the plurality of nodes, a peering relationship between the first interface and the second interface, wherein each of the first interface of the first node and the second interface of the second node is configured with an attribute specifying a neighborhood of a plurality of neighborhoods within an administrative domain, wherein the first communication role is within a network topology formed by nodes of the plurality of nodes configured with the attribute specifying the neighborhood, wherein the plurality of nodes are configured to form a Layer 3 network, and wherein the second communication role is within the network topology; and

in response to forming the peering relationship between the first interface and the second interface, forward, via the first interface, network traffic to the second interface of the second node.

* * * * *