

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.
Petitioner

v.

PORTSMOUTH NETWORK CORP.
Patent Owner

IPR2024-00506
U.S. Patent No. 8,014,394

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. § 312 AND 37 C.F.R. § 42.104**

Mail Stop Patent Board
Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. MANDATORY NOTICES	2
A. Real Party-in-Interest (37 C.F.R. § 42.8(b)(1)).....	2
B. Related Matters (37 C.F.R. § 42.8(b)(2)).....	2
C. Lead and Back-up Counsel (37 C.F.R. § 42.8(b)(3)).....	3
D. Service Information (37 C.F.R. § 42.8(b)(4))	3
III. FEES	4
IV. GROUNDS FOR STANDING.....	4
V. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED.....	4
A. Prior Art References.....	4
B. Grounds For Challenge	5
VI. TECHNOLOGY BACKGROUND.....	6
A. Networking, Multicasting, and IGMP.....	6
B. Distributed Router	8
C. Port Bandwidth Control	8
VII. OVERVIEW OF THE '394 PATENT	9
A. Purported Invention.....	9
B. Prosecution History	11
VIII. CLAIM CONSTRUCTION	12
IX. LEVEL OF ORDINARY SKILL IN THE ART.....	12

X.	OVERVIEW OF PRIMARY PRIOR ART REFERENCE.....	13
A.	Weyman.....	13
XI.	SPECIFIC GROUNDS OF INVALIDITY	17
A.	Ground 1: Weyman and Li Render Obvious Claims 1, 4, 6-7, 11, 14, and 17.	17
1.	Weyman In View Of Li	17
2.	Independent Claim 1	25
3.	Claim 4.....	49
4.	Claim 6.....	51
5.	Claim 7.....	51
6.	Independent Claim 11	51
7.	Claim 14.....	56
8.	Claim 17.....	56
B.	Ground 2: Weyman, Li, And Rooholamini Render Obvious Claims 5 and 15.....	56
1.	Weyman/Li in view of Rooholamini	56
2.	Claim 5.....	61
3.	Claim 15.....	64
C.	Ground 3: Weyman, Li, And Pitcher Render Obvious Claim 16.	65
1.	Weyman/Li In View Of Pitcher.....	65
2.	Claim 16.....	67
D.	Ground 4: Weyman, Jain, And LeMaire Render Obvious Claims 8 And 18.....	68
1.	Weyman In View Of Jain And LeMaire.....	68

2.	Independent Claim 8	75
3.	Independent Claim 18	81
E.	Ground 5: Weyman, Jain, LeMaire, And Li Render Obvious Claim 9.	83
1.	Weyman/Jain/LeMaire In View Of Li.....	83
2.	Claim 9	84
F.	Ground 6: Weyman, Jain, LeMaire, And Pitcher Render Obvious Claim 10.....	85
1.	Weyman/Jain/LeMaire In View Of Pitcher	85
2.	Claim 10	86
G.	Ground 7: Weyman, Li, and Ding Render Obvious Claims 2 and 12.	87
1.	Claim 2	87
2.	Claim 12	89
XII.	DISCRETIONARY DENIAL IS NOT WARRANTED.....	89
XIII.	CONCLUSION.....	90

TABLE OF AUTHORITIES

Pages(s)

Cases

<i>Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.</i> , 868 F.3d 1013 (Fed. Cir. 2017)	12
---	----

Statutes

35 U.S.C. § 102	4
35 U.S.C. § 102(b)	4, 5
35 U.S.C. § 103	4
35 U.S.C. § 314(a)	6
37 C.F.R. § 42.104(b)(1)-(2)	4
37 C.F.R. §§ 42.22(a)(1)	4

PETITIONER’S EXHIBIT LIST

Exhibit No.	Description
1001	U.S. Patent No. 8,014,394 to Ram et al. (“’394 Patent”)
1002	Declaration of Dr. Henry Houh
1003	<i>Curriculum Vitae</i> of Dr. Houh
1004	Prosecution History of U.S. Patent No. 8,014,394
1005	U.S. Patent Publication No. 2005/0041665 (“Weyman”)
1006	W.O. Patent Publication No. 2007/140702 (“Li”)
1007	Certified Translation of W.O. Patent Publication No. 2007/140702 (“Li”)
1008	U.S. Patent Publication No. 2007/0297406 (“Rooholamini”)
1009	U.S. Patent No. 6,370,142 (“Pitcher”)
1010	U.S. Patent Publication No. 2006/0224877 (“Jain”)
1011	U.S. Patent No. 5,999,530 (“LeMaire”)
1012	U.S. Patent Publication No. 2006/0146823 (“Ding”)
1013	U.S. Patent Publication No. 2008/0151890 (“Zelig”)
1014	U.S. Patent Publication No. 2006/0200575 (“Sherer”)
1015	RFC 3376 (“IGMPv3”), Internet Group Management Protocol, Version 3
1016	U.S. Patent Publication No. 2005/0238018 (“Devi”)
1017	U.S. Patent Publication No. 2006/0146857 (“Naik”)
1018	U.S. Patent Publication No. 2003/0147412 (“’412 Weyman”)
1019	U.S. Patent No. 6,034,957 (“Haddock”)

Exhibit No.	Description
1020	U.S. Patent Publication No. 2006/0067213 (“Evans”)
1021	U.S. Patent Publication No. 2002/0024964 (“Baum”)
1022	U.S. Patent Publication No. 2002/0143951 (“Khan”)
1023	U.S. Patent Publication No. 2008/0298236 (“Ervin”)
1024	U.S. Patent Publication No. 2010/0020680 (“Salam”)
1025	U.S. Patent No. 8,238,344 (“Chen”)
1026	U.S. Patent No. 7,525,965 (“St. Pierre”)
1027	U.S. Patent Publication No. 2005/0041680 (“Tanaka”)
1028	Docket Control Order, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Dec. 18, 2023)
1029	Notice of Disclosures, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Jan. 8, 2024)
1030	Order Granting Unopposed Motion for Leave to File Amended Complaint for Patent Infringement, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Mar. 12, 2024)
1031	Director’s June 21, 2022, Memorandum on Discretionary Denials

Citations to specific pages of exhibits follow the pagination added to those exhibits per 37 C.F.R. § 42.63(d)(2)(i).

All emphasis has been added unless noted otherwise.

All annotations in Figures have been added unless noted otherwise.

CLAIM LISTING

Claim 1

[1pre] A method for communication, comprising:

[1a] operating a network element,

[1b] wherein the network element further comprises multiple interconnected processing units and multiple ports,

[1c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

[1d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned;

[1e] storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element, and further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to;

[1f] receiving from a client a request relating to reception of a given multicast packet stream;

[1g] distributing the request among the multiple processing units within the network element, and

[1h] updating the respective list in each processing unit responsively to the request; and

[1i] selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists,

[1j] wherein operating the network element comprises aggregating two or more of the ports to form a high-capacity port, and

[1k] wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port,

[1l] wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and

[1m] wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

Claim 2

[2] The method according to claim 1, wherein storing the list, distributing the request and selectively forwarding the multicast packets are performed by dedicated hardware logic.

Claim 4

[4a] The method according to claim 1, wherein distributing the request comprises distributing among the multiple processing units an identity of an ingress port over which the request was received at the network element, and

[4b] wherein updating the list comprises causing the list to indicate that the given multicast packet stream is permitted for forwarding over the ingress port.

Claim 5

[5] The method according to claim 1, wherein the ports forming the high-capacity port are assigned to at least two of the processing units, and wherein updating the list comprises holding respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, updating the local copies based on the distributed request, and selectively updating the list in each of the at least two processing units based on the respective local copies.

Claim 6

[6] The method according to claim 1, wherein the request specifies a requested action selected from a group of actions consisting of receiving the given multicast packet stream and stopping to receive the given multicast packet stream.

Claim 7

[7] The method according to claim 1, wherein the request conforms to one of an Internet Group Management Protocol (IGMP) and a Multicast Listener Discovery (MLD) protocol.

Claim 8

[8pre] A method for communication, comprising:

[8a] operating a network element,

[8b] wherein the network element further comprises multiple interconnected processing units and multiple ports,

[8c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

[8d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned;

[8e] storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element, and further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to;

[8f] receiving from a client a request relating to reception of a given multicast packet stream;

[8g] distributing the request among the multiple processing units within the network element, and

[8h] updating the respective list in each processing unit responsively to the request; and

[8i] selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists,

[8j] wherein operating the network element comprises initializing a new processing unit added to the network element by:

[8k] deactivating the ports to which the new processing unit is assigned;

[8l] while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request, receiving multicast membership

reports sent by multicast clients in response to the multicast status request, constructing the list by the new processing unit based on the received multicast membership reports; and

[8m] after constructing the list, activating the ports to which the new processing unit is assigned.

Claim 9

[9] The method according to claim 8, wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over a given port only when a number of the multicast packet streams that are permitted for forwarding over the given port, including the given multicast packet stream, does not exceed a maximum value specified for the given port.

Claim 10

[10] The method according to claim 8, wherein updating the list comprises tracking a use of a permission to forward a certain multicast packet stream over a certain port by monitoring requests relating to the certain multicast packet stream that are received at the network element, and removing the permission from the list responsively to determining that the permission is aged.

Claim 11

[11pre] A network element, comprising;

[11a] multiple ports; and

[11b] multiple interconnected processing units,

[11c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

[11d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned,

[11e] wherein the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,

[11f] to store respective lists, each of which indicates all of one or more multicast packet streams that are permitted for forwarding by the network element and further indicates which of the one or more ports multicast packet stream is permitted to be forwarded to,

[11g] to receive from a client a request relating to reception of a given multicast packet stream, and

[11h] to distribute the request among the processing units within the network element so as to update the lists responsively to the distributed request and to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists,

[11i] wherein two or more of the ports are aggregated to form a high-capacity port, and

[11j] wherein an upper limit is set on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port,

[11k] wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and

[11l] wherein the processing units are configured to update the lists to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

Claim 12

[12] The network element according to claim 11, wherein the processing units comprise dedicated hardware logic for storing the lists, distributing the request and selectively forwarding the multicast packets.

Claim 14

[14] The network element according to claim 11, wherein a processing unit receiving the request is configured to distribute among the multiple processing units an identity of an ingress port over which the request was received, and wherein the processing units are configured to cause the respective lists to indicate that the given multicast packet stream is permitted for forwarding over the ingress port.

Claim 15

[15] The network element according to claim 11, wherein the ports forming the high-capacity port are assigned to at least two of the processing units, and wherein the at least two processing units are configured to hold respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, to update the local copies based on the distributed request, and to selectively update their respective lists based on the respective local copies.

Claim 16

[16a] The network element according to claim 11, wherein the processing units are configured to track a use of a permission to forward a certain multicast packet stream over a certain port by monitoring requests relating to the certain multicast packet stream that are received at the network element, and

[16b] to remove the permission from the respective lists responsively to determining that the permission is aged.

Claim 17

[17] The network element according to claim 11, wherein the request conforms to one of an Internet Group Management Protocol (IGMP) and a Multicast Listener Discovery (MLD) protocol.

Claim 18

[18pre] A network element, comprising;

[18a] multiple ports; and

[18b] multiple interconnected processing units,

[18c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

[18d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned,

[18e] wherein the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,

[18f] to store respective lists, each of which indicates all of one or more multicast packet streams that are permitted for forwarding by the network element and further indicates which of the one or more ports multicast packet stream is permitted to be forwarded to,

[18g] to receive from a client a request relating to reception of a given multicast packet stream, and

[18h] to distribute the request among the processing units within the network element so as to update the lists responsively to the distributed request and to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists,

[18i] wherein a new processing unit added to the network element is configured to initialize by:

[18j] deactivating the ports to which the new processing unit is assigned;

[18k] while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request, receiving multicast membership reports sent by multicast clients in response to the multicast status request, constructing the list by the new processing unit based on the received multicast membership reports; and

[181] after constructing the list, activating the ports to which the new processing unit is assigned.

I. INTRODUCTION

U.S. Patent No. 8,014,394 (“’394 Patent”) uses well-known and routine operations for delivering multicast content to clients. Although the ’394 Patent purports to claim a novel way of providing multicast streams using a distributed router with a high-capacity port, the claimed subject matter was obvious at the time the ’394 Patent was filed.

For example, claim 1 of the ’394 Patent recites “[a] method” that includes “operating a network element ... [that] comprises multiple interconnected processing units ...,” “distributing the request [for a multicast packet stream] among the multiple processing units ... and updating the respective list,” “aggregating two or more ports to form a high-capacity port,” and “modifying the list to indicate that ... [the multicast packet stream] is permitted for forwarding over the high-capacity port only when the number of [multicast packet streams] for forwarding over the high-capacity port ... does not exceed the upper limit.”

None of these claim requirements was new. U.S. 2005/0041665 to Weyman discloses operating a distributed router that includes multiple interconnected units, distributing a request for a multicast stream among the multiple units and updating the respective lists, and aggregating ports to form a high-capacity port. W.O. 2007/140702 to Li, on the other hand, discloses limiting the number of permitted multicast streams on ports of a multicast router. It would have been

obvious to combine Weyman’s distributed router with Li’s port bandwidth control method, as detailed below.

Accordingly, Cisco Systems, Inc. (“Petitioner”) requests cancellation of claims challenged in this Petition.

II. MANDATORY NOTICES

A. Real Party-in-Interest (37 C.F.R. § 42.8(b)(1))

Petitioner certifies that the real party-in-interest is Cisco Systems, Inc.

B. Related Matters (37 C.F.R. § 42.8(b)(2))

On September 26, 2023, Portsmouth Network, Corp. (“Patent Owner”) filed a complaint against Cisco Systems, Inc. alleging infringement of U.S. Patent Nos. 6,917,986, 7,035,279, 7,596,088, and 8,014,394. *Portsmouth Network, Corp. v. Cisco Sys., Inc.*, Case No. 2:23-cv-00441 (E.D. Tex).¹ Petitioner has filed petitions for *inter partes* review of U.S. Patent Nos. 6,917,986 and 7,035,279.

Patent No.	IPR Case No.
U.S. 6,917,986	IPR2024-00505
U.S. 7,035,279	IPR2024-00503

¹ Patent Owner has since dropped infringement allegations based on U.S. Patent No. 7,596,088, and added infringement allegations based on U.S. Patent No. 8,199,637.

C. Lead and Back-up Counsel (37 C.F.R. § 42.8(b)(3))

Petitioner is filing a power of attorney appointing the practitioners associated with Customer Number 132,593. Petitioner designates the following lead and first backup counsel:

Lead Counsel

Taeg Sang Cho
Desmarais LLP
230 Park Ave
New York, NY 10023

Phone: (212) 351-3400
tcho@desmaraisllp.com
USPTO Reg. No. 69,618

First Back-up Counsel

Jonathan V. Lewis
Desmarais LLP
1899 Pennsylvania Ave, NW
Washington, DC 20006

Phone: (202) 451-4900
jlewis@desmaraisllp.com
USPTO Reg. No. 78,394

D. Service Information (37 C.F.R. § 42.8(b)(4))

Post and hand delivery: Desmarais LLP, 230 Park Avenue, New York, NY 10169

Telephone: 212-351-3400

Email: CiscoPortsmouthIPRService@desmaraisllp.com

Please address all correspondence to lead and back-up counsel. Petitioner consents to service in this proceeding by email at the addresses above and at

CiscoPortsmouthIPRService@desmaraisllp.com.

III. FEES

Petitioner concurrently submits required fees for this Petition. The Board is authorized to charge Desmarais LLP's deposit account, No. 50-6822, for any fee deficiency.

IV. GROUNDS FOR STANDING

Petitioner certifies that the '394 Patent is available for *inter partes* review and that Petitioner is not estopped or barred from requesting *inter partes* review.

V. OVERVIEW OF CHALLENGE AND RELIEF REQUESTED

Under 37 C.F.R. §§ 42.22(a)(1) and 42.104(b)(1)-(2), Petitioner requests cancellation of claims 1-2, 4-12 and 14-18 ("Challenged Claims") of the '394 Patent.

A. Prior Art References

The '394 Patent was filed on February 11, 2009, and thus, the '394 Patent is subject to pre-AIA provisions of 35 U.S.C. §§ 102, 103. The following references are pertinent to the grounds of unpatentability provided below:

1. U.S. 2005/0041665 to Weyman *et al.* ("Weyman")(EX1005), published February 24, 2005, is prior art under 35 U.S.C. § 102(b).

2. W.O. 2007/140702 to Li *et al.* (“Li”)(EX1007), published December 13, 2007, is prior art under 35 U.S.C. § 102(b).²
3. U.S. 2007/0297406 to Rooholamini (“Rooholamini”)(EX1008), published December 27, 2007, is prior art under 35 U.S.C. § 102(b).
4. U.S. 6,370,142 to Pitcher *et al.* (“Pitcher”)(EX1009), issued April 9, 2002, is prior art under 35 U.S.C. § 102(b).
5. U.S. 2006/0224877 to Jain *et al.* (“Jain”)(EX1010), published October 5, 2006, is prior art under 35 U.S.C. § 102(b).
6. U.S. 5,999,530 to LeMaire *et al.* (“LeMaire”)(EX1011), issued December 7, 1999, is prior art under 35 U.S.C. § 102(b).
7. U.S. 2006/0146823 to Ding (“Ding”)(EX1012), published July 6, 2006, is prior art under 35 U.S.C. § 102(b).

B. Grounds For Challenge

² Li was originally published in Chinese. EX1006 includes the original publication, and EX1007—Exhibit B includes the certified English translation of EX1006. This Petition relies on the certified English translation of Li (referenced as EX1007 in short), and the page/line cites correspond to that of EX1007—Exhibit B.

Ground	Claims	Basis	References
#1	1, 4, 6-7, 11, 14, 16-17	103	Weyman, Li
#2	5, 15	103	Weyman, Li, Rooholamini
#3	16	103	Weyman, Li, Pitcher
#4	8, 18	103	Weyman, Jain, LeMaire
#5	9	103	Weyman, Jain, LeMaire, Li
#6	10	103	Weyman, Jain, LeMaire, Pitcher
#7	2, 12	103	Weyman, Li, Ding

This Petition, supported by the declaration of Dr. Houh (EX1002), demonstrates there is a reasonable likelihood that the Petitioner would prevail with respect to at least one challenged claim. 35 U.S.C. § 314(a). Petitioner respectfully requests institution.

VI. TECHNOLOGY BACKGROUND

A. Networking, Multicasting, and IGMP

At the time of the '394 Patent, it was well-known that communication networks deliver packets from a source to a client. EX1017, [0005]. Some packets are delivered in a “unicast stream”—e.g., from a single source to a single client. EX1017, [0005]. Some packets, however, are delivered in a “multicast stream”—from a single source to multiple clients. EX1017, [0005]. EX1002, ¶34.




FIG.1

EX1014, FIG. 1. EX1002, ¶35.

Multicast streams may be set up using the Internet Group Management Protocol (IGMP). EX1015; EX1001, 1:25-30 (background). A client can “join” a multicast stream by sending an “IGMP Membership Report” that includes the identification of the multicast stream; a client can “leave” a multicast stream by sending a different IGMP Membership Report or an “IGMP Leave Group” message. EX1015, 4-8. EX1002, ¶¶36-37.

The network typically includes a network device—such as a router, switch, or bridge—that can route or forward a multicast stream to clients using a forwarding table. EX1017, [0050]-[0052]. A network device can populate its forwarding table by “snooping” IGMP messages. For example, a switch can snoop (e.g., detect) an IGMP Membership Report on a port to determine which multicast streams should

be sent to that port, and add that information to the forwarding table. EX1017, [0066]. EX1002, ¶¶38-41.

B. Distributed Router

A router may be implemented as a distributed router that includes multiple modules 110, 120 connected by an interconnection fabric 140. EX1012, [0012], [0036]; EX1018, [0002].




FIG. 1 100

EX1012, FIG. 1. EX1002, ¶42.

C. Port Bandwidth Control


Typically, ports have limited bandwidth. EX1019, 1:25-2:36. If a port has to transfer data at a rate higher than its maximum bandwidth, the port may drop packets and cause communication failures. EX1020, [0003]. EX1002, ¶44.

To avoid overloading ports, network devices can implement port bandwidth control. For example, a network device can count the number of data streams (e.g., multicast streams) transmitted over a port, and if the port is already carrying the maximum allowable number of streams, the network device can deny the request to send another data stream over that port. EX1017, [0055]. EX1002, ¶¶45-46.

VII. OVERVIEW OF THE '394 PATENT

A. Purported Invention


The '394 Patent discloses a network element that delivers multicast streams. As shown in Figure 1 (below), the “network element 32 comprises multiple processing units, such as line cards 36,” which are “interconnected by an interconnection switch fabric 40.” EX1001, 6:23-25. “[E]ach line card 36 is assigned to process the packets communicated over a given subset of the ports.” EX1001, 6:25-28. EX1002, ¶47.



EX1001, FIG. 1. EX1002, ¶48.

Each line card includes a forwarding table, such as “a local filtering database [DB].” EX1001, 5:4-8. The line cards construct or update their forwarding tables “by monitoring [multicast] request messages.” EX1001, 7:23-26. Figure 2 (below) shows that when a line card receives a multicast request from a client (step 60), the line card forwards a copy of the request to every other line card in the network element (step 72). EX1001, 7:30-58. Each line card then “updates its respective database 52” according to the received request (step 76). EX1001, 7:63-65. EX1002, ¶49.

FIG. 2



EX1001, FIG. 2.

The '394 Patent discloses “the network element [may] allow[] only a limited number of multicast streams to be forwarded over any given port, so as not to exceed the maximum bandwidths of the ports.” EX1001, 8:36-39. The '394 Patent uses the same concept to control a high-capacity port that is formed by aggregating physical ports as “a Link Aggregation Group (LAG).” EX1001, 6:61-7:1, 9:64-10:14. EX1002, ¶¶50-55.

B. Prosecution History

The '394 Patent was filed on February 11, 2009. EX1001, Cover; EX1004, 263-272. On October 29, 2010, the Examiner issued a non-final rejection, noting that independent claims generally directed to a distributed router and a high-capacity

port was obvious. EX1004, 106-131. But the Examiner noted that some dependent claims that further required controlling the number of multicast streams over the high-capacity port were allowable, and that some other dependent claims that further required initialization of a newly added module were also allowable. EX1004, 106-131. After a round of exchanges, Applicant amended the corresponding independent claims to incorporate the allowable subject matter, and that led to the '394 Patent. EX1004, 160-179, 232-241. EX1002, ¶¶56-58.

As explained below, however, the allowable subject matter would have been obvious in view of prior art that the Examiner did not consider. EX1002, ¶59.

VIII. CLAIM CONSTRUCTION

For the grounds presented herein, no claim term requires express construction. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017). EX1002, ¶¶60-61.

IX. LEVEL OF ORDINARY SKILL IN THE ART

A POSITA would have had a bachelor's degree in computer science, computer engineering, electrical engineering, or an equivalent field, or approximately two years of experience working in the field of information technology and networking by February 11, 2009. Lack of professional experience can be substituted by additional education, and vice versa. EX1002, ¶¶62-63.

X. OVERVIEW OF PRIMARY PRIOR ART REFERENCE

A. Weyman

Weyman describes a “distributed router composed of individual units ... which appears to be a single network entity.” EX1005, Abstract, [0024]. As shown in Figure 5 (below), the distributed router includes multiple units 10, 20. EX1005, [0060], FIG. 5. EX1002, ¶64.





FIG.5

Each unit has: (1) user ports (e.g., P1-P4) facing other devices in the network (e.g., router, client); and (2) one or more fabric ports (e.g., F1) connected to the internal fabric link 30 that is also connected to other units within the distributed router. EX1005, [0029], [0036], [0064]. Each unit also includes a bridge (e.g.,

Vbridge 21), the bridge's forwarding table (e.g., FDB (forwarding database) 42), a router (e.g., router 32), and the router's forwarding table (e.g., RIB (routing information base) 52). EX1005, [0064]. All units have the "identical" "router configuration" so that the units collectively "operate as a single distributed router." EX1005, [0069], [0100]. EX1002, ¶¶65-67.

Weyman's distributed router can handle multicast streams comprising packets. EX1005, [0091]-[0093]. As shown in Figure 11 (below), when unit 10 receives a multicast stream from a multicast source S, unit 10 uses Vbridge 11 (and its forwarding table, not shown) and router 31 (and its forwarding table, not shown) to route the stream to only the select user ports via which the stream can be forwarded to clients consuming the multicast stream. EX1005, [0103]-[0104], [0122]-[0124], [0127], FIG. 11. EX1002, ¶¶68-70.

**FIG.11**

For example, if a destination port is in the same unit, then unit 10 forwards the multicast stream directly to the local destination port. If the destination port is in a different unit, then unit 10 forwards the multicast stream to a fabric port, which then provides the stream to the fabric link, which then provides the stream to the other unit, which can subsequently forward the packet to its local destination port. EX1005, [0103]-[0104], [0122]-[0124], [0127]. *See also* EX1005, [0065]-[0067]. EX1002, ¶¶70-72.

Such selective forwarding is enabled by Vbridge's forwarding table and router's forwarding table that collectively identify the user port and/or fabric port to which the stream should be forwarded. EX1005, [0112] (A router forwarding table

includes, for each multicast stream, “the router interface on which to accept the stream and the router interfaces to which to forward the stream.”); [0115] (A Vbridge forwarding table includes “for each Vbridge and for each multicast group within each Vbridge[,], the ports to which to forward traffic for that group.”). EX1002, ¶73.

Weyman explains that the unit’s forwarding tables may be built using IGMP messages. EX1005, [0130], [0138]. For example, Weyman teaches learning the router’s forwarding table from IGMP messages. When a client C in Figure 11 (above) wants to join a multicast stream, the client “send[s] an IGMP host report” to unit 20. EX1005, [0137], [0162]. Unit 20 then “forward[s] this report to the unit’s [internal] multicast router [32] and also the fabric port” to send the report to the router 31 in unit 10. EX1005, [0137], [0162]. Then “the router[s] in *both units* would snoop the IGMP report and would use the report in the construction of the multicast forwarding tables.” EX1005, [0139]. This way, all router forwarding tables within the distributed router share the same entries. EX1005, [0065], [0162]. Similarly, Weyman teaches learning the Vbridge forwarding table from “MAC-addressed traffic,” which includes frames encapsulating IGMP messages. EX1005, [0045]-[0046]. EX1002, ¶¶74-77.

Weyman also explains that a unit’s ports can be aggregated to form an aggregated link to provide “resilience in a distributed router.” EX1005, [0050], [0086]. An aggregated link “comprises a multiplicity of physical connections [e.g.,

ports] constituting a single logical connection,” as shown in Figure 6 (below).

EX1005, [0050], [0088]-[0089]. EX1002, ¶¶78-82.




FIG.6

XI. SPECIFIC GROUNDS OF INVALIDITY

A. Ground 1: Weyman and Li Render Obvious Claims 1, 4, 6-7, 11, 14, and 17.


1. Weyman In View Of Li

As explained in Section X.A, Weyman discloses a distributed router with multiple units, each of which can selectively send and/or receive multicast streams using its own forwarding table. Weyman’s distributed router can also forward a multicast stream over an aggregated link that “comprises a multiplicity of physical connections constituting a single logical connection.” EX1005, [0050], [0148]. EX1002, ¶83.

**i. Controlling Number of Multicast Streams For a Port
In View of Li**

As explained in Section VI.C, it was understood that controlling the amount of data being transferred over a particular interface (such as a port or an aggregated link) was important for reliability. Indeed, Weyman recognizes that “[w]hen several individual data switching units are connected to form a fabric [e.g., a distributed router], ... *sharing of the workload* between the units” is one of the “[a]ncillary requirements.” EX1005, [0024]; [0050]. Accordingly, a POSITA would have been motivated to share the workload between units to control the bandwidth of ports, including any aggregated links. EX1002, ¶¶84-86.

A POSITA would have understood that Li discloses a port bandwidth control technique. Li discloses an IGMP-based multicast system that includes an access device. EX1007, 1:20-2:11. “[A]n access device 91 of a multicast service is a Digital Subscriber Line Access Multiplexer (DSLAM) functioning as a multicast router, and provides the user with a Digital Subscriber Line (xDSL) service port,” as shown in Figure 9 (below). EX1007, 26:13-15. EX1002, ¶87.



EX1007, FIG. 9. EX1002, ¶87.

When a client—such as a set-top box (STB)—sends a multicast join request to the multicast router 91, the multicast router 91 determines “whether a multicast channel requesting to join is permitted to join at a current XDSL port.” EX1007, 27:10-13. EX1002, ¶88.

In deciding whether to permit a client to join a multicast channel, the multicast router 91 looks to the maximum number of multicast channels allowable on its service port. Specifically:

“[E]ach xDSL service port is configured with the *maximum number of accessed multicast channels allowable*, which can be set dependent upon a *bandwidth of the xDSL port or as required for management*. It is determined whether the current xDSL service port allows joining the new channel, that is, *whether the number of channels being forwarded over the port is below the maximum number of accessed channels allowable, which is set for the port*, and if so, the port allows an access to the new channel for joining; otherwise, the new channel cannot be added.”

EX1007, 28:23-29:1. *See also* EX1007, 29:13-19, 30:25-31:3. EX1002, ¶89.

Such port control is performed by the “channel joining processing unit 122” in the multicast router, illustrated in Figure 12 (below). EX1007, 30:8-9. When the multicast router receives a multicast join request, the channel-number comparison unit 126 in the channel joining processing unit 122 “determine[s] whether the number of channels being forwarded at the current service port is below the maximum number of accessed channels allowable, which is set for the port; and, if so, the channel-number comparison unit instructs the multicast forwarding unit to start forwarding data of the new channel; otherwise, it is preliminarily determined that the channel is rejected.” EX1007, 31:4-12. When the channel-number comparison unit 126 preliminarily determines that the multicast join request should be rejected, the channel joining processing unit 122 checks to see if any of the port’s existing multicast streams is being terminated, and if not, the channel joining

processing unit 122 rejects the multicast join request. EX1007, 31:13-19. EX1002, ¶¶90-92.





FIG. 12

EX1007, FIG. 12.

A POSITA would have found it obvious to combine Weyman's distributed router with Li's channel joining processing unit to provide Li's channel joining processing unit in each of Weyman's units (as shown below). The Weyman-Li combination ("Weyman/Li") would have used the channel joining processing unit to control the workload of individual ports in each unit as well as any aggregated link in order not to overload the ports/aggregated link. EX1002, ¶¶93-94.

Weyman in combination with Li



A POSITA would have been motivated to incorporate Li's channel joining processing unit into Weyman's distributed router so that Weyman's distributed router does not overload a particular port and degrade the router's reliability, which was a known problem in the art. *See* Section VI.C [Bandwidth Control]. *See also* EX1020, [0003], [0006]. Indeed, Li expressly teaches that "the maximum number of multicast channels that can be selected of a service port *needs to be controlled*." EX1007, 5:5-23. EX1002, ¶95.

Also, in making the combination, a POSITA would have been motivated to incorporate Li's channel joining processing unit into individual units within Weyman's distributed router, as shown above, for two reasons. **First**, Weyman

teaches processing multicast join requests (IGMP Report) locally at the unit (EX1005, [0137]-[0139], [0162]), so a POSITA would have been motivated to control the assignment of multicast streams to ports locally at the unit as well. Further, a POSITA would have understood that the local control of port's workload is desirable because the alternative option—the global control of port's workload—would have required the use of internal fabric link to communicate with all units in the distributed router, and Weyman teaches that “a multicast router needs to *minimise* its use of the fabric link or links.” EX1005, [0119]. *See also* EX1005, [0076], [0098]. EX1002, ¶¶96-97.

Additionally, a POSITA would have been motivated to use Li's channel joining processing unit to control the workload of Weyman's aggregated link because: (1) a POSITA would have understood that the bandwidth control is desirable for both individual ports and aggregated links because even aggregated links have a limited bandwidth; and (2) Li's channel joining processing unit is equally applicable to both individual ports and aggregated links. EX1002, ¶¶98, 100.

A POSITA would have had a reasonable expectation of success in incorporating Li's channel joining processing unit into Weyman's unit. Both Weyman and Li are directed to an IGMP based multicast router, and the ports (and any aggregated link) in Weyman and the service port in Li are both carrying

multicast streams. EX1005, [0104]-[0105]; EX1007, 33:16-35:23. Furthermore, Li's channel joining processing unit is configured to control the number of multicast streams on a particular port based on a straightforward determination of the number of multicast streams being carried by the particular port. Accordingly, a POSITA would have understood that Li's channel joining processing unit would have provided a predictable result of limiting the number of multicast streams carried on Weyman's ports and aggregated links. EX1002, ¶¶99-100.

As described above, the combination of Weyman's distributed router with Li's channel joining processing unit would have merely amounted to combining prior art elements (e.g., Weyman's units and Li's channel joining processing unit) according to known methods (e.g., implementing logical operations in each of Weyman's units) to yield predictable results (e.g., limiting the number of multicast streams carried on Weyman's ports to avoid overloading Weyman's ports). EX1002, ¶101.


Also, the combination of Weyman's distributed router with Li's channel joining processing unit would have merely amounted to the use of known technique (e.g., Li's channel-number comparison technique) to improve similar devices (e.g., Weyman's units in the distributed router) in the same way (e.g., limiting the number of multicast streams carried on Weyman's ports and aggregated link to avoid overloading Weyman's ports). EX1002, ¶¶101-102.

At a minimum, it would have been obvious to try Li's channel joining processing unit in Weyman's distributed router. As explained in Section VI.C [Port Bandwidth Control], there was a design or market need to control a port's bandwidth, and there were only a finite number of ways to control the port's bandwidth (such as measuring the port's bandwidth real-time and throttling the bandwidth, or limiting the number of multicast streams assigned to the port). EX1013, [0036]; EX1021, [0225]; EX1022, [0024]; EX1019, 2:37-3:14. Each of those solutions would have had a reasonable expectation of success in Weyman's router because those techniques would have been implemented as a conventional logic operation in Weyman's units, yielding predictable results of controlling the port's bandwidth. EX1002, ¶103.


2. Independent Claim 1

i. [1pre] A method for communication, comprising:

To the extent limiting, Weyman/Li teaches the preamble. For example, Weyman discloses a "distributed router composed of individual units connected by links in a routing fabric." EX1005, Abstract. Figure 11 of Weyman illustrates how the first unit 10 in the distributed router provide a multicast stream to a user port in the second unit 20. EX1005, [0136]-[0138]. EX1002, ¶¶104-105.

Weyman in combination with Li**ii. [1a] operating a network element,**

Weyman/Li teaches limitation [1a]. Weyman teaches operating a distributed router (“network element”) to provide a multicast stream from a multicast source to a multicast client. EX1005, Abstract. *See also* EX1005, [0136]-[0137], [0148], [0151]-[0153]. Weyman discloses that operating the distributed router involves executing “routing protocols” to forward a multicast stream between the source and clients. EX1005, [0006], [0093], [0107]. EX1002, ¶¶106-107.


Weyman in combination with Li

- iii. [1b] wherein the network element further comprises multiple interconnected processing units and multiple ports,


Weyman/Li teaches limitation [1b]. Weyman discloses that the distributed router (“network element”) includes multiple units that are interconnected via a fabric link. EX1005, [0122] (“FIG. 10 represents one embodiment of a distributed multicast router according to the invention and demonstrates the forwarding of a particular multicast stream through a distributed fabric consisting in this example of the two units 10 and 20 and link 30.”). Each unit is a “processing unit” because the unit processes IGMP messages and multicast packet streams to forward the streams

to their destinations. EX1005, [0122]-[0124], [0162] (“the IGMP message could be delivered to every unit in the fabric for processing”). Each unit is also an “interconnected ... unit” because the unit is interconnected to other units via a fabric link. See EX1005, [0024] (“several individual data switching units are connected to form a fabric”), [0025] (“several units are connected through their fabric ports (i.e. ports connected only to other units within the fabric and not to the external network)”), [0122]-[0124], [0140], [0151], [0158] (“fabric interconnect”). EX1002, ¶¶108-110.

Weyman in combination with Li




Additionally, Weyman discloses that the distributed router has “multiple ports.” For instance, Weyman describes that each unit in the distributed router includes user ports. EX1005, [0029] (“unit 10 has three ‘Vbridges’ 11, 12 and 13 ... [and] [e]ach Vbridge 11, 12 and 13 is connected to at least one and usually more than one of a multiplicity of user ports 15.”), [0060], [0062] (“user ports P1 to P24”), [0064], [0122]-[0125]. Indeed, Weyman’s Figure 10 shows that unit 10 includes a first set of user ports, and unit 20 includes a second set of user ports. EX1002, ¶¶111-112.



- iv. [1c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

Weyman/Li teaches limitation [1c]. As explained in Section XI.A.2.iii [limitation 1[b]], unit 10 includes a first set of user ports, and unit 20 includes a

second set of user ports. Accordingly, each of unit 10 and unit 20 (“each processing unit”) is assigned to a different subset of user ports (“the multiple ports”) of the distributed router (“the network element.”). *See also* EX1005, [0064]; FIG. 8. EX1002, ¶¶113-114.



- v. [1d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned;


Weyman/Li teaches or renders obvious limitation [1d]. EX1002, ¶115.

First, each unit (“each processing unit”) includes one or more Vbridges and a router that “process[] [multicast streams] communicated over the different subset of the multiple ports to which the processing unit is assigned.” For example, when a unit (such as unit 10 or unit 20) receives a multicast stream communicated over

its user ports (a “subset of the multiple ports to which the processing unit is assigned”), the Vbridges and the router in the unit determine the egress port(s) for the multicast stream and forwards the multicast stream to the egress port(s).

EX1005, [0122]-[0124]; [0104]. Similarly, when a unit (such as unit 10 or unit 20) receives a multicast stream from its fabric port, the Vbridges and the router in the unit can determine the egress port(s) for the multicast stream and forwards the multicast stream to the egress port(s) (a “subset of the multiple ports to which the processing unit is assigned”). EX1005, [0124], [0104]. *See also* EX1005, [0025]. EX1002, ¶¶116-117.

Weyman in combination with Li



Second, the multicast streams processed by the Vbridges and a router of the units comprise packets. EX1005, [0093] (“each unit within the fabric would be able to route *multicast packets that it received*”), claim 1 (“distributed router for multicast packets”), [0127]. EX1002, ¶118.

Third, a POSITA would have understood that the one or more Vbridges and the router collectively form a “processor” because they “process” the packets in the multicast streams to forward them to the appropriate egress port(s). EX1005, [0097]-[0098], [0124]-[0125]. *See also* EX1017, [0040]. Indeed, it was well known that a unit (e.g., a line card) within a distributed router includes a processor, such as a CPU. *See* EX1008, [0042] (“The line card(s) 820 may include ... CPU(s) 828”), FIG. 8; EX1017, [0040]. EX1002, ¶119.


To the extent Patent Owner contends that the one or more Vbridges and the router in each unit are not a “processor,” it would have been obvious to implement the one or more Vbridges and the router using a processor. A POSITA would have understood that a processor would be programmed to perform the functions of the Vbridges and the router in Weyman—forwarding packets to destination ports. *See* EX1008, [0042], FIG. 8; EX1017, [0040]. It would have been obvious to combine Weyman/Li with a processor to implement the one or more Vbridges and the router using a processor because using a processor (instead of other processing mechanisms) is merely a design choice with the expected result of processing

packets. Further, implementing the one or more Vbridges and the router using a processor would have merely amounted to a simple substitution of one known element (e.g., Weyman's Vbridges and router) for another (e.g., processor) to obtain predictable results of processing packets. EX1002, ¶¶120-121.

- vi. **[1e] storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element, and further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to;**

Weyman/Li teaches limitation [1e]. EX1002, ¶122.

First, Weyman discloses “storing in each processing unit a respective list.” As described in Section XI.A.1. [Weyman/Li], each unit in Weyman/Li includes a Vbridge forwarding table and a router forwarding table. Those forwarding tables collectively include rules for forwarding multicast streams of the entire distributed router, and are thus the claimed “respective list.” See EX1005, [0031], [0046], [0069]-[0070]. EX1002, ¶123.

Weyman in combination with Li

Specifically, Weyman explains that the unit has a vbridge forwarding table and a router forwarding table (EX1005, [0105], [0115]), and “the **bridge multicast forwarding tables and the router multicast forwarding tables** should be used to control the forwarding of multicast traffic to both user ports and fabric ports” (EX1005, [0128]). A POSITA would have understood the Vbridge forwarding table and the router forwarding table to operate collectively as a single forwarding table because Weyman teaches that each unit “combine[s]” the Vbridge forwarding table and the router forwarding table to decide whether to forward a multicast stream to a particular port. EX1005, [0127] (“**the unit must combine the router’s forwarding**

tables and the bridge forwarding tables and only forward to a particular port if both tables indicate that the packet should be so forwarded.”). EX1002, ¶¶124-125.

A POSITA would have understood that the forwarding tables in the Vbridges and router have the structure of a “list” because Figure 2 of Weyman (below) provides the content of a Vbridge forwarding table as a list. Also, when referring to forwarding rules, the POSITA would have recognized that a table includes a list. EX1017, [0051]; EX1009, 6:59-64. *See* EX1008, [0030], FIG. 4. EX1002, ¶¶126-127.




FIG.2

EX1005, FIG. 2. EX1002, ¶¶126-127.

Second, Weyman/Li teaches “each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element.” For

example, Weyman teaches that the forwarding tables (“list”) “will provide, for each multicast source and destination group pair, the rule by which the router’s data plane must forward or filter the stream.” EX1005, [0111]. Each unit forwards a multicast packet stream to a port only if the forwarding tables “indicate that the packet should be so forwarded.” EX1005, [0127]. EX1002, ¶128.

Weyman also teaches that the forwarding tables maintain forwarding rules for the *entire* distributed router (*i.e.*, for “all of one or more multicast packet streams [for] the network element”). For example, Weyman teaches that the forwarding tables in each unit are “synchronised.” EX1005, [0046] (“The resultant forwarding databases [for Vbridges] in *each unit* would contain entries that are effectively ‘synchronised’ with each other to *allow traffic to be forwarded through the fabric as if the fabric were a single unit.*”), [0065] (“the forwarding databases on the units are *synchronized.*”), [0100] (“the multicast router will also have synchronised router interfaces [router forwarding table]”), [0162]. EX1002, ¶¶129-130.

Third, Weyman/Li teaches “indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to.” Weyman discloses that “the bridge multicast forwarding tables and the router multicast forwarding tables ... control the forwarding of *multicast traffic* to both *user ports* and fabric ports.” EX1005, [0128]. *See also* EX1005, [0131]. EX1002, ¶131.

Weyman illustrates the process in Figures 9-10. When Vbridge 11 of unit 10 receives a multicast stream, Vbridge 11 looks up the Vbridge forwarding table to determine the destination of the multicast stream. In this case, the Vbridge forwarding table of Vbridge 11 provides that the multicast stream should be provided to three destinations: (1) a local port connected to the client C1; (2) an internal router 31; and (3) a fabric port 16 connected to unit 20 that serves destination clients. EX1005, [0105], [0115], [0128], FIG. 10. Thus, Vbridge 11 forwards the multicast stream to the local port, the router 31, and the fabric port 16 towards unit 20. EX1005, [0105], [0122]-[0124]. EX1002, ¶132.




FIG. 10

EX1005, FIG. 10. EX1002, ¶132.


When the router 31 receives the multicast stream, the router 31 uses the router forwarding table to determine that Vbridge 12, connected via a router interface, will subsequently handle the multicast stream, and thus sends the stream to Vbridge 12. EX1005, [0112], [0123]. The Vbridge 12 then uses its Vbridge forwarding table to route the multicast stream to client C2. EX1005, [0105], [0123]. EX1002, ¶133.

Similarly, when unit 20 receives the multicast stream from unit 10 via the fabric port 16, the Vbridges and router in unit 20 use their forwarding tables to forward the multicast stream to the appropriate user ports. EX1005, [0124], [0131]. EX1002, ¶134.

As explained, each unit uses the Vbridge forwarding table(s) and the router forwarding table (collectively, the “list”) to “forward a multicast stream only to the ports that require it.” EX1005, [0127]. Thus, the “list” indicates “which of the one or more ports each multicast packet stream is permitted to be forwarded to.” EX1002, ¶135.

vii. [1f] receiving from a client a request relating to reception of a given multicast packet stream;

Weyman/Li teaches limitation [1f]. Weyman explains that “[t]he IGMP protocol enables multicast clients on a sub-network to register with their local routers in order to receive IP multicast streams.” EX1005, [0160]. For example, Weyman teaches that “[i]n order to join the destination group for the stream, the client would need to send an IGMP host report” (EX1005, [0137]), as illustrated in Figure 11 (below). EX1002, ¶¶136-137.

**FIG.11**

EX1005, FIG. 11. EX1002, ¶137.

viii. [1g] distributing the request among the multiple processing units within the network element, and

Weyman/Li teaches limitation [1g]. Weyman teaches that when Vbridge 22 of unit 20 receives an IGMP host report (“request”), “Vbridge 22 in device 20 would forward this report to the unit’s multicast router and also the fabric port. Vbridge 12 in unit 10 would then receive the report and forward it to the router

31,” as shown in Figure 11 (below). EX1005, [0137]. EX1002, ¶138.




FIG.11

EX1005, FIG. 11. EX1002, ¶139.

Although Figure 11 shows forwarding the IGMP host report to only one unit 10 because Figure 11 shows only two units, Weyman expressly teaches distributing the IGMP host report on all other units so that “[a]ny changes to the multicast distribution tree [is] kept synchronised for *all the units in the fabric.*” EX1005, [0162] (“the IGMP message could be delivered to *every unit* in the fabric for processing.”). EX1002, ¶140.

ix. [1h] updating the respective list in each processing unit responsively to the request; and

Weyman/Li teaches limitation [1h]. Weyman explains that once the IGMP host report is distributed to each unit, each unit “would use the report in the *construction of the respective forwarding table* in the router. As a result, the forwarding table of each multicast router would include the specification that the stream should be replicated and routed from VLAN 1 to VLAN 2.” EX1005, [0138]. Also, Weyman explains that the forwarding rules “must be kept synchronised for all the units in the fabric” in response to an IGMP Report (EX1005, [0162]), teaching that the forwarding tables are updated in response to the IGMP Report. Accordingly, Weyman teaches that each unit updates its forwarding tables (“the respective list”) responsively to the IGMP Report. EX1002, ¶¶141-143.

x. [1i] selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists,

Weyman/Li teaches limitation [1i]. EX1002, ¶144

As discussed in Section XI.A.2.vi [1e], Weyman teaches “selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective [] lists”:

“A unit needs to configure its forwarding engine to forward a multicast stream only to the ports that require it. *It must not forward the stream*

to any other port. In order to achieve this end the unit must combine the router’s forwarding tables and the bridge forwarding tables and only forward to a particular port *if both tables indicate that the packet should be so forwarded.”*

EX1005, [0127]. EX1002, ¶145.

Furthermore, Weyman teaches that the “router’s forwarding tables and the bridge forwarding tables” is updated during operation based on the IGMP host report, as discussed in Section XI.A.2.xi [limitation 1[j]]. Thus, the “router’s forwarding tables and the bridge forwarding tables” is an “updated list.” EX1002, ¶146.

Accordingly, Weyman/Li teaches that the units “selectively forward[]” the multicast stream (which comprises packets) in accordance with the router’s forwarding tables and the bridge forwarding tables (“respective updated list”). EX1002, ¶147.

- xi. [1j] wherein operating the network element comprises aggregating two or more of the ports to form a high-capacity port, and**


Weyman/Li teaches limitation [1j]. EX1002, ¶148.

Weyman discloses that “operating the network element comprises aggregating two or more of the [physical connections] to form [an aggregated link].”

Weyman discloses that the distributed router can forward “a multicast stream ... to

an ***aggregated link***.” EX1005, [0148]. Weyman explains that “an aggregated link ... comprises a ***multiplicity of physical connections constituting a single logical connection***.” EX1005, [0050]. *See also* EX1005, [0050] (“the protocol considers the aggregated link to be a single logical link.”). EX1002, ¶¶148-149.

A POSITA would have understood that the physical connections aggregated in Weyman are physical “ports,” and that the aggregated link is a “high-capacity port.” As shown in Figure 6 (below), Weyman shows that one physical connection in unit 10 and another physical connection in unit 20 are aggregated to form an aggregated link with the switch 91. EX1005, [0013], [0088]. Weyman teaches that the aggregated link can be formed using “a link aggregation control protocol” or “LACP” (EX1005, [0025], [0088]), which was well-understood to bundle ***physical ports*** together (EX1023, [0002]). Accordingly, a POSITA would have understood that the aggregated link in Weyman bundles the port in unit 10 and the port in unit 20 together. Also, the aggregated link formed by LACP was understood to be a link aggregation group (LAG) (EX1024, [0028]-[0029]), which is precisely the “high-capacity port” disclosed by the ’394 Patent (EX1001, 6:61-7:1). *See also* EX1025, 5:49-53. EX1002, ¶¶150-152.

**FIG.6**

EX1005, FIG. 6. EX1002, ¶150.

- xii. [1k] wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port,

Weyman/Li teaches limitation [1k]. For example, Li teaches that “each xDSL service port is configured with the *maximum number of accessed multicast channels allowable*.” EX1007, 28:23-29:1. When the multicast router receives a multicast join request, the Weyman/Li’s channel joining processing unit determines whether the number of channels being forwarded at the port or aggregated link is below the maximum number of accessed channels allowable, and if not, the channel joining processing unit rejects the multicast join request. EX1007, 31:4-19. Accordingly, the distributed router in Weyman/Li sets an upper limit on a number

of multicast streams that are permitted for forwarding over Weyman's ports as well as the aggregated link ("high-capacity port"), as described in Section XI.A.1. [Weyman/Li]. EX1002, ¶¶153-155.

xiii. [11] wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and

Weyman/Li teaches limitation [11]. On its face, limitation [11] is ambiguous as to whether "the request" refers to the "request" recited in limitation [1f], or the "request" recited in limitation [1g]. But the '394 Patent discloses that "the distributed copy [of the request] comprises an indication of the ingress port, i.e., the port via which the request arrived at the network element." EX1001, 8:3-5. Thus, a POSITA would have understood the "request" to be a "distributed copy of the request." And Weyman discloses that the distributed copy of the IGMP report ("request") (the frame encapsulating the IGMP report sent via the fabric link) satisfies limitation [11]. EX1002, ¶¶156-158.

Specifically, Weyman teaches distributing a copy of the IGMP report to other units by encapsulating the IGMP report in a frame that also includes an identity of an ingress port (including the aggregated link ("high-capacity port")). For example, Weyman teaches that a fabric link communicates messages (such as an IGMP host report (EX1005, [0137], [0162])) using a "frame." EX1005, [0042]. The fabric link's frame includes "a special header which among other things identifies *an*

ingress port and an egress port ... so that the special header can identify uniquely within the fabric both a *source port and unit* and a destination port and unit.” EX1005, [0042]. *See also* EX1005, [0043]. Thus, when a unit distributes via the fabric link an IGMP host report received from an aggregated link (“high-capacity port”) (*see* Section XI.A.2.viii [1g], EX1005, [0137], [0162]), a POSITA would have understood that the frame carrying the IGMP host report (the “distributed copy of the request”) includes an identity of the high-capacity port over which the IGMP host report was received at the distributed router. EX1005, [0042]-[0043] (“the fabric header ... comprise[s] ... a source port identification.”). EX1002, ¶¶159-161.

Weyman uses such identification of “an ingress port”/“source port identification” (EX1005, [0042], [0043]) as “the port to which to forward any traffic destined for that MAC address” from which the IGMP host report was received (EX1005, [0041], [0137], [0162]). Accordingly, Weyman teaches that the frame carrying the IGMP host report (the “distributed copy of the request”) indicates that the multicast stream requested by the IGMP host report “is to be forwarded over the high-capacity port.” EX1002, ¶¶159-162.

- xiv. **[1m] wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.**

Weyman/Li teaches limitation [1m]. EX1002, ¶163.

As explained in Sections XI.A.2.xii [1k], XI.A.2.xiii [1l], and XI.A.1 [Weyman/Li], when a unit receives the frame including an IGMP host report, the channel joining processing unit in the unit determines (e.g., using the identity of the aggregated link in the header of the received frame) whether the number of multicast streams being forwarded at the aggregated link (“high-capacity port”) is below the maximum number of multicast channels for the aggregated link. EX1007, 31:4-12. If not, the channel joining processing unit rejects the IGMP host report. EX1007, 31:4-31:24. But if it is (i.e., “only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit”), the channel joining processing unit accepts the IGMP host report. EX1007, 30:21-24. EX1002, ¶¶164-165.

Then, the unit processes the IGMP host report (along with the identity of the aggregate link in the header of the received frame) to modify the unit’s forwarding tables (“list”) (EX1005, [0137]-[0139], [0162], Section XI.A.2.ix [1h]) to indicate that the requested multicast stream is permitted for forwarding over the port corresponding to the aggregated link (“high-capacity port”) (EX1005, [0127], [0128] (“the bridge multicast forwarding tables and the router multicast forwarding tables should be used to control the forwarding of multicast traffic to both *user ports*

and fabric ports”), [0148] (“Each replica of a multicast stream should be forwarded to an ***aggregated link***...”); Section XI.A.2.vi [1e]). *See also* EX1005, [0042] (teaching that the forwarding table identifies ports at which the packet (such as the packet comprising the IGMP message) was received). EX1002, ¶166.

3. Claim 4

- i. **[4a] The method according to claim 1, wherein distributing the request comprises distributing among the multiple processing units an identity of an ingress port over which the request was received at the network element, and**

Weyman/Li teaches the additional limitation of [4a]. As explained in Section XI.A.2.xiii [11], Weyman teaches distributing a copy of the IGMP report (“request”) to other units by encapsulating the IGMP report in a frame that also includes “an identity of an ingress port” over which the IGMP report was received at the network element. *See also* EX1005, [0041]-[0043], [0137]-[0138], [0162]. EX1002, ¶¶167-168.

- ii. **[4b] wherein updating the list comprises causing the list to indicate that the given multicast packet stream is permitted for forwarding over the ingress port.**

Weyman/Li teaches limitation [4b]. As explained in Section XI.A.2.xiii [11], when a unit receives via the fabric link a frame including the IGMP host report (along with the identity of the aggregate link in the header of the received frame), the unit modifies its forwarding tables (“list”) (EX1005, [0137]-[0139], [0162])

indicating that the requested multicast stream is “permitted for forwarding over the ingress port” (EX1005, [0042], [0127], [0128]). EX1002, ¶¶169-170.

Specifically, ***first***, Weyman teaches that the forwarding table in a unit may identify a user port (“ingress port”) from which a packet (e.g., a medium using which the IGMP message is communicated) was received, regardless of the unit to which the user port is assigned. EX1005, [0042] (“[I]t would not be necessary to reference the fabric ports in the forwarding databases. ***An identifier for the user port may be used instead***, this identifier being unique within the entire distributed fabric.”). EX1002, ¶171.

Second, as discussed in Section XI.A.2.xiii [1l], the frame carrying the IGMP host report would have included the unique identifier of the “ingress port.” EX1005, [0042]-[0043]. EX1002, ¶172.

Third, Weyman teaches using the IGMP host report to update the forwarding table in the unit. EX1005, [0138]-[0139], [0162]. As discussed in Sections XI.A.2.vi [1e] and XI.A.2.xiv [1m], the unique identifier for the “ingress port” in the frame of the IGMP host report is used to update the forwarding table in the unit to indicate that the multicast stream requested by the IGMP host report is permitted for forwarding over the ingress port, thereby satisfying limitation 4[b]. EX1002, ¶173.

4. Claim 6

The method according to claim 1, wherein the request specifies a requested action selected from a group of actions consisting of receiving the given multicast packet stream and stopping to receive the given multicast packet stream.

Weyman/Li teaches the additional limitation of claim 6. Weyman discloses that “[i]n order to join the destination group for the stream, the client would need to send an IGMP host report.” EX1005, [0137]. The IGMP host report is a “request” to receive the given multicast packet stream of the destination group. EX1002, ¶¶174-175.

5. Claim 7

The method according to claim 1, wherein the request conforms to one of an Internet Group Management Protocol (IGMP) and a Multicast Listener Discovery (MLD) protocol.

Weyman/Li teaches the additional limitation of claim 7. Weyman discloses that “[i]n order to join the destination group for the stream, the client would need to send an IGMP host report.” EX1005, [0137]. IGMP stands for “internet group management protocol.” EX1005, [0025]. EX1002, ¶¶176-177.


6. Independent Claim 11

i. [11pre] A network element, comprising:

To the extent limiting, Weyman/Li teaches the preamble for the same reasons as in Section XI.A.2.i. [element [1pre]]. EX1002, ¶178.


ii. [11a] multiple ports; and

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.iii [1b]. EX1002, ¶179.




iii. [11b] multiple interconnected processing units,

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.iii [1b]. EX1002, ¶180.



- iv. [11c] wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.iv [1c]. EX1002, ¶181.



- v. **[11d] wherein each processing unit contains a processor for processing packets communicated over the different subset of the multiple ports to which the processing unit is assigned,**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.v

[1d]. EX1002, ¶182.

- vi. **[11e] wherein the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.v

[1d]. EX1002, ¶183.

- vii. **[11f] to store respective lists, each of which [respective lists] indicates all of one or more multicast packet streams that are permitted for forwarding by the network element and further indicates which of the one or more ports multicast packet stream is permitted to be forwarded to**

Weyman/Li teaches this limitation for the same reasons as in Section

XI.A.2.vi [1e]. EX1002, ¶184.

- viii. **[11g] to receive from a client a request relating to reception of a given multicast packet stream, and**

Weyman/Li teaches this limitation for the same reasons as in Section

XI.A.2.vii [1f]. EX1002, ¶185.

- ix. **[11h] to distribute the request among the processing units within the network element so as to update the lists responsively to the distributed request and to selectively forward multicast packets associated with**

the given multicast packet stream in accordance with the updated lists,

Weyman/Li teaches this limitation for the same reasons as in Sections XI.A.2.viii [1g]-XI.A.2.x [1i]. EX1002, ¶186.

- x. [11i] wherein two or more of the ports are aggregated to form a high-capacity port, and**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.xi [1j]. EX1002, ¶187.

- xi. [11j] wherein an upper limit is set on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, and**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.xii [1k]. EX1002, ¶188.

- xii. [11k] wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.xiii [1l]. EX1002, ¶189.

- xiii. [11l] wherein the processing units are configured to update the lists to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.**

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.2.xiv [1m]. EX1002, ¶190.

7. Claim 14

The network element according to claim 11, wherein a processing unit receiving the request is configured to distribute among the multiple processing units an identity of an ingress port over which the request was received, and wherein the processing units are configured to cause the respective lists to indicate that the given multicast packet stream is permitted for forwarding over the ingress port.

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.3. [Claim 4]. EX1002, ¶191.

8. Claim 17

The network element according to claim 11, wherein the request conforms to one of an Internet Group Management Protocol (IGMP) and a Multicast Listener Discovery (MLD) protocol.

Weyman/Li teaches this limitation for the same reasons as in Section XI.A.5. [Claim 7]. EX1002, ¶192.

B. Ground 2: Weyman, Li, And Rooholamini Render Obvious Claims 5 and 15.

1. Weyman/Li in view of Rooholamini

As explained in Section XI.A.1. [Weyman/Li], the channel joining processing unit accepts or rejects a request for a multicast stream at a particular port based on the number of multicast streams being carried by the particular port. Accordingly,

a POSITA would have been motivated to keep a record of the number of multicast streams being carried by ports in the Weyman/Li's distributed router. EX1002, ¶193.

And that was well-known in the art. For example, Rooholamini teaches a switch fabric (a distributed router) that manages multicast streams by maintaining a "port path count table," which tracks the number of multicast streams that pass through a particular port. EX1008, [0029]-[0030]. Rooholamini provides an example of the port path count table in Figure 4 (below). The rows in the table correspond switches S1, ..., Sn; the columns correspond to ports P0 ... Pm within those switches; and the numbers correspond to the number of multicast streams handled by the particular port of the particular switch. EX1008, [0030]. Rooholamini teaches that the port path count table is updated during the operation of the router such that it maintains a real-time count of the number of multicast streams passing through a particular port of a particular switch. EX1008, [0029], [0035]. EX1002, ¶¶194-195.

	P_0	P_1	...	P_m
S_1	2	1	0	0
S_2	0	3	0	1
...	...	...	...	...
S_n	0	0	0	0


434

FIG. 4

EX1008, FIG. 4. EX1002, ¶194.

A POSITA would have found it obvious to combine Weyman/Li with Rooholamini such that the unit includes a port count table that identifies the number of streams handled by the ports in the distributed router (including the aggregated link). The combination is illustrated in the figure below. This combination would have allowed Li's channel joining processing unit to control the number of multicast streams for a particular port using the port count available locally in the unit.

EX1002, ¶196.

Weyman in combination with Li and Rooholamini

EX1002, ¶196.

A POSITA would have been motivated to combine Weyman/Li with Rooholamini as illustrated above because Li's channel joining processing unit needs to track and rely on the number of streams handled by a port. EX1007, 5:5-23 ("In the controlled multicast service, the maximum number of multicast channels that can be selected of a service port needs to be controlled...."). And indeed, Rooholamini also contemplates that the port count table is used to manage the multicast streams served by the distributed router's ports. See EX1008, [0029] ("By tracking the number of paths that go through the ingress and egress ports ..., the fabric manager 230 may determine if and when *the switch multicast tables 280-1 to 280-3 need to*

be configured or reconfigured when there is a change in multicast group membership.”). A POSITA would have been motivated to use the port count table to also keep track of the number of multicast streams handled by the aggregated link because that information would have allowed the channel joining processing unit to manage the bandwidth of the aggregated link. EX1002, ¶197.

A POSITA would have been motivated to provide the port count table locally in the unit in order to reduce the use of the fabric link as taught by Weyman and also to reduce the latency associated with pulling the port count information. EX1005, [0119] (“a multicast router needs to *minimise* its use of the fabric link or links.”). *See also* EX1005, [0076], [0098]. EX1002, ¶198.

A POSITA would have had a reasonable expectation of success in combining Weyman/Li’s distributed multicast router with Rooholamini’s count table because, as described above, because Weyman, Li, and Rooholamini all relate to routing multicast streams using a distributed multicast router (e.g., switch fabric). EX1005, [0004]; EX1007, 48:5-18; EX1008, [0004]. Accordingly, the combination would have merely amounted to combining prior art elements (e.g., Weyman/Li’s forwarding tables and Rooholamini’s count table) according to known methods (e.g., implementing the tables in memory on the unit) to yield predictable results (e.g., maintaining the count information in Weyman/Li’s forwarding table such that Li’s

channel-number comparison unit can use the table to determine whether to accept a multicast join request for a particular port.) EX1002, ¶199.

2. Claim 5

[5a] The method according to claim 1, wherein the ports forming the high-capacity port are assigned to at least two of the processing units, and

Weyman/Li/Rooholamini teaches the additional limitation of 5[a]. As explained in Section XI.A.2.xi [limitation 1[j]], the Weyman/Li router includes an aggregated link (“high-capacity port”) that aggregates one port in unit 10 and another port in a different unit 20. Accordingly, “the ports forming the high-capacity port are assigned to at least two of the processing units.” EX1002, ¶¶201-202.





FIG. 6

EX1005, FIG. 6. EX1002, ¶203.

[5b] wherein updating the list comprises holding respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, updating the local copies based on the distributed request, and selectively updating the list in each of the at least two processing units based on the respective local copies.

Weyman/Li/Rooholamini teaches limitation 5[b]. EX1002, ¶204.

First, Weyman/Li/Rooholamini teaches “holding respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port.” For example, as explained in Section XI.B.1. [Weyman/Li/Rooholamini], each unit holds a port count table, which is a copy local to the unit as shown below. The port count table includes the number of the multicast packet streams that are permitted for forwarding over the ports in the distributed router, including the aggregated link. EX1002, ¶¶205-207.

Weyman in combination with Li and Rooholamini

EX1002, ¶206.

Second, Weyman/Li/Rooholamini teaches “updating the local copies based on the distributed request.” Rooholamini teaches that “[f]or each of the multicast groups, port path count tables 234 *track* the number of paths through the ports of the switches configured to route multicast packets.” EX1008, [0029]. Thus, the port count table in each unit is updated whenever there is a change in the number of multicast streams carried by a particular port (such as the aggregated link) of the corresponding unit. And such change is triggered by the IGMP host report (“request”) when it is distributed and processed by each of the units in the distributed

router. *See* EX1005, [0137]-[0139], [0162] (“IGMP message could be delivered to every unit in the fabric for processing.”). EX1002, ¶¶208-209.

Third, Weyman/Li/Rooholamini teaches “selectively updating the list in each of the at least two processing units based on the respective local copies.” For example, in Weyman/Li/Rooholamini, the channel joining processing unit in each unit can use the updated local port count table to determine that the newly requested multicast stream is permitted for forwarding. EX1007, 28:23-29:1, 31:4-12. The unit then can selectively update its forwarding table to note that the requested multicast stream is permitted forwarding. EX1005, [0138]-[0139], [0162]. EX1002, ¶210.

3. Claim 15

The network element according to claim 11, wherein the ports forming the high-capacity port are assigned to at least two of the processing units, and wherein the at least two processing units are configured to hold respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, to update the local copies based on the distributed request, and to selectively update their respective lists based on the respective local copies.

Weyman/Li teaches this limitation for the same reasons as in Section XI.B.2.

[Claim 5]. EX1002, ¶211.

C. Ground 3: Weyman, Li, And Pitcher Render Obvious Claim 16.

1. Weyman/Li In View Of Pitcher

As discussed in Sections X.A. [Weyman] and XI.A.1. [Weyman/Li], each unit in Weyman/Li's distributed router includes forwarding tables. Those forwarding tables include, for each multicast stream, the ports (including any aggregated link) to which to forward the multicast stream. EX1005, [0112], [0115]. EX1002, ¶212.

A POSITA would have understood, however, that certain entries in the forwarding tables become stale, i.e., not used to route multicast streams, and that stale entries are undesirable because they unnecessarily consume memory space. EX1005, [0156] (teaching that the forwarding table is stored "in non-volatile storage"); EX1026, 5:64-6:4. Accordingly, a POSITA would have been motivated to identify stale entries in forwarding tables and remove them. EX1002, ¶213.

A POSITA would have understood that stale entries can be removed using a timer. For example, Pitcher teaches that each entry of the multicast group destination lists (forwarding tables) are associated with a port report expiration timer. EX1009, 7:56-62. When a switch (e.g., a unit in Weyman/Li's distributed router) receives an IGMP membership report (e.g., a multicast join request), the switch adds the "receiving port" to the appropriate forwarding table "if it is not already a current member" and resets the corresponding port report timer. EX1009, 7:62-66. And

“[i]f the port report timer expires, the port is removed from the multicast group destination list [forwarding table].” EX1009, 7:66-67. EX1002, ¶214.

A POSITA would have found it obvious to combine Weyman/Li’s unit with Pitcher’s timer-based entry removal technique to provide a “port report timer” in each entry of the forwarding table in each unit, so that the Weyman/Li’s unit can remove stale entries in the unit’s forwarding tables. A POSITA would have understood that stale entries cause inefficient use of memory in the unit, and would have been motivated to remove such stale entries. EX1026, 5:64-6:4. EX1002, ¶215.

A POSITA would have had a reasonable expectation of success in combining Weyman/Li’s unit with Pitcher’s timer-based entry removal technique because the combination would have merely amounted to using a known technique (Pitcher’s timer-based entry removal technique) to improve a similar device (Weyman/Li’s unit, which performs substantially similar function as a switch) in the same way (remove stale entries from the unit’s forwarding tables). EX1002, ¶216.

The combination of Weyman/Li’s unit and Pitcher’s timer-based entry removal technique would have entailed merely applying a known technique (Pitcher’s timer-based entry removal technique) to a known device (Weyman/Li’s unit) ready for improvement (the stale entries in forwarding tables unnecessarily

consuming memory space) to yield predictable results (removing stale entries from the unit's forwarding tables). EX1002, ¶217.

2. Claim 16

- i. **[16a] The network element according to claim 11, wherein the processing units are configured to track a use of a permission to forward a certain multicast packet stream over a certain port by monitoring requests relating to the certain multicast packet stream that are received at the network element, and**

Weyman/Li/Pitcher teaches the additional limitation of [16a]. As explained in Section XI.C.1 [Weyman/Li/Pitcher], each unit tracks the use of entries in the forwarding tables using a timer. EX1009, 7:56-62. Entries in the forwarding tables (e.g., the Vbridge forwarding table) indicates “a permission to forward a certain multicast packet stream over a certain port.” EX1005, [0112], [0115]. And the timer is updated by monitoring IGMP membership reports (“requests”) relating to the certain multicast stream that are received at the Weyman/Li/Pitcher's unit. EX1009, 7:62-67. EX1002, ¶218.

- ii. **[16b] to remove the permission from the respective lists responsively to determining that the permission is aged.**

Weyman/Li/Pitcher teaches limitation [16b]. For example, “[i]f the port report timer expires [i.e., determining that the permission is aged], the port is removed from the multicast group destination list [i.e., the permission is removed

from the respective forwarding tables].” EX1009, 7:66-67. Since each unit has its own port report timer for each entry of the forwarding table, the “processing units are configured ... to remove the permission from the respective lists” as required by limitation [16b]. EX1002, ¶¶219-220.

D. Ground 4: Weyman, Jain, And LeMaire Render Obvious Claims 8 And 18.

1. Weyman In View Of Jain And LeMaire

As explained in Section X.A [Weyman], Weyman discloses a distributed multicast router with multiple units, each of which can selectively forward multicast streams using its own forwarding tables. Weyman also discloses that “the system can accommodate the *addition* or removal of unit” in the distributed router. EX1005, [0069]. EX1002, ¶221.

Weyman explains that “[i]f a unit with a default configuration joins the fabric[,] it needs to update its configuration to that of the distributed multicast router.” EX1005, [0156]. But Weyman does not specifically explain how the new unit would “update its configuration to that of the distributed multicast router,” including the forwarding tables in the unit. Accordingly, a POSITA would have been motivated to look for ways to update the initial configuration of the newly added unit. As explained below, various features of the initialization process of a switch were well-known in the art, and it would have been obvious to combine those

features with the newly added unit in Weyman's distributed router. EX1002, ¶¶222-223.

i. Disabling Ports During Initialization In View Of Jain

Jain explains that when a node in a network (e.g., a switch) is added to the network after being down (e.g., “reset, power down, rebooting”), a node executes a “controlled boot-up process” to configure the node. EX1010, [0019]. As shown in Figure 2 (below), the controlled boot-up process involves first issuing a “Link Down” command on at least some of the node's ports to disable them. EX1010, [0020] (“a ‘Link Down’ condition is effectuated on at least a portion of the link ports in the network node. ... When the Link Down condition is sensed by the entities (i.e., devices and nodes) connected to the network node via respective links, network traffic on the links is *discontinued* (block 206).”). And then “[w]hile the Link Down (or, disabled) condition is maintained on the link ports, the system software of the network node continues with its initialization and configuration process until it is complete and operational.” EX1010, [0021]. “Thereafter, ... a ‘Link Up’ condition is effectuated on the ports.” EX1010, [0021]. EX1002, ¶¶224-226.




FIG. 2

EX1010, FIG. 2. EX1002, ¶225.

A POSITA would have been motivated to combine Weyman's distributed router with Jain's controlled boot-up process to disable or deactivate the user ports³

³ To be clear, a POSITA would not have disabled the fabric port on the newly added unit because Weyman specifically teaches that "[i]f a unit with a default configuration joins the fabric it *needs to update* its configuration to that of the distributed multicast router" (EX1005, [0156]), and a POSITA would have

(the ports providing multicast streams to the network) on the newly added unit of Weyman's distributed router during the initialization of the newly added unit. A POSITA would have been motivated to do so because when the newly added unit is not properly initialized, the "ports [may] exhibit spurious logic conditions that may cause the interfacing entities (i.e., remote devices as well as other network nodes) to commence network traffic while the switch software (i.e., system software) is still in an incomplete operational state." EX1010, [0019]. A POSITA would have understood that commencing communications without proper initialization would have caused "various deleterious effects with respect to the integrity of the traffic received at the network node." EX1010, [0019]. EX1002, ¶¶227-228.

A POSITA would have had a reasonable expectation of success in making the combination because it merely entails using a known technique (e.g., disabling user ports of the newly added network node during initialization, per Jain) to improve a similar device (e.g., a newly added unit in Weyman's distributed router) in the same way (e.g., temporarily disabling the ports so that the new network node does not inadvertently cause disruptions). EX1002, ¶229.

understood that the new unit needs to use the fabric port to synchronize the new unit's configuration to that of the distributed router (and other units connected to the fabric link). EX1002, ¶227.

ii. Initialization of New Unit In View Of LeMaire

LeMaire discloses a bridge for handling multicast streams. EX1011, Abstract. As shown in Figure 5 (below), the bridge 62 includes multiple ports for communicating with other devices in the network. EX1011, 7:64-8:2. The bridge 62 also includes a “Bridge Filtering Table 88” in which “[e]ach entry specifies the LAN address from which a packet is received and includes a means for indicating a port to which that address is connected.” EX1011, 8:8-15. EX1002, ¶¶230-231.




FIG. 5

EX1011, FIG. 5. EX1002, ¶231.

When the bridge is first added to the network (i.e., “power up”), the bridge transmits IGMP queries requesting the multicast membership status from clients coupled to the ports. EX1011, 9:12-19. “In response to the IGMP Queries, [clients] attached to the bridge 62 that wish to receive [multicast streams] will transmit IGMP Reports” to the bridge 62. EX1011, 9:21-23. The bridge then uses the reports to populate the bridge filtering table 88. EX1011, 9:23-34. EX1002, ¶¶231-232.

A POSITA would have found it obvious to combine Weyman/Jain’s distributed router with LeMaire’s bridge initialization technique such that when a new unit is added to the distributed router, the distributed router uses LeMaire’s initialization technique to update its forwarding table while the new unit’s user ports are disabled. For example, when a new unit is added to the distributed router, the distributed router can disable its user ports as discussed in Section XI.D.1.i. [Weyman/Jain], and it can then transmit IGMP queries to its clients, as taught by LeMaire, using other units in the distributed router. EX1011, 9:12-19. In response to the IGMP queries, the other units can receive IGMP reports from the clients and forward those IGMP reports to the new unit via the fabric link and the new unit’s fabric port, as taught by LeMaire and Weyman. EX1011, 9:21-23; EX1005, [0137]-[0139], [0162]. The new unit can then update its forwarding table in accordance with the received IGMP reports, as taught by both Weyman and LeMaire. EX1005,

[0138]-[0139]; EX1011, 9:23-34. Once such initialization process is completed, the unit activates the previously disabled user ports to initiate communication. EX1008, [0021], [0007]. EX1002, ¶¶233-235.

A POSITA would have been motivated to combine Weyman/Jain's distributed router with LeMaire's initialization technique as described above because, although Weyman teaches adding a new unit to the distributed router and teaches that initialization is needed (EX1005, [0069], [0156]), Weyman does not provide details of such initialization process, and LeMaire provides the missing details. EX1002, ¶236.

A POSITA would have had a reasonable expectation of success in combining Weyman/Jain's distributed router with LeMaire's initialization technique because both Weyman/Jain and LeMaire are both directed to a network element that supports multicast streams, and use IGMP to set up the forwarding tables. EX1005, [0104]-[0105]; EX1011, 4:42-5:4. EX1002, ¶237.

Besides, the combination would have been obvious because it would have merely amounted to using a known technique (e.g., LeMaire's initialization technique) to improve a similar device (e.g., a newly added unit in Weyman/Jain's distributed router, which performs substantially similar function as the bridge in LeMaire in that the newly added unit routes multicast streams) in the same way (e.g.,

the newly added unit initializing the forwarding table and functioning as a component of Weyman/Jain's distributed router). EX1002, ¶238.

The combination would have been obvious also because it would have merely amounted to combining prior art elements (e.g., a newly added unit in Weyman/Jain's distributed router and LeMaire's bridge initialization technique) according to known methods (e.g., implementing LeMaire's bridge initialization technique as part of Weyman's initialization process) to yield predictable results (e.g., the newly added unit initializing the forwarding table and functioning as a component of Weyman/Jain's distributed router). EX1002, ¶239.

2. Independent Claim 8

i. [8pre]⁴

To the extent the preamble is limiting, Weyman/Jain/LeMaire teaches the preamble for the same reasons as in Section XI.A.2.i [element [1pre]]. *See* EX1005, Abstract. EX1005, [0136]-[0138], FIG. 11. EX1002, ¶¶240-241.

⁴ Claim elements [8pre]-[8i] are identical to claim elements [1pre]-[1i].

Accordingly, claim elements [8pre]-[8i] are not repeated in this section. *See also* Claim Listing, *supra*.

ii. [8a]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.ii [1a]. *See* EX1005, Abstract, [0006], [0093], [0107], [0136]-[0137], [0148], [0151]-[0153], FIG. 11. EX1002, ¶242.

iii. [8b]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.iii [1b]. *See* EX1005, [0024]-[0025], [0029], [0060], [0062], [0064], [0122]-[0125], [0140], [0151], [0158], [0162], FIG. 11. EX1002, ¶243.

iv. [8c]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.iv [1c]. *See* EX1005, [0024]-[0025], [0029], [0060], [0062], [0064], [0122]-[0125], [0140], [0151], [0158], [0162], FIGS. 8, 11. EX1002, ¶244.

v. [8d]

Weyman/Jain/LeMaire teaches or renders obvious this limitation for the same reasons as in Section XI.A.2.v [1d]. *See* EX1005, [0025], [0093], [0097]-[0098], [0104], [0122]-[0125], [0127], FIGS. 8, 11. EX1002, ¶245.

vi. [8e]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.vi [1e]. *See* EX1005, [0046], [0065], [0100], [0105], [0111]-[0112], [0115], [0122]-[0124], [0127]-[0128], [0131], [0162], FIGS. 2, 10. EX1002, ¶246.

vii. [8f]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.vii [1f]. *See* EX1005, [0137], [0160], FIG. 11. EX1002, ¶247.

viii. [8g]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.viii [1g]. *See* EX1005, [0137], [0160], [0162], FIG. 11. EX1002, ¶248.

ix. [8h]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.ix [1h]. *See* EX1005, [0138], [0162]. EX1002, ¶249.

x. [8i]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.A.2.x [1i]. *See* EX1005, [0127]; Sections XI.D.2.v-XI.D.2.vi, XI.D.2.ix [[8d], [8e], [8h]]. EX1002, ¶250.

xi. [8j] wherein operating the network element comprises initializing a new processing unit added to the network element by:

Weyman/Jain/LeMaire teaches this limitation. Weyman discloses that the distributed router “can accommodate the *addition* or removal of unit” (EX1005, [0069]), and “[i]f a unit with a default configuration joins the fabric[,] it needs to *update its configuration* to that of the distributed multicast router” (EX1005,

[0156]). Indeed, as explained in Section XI.D.1.ii. [Weyman/Jain/LeMaire], upon addition of the unit, the new unit is configured to initialize the new unit by sending and receiving IGMP messages and updating the forwarding table. *See also* Sections XI.D.2.xii.-XI.D.2.xiv [limitations [8k]-[8m]]. EX1002, ¶¶251-253.

xii. [8k] deactivating the ports to which the new processing unit is assigned;

Weyman/Jain/LeMaire teaches this limitation. As explained in Section XI.D.1. [Weyman/Jain/LeMaire], when a new unit is added to the distributed router, the distributed router can disable or discontinue (“deactivat[e]”) the user ports of the new unit. EX1010, [0020] (“a ‘**Link Down**’ condition is effectuated on at least a portion of the link ports in the network node. ... When the Link Down condition is sensed by the entities (i.e., devices and nodes) connected to the network node via respective links, network traffic on the links is **discontinued** (block 206).”). EX1002, ¶¶254-255.

xiii. [8l] while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request, receiving multicast membership reports sent by multicast clients in response to the multicast status request, constructing the list by the new processing unit based on the received multicast membership reports; and

Weyman/Jain/LeMaire teaches this limitation. EX1002, ¶256.

First, Weyman/Jain/LeMaire teaches initializing a newly added unit “while the ports to which the new processing unit is assigned are deactivated.” As explained in Sections XI.D.1 [Weyman/Jain/LeMaire], when a new unit is added to the Weyman/Jain/LeMaire’s distributed router, the distributed router disables its user ports (*see also* EX1010, [0020]), and ***while the ports are disabled, the distributed router performs the initialization process.*** EX1010, [0021] (“***While the Link Down (or, disabled) condition*** is maintained on the link ports, the system software of the network node ***continues with its initialization and configuration process*** until it is complete and operational.”). EX1002, ¶¶257-258.

Second, Weyman/Jain/LeMaire teaches that the initialization step includes “sending a multicast status request.” As explained in Section XI.D.1 [Weyman/Jain/LeMaire], while the ports are disabled, the distributed router transmits IGMP queries to its clients, as taught by LeMaire, using other units in the distributed router. EX1011, 9:12-19. A POSITA would have understood that a IGMP query is a “multicast status request,” asking for the multicast “group membership” of clients. EX1005, [0161]. *See also* EX1017, [0061]. EX1002, ¶¶259-261.

Third, Weyman/Jain/LeMaire teaches that the initialization step includes “receiving multicast membership reports sent by multicast clients in response to the multicast status request.” As explained in Section XI.D.1 [Weyman/Jain/LeMaire],

while the ports are disabled, “[i]n response to IGMP Queries,” the other units can receive IGMP reports from the clients and forward those IGMP reports to the new unit via the fabric link and the new unit’s fabric port. *See* EX1011, 9:21-23; EX1005, [0137]-[0139], [0162]. EX1002, ¶¶262-263.

Lastly, Weyman/Jain/LeMaire teaches that the initialization step includes “constructing the list by the new processing unit based on the received multicast membership report.” As explained in Section XI.D.1 [Weyman/Jain/LeMaire], while the ports are disabled, the new unit can use the received IGMP reports to update its forwarding table. EX1005, [0138] (“The control plane of the multicast routers in both units would use the report in the construction of the respective forwarding table in the router.”), [0139], [0162]; EX1011, 9:23-34. EX1002, ¶¶264-265.

xiv. [8m] after constructing the list, activating the ports to which the new processing unit is assigned.

Weyman/Jain/LeMaire teaches this limitation. As explained in Section XI.D.1 [Weyman/Jain/LeMaire], once the new unit completes the initialization process—including the construction of the forwarding table—the unit activates the previously disabled user ports. EX1010, [0021] (“*[U]pon a suitable indication to the port hardware that the system software has completed the requisite configuration process, a ‘Link Up’ condition is effectuated on the ports (block*

212).”); EX1010, [0007] (“a link up condition is effectuated on the link ports, whereupon network traffic on the link ports may recommence.”). EX1002, ¶¶266-267.

3. Independent Claim 18

i. [18pre]⁵

To the extent limiting, Weyman/Jain/LeMaire teaches the preamble for the same reasons as in Section XI.D.2.i. [element [8pre]]. EX1002, ¶¶268-269.

ii. [18a]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.iii. [8b]. EX1002, ¶270.

iii. [18b]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.iii. [8b]. EX1002, ¶271.

iv. [18c]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.iv. [8c]. EX1002, ¶272.

⁵ Claim elements [18pre]-[18h] are identical to claim elements [11pre]-[11h]. *See also* Claim Listing, *supra*.

v. [18d]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.v. [8d]. EX1002, ¶273.

vi. [18e]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.v. [8d]. EX1002, ¶274.

vii. [18f]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.vi. [8e]. EX1002, ¶275.

viii. [18g]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.vii. [8f]. EX1002, ¶276.

ix. [18h]

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.viii.-XI.D.2.x. [[8g]-[8i]]. EX1002, ¶277.

x. [18i] wherein a new processing unit added to the network element is configured to initialize by:

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.xi [8j]. EX1002, ¶278.

- xi. [18j] deactivating the ports to which the new processing unit is assigned;**

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.xii [8k]. EX1002, ¶279.

- xii. [18k] while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request, receiving multicast membership reports sent by multicast clients in response to the multicast status request, constructing the list by the new processing unit based on the received multicast membership reports; and**

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.xiii [8l]. EX1002, ¶280.

- xiii. [18l] after constructing the list, activating the ports to which the new processing unit is assigned.**

Weyman/Jain/LeMaire teaches this limitation for the same reasons as in Section XI.D.2.xiv [8m]. EX1002, ¶281.

E. Ground 5: Weyman, Jain, LeMaire, And Li Render Obvious Claim 9.

1. Weyman/Jain/LeMaire In View Of Li

For the reasons provided in Section XI.A.1. [Weiman/Li], a POSITA would have been motivated to look for ways to share the workload between units to, for example, control the bandwidth of the ports, including the aggregated link, in Weyman/Jain/LeMaire's distributed router. EX1002, ¶282.

As described in Section XI.A.1 [Weiman/Li], a POSITA would have understood that Li discloses a port bandwidth control technique. EX1002, ¶283.

A POSITA would have found it obvious to combine Weyman/Jain/LeMaire's distributed multicast router with Li's channel joining processing unit to provide Li's channel joining processing unit in each of Weyman/Jain/LeMaire's unit. The reasons for combining Weyman and Li, provided in Section XI.A.1. [Weiman/Li], are equally applicable here because the initialization process of Jain and LeMaire are independent of the port bandwidth control process of Li, and the combination Weyman/Jain/LeMaire's unit with Li's channel joining processing unit would have yielded predictable results provided in Section XI.A.1 [Weiman/Li]. EX1002, ¶284.

2. Claim 9

The method according to claim 8, wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over a given port only when a number of the multicast packet streams that are permitted for forwarding over the given port, including the given multicast packet stream, does not exceed a maximum value specified for the given port.

Weyman/Jain/LeMaire/Li teaches this limitation for the same reasons as in Section XI.A.2.xiv [1m]. EX1002, ¶285.

For example, as explained in Sections XI.A.2.xiv [1m], XI.A.1. [Weyman/Li], and XI.E.1. [Weyman/Jain/LeMaire/Li], when the multicast router receives a multicast join request, the channel joining processing unit determines whether the

number of multicast streams being forwarded at a particular port (including a high-capacity port) is below the maximum number of accessed channels allowable, and if not, the channel joining processing unit rejects the multicast join request. EX1007, 31:4-24. EX1002, ¶286.

But if it is, the channel joining processing unit accepts the multicast join request. EX1007, 30:21-24. Then, the unit would process the join request to modify the unit's forwarding tables ("list") (EX1005, [0137]-[0139], [0162]) to indicate that the requested multicast stream is "permitted for forwarding over a given port." EX1002, ¶287.

F. Ground 6: Weyman, Jain, LeMaire, And Pitcher Render Obvious Claim 10.

1. Weyman/Jain/LeMaire In View Of Pitcher

For the reasons provided in Section XI.C.1 [Weiman/Li/Pitcher], a POSITA would have understood that stale entries in forwarding tables are undesirable because they unnecessarily consume memory space. EX1026, 5:64-6:4; EX1027, [0015]; EX1005, [0156]. Thus, a POSITA would have been motivated to find a way to identify stale entries in forwarding tables of Weyman/Jain/LeMaire's distributed router and remove them. EX1002, ¶288.

For the reasons provided in Section XI.C.1. [Weyman/Li/Pitcher], a POSITA would have understood that Pitcher discloses a timer-based method to remove stale entries from forwarding tables. EX1002, ¶289.

A POSITA would have found it obvious to combine Weyman/Jain/LeMaire's unit with Pitcher's timer-based entry removal technique to remove stale entries in the unit's forwarding tables. The reasons for combining Weyman, Li, and Pitcher, provided in Section XI.C.1. [Weiman/Li/Pitcher], are equally applicable here because Jain and LeMaire's initialization technique is independent of Pitcher's timer-based stale entry removal method and the combination of Weyman/Jain/LeMaire's unit with Pitcher's timer-based stale entry removal method would have yielded predictable results provided in Section XI.C.1. [Weiman/Li/Pitcher]. EX1002, ¶¶290-291.

2. Claim 10

The method according to claim 8, wherein updating the list comprises tracking a use of a permission to forward a certain multicast packet stream over a certain port by monitoring requests relating to the certain multicast packet stream that are received at the network element, and removing the permission from the list responsively to determining that the permission is aged.

Weyman/Jain/LeMaire/Pitcher teaches this limitation for the same reasons as in Section XI.C.2. [Claim 16]. EX1002, ¶292.

For example, as explained in Section XI.C.2. [Claim 16] and in Sections XI.C.1. [Weyman/Li/Pitcher] and XI.F.1. [Weyman/Jain/LeMaire/Pitcher], each unit of the distributed router tracks the use of entries in the forwarding tables using a timer. EX1009, 7:56-62. Entries in the forwarding tables (e.g., the Vbridge forwarding table) indicates “a permission to forward a certain multicast packet stream over a certain port.” EX1005, [0112], [0115]. And the timer is updated by monitoring IGMP membership reports (“requests”) relating to the certain multicast stream that are received at the Weyman/Jain/LeMaire/Pitcher’s distributed router. EX1009, 7:62-67. Then “[i]f the port report timer expires [i.e., determining that the permission is aged], the port is removed from the multicast group destination list [i.e., the permission is removed from the respective forwarding tables].” EX1009, 7:66-67. EX1002, ¶¶293-294.

G. Ground 7: Weyman, Li, and Ding Render Obvious Claims 2 and 12.

1. Claim 2

The method according to claim 1, wherein storing the list, distributing the request and selectively forwarding the multicast packets are performed by dedicated hardware logic.

Weyman/Li in view of Ding (“Weyman/Li/Ding”) would have rendered obvious the additional limitation of claim 2. EX1002, ¶295.

In Weyman/Li, the steps of storing the forwarding tables, distributing the IGMP host report, and selectively forwarding the multicast streams are performed by the unit's Vbridge(s), router, and channel joining processing unit. *See* Sections XI.A.2.vi-XI.A.2.viii, XI.A.2.x [limitations [1e], [1g], [1i]]. It was well-known in the art that such functionalities may be embodied in dedicated hardware logic. For example, Ding teaches that its “method for forwarding multicast packets by a module” may be implemented in hardware logic such as “*integrated circuitry, ... a Field Programmable Gate Array (FPGA)*.” EX1012, Claim 1, [0086]. EX1002, ¶¶296-297.

It would have been obvious to implement the “storing,” “distributing,” and “selectively forwarding” functionalities of the unit in Weyman/Li using Ding's dedicated hardware logic because that would have merely amounted combining prior art elements (Weyman/Li's unit and Ding's dedicated hardware logic) according to a known method (e.g., integration of the hardware logic into the distributed router) to yield a predictable result of operating the distributed router using Ding's dedicated hardware logic. Implementing Weyman/Li's unit would have been a routine design choice—for example, between a generic processor and a dedicated hardware logic—and that would have been obvious. EX1002, ¶298.

2. Claim 12

The network element according to claim 11, wherein the processing units comprise dedicated hardware logic for storing the lists, distributing the request and selectively forwarding the multicast packets.

Weyman/Li/Ding teaches this limitation for the same reasons as in Section XI.G.1[Claim 2]. EX1002, ¶299.

XII. DISCRETIONARY DENIAL IS NOT WARRANTED

Discretionary denial under the *General Plastic* factors is not warranted because the '394 Patent has never been challenged in an IPR proceeding.

Discretionary denial under the *Fintiv* factors is not warranted because Petitioner provides a *Sotera* stipulation that, if the instant IPR petition is instituted, Petitioner will not pursue in the district court any grounds of invalidity against the challenged claims that were or reasonably could have been asserted in the instant IPR petition. Furthermore, the pending litigation is in early stages (*see* EX1028-EX1030) and claim construction briefing has not started (EX1028, 4-5).

Lastly, discretionary denial under Section 325(d) is also not warranted because the prior art references cited in the grounds were not previously considered during prosecution, and are not cumulative over the references already considered. EX1031.

XIII. CONCLUSION

Petitioner respectfully requests the Board to institute the IPR and cancel all challenged claims.

Dated: March 31, 2024

Respectfully submitted,

/Taeg Sang Cho/

Taeg Sang Cho

(Reg. No. 69,618)

tcho@desmaraisllp.com

DESMARAIS LLP

230 Park Ave

New York, NY 10169

212-351-3400 (telephone)

212-351-3401 (facsimile)

Counsel for Petitioner

CERTIFICATION UNDER 37 C.F.R. § 42.24(d)

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and in reliance on the word count provided by the word-processing system used to prepare this Petition, that the number of words in this paper is 13,943. Pursuant to 37 C.F.R. § 42.24(d), this word count excludes the table of contents, table of authorities, mandatory notices under § 42.8, certificate of service, certificate of word count, appendix of exhibits, and any claim listing.

Dated: March 31, 2024

Respectfully submitted,

/Taeg Sang Cho/

Taeg Sang Cho

(Reg. No. 69,618)

tcho@desmaraisllp.com

DESMARAIS LLP

230 Park Ave

New York, NY 10169

212-351-3400 (telephone)

212-351-3401 (facsimile)

Counsel for Petitioner

CERTIFICATE OF SERVICE

Pursuant to 37 C.F.R. §42.6(e), the undersigned certifies that a complete copy of the foregoing materials:

- Petition for Inter Partes Review Under 35 U.S.C. §312 and 37 C.F.R. §42.104 of U.S.;
- Exhibit List;
- Exhibits 1001-1031;
- Word Count Certification Under 37 CFR §42.24(d); and
- Petitioner's Power of Attorney

was caused to be served on counsel of record for Patent Owner by filing this document through P-TACTS and by sending this document via Federal Express Priority Overnight shipping, which is at least as fast and reliable as U.S. Priority Mail Express, to counsel of record for Patent Owner at the correspondence address indicated in Patent Center:

Yehuda Binder
131926 - May Patents Ltd. c/o Dorit Shem-Tov
P.O.B 7230
Ramat-Gan, - 5217102
ISRAEL

Dated: March 31, 2024

Respectfully submitted,

/Taeg Sang Cho/

Taeg Sang Cho

(Reg. No. 69,618)

tcho@desmaraisllp.com

DESMARAIS LLP

230 Park Ave

New York, NY 10169

212-351-3400 (telephone)

212-351-3401 (facsimile)

Counsel for Petitioner