

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.

Petitioner

v.

PORTSMOUTH NETWORK CORP.

Patent Owner

Case IPR2024-00506

U.S. Patent No. 8,014,394

DECLARATION OF HENRY H. HOUH, PH.D.

ON BEHALF OF PETITIONER

TABLE OF CONTENTS

	<u>Page</u>
I. Background.....	1
II. Relevant Law	9
A. Claim Construction.....	9
B. Obviousness.....	10
C. Ordinary Skill in the Art.....	12
III. Summary of Opinions.....	13
IV. Technology Background.....	13
A. Networking, Multicasting, and IGMP.....	13
B. Distributed Router	18
C. Link Aggregation	19
D. Port Bandwidth Control	20
V. Overview of the '394 Patent.....	21
A. Purported Invention Of The '394 Patent.....	21
B. Prosecution History	27
VI. Claim Construction.....	28
VII. Level of Ordinary Skill in the Art	29
VIII. Overview of the Prior Art.....	29
A. Weyman (“US 6,032,194”).....	29
IX. Specific Grounds for Challenge	46

A.	Ground 1: Weyman and Li Render Obvious Claims 1, 4, 6-7, 11, 14, and 17.....	46
1.	Weyman In View Of Li	46
2.	Independent Claim 1	58
3.	Claim 4.....	97
4.	Claim 6.....	99
5.	Claim 7.....	100
6.	Independent Claim 11	100
7.	Claim 14.....	104
8.	Claim 17.....	105
B.	Ground 2: Weyman, Li, And Rooholamini Render Obvious Claims 5 and 15.....	105
1.	Weyman/Li in view of Rooholamini	105
2.	Claim 5.....	110
3.	Claim 15	114
C.	Ground 3: Weyman, Li, And Pitcher Render Obvious Claim 16.	114
1.	Weyman/Li In View Of Pitcher.....	114
2.	Claim 16.....	117
D.	Ground 4: Weyman, Jain, And LeMaire Render Obvious Claims 8 And 18.	118
1.	Weyman In View Of Jain And LeMaire.....	118
2.	Independent Claim 8	129

3.	Independent Claim 18	136
E.	Ground 5: Weyman, Jain, LeMaire, And Li Render Obvious Claim 9.	138
1.	Weyman/Jain/LeMaire In View Of Li.....	139
2.	Claim 9	139
F.	Ground 6: Weyman, Jain, LeMaire, And Pitcher Render Obvious Claim 10.	141
1.	Weyman/Jain/LeMaire In View Of Pitcher	141
2.	Claim 10	142
G.	Ground 7: Weyman, Li, and Ding Render Obvious Claims 2 and 12.	143
1.	Claim 2	143
2.	Claim 12	145
X.	Conclusion	145
XI.	Availability for Cross-Examination	145
XII.	Right to Supplement	146
XIII.	Jurat.....	146

I, Henry H. Houh, declare as follows:

1. My name is Henry H. Houh.
2. I have been retained as an expert witness on behalf of Petitioner Cisco Systems, Inc. to provide expert opinions on the validity of U.S. Patent No. 8,014,394 (the “394 patent”).

I. Background

3. I received a Ph.D. in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (“MIT”) in 1998. I also received a Master of Science degree in Electrical Engineering and Computer Science in 1991, a Bachelor of Science degree in Electrical Engineering and Computer Science in 1989, and a Bachelor of Science degree in Physics in 1990, all from MIT.

4. I am currently self-employed as an independent technical consultant. I am also president of a company, Einstein’s Workshop, that provides supplemental science, technology, engineering, and mathematics (“STEM”) education to children of all ages.

5. I first entered telecommunications in 1987 when I worked as a summer intern at AT&T Bell Laboratories as part of a five-year dual degree program at MIT. I continued to work at AT&T Bell Laboratories as part of this MIT program. While at MIT, I was a teaching assistant (“TA”) in the Electrical Engineering and Computer Science Department’s core Computer Architectures

course. I first was a TA as a senior for a role typically reserved for graduate students. I later became head TA. The course covered various topics in computer architectures, as well as programming. As a TA, I helped write homework assignments, lab assignments (including those that involved programming), and exams. I also taught in the recitation sections.

6. Later, as part of my doctoral research at MIT from 1991-1998, I was a research assistant in the Telemedia Network Systems (“TNS”) group at the Laboratory for Computer Science. The TNS group built a high-speed gigabit network and created applications that ran over the network. Example applications included ones for remote video capture, processing, and display of video on computer terminals. In addition to working on the design of core network components, designing and building the high-speed links, and designing and writing the device drivers for the interface cards, I also set up the group’s web server. Also, I helped to maintain, install, and upgrade the networking devices used within the group, along with other graduate students. Being a networking group, we regularly discussed topic of concern in networking, including network security and the kinds of attacks that were carried out over a network. I used several network scanning applications, including Nmap, to assess the level of security in our own network.

7. I also helped to build the web pages that initiated the above-mentioned video sessions via a web interface. Vice President Al Gore visited our group in 1996 and received a demonstration of—and remotely drove—a radio-controlled toy car with a wireless video camera mounted on it that was built by our group. This toy car device received commands transmitted over a network from a remote computer, and video data from the toy car was transmitted wirelessly then over a computer network back to the user controller. On occasion, we allowed users visiting our web site to drive the toy car from their remote computer while they watched the video on their computer. The video stream was encoded by TNS-designed hardware, streamed over the TNS-designed network, and displayed using TNS-designed software. We also installed a multicast backbone (“MBONE”) router within our group around the mid 1990’s, and I personally accessed multicast video at that time. As part of my consulting work, I have also studied commercial deployments of IPTV and their use of multicast, including the AT&T U-Verse system first deployed around 2006.

8. I defended and submitted my Ph.D. thesis, titled “Designing Networks for Tomorrow’s Traffic,” in January 1998. As part of my thesis research, I analyzed local area and wide area flows to show a more efficient method for routing packets in a network, based on traffic patterns at the time. The traffic flow

data included Ethernet, IP, TCP or UDP, and RTP header information, which I analyzed to come to the conclusions in my thesis.

9. From 1997 to 1999, I was a Senior Scientist and Engineer at NBX Corporation, a start-up that made business telephone systems for streaming packetized audio over data networks instead of using traditional telephone lines. NBX was later acquired by 3Com Corporation.

10. As part of my work at NBX, I designed the core audio reconstruction software algorithms for the telephones, as well as the packet transmission algorithms. I also designed and validated the core packet transport protocol used by the phone system. The protocol was used for all signaling in the phone system, including for the setup of conference calls. The NBX system also featured a computer interface for initiating phone calls, which could also initiate conference calls. The NBX system also supported the Telephony Application Programming Interface (“TAPI”) that allowed other computer programs to integrate with our system telephony features. I also programmed the first prototype of our phone which communicated using IP, and I demonstrated our IP phones working over the Internet when we attended a trade show in California. The phone connected over the Internet to our headquarters in Andover, MA. I also programmed a data capture and analysis system that provided data on the operation of the system.

11. From 1999-2004, I was employed by Empirix or its predecessor company, Teradyne. Empirix was a leader in test tools for telecommunications protocols and systems, providing functional testing tools as well as load testing tools. From 2000-2001, I conceived and built a test platform for testing Voice-over-IP (VoIP). The first application on this new test platform was a cloud emulator for simulating the effects of transmitting VoIP over a busy network. The test platform was based on a network processor chip, which could be programmed to cut-through packets while processing packet data such as various protocol layer headers including addresses included therein and even packet data contents. I also designed a protocol analyzer built on the same platform. The application captured and performed protocol decoding at various layers in the protocol stacks of captured packets, including detailed Ethernet header decoding, IP header decoding, TCP header decoding, UDP header decoding, RTP header decoding, and many other specified protocols. The application was also designed to reconstruct entire conversations that spanned multiple packets.

12. At Empirix, I also rearchitected the design of the Web Test division's core product, e-Test Suite. E-Test Suite was a software program designed to perform functional and load testing of web applications. I was technical lead on the project, and also worked on programming a Javascript interpreter for the product. Javascript is the flexible programming layer that allows users to interact with web

pages through actions such as clicking and dragging that allows web page programmers to extend the functionality of a web page beyond clicking on links. I also designed and prototyped a web server security testing product, which scanned web servers for specific types of vulnerabilities, including looking for administrative logins which could be attacked with password attacks and buffer overflow attacks, and attacks such as SQL query injection attacks. We used the prototype on some customer web sites, with their permission, and successfully detected vulnerabilities.

13. In 2006, as part of my role at BBN Technologies, I helped found PodZinger Inc., now known as RAMP Inc. PodZinger utilized BBN's speech recognition algorithms to search through the spoken words in audio and video segments. While I was Vice President of Operations and Technology, PodZinger followed its initial prototype with a full streaming audio and video search solution. I also created a social networking web site, which BBN sold to a venture-funded startup company. In the process of creating the web site, I designed and specified the authentication and authorization protocols.

14. Until October 2022, I was President of Einstein's Workshop, a STEM education company for kids. At Einstein's Workshop, we taught programming to children of all ages, starting with Scratch programming for young elementary students and advancing through programming in Python and for game

development for high school students. I conceived many of the programming classes and have taught programming to novice and more advanced students. We also created an educational computer-aided design tool, called BlocksCAD. We received a grant, on which I was the principal investigator, to further develop BlocksCAD and successfully met our grant objectives. BlocksCAD was spun-off from Einstein's Workshop and has been part of various incubator/accelerator programs in education including LearnLaunch and Mass Challenge as well as Play Labs at the MIT Media Lab. Today, BlocksCAD is used worldwide in makerspaces as well as in schools throughout the United States.

15. I have reviewed all of the exhibits cited in this declaration, including the following:

Exhibit No.	DESCRIPTION
1001	U.S. Patent No. 8,014,394 to Ram et al. (the "'394 patent")
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

Exhibit No.	DESCRIPTION
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”)
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

16. I understand the ’394 patent has a filing date of February 11, 2009.

Thus, I have considered certain issues from the perspective of a person of ordinary skill in the art (“POSITA”) as described below as of the February 11, 2009, filing date of the ’394 patent.

17. I have been retained by the Petitioner as an expert in the field of information technology and networking. I am working as an independent consultant in this matter and am being compensated at my normal consulting rate

of \$680 per hour for my time. My compensation is not dependent on and in no way affects the substance of my statements in this Declaration.

18. I have no financial interest in the Petitioner. I similarly have no financial interest in the '394 patent and have had no contact with any of the named inventors of the '394 patent.

II. Relevant Law

19. I am not an attorney. For the purposes of this declaration, I have been informed about certain aspects of the law that are relevant to my opinions. My understanding of the law is as follows:

A. Claim Construction

20. I have been informed that claim construction is a matter of law and that the final claim construction will ultimately be determined by the Board. For the purposes of my analysis in this proceeding and with respect to the prior art, I have been informed that IPRs are currently reviewed under “the *Phillips* standard.”

21. I have been informed that under the *Phillips* standard, claim terms are given their plain and ordinary meaning as understood by a person of ordinary skill in the art at the time of the invention in light of the claim language and the patent specification.

22. I have been informed that embodiments described in the specification are generally encompassed by the claims.

23. I have been informed that the patentee can serve as their own lexicographer. As such, if a claim term is provided with a specific definition in the specific, I should interpret that claim term in light of the particular definition provided by the patentee.

B. Obviousness

24. I have been informed and understand that a patent claim can be considered to have been obvious to a person of ordinary skill in the art at the time the application was filed. This means that, even if all of the requirements of a claim are not found in a single prior art reference, the claim is not patentable if the differences between the subject matter in the prior art and the subject matter in the claim would have been obvious to a person of ordinary skill in the art at the time the application was filed.

25. I have been informed and understand that a determination of whether a claim would have been obvious should be based upon several factors, including, among others:

- the level of ordinary skill in the art at the time the application was filed;
- the scope and content of the prior art; and
- what differences, if any, existed between the claimed invention and the prior art.

26. I have been informed and understand that a claim can be unpatentable in light of a single prior art reference if it would have been obvious to modify that reference to come up with the claimed subject matter.

27. I have been informed and understand that the teachings of two or more references may be combined in the same way as disclosed in the claims, if such a combination would have been obvious to one having ordinary skill in the art.

28. In determining whether a combination based on either a single reference or multiple references would have been obvious, it is appropriate to consider, among other factors:

- whether the teachings of the prior art references disclose known concepts combined in familiar ways, which, when combined, would yield predictable results;
- whether a person of ordinary skill in the art would have implemented a predictable variation, and would have seen the benefit of doing so;
- whether the claimed elements represent one of a limited number of known design choices, and would have a reasonable expectation of success by those skilled in the art;
- whether a person of ordinary skill would have recognized a reason to combine known elements in the manner described in the claim;

- whether there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; and
- whether the claim applies a known technique that had been used to improve a similar device or method in a similar way.

29. I understand that one of ordinary skill in the art has ordinary creativity and is not an automaton.

30. I understand that in considering obviousness, it is important not to determine obviousness using the benefit of hindsight derived from the patent being considered.

C. Ordinary Skill in the Art

31. I understand that, in the context of an invalidity analysis, a POSITA is a hypothetical person who looks to prior art at the time of the invention.

32. I further understand that the factors that may be considered in determining the level of ordinary skill in the art include: (1) the problems encountered in the art; (2) the prior art solutions to the problems encountered in the art; (3) the rapidity of innovation; (4) the sophistication of the technology; and (5) the education level of active workers in the field. I understand that these factors need not all be taken into account for the analysis and that one or more of these factors may control.

III. Summary of Opinions

33. It is my conclusion that claims 1-2, 4-12 and 14-18 (the “Challenged Claims”) of the ’394 patent are rendered obvious by the prior art identified below:

Ground	References	Basis (35 U.S.C.)	Claims
1	Weyman and Li	§ 103 (Obviousness)	1, 4, 6-7, 11, 14, 16-17
2	Weyman, Li, and Rooholamini	§ 103 (Obviousness)	5, 15
3	Weyman, Li, and Pitcher	§ 103 (Obviousness)	16
4	Weyman, Jain, and LeMaire	§ 103 (Obviousness)	8, 18
5	Weyman, Jain, LeMaire, and Li	§ 103 (Obviousness)	9
6	Weyman, Jain, LeMaire, and Pitcher	§ 103 (Obviousness)	10
7	Weyman, Li, and Ding	§ 103 (Obviousness)	2, 12

IV. Technology Background

A. Networking, Multicasting, and IGMP

34. At the time of the filing of the ’394 Patent, communication networks that could deliver packets from a source to a client device were well-known. EX1017, [0005]. In some cases, packets are delivered in a “unicast stream” (e.g., from a single source to a single client). EX1017, [0005]. In other cases, packets are delivered in a “multicast stream” (e.g., from a single source to multiple clients). EX1017, [0005].

35. In one example, Sherer discloses a “technique for managing the streaming of digital video content involves providing a unicast stream to a client in response to the playout status of the unicast stream at the client.” EX1014, Abstract. Sherer discloses that “FIG. 1 depicts a network that includes a stream server configured to manage the streaming of digital video content in accordance with an embodiment of the invention.” EX1014, [0010]. Specifically, Sherer describes that “FIG. 1 depicts a network 100 that includes a stream server 102 that is configured to manage the streaming of digital video content in accordance with an embodiment of the invention.” EX1014, [0026]. In this example, Sherer discloses that the “stream server 102 provides digital video content simultaneously to multiple clients 110 via multicasting (e.g., multicast stream 114) and provides digital video content to a particular single client via unicasting (e.g., unicast stream 116).” EX1014, [0028].




FIG.1

EX1014, FIG. 1.

36. Indeed, it was known that multicast groups may be managed using communication protocols, such as Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD). EX1015; EX1001, 1:25-35 (Background, “IGMP is specified, for example, by Fenner in Request For Comments (RFC) 2236 of the Internet Engineering Task Force (IETF), entitled “Internet Group Management Protocol, Version 2,” November, 1997, and by Cain et al. in “Internet Group Management Protocol, Version 3,” RFC 3376, October, 2002, which are incorporated herein by reference. MLD is specified, for example, by Deering et al. in “Multicast Listener Discovery (MLD) for IPv6,” RFC 2710, October, 1999, and by Vida and Costa in “Multicast Listener Discovery Version 2

(MLDv2) for IPv6,” RFC 3810, June, 2004, which are incorporated herein by reference”).

37. In some examples, a client device may request to “join” a multicast stream by sending an “IGMP Membership Report” (sometimes referred to as an “IGMP host report” or simply “IGMP report”). *See, e.g.*, EX1015, 4-8. The IGMP Membership Report may include an identification of a multicast stream. EX1015, 4-8. In addition, a client can “leave” a multicast stream by sending a different IGMP Membership Report or an “IGMP Leave Group” message that includes the identification of the multicast stream. EX1015, 4-8.

38. Traditionally, communication networks include network devices, such as routers, switches, and bridges. Those network devices can facilitate network communication, such as multicast packet streams.

39. Typically, a network device can route or forward a multicast stream using a forwarding table. For example, the forwarding table can identify, for each multicast stream, ports or interfaces that are coupled to the multicast clients and that would receive the particular multicast stream. *See, e.g.*, EX1017, [0050]-[0052]. So when a network device receives a multicast stream through a port (e.g., an ingress port), the network device can look up the forwarding table and determine the ports to which the multicast stream should be forwarded (also known as the egress port(s)).

40. Various implementations of forwarding tables have been employed in the art. One implementation of a multicast forwarding table is reproduced below. In this example, each entry of the multicast forwarding table identifies: (1) a multicast stream, which is the combination of a “Source Address 414” and a “Group Address 412” and (2) an outgoing port to which the multicast stream should be forwarded over (e.g., the “Egress Interface 430”).

400 ↙

ENTRY 410								
GROUP ADDRESS 412	SOURCE ADDRESS 414	MASK 416	INGRESS INTERFACE 420	FLAGS 422	EGRESS INTERFACE 430	FLAGS 432		
						F 434	SP 436	NS 438
ENTRY 410								
• • •								
ENTRY 410								

EX1017, FIG. 4.

41. In some examples, a network device may populate its multicast forwarding table by “snooping” incoming messages, such as IGMP messages. Specifically, some examples contemplate a switch that can “snoop” (e.g., detect) an IGMP Membership Report that arrives on a particular port; and in doing so, the

switch can (1) determine which multicast streams should be sent to that port and (2) add that information to the multicast forwarding table. EX1017, [0066].

B. Distributed Router

42. Exemplary network devices includes routers. Routers may be implemented in a number of ways, including as a distributed router. One example distributed router includes multiple modules connected by an interconnection fabric. *See, e.g.*, EX1012, [0012], [0036]; EX1018, [0002].




FIG. 1 100

EX1012, FIG. 1.

C. Link Aggregation

43. Network devices are connected by data links between endpoints called ports. Link aggregation is a process that combines (e.g., aggregates) multiple physical ports into a single logical connection called a link aggregation group

(LAG). EX1016, [0003]-[0004], EX1013, [0003]. By combining multiple physical ports, LAGs provide high-capacity connections (e.g., aggregated links) that can be used to increase network efficiencies, for example, by load balancing. EX1016, [0006]; EX1013, [0036], [0065].

D. Port Bandwidth Control

44. It is well known that ports have limited bandwidth. EX1019, 1:25-2:36. For example, a standard Ethernet port has a maximum bandwidth of 10 Megabits per second (Mbps); a Fast Ethernet port has a maximum bandwidth of 100 Mbps; a Gigabit Ethernet port has a maximum bandwidth of 1 Gigabits per second (Gbps). EX1019, 1:25-2:36. And when a data transfer request for a particular port requires transferring data at a rate higher than the port's maximum bandwidth, the port can drop packets, which can ultimately cause communication failures. *See* EX1020, [0003].

45. Network devices can implement port bandwidth control to avoid overloading ports and control the amount of data transfer on ports. For instance, a network device can track an amount of data that is transferred on a port in real-time; and if the data transfer rate approaches the maximum bandwidth, the network device can divert one or more data streams to other ports. *See, e.g.*, EX1016, [0009] ("Other important technical advantages of certain embodiments of the present invention include load balancing. *The tables or other information used by*

the ingress card to determine a destination port for packets may also include additional usage information that allows the ingress card to determine a destination port from the egress ports in an aggregated link. This means that rather than using one link to the exclusion of others, loads may be distributed among several links in an aggregated link in a more balanced fashion. Particular embodiments of the present invention may have some, all or none of the enumerated technical advantages. Still other important technical advantages will be apparent to one skilled in the art from the following figures, description, and claims.’’).


46. Further, in some examples, a network device may count a number of data streams, such as multicast packet streams, that are transmitted over a given port. In this, if the particular port is already carrying a maximum number of allowable data streams, the network device may deny any request to send one or more additional data streams over that particular port. *See, e.g., EX1017, [0055].*

V. Overview of the ’394 Patent

A. Purported Invention Of The ’394 Patent

47. The ’394 Patent discloses a network element that purports to deliver multicast streams to clients. I have reproduced Figure 1 of the ’394 Patent below. As shown in Figure 1, the “network element 32 comprises multiple processing units, such as line cards 36. The line cards are interconnected by an interconnection

switch fabric 40.” EX1001, 6:23-25. Figure 1 also shows that “[t]he network element comprises multiple Input/Output (I/O) ports 44, and each line card 36 is assigned to process the packets communicated over a given subset of the ports.” EX1001, 6:25-28.




EX1001, FIG. 1.

48. Further, Figure 1 depicts that each line card includes a multicast forwarding table, which is described as “a local filtering database [DB], which stores a list of one or more multicast streams that are permitted for forwarding,” and “[t]he filtering database also stores the identities of the specific ports over which each of the permitted streams is to be forwarded.” EX1001, 5:4-8. The line

cards construct/update their filtering databases, and the '394 Patent notes that “[t]he filtering databases in the different line cards are updated by monitoring [multicast] request messages.” EX1001, 7:23-26.

49. I have reproduced Figure 2 of the '394 Patent below. Figure 2 shows a line card that receives a multicast request from a multicast client in step 60, and then, the line card forwards a copy of the request to every other line card in the network element at step 72. EX1001, 7:30-58. The '394 Patent discloses that at step 76, each line card “updates its respective database 52” according to the received request. EX1001, 7:63-65.

FIG. 2



EX1001, FIG. 2.

50. The '394 Patent further discloses that “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. For instance, the '394 Patent describes that the “filtering databases 52 hold port

counters, which are indicative of the number of multicast streams that are permitted for forwarding over each port.” EX1001, 8:41-45.

51. In addition, the '394 Patent discloses that its line cards can look up a count in a particular port counter that is associated with a requesting port; and if that particular port counter indicates that the “current number of multicast streams” has already reached its maximum value (e.g., count), then the line card will reject the request for the multicast stream. EX1001, 8:46-9:4. Further, the '394 Patent describes using the same concept to control a high-capacity port that is formed by aggregating physical ports as “a Link Aggregation Group (LAG).” *See, e.g.,* EX1001, 6:61-7:4 (“In some embodiments, two or more ports 44 are aggregated to form a high-capacity port 56, in order to provide high traffic capacity and redundancy to a certain content client or content source. For example, *ports can be aggregated to form a Link Aggregation Group (LAG)*, as specified in the IEEE 802.1AX standard, entitled “IEEE Standard for Local and Metropolitan Area Networks—Link Aggregation,” Nov. 3, 2008, which is incorporated herein by reference. Each individual port in a LAG is referred to as a group member. *A LAG may comprise any desired number of member ports, which may reside in the same line card or in different line cards.*”), 9:64-10:14 (“Typically, the total number of multicast streams forwarded over the entire LAG is limited. For this purpose, *the network element maintains a counter that limits the number of*

multicast streams that are permitted for forwarding over the LAG. As noted above, member ports of a given LAG may reside in different line cards. In these embodiments, the different NPs in the line cards that host the different LAG members hold respective local copies of the LAG's limit counter. When one of these NPs receives a multicast content request, the request typically identifies the entire LAG (and not the individual member over which it arrived) as the ingress port. The receiving NP duplicates the request and distributes one copy to all NPs, as described above. In particular, the NPs that host the different LAG members receive the distributed copy, and update their local copies of the LAG limit counter accordingly. This mechanism allows the NPs to manage local copies of the LAG limit counter without a need for additional communication or synchronization among the LAG members.”).

52. The '394 Patent explains that “two or more ports 44 [in the network element] are aggregated to form a high-capacity port 56, in order to provide high traffic capacity and redundancy to a certain content client or content source.” EX1001, 6:61-64. “For example, ports can be aggregated to form a Link Aggregation Group (LAG).” EX1001, 6:64-65.

53. “Typically, the total number of multicast streams forwarded over the entire LAG is limited.” EX1001, 9:64-67. Similar to individual ports, “the network element maintains a counter that limits the number of multicast streams [over] the

LAG.” And when a line card receives a request for a multicast stream over the LAG, but the LAG is already carrying the maximum number of streams, then the line card rejects the request. EX1001, 10:1-14.

54. The '394 Patent also discloses updating a multicast forwarding table when a new line card is installed. Specifically, the '394 Patent describes “[w]hen a new line card is inserted into an active network element, [] the new line card runs a process that updates its filtering database.” EX1001, 9:41-51.

55. “Typically, the NP of the new line card disables its local ports 44 for a certain period of time (e.g., by setting the ADMIN_STATUS of these ports to DOWN), during which a multicast status update request (e.g., an IGMP general query) is transmitted through all the ports of the network element.” EX1001, 9:43-48. “The multicast content clients respond to the status update request, thus allowing the newly-added NP to update its local filtering database, port counters and aging counters.” EX1001, 9:48-51. According to the '394 Patent, “[w]hen the database and counters are up-to-date (or after a certain time out), the NP activates the ports of the newly-inserted line card, and may start normal operation.” EX1001, 9:51-53.

B. Prosecution History

56. The '394 Patent was originally filed on February 11, 2009. EX1001, Cover (22); EX1004, 263-272.

57. I understand that on October 29, 2010, the Examiner issued a first, non-final rejection, noting that the independent claims were generally directed to a distributed router and multicasting over a high-capacity port would have been obvious to an ordinarily skilled artisan. EX1004, 106-131. Further, I understand that the Examiner noted that certain dependent claims, which also required controlling a maximum number of multicast streams over that high-capacity port would have been allowable if rewritten in independent form, and that some other dependent claims, which additionally required initialization of a new processing unit would have also been allowable if rewritten in independent form. EX1004, 106-131.

58. Eventually, Applicant amended the independent claims to incorporate the previously indicated allowable subject matter, which led to the allowance of the claims recited in the '394 Patent. EX1004, 160-179, 232-241.

59. However, as I explain below, the purported allowable subject matter would have been obvious to a POSITA, at the time of the alleged invention, in view of certain prior art references that the Examiner never considered.

VI. Claim Construction

60. I have reviewed the claims, specification, and file history of the '394 patent.

61. I did not find any explicit definition of claim terms. I have applied the ordinary and customary meaning of the claim terms as understood by a POSITA in my analysis provided herein.

VII. Level of Ordinary Skill in the Art

62. In my opinion, a person of ordinary skill in the art (“POSITA”) in the field of the ’394 patent as of its effective filing date of February 11, 2009, 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. In my opinion, additional education could make up for less practical experience, and vice versa.

63. I met these requirements for POSITA at the time of the effective filing date of the ’394 patent.

VIII. Overview of the Prior Art

A. Weyman (“US 6,032,194”)

64. Weyman describes a “distributed router composed of individual units connected by links in a routing fabric which appears to be a single network entity.” EX1005, Abstract; *see also* EX1005, [0024] (“*When several individual data switching units are connected to form a fabric, their functionalities may be distributed among the individual units in the fabric. It is generally important to allow this distribution to occur so that the entire fabric appears to other entities*”).

as if it were a single entity. Ancillary requirements are intended to facilitate the sharing of the workload between the units, to maximise the data forwarding performance and to minimise the dependence on any single unit or connection within the fabric.”). I have reproduced Figure 5 of Weyman below. As shown in Figure 5, Weyman’s distributed router includes multiple units 10, 20. EX1005, [0064] (“FIG. 5 illustrate this operation in greater detail. *The fabric in this example consists of units 10 and 20 and link 30*. The unit 10 has fabric ports P1 to P4, a fabric port F1, a router 31, forwarding database 41 and a routing information base 51. Unit 2 has user ports P1 to P4, a fabric port F2, bridges 21, a router 32, forwarding database 42 and a routing information base 52.”); *see also* EX1005, [0060] (“FIG. 4 illustrates a unit 10 which has a multiplicity of Vbridges, such as the Vbridges previously mentioned, user ports 15 to which the Vbridges 11 are variously connected, a fabric port 16, to which all the Vbridges are connected, and a router 31 having a respective router interface coupled to each of the Vbridges.”).




FIG.5

EX1005, FIG. 5.

65. Weyman discloses that each unit includes: (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] (“FIG. 1 illustrates a single unit, operating as a bridge, for use in a network ‘fabric’. In this example, the unit 10 has three ‘Vbridges’ 11, 12 and 13, each with a forwarding database FDB 1, FDB 2, etc. Each Vbridge 11, 12 and 13 is connected to at least one and usually more than one of a multiplicity of user ports 15 and all have logical connections to each of at least one and in generally more than one fabric port 16. *The term ‘fabric port’ is used to refer to a port which is only connected to a link internal of ‘fabric’ of units.*”), [0036] (“As is used herein, the term ‘fabric port’ refers to a visible port that is dedicated to the task of connecting a unit to other units within the fabric so as to constitute the distributed fabric of units. A fabric port may not be connected to anything other than a fabric port on another unit within the same fabric. The physical implementation of a fabric port and the number of them available on each unit will vary from product to product. For example, some products may have physically dedicated fabric ports and some of these may be physically unable to be connected to standard user ports. Alternatively, products may have dual-purpose ports that may be configured to be either user ports or

fabric ports. The method used to configure dual purpose ports may be by manual specification or may be automatic using some discovery protocol.”), [0064].

66. Further, Weyman discloses that each unit also includes a bridge (e.g., Vbridge 21), the bridge’s forwarding table, which are stored in the forwarding database(e.g., FDBs 41, 42), an internal router (e.g., routers 31, 32), and the routers’ forwarding table (e.g., RIBs [routing information bases] 51, 52). EX1005, [0064] (“FIG. 5 illustrate this operation in greater detail. The fabric in this example consists of **units 10 and 20** and link 30. The unit 10 has fabric ports P1 to P4, a fabric port F1, **a router 31, forwarding database 41 and a routing information base 51**. Unit 2 has user ports P1 to P4, a fabric port F2, bridges 21, **a router 32, forwarding database 42 and a routing information base 52**.”).

67. According to Weyman, all of the units must have an “identical” “router configuration,” which allows the units to collectively “operate as a single distributed router.” EX1005, [0069] (“**For the routers to operate as a single distributed router every defined router configuration should be identical on all units in the fabric**. Preferably the configuration is stored in non-volatile storage in all the units. If the system can accommodate the addition or removal of units, then when a unit with a default configuration joins the fabric, it needs to update its configuration to that of the distributed router. If a unit with a non-default configuration joins the fabric then there should be a selection mechanism which

determines the distributed router's configuration. This may be as simple as selecting the lowest numbered unit to provide the configuration. Alternatively, the fabric could be configurable through a priority system.”), [0100] (“If the unicast router is a distributed router as described in the foregoing, the router interfaces must be synchronised among the devices in the fabric. *Since the multicast router shares those same router interfaces with the unicast router, the multicast router will also have synchronised router interfaces.*”).

68. Weyman describes that when a unit receives a packet from a source via a user port, the unit uses the bridge / forwarding database and the router / RIB to determine the destination user port to which the packet should be forwarded. EX1005, [0065]-[0067]. Once the unit determines the correct user port, it forwards the packet to that destination user port. *Id.* Weyman discloses that, if the destination port is in another unit within the distributed router, Weyman’s unit is able to forward the packet to the other unit via a fabric link, for example, such that the other unit can subsequently forward the packet to the destination port associated with the requesting client. EX1005, [0065]-[0067].

69. Further, Weyman discloses that its distributed router can handle multicast streams comprising multicast packets. EX1005, [0091]-[0093]. I have reproduced Figure 11 of Weyman below.




FIG.11

EX1005, FIG. 11.

70. Figure 11 of Weyman depicts that a unit 10 receives a multicast stream from a multicast source S, and then unit 10 uses its Vbridge 11 (and its forwarding table, not shown) and its router 31 (and its forwarding table, not shown) to transmit a multicast stream, for example, by only routing the multicast to selectively chosen user ports. EX1005, [0103]. In doing so, unit 10 determines that the requested multicast stream may only be forwarded to multicast clients that have requested to receive the particular multicast packet stream. *See, e.g.*, EX1005, [0103] (“Mere routing of multicast traffic controls the forwarding of traffic

between a router's interfaces.”), [0104] (“Many modern bridges have IGMP (internet group management protocol) snooping functionality that controls the bridge's forwarding of IP multicast address traffic within a VLAN. It is preferable, in order to control traffic within a VLAN, that a Vbridge monitors the IGMP traffic in order to determine which user ports connect to a client of a particular multicast stream. Accordingly the snooper must also determine which user ports are connected to other multicast routers.”), [0127] (“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.***”).

71. In one example, Weyman describes that the distributed router only routes the requested multicast packet stream to only the select user ports via which the stream can be forwarded to clients consuming the multicast stream. For example, Weyman discloses:

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. A source

S of the multicast stream is directly connected by way of a user port on unit 10 to Vbridge 11 in unit 10. *A multicast client C1 is also directly connected to the Vbridge 11 and accordingly the stream is directly bridged to the user C1 from the multicast source S.* The stream is also bridged to the fabric port 16 and to the multicast router 31 in unit 10.

The multicast router within unit 10 replicates (in a manner known per se) and routes the stream to Vbridge 12 for forwarding to multicast client C2.

The stream having been bridged to the fabric port 16 by Vbridge 11 in unit 10, the Vbridge 21 within unit 20 receives the multicast stream. Vbridge 21 serves VLAN 1 and so the stream is carried on VLAN 1 between the two devices. *Vbridge 21 within unit 2 forwards the stream to the directly connected client C3* and also to the multicast router within unit 20 for replication of the stream and its routing to Vbridge 22. *Vbridge 22 then forwards the stream to the directly connected clients C4 and C5.*

EX1005, [0122]-[0124].

72. In some examples, Weyman discloses that if a destination port is in the same unit, then Weyman's unit 10 can forward a requested multicast packet stream directly to a particular local destination port. Specifically, Weyman

discloses that if the destination port is disposed in a different unit, then Weyman's unit 10 may (1) forward the requested multicast packet stream to a fabric port, (2) then provide the multicast packet stream to the fabric link (e.g., via the fabric port), and (3) next, provide the multicast packet stream to the other unit, and (4) lastly, forward the multicast packet to the appropriate local destination port. EX1005, [0103]-[0104], [0122]-[0124], [0127]; *see also* EX1005, [0065]-[0066], [0067] ("After the routing operation, the outgoing router interface for the packet will be 158.101.58.254. Router 32 will deliver the packet to the bridge attached to this router interface. Since the MAC address of client 62 is being learned on unit 10, the packet will be forwarded by bridging to the fabric link 30. ***Unit 10 will receive the packet on its fabric port F1 and, based on the destination MAC address of client 62, it will forward the packet by bridging it to port P1 connected to client 62.***").

73. Weyman further discloses that such selective forwarding may be enabled by the Vbridge's forwarding table and the router's forwarding table, which can collectively identify a particular user port and/or fabric port that the multicast stream is permitted to be forwarded over. *See, e.g.*, EX1005, [0111] ("The actual tables maintained by the control plane would depend on whichever multicast routing protocol is in use. In general, the control plane must present a composite forwarding table to the data plane of the multicast router. This composite

forwarding table 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.”); [0112] (Each router forwarding table includes “the router interface on which to accept the stream and the router interfaces to which to forward the stream,” for each multicast stream); [0115] (Each Vbridge forwarding table also includes, “for each Vbridge and for each multicast group within each Vbridge, the ports to which to forward traffic for that group.”); *see also* EX1005, [0105] (“The control of forwarding within an VLAN using a forwarding table derived by the IGMP snooper prevents Vbridge 11 from forwarding the stream to any other user ports. The router has forwarding tables populated by its control plane.”), [0127] (“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.”), [0150].

74. In addition, Weyman goes on to explain that each of the unit’s multicast forwarding tables may be constructed from IGMP messages (e.g., IGMP Membership Reports or IGMP host reports). *See, e.g.*, EX1005, [0130] (“The rules according to which multicast traffic must be forwarded to a user port are no different for a distributed fabric than for a single unit. They may be directly

derived from the requirements of multicast routing and multicast bridging with an associated IGMP snooper.”), [0138] (“The control plane of the multicast routers in both units would use the [IGMP host] report in the construction of the respective forwarding table in the router.”).

75. In some examples, Weyman discloses that the unit may learn the router’s forwarding table from IGMP messages. I have reproduced Figure 11 of Weyman below.




FIG.11


EX1005, FIG. 11.

76. As Weyman shows in Figure 11, when a multicast client C requests to join a multicast stream, the multicast client “send[s] an IGMP host report” to unit 20. EX1005, [0137], [0162]. In this example, Weyman discloses that unit 20 then “forward[s] this report to the unit’s [internal] multicast router [32] and also the fabric port” to send the IGMP membership report to the router 31 in unit 10. EX1005, [0137], [0162].

77. Subsequently, Weyman explains that “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]. In this way, all router forwarding tables within Weymans’s distributed router are able to share the same entries. *See* EX1005, [0065] (“It is assumed that the routing tables contained in the routers 31 and 32 have been synchronised. It is also assumed that the forwarding databases on the units are synchronised and the port connected to client 62 has been learned in forwarding database 41.”), [0162] (“Changes in the multicast membership reported may affect the multicast distribution tree that the router maintains. Any changes to the multicast distribution tree must be kept synchronised for all the units in the fabric. One possibility is to cause merely one unit to process the IGMP report and to send updates to the remainder of the units in the fabric. Alternatively the IGMP message could be delivered to every unit in the fabric for processing.”). Likewise, Weyman discloses that learning the Vbridge forwarding table from “MAC-

addressed traffic” includes frames encapsulating IGMP messages. EX1005, [0045]-[0046], [0127] (“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.*”).

78. Weyman further explains that each unit’s ports can be aggregated to form an aggregated link, and each aggregated link is able to provide “resilience in a distributed router.” EX1005, [0050]; *see also* EX1005, [0086] (“In the foregoing it has been described how a distributed router maintains resilience for its active routing protocols, with reference to the earlier co-pending applications which are incorporated by reference herein. This resilience allows a distributed router to continue routing. To enable resilience throughout the network the network units connected to the distributed router should preferably be configured with redundant physical links to the network that form the distributed router.”). Weyman discloses that an aggregated link “comprises a multiplicity of physical connections [e.g., ports] constituting a single logical connection.” EX1005, [0050]. I have reproduced Figure 6 of Weyman below.



EX1005, FIG. 6.

79. As depicted above, Figure 6 of Weyman shows that two or more ports on the distributed router are aggregated to form a single logical link to a switch (e.g., 91). EX1005, [0088] (“In FIG. 6, a distributed router is formed as before by unit 10 and unit 20 connected by a link 30. *An external switch 91 has a link aggregation to both unit 10 and unit 20. An external switch 92 has an aggregated link connected to unit 10 and unit 20 and a switch 93 has just a simple link to unit 2.* It will be supposed that switch 10 uses a link aggregation control protocol to create an aggregation; the links in switch 2 have been configured manually.”).

80. In this example, Weyman discloses that the aggregated links may provide redundancy in the event of a failure. EX1005, [0089] (“In FIG. 6, a distributed router is formed as before by unit 10 and unit 20 connected by a link 30. An external switch 91 has a link aggregation to both unit 10 and unit 20. An external switch 92 has an aggregated link connected to unit 10 and unit 20 and a switch 93 has just a simple link to unit 2. It will be supposed that switch 10 uses a link aggregation control protocol to create an aggregation; the links in switch 2 have been configured manually.”).

81. Further, Weyman explains that when forwarding to aggregated links:


Each replica of *a multicast stream should be forwarded to an aggregated link only once regardless of how many units over which the aggregated link is distributed.*

Because the bridge controls the aggregated links, this requirement is strictly speaking not directly relevant to the invention. There are various options. For example, *it may be provided that a particular unit is nominated to forward multicast traffic to a specified aggregated link.*

If this option is adopted, any other unit within the fabric port must not forward multicast traffic to user ports that are members of the aggregated link. Alternatively, the forwarding of multicast traffic may be distributed to aggregations according to the unit through which the stream enters the fabric. This option may offer more

opportunities for reducing unnecessary usage of fabric links and so would generally be preferable to the previous option.

EX1005, [0148].



EX1005, FIG. 6.

82. According to Weyman, the distributed router shown in Figure 6 may be implemented to avoid losses of connectivity. EX1005, [0090] (Noting that “it is preferable that manual (non-LACP) link aggregation should not be used on the links between an external network unit and multiple units within the fabric. It should be said that such links do not prevent the network from working normally; the significance is a loss of redundancy which provides resilience to failure.”).

Weyman explains that a “multicast router may use available multicast protocols to organise the distribution of multicast streams to those parts of the network that require it.” EX1005, [0150].

IX. Specific Grounds for Challenge

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

1. Weyman In View Of Li

83. As I explained in Section VIII.A, Weyman discloses a distributed multicast router that includes multiple units and that each of the multiple units can selectively send and/or receive multicast streams using its own multicast forwarding table. In addition, Weyman discloses that its distributed router is capable of forwarding a multicast stream over an aggregated link. For example, Weyman discloses that the aggregated link “comprises a multiplicity of physical connections constituting a single logical connection.” EX1005, [0050], [0148].

Weyman also discloses:

Further, *there needs to be coordination between the units so that where the aggregated link has member ports on more than one unit in the fabric*, there is coordination between the units so that each traffic stream is forwarded to the aggregated link consistently regardless of which unit is forwarding to it. In particular, *a frame should be forwarded to only one member port*

of the aggregated link and where traffic has arrived from the aggregated link it should not be forwarded to another port on the same aggregated link.

EX1005, [0050].


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

84. As I explained in Section IV.D, a POSITA would have understood the importance of controlling the amount of data being transferred over a particular interface (such as a port or an aggregated link), especially with respect to reliability.

85. To that end, Weyman discloses that “[w]hen several individual data switching units are connected to form a fabric [i.e., a distributed router], their functionalities may be distributed among the individual units in the fabric.” EX1005, [0024]. According to Weyman, “[i]t is generally important to allow this distribution to occur so that the entire fabric appears to other entities as if it were a single entity.” EX1005, [0024]. Weyman explains that “*ancillary requirements* are intended to facilitate the *sharing of the workload* between the units.” EX1005, [0024]. Weyman further explains that “there needs to be coordination between the units so that where the aggregated link has member ports on more than one unit in the fabric.” EX1005, [0050].

86. Consequently, a POSITA would have been motivated to find ways to distribute an amount of processing workload amongst switching units to better control bandwidth usage across ports, including any aggregated links.

87. A POSITA would have readily recognized that Li discloses a technique for port bandwidth control. Specifically, Li discloses that an IGMP-based multicast system can include a multicast 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.” EX1007, 26:13-15. I have reproduced Figure 9 of Li below.



EX1007, FIG. 9.

88. As shown above, in Figure 9 of Li, a multicast client, such as one of the set-top boxes (STBs) 92, may send a multicast join request to the multicast router 91. In this example, Li discloses that 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. Thus, Li discloses that the multicast router 91 sends the requested multicast packet stream to the client STB. *Id. See also id.*, FIG. 9.

89. Li further discloses that the multicast router 91 looks to the maximum number of multicast channels allowable on its service port in order to determine whether a requesting multicast client has permission to join a particular multicast channel. Indeed, Weyman expressly discloses that:

[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.

90. Li continues by describing that its port control is performed by the “channel joining processing unit 122” in the multicast router. EX1007, 30:8-9. I have reproduced Figure 12 below.



EX1007, FIG. 12.


91. Li discloses that when the multicast router receives a multicast join request, the channel joining processing unit 122 uses the channel-number comparison unit 126 to “determine 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.

92. Further, Li explains that when the channel-number comparison unit 126 makes a preliminarily determination to reject the multicast join request, the channel joining processing unit 122 can check to determine whether any of the port’s existing multicast streams is currently being terminated, and if not, the multicast join request is rejected by the channel joining processing unit 122. EX1007, 31:13-19.

93. Thus, a POSITA would have found it obvious to combine the distributed multicast router disclosed in Weyman with the channel joining processing unit described in Li, to provide Li’s channel joining processing unit in each of Weyman’s unit (as depicted below).

Weyman in combination with Li



94. The combination of Weyman in view of Li (e.g., “Weyman/Li”) would have controlled the workload of each of individual ports in each unit as well as any aggregated link, for example, using the channel joining processing unit in Li to avoid overloading any port (including any aggregated link). The combination would have controlled the workload of an aggregated link by, for example, treating the aggregated link as a virtual port and limiting the number of multicast streams sent over the virtual port (e.g., the total number of multicast streams sent over the physical ports that comprise the aggregated link), as taught by the channel joining processing unit in Li.

95. A POSITA would have been motivated to modify Weyman's distributed router to include Li's channel joining processing unit to ensure that Weyman's distributed router does not overload any particular port (including the aggregated link in Weyman), thereby degrading the distributed router's reliability, which was a well-known problem in the prior art. *See, e.g.*, Section IV.D [Port Bandwidth Control]; *see also* EX1020, [0003], [0006]. Indeed, Li expressly discloses: "the maximum number of multicast channels that can be selected of a service port ***needs to be controlled***." EX1007, 5:5-23.

96. Moreover, a POSITA would have been motivated to integrate Li's channel joining processing unit into each of the units within Weyman's distributed router for two additional reasons. As an initial matter, Weyman discloses that the processing of each of the multicast join requests (e.g., IGMP Membership Reports) is performed locally, for example, by each of the units. *See, e.g.*, EX1005, [0137]-[0139], [0162] ("Any changes to the multicast distribution tree must be kept synchronised for all the units in the fabric."). Consequently, a POSITA would have been motivated to use Li's channel joining processing unit to control the assignment of multicast streams of Weyman's ports locally at the unit.

97. Secondly, a POSITA would have known that it would have been desirable to control the workload of each unit's ports locally because the alternative option (e.g., controlling the workload of each units' ports globally)

would have required using the internal fabric link to communicate with all of the units in the distributed router. In fact, Weyman expressly discloses that “a multicast router needs to *minimise* its use of the fabric link or links.” EX1005, [0119]; *see also* EX1005, [0076] (“It will however be appreciated that any synchronisation between units of the routing protocols will occupy bandwidth in the fabric links. To minimise the bandwidth used for synchronisation over a fabric link, the routing information base used for routing can be kept synchronised by allowing the routing protocols to run on each device.”), [0098] (“A distributed multicast router should behave similarly in that it should not forward traffic across any fabric link unnecessarily.”).

98. In addition, a POSITA would have been motivated to control the workload of Weyman’s individual ports, as well as any aggregated link, using Li’s channel joining processing unit because the POSITA would have recognized that implementing bandwidth control is desirable across individual ports and across aggregated links because both individual ports and aggregated links have a limited amount of bandwidth. Further, the POSITA would have also recognized that implementing bandwidth control is desirable across individual ports and across aggregated links because Li’s channel joining processing unit would have been just as easy to use with both individual ports and aggregated links.

99. A POSITA would have had a reasonable expectation of success in modifying each of Weyman's units to include Li's channel joining processing unit. For instance, both of Weyman and Li are directed to multicast routing across an access network to deliver multicast streams (e.g., IGMP multicast packets), and both the ports (including aggregated links) described in Weyman and the service ports disclosed in Li are used to deliver these multicast streams. *See, e.g.*, EX1005, [0104]-[0105]; EX1007, 33:16-35:23.

100. Further still, Li discloses that its channel joining processing unit can be configured to control a maximum number of multicast packet streams that are permitted to be transmitted over a particular port, for example, based on a straightforward determination of the number of multicast streams being carried by the particular port. As a result, the POSITA would have also understood that implementing Li's channel joining processing unit in Weyman's distributed router would have merely provided the POSITA with the predictable result of limiting the maximum number of multicast streams that are permitted to be carried over each port disclosed in Weyman (including the aggregated link forming a virtual port). Implementing the bandwidth control would have merely entailed counting the number of multicast streams handled by the aggregated link (e.g., by summing up the number of multicast streams handled by the physical ports that are members of

the aggregated link) to limit the number of multicast streams handled by the aggregated link.

101. Moreover, a POSITA would have recognized that combining each of Weyman's units with one of Li's channel joining processing units would merely have required combining known prior art elements (e.g., Weyman's units and Li's channel joining processing unit) using known methods of implementing logical operations (e.g., software routines) in each of Weyman's units to yield predictable results that include limiting a total number of multicast streams that are permitted to be carried over Weyman's ports (including any aggregated link) to avoid overloading any given one of Weyman's ports (including any aggregated link).

102. Furthermore, combining each of Weyman's units with Li's channel joining processing unit would have merely amounted using a known technique—performing Li's channel-number comparison—to improve similar devices (e.g., each of Weyman's units in Weyman's distributed multicast router) in the same way: by limiting the maximum number of multicast streams that are permitted to be carried over one of Weyman's ports to avoid overloading any particular one of Weyman's ports. Further, the combination would have merely amounted using a known technique—performing Li's channel-number comparison—to improve similar devices (e.g., each of Weyman's units in Weyman's distributed multicast router) in the same way: by limiting the maximum number of multicast streams

that are permitted to be carried over one of Weyman's ports, including the aggregated link.

103. Thus, at a minimum, it would have been obvious to a POSITA to try to implement Li's channel joining processing unit within Weyman's units. Specifically, and as I explained in Section IV.D [Port Bandwidth Control], there was a well-known design need and/or market need to control a port's limited amount of bandwidth, and there were only a finite number of ways to control the port's bandwidth, including for example, measuring a particular port's bandwidth in real-time and responsively throttling the bandwidth, or limiting a maximum number of multicast streams that may be assigned to the particular port. EX1013, [0036]; EX1021, [0225]; EX1022, [0024]; EX1019, 2:37-3:14. And a POSITA would have readily recognized that each of those solutions would have had a reasonable expectation of success, for example, in Weyman's units within the distributed router because each of those techniques would have been implemented as a conventional logic operation (e.g., admission control, software routing, among others) in each of Weyman's units, which would have yielded the predictable result of controlling the port's bandwidth for each of Weyman's units within its distributed router.


2. Independent Claim 1

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

104. To the extent the preamble may be interpreted to be a claim limitation, Weyman/Li teaches the preamble.

105. For instance, Weyman discloses a “distributed router composed of individual units connected by links in a routing fabric.” EX1005, Abstract. Further, Weyman depicts, in Figure 11, that the distributed router includes a first unit 10, which provides a multicast stream to a user port in a second unit 20 in the distributed router. *See, e.g.,* EX1005, [0136]-[0138].

Weyman in combination with Li




ii. [1a] operating a network element,

106. Weyman/Li teaches claim limitation [1a].

107. Specifically, Weyman discloses operating a network element (e.g., a distributed router) that provides connectivity between a multicast source and a multicast client that allows for the delivery of a multicast stream. *See, e.g.*, EX1005, Abstract; *see also* EX1005, [0136]-[0137], [0148], [0151]-[0153]. To do so, Weyman discloses that the distributed router is capable of executing “routing protocols” for multicasting, enabling the distributed router to forward a multicast stream between the multicast source and multicast clients. EX1005, [0006] (“The invention is based on the bridging of multicast traffic to each unit, for example on an ingress VLAN, and thereafter the replication by each unit and local routing so that no replica packets cross the ‘cascade’, that is to say to the system of links between the units. It is preferable for a ‘master’ or ‘lead’ unit in the fabric to run routing protocols and for the other units to synchronise their operation by the action of ‘snooping’.”), [0093], [0107] (“The control plane of the multicast router, including the multicast routing protocols, is used to create and maintain various tables. These are likely to include: (1) For each routing interface, a list of neighbouring multicast routers; (2) for each multicast source, the interface on which streams from this source should be accepted; (3) for each source/group pair, the router interfaces for which this is the designated router, the number of routers that are downstream of each interface for which this is the designated router and

whether each of these routers is pruned from the forwarding tree; and whether there is any locally attached member of the group.”).

Weyman in combination with Li




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

108. Weyman/Li teaches claim limitation [1b].

109. For instance, Weyman discloses that its distributed router (e.g., a “network element”) may include multiple units (e.g., “processing units”) that are interconnected via a fabric link. *See, e.g., 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.”).

110. It is my conclusion that each of Weyman’s units are a “processing unit” because each unit is configured to process multicast requests, such as IGMP host reports, and forward multicast packets streams to their appropriate destinations (e.g., multicast client). EX1005, [0122]-[0124], [0162] (“Report messages can be generated by the client units in order to route query messages or a change in an applicant's multicast membership. Changes in the multicast membership reported may affect the multicast distribution tree that the router maintains. Any changes to the multicast distribution tree must be kept synchronised for all the units in the fabric. One possibility is to cause merely one unit to process the IGMP report and to send updates to the remainder of the units in the fabric. *Alternatively the IGMP message could be delivered to every unit in the fabric for processing.*”).




EX1005, FIG. 10.

111. Each unit is also an “interconnected ... unit” because the unit is interconnected to other units via a fabric link. *See EX1005, [0024] (“When several individual data switching units are connected to form a fabric, their functionalities may be distributed among the individual units in the fabric.* It is generally important to allow this distribution to occur so that the entire fabric appears to other entities as if it were a single entity. Ancillary requirements are intended to facilitate the sharing of the workload between the units, to maximise the data forwarding performance and to minimise the dependence on any single unit or connection within the fabric”), [0025] (“The functionalities required of a


network unit may be divided into two broad categories. They may be characterised as ‘data plane’ and ‘control plane’. *The data plane relates directly to the forwarding of the data traffic, namely the packets, and the control plane relates to the ‘overhead’ activities of establishing the topology of the network in which the unit is deployed.* More specifically the data plane includes the Vbridges, fabric ports, user ports, forwarding databases and forwarding engines. The control plane includes the protocol entities which include, for example, LACP (link aggregation control protocol), STAP (spanning tree algorithm and protocol), IGMP (internet group management protocol), routing protocols and network login. *In the following description, there will be an explanation of the requirements which must be met in practice by the data plane and control plane in each device so that when 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), they will appear from outside the fabric as if they were a single entity.”), [0122]-[0124], [0140], [0151], [0158]* (“It should be noted that any synchronisation between units of the routing protocols will consume bandwidth *from the fabric interconnect.*”).

Weyman in combination with Li



112. Weyman further discloses that its distributed router includes “multiple ports.” Specifically, Weyman describes that each unit within its distributed router includes Vbridges that are connected to one or more user ports. EX1005, [0029] (“In this example, the unit 10 has three ‘Vbridges’ 11, 12 and 13, each with a forwarding database FDB 1, FDB 2, etc. Each Vbridge 11, 12 and 13 is connected to at least one and usually more than one of *a multiplicity of user ports 15* and all have logical connections to each of at least one and in generally more than one fabric port 16.”), [0060] (“FIG. 4 illustrates a unit 10 which has a multiplicity of Vbridges, such as the Vbridges previously mentioned, *user ports 15 to which the*

Vbridges 11 are variously connected, a fabric port 16, to which all the Vbridges are connected, and a router 31 having a respective router interface coupled to each of the Vbridges.”), [0062] (“Thus with reference to FIG. 4, *a packet received on one of the user ports P1 to P24* and needing to be routed would be delivered to the router by the receiving Vbridge on the router interface attached to that bridge”), [0064] (“The unit 10 has *fabric ports P1 to P4*, a fabric port F1, a router 31, forwarding database 41 and a routing information base 51. Unit 2 has *user ports P1 to P4*, a fabric port F2, bridges 21, a router 32, forwarding database 42 and a routing information base 52.”), [0122] (“A source S of the multicast stream is directly connected by way of *a user port on unit 10* to Vbridge 11 in unit 10.”), [0123]-[0125]. Indeed, Figure 10 of Weyman shows that the distributed router’s unit 10 includes a first set of user ports, and unit 20 includes a second set of user ports.




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

113. Weyman/Li teaches claim limitation [1c].

114. As I explained in Section IX.A.2.iii [claim limitation 1[b]], Weyman discloses that its distributed router includes a unit 10 having a first set of user ports, and a unit 20 having a second set of user ports. Thus, Weyman discloses that each of the units 10 and 20 (e.g., “each processing unit”) may be assigned to a different subset of user ports (e.g., “the multiple ports”) of the distributed router (e.g., “the network element.”); *see also* EX1005, [0064] (“FIG. 5 illustrate this operation in greater detail. The fabric in this example consists of units 10 and 20 and link 30. The unit 10 has fabric ports P1 to P4, a fabric port F1, a router 31, forwarding database 41 and a routing information base 51. Unit 2 has user ports P1

to P4, a fabric port F2, bridges 21, a router 32, forwarding database 42 and a routing information base 52.”); FIG. 8.



- 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;

115. Weyman/Li teaches or renders obvious claim limitation [1d].


116. **First**, Weyman discloses that each unit (e.g., “each processing unit”) may include one or more Vbridges, as well as a router that “process[] [multicast streams] communicated over the different subset of the multiple ports to which the processing unit is assigned.” For instance, Weyman discloses that the Vbridges and the router in unit 10 may determine an egress port(s) for a particular multicast stream, for example, when its processing unit (e.g., unit 10) receives a multicast

packet stream that is communicated over one its user ports (a “subset of the multiple ports to which the processing unit is assigned”), and can then forward the multicast packet stream to the appropriate egress port(s) of unit 10 (e.g., including its user ports). EX1005, [0122]-[0124]; [0104] (“Many modern bridges have IGMP (internet group management protocol) snooping functionality that controls the bridge's forwarding of IP multicast address traffic within a VLAN. It is preferable, in order to control traffic within a VLAN, that a Vbridge monitors the IGMP traffic in order to determine which user ports connect to a client of a particular multicast stream. Accordingly the snooper must also determine which user ports are connected to other multicast routers. The Vbridge must then forward each multicast stream to each user port that connects either to another multicast router or to a client of that stream, and the Vbridge should filter the multicast stream from all other user ports.”).

117. Likewise, Weyman discloses that unit 20 may receive a multicast packet stream from another interconnected unit (e.g., unit 10) via the fabric port 16, and in response, unit 20's Vbridges and router can determine the appropriate egress port(s) for the multicast packet stream and forward the multicast stream to the multicast client via that egress port(s), which may also be a “subset of the multiple ports to which the processing unit is assigned.” EX1005, [0124] (“The stream having been bridged to the fabric port 16 by Vbridge 11 in unit 10, the

Vbridge 21 within unit 20 receives the multicast stream. Vbridge 21 serves VLAN 1 and so the stream is carried on VLAN 1 between the two devices. Vbridge 21 within unit 2 forwards the stream to the directly connected client C3 and also to the multicast router within unit 20 for replication of the stream and its routing to Vbridge 22. Vbridge 22 then forwards the stream to the directly connected clients C4 and C5.”), [0104]; *see also* EX1005, [0025].

Weyman in combination with Li



118. **Second**, Weyman describes that the multicast streams which are processed by the Vbridges and router of the units (e.g., within the distributed router) comprise packets. EX1005, [0093] (“Furthermore, it will be presumed that

the time to live count of each data packet will be decremented no more than once in the distributed fabric, the collection of routers is managed as a single multicast router, and each unit within the fabric would be able to route *multicast packets* that it received.”); claim 1 (“distributed router for *multicast packets*”); [0127] (“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.*”).

119. *Third*, a POSITA would have understood that each of the units (e.g., processing units) include one or more Vbridges and a router that collectively form a “processor” because the Vbridges and router collectively “process” the packets in the multicast streams, in order to forward the multicast packet streams to the appropriate egress port(s) associated with the multicast client. EX1005, [0097]-[0098], [0124]-[0125]; *see also* EX1017, [0040]. Further, a POSITA would have recognized that it was well known for a unit (e.g., a processing unit or line card) within a distributed router would include a processor, such as a CPU. *See, e.g.*, EX1008, [0042] (“The line card(s) 820 may include a local ASI switch component 822 that is linked to local framer/media access control (MAC)/physical layer

(PHY) component(s) 824, NPU(s) 826, and/or **CPU(s) 828**”), FIG. 8; EX1017, [0040].

120. At a minimum, it would have been obvious to an ordinarily skilled artisan to implement the one or more Vbridges and router using a processor. A POSITA would have readily recognized that the functions and/or routines of the Vbridges and router—e.g., forwarding packet streams to their appropriate destination ports, including unicast packet streams and multicast packet streams—may be implemented using a processor. *See, e.g.*, EX1005, [0099], [0100]; EX1008, [0042], FIG. 8; EX1017, [0040].


121. It would also have been obvious to a POSITA to combine Weyman/Li’s unit with a processor, in order to implement the functions performed by the one or more Vbridges and the router with a processor because implementing the Vbridges and the router using a processor is merely a design choice with the predictable results of processing packets. In addition, implementing the Vbridges and the router using a processor would have merely amounted to a simple substitution of one known element (e.g., Weyman’s unit (including the Vbridges and router)) for another (e.g., processor) to obtain predictable results of processing packets, including unicast packets and multicast packets. *See, e.g.*, EX1005, [0099], [0100]; EX1008, [0042], FIG. 8; EX1017, [0040].

- 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;**

122. Weyman/Li teaches claim limitation [1e].

123. *First*, Weyman discloses “storing in each processing unit a respective list.” As I described above, in Section IX.A.1 [Weyman/Li], Weyman/Li discloses that each unit includes a Vbridge forwarding table (e.g., in the FDB) and a router forwarding table (e.g., RIB). The Vbridge forwarding table and the router forwarding table collectively include various rules that are implemented for forwarding the multicast packet streams to multicast clients across the entire distributed router, and are thus the claimed “respective list.” *See* EX1005, [0031], [0046], [0069]-[0070].

Weyman in combination with Li



124. Indeed, Weyman explains that each unit within the distributed includes a Vbridge forwarding table (FDB) and a router forwarding table (RIB) (EX1005, [0105], [0115]) and that “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]).///

125. A POSITA would have understood that Weyman/Li’s forwarding table (e.g., the Vbridge forwarding table and the router forwarding table, collectively) may be configured to operate collectively and cooperatively as a single multicast forwarding table because Weyman expressly discloses that each of

its units can “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] (“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.”).

126. It is my conclusion that a POSITA would have also understood that the multicast forwarding tables in the Vbridges and router have the structure of a “list” because Figure 2 of Weyman, which I have reproduced and annotated below, provides the FDB content of a Vbridge forwarding table as addressable entries in a list.





FIG.2

EX1005, FIG. 2.

127. In addition, such a POSITA would have understood that when referring to the packet forwarding rules, the “table” includes the “list”. EX1017, [0051]; EX1009, 6:59-64; *see also* EX1008, [0030], FIG. 4 (below).



EX1008, FIG. 4.

128. **Second**, Weyman/Li discloses “each list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element.” For instance, Weyman describes that the distributed router’s forwarding tables (e.g., the “respective 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]. Weyman further discloses that each unit can forward a multicast packet stream to a port if, and only if, the forwarding tables “indicate that the packet should be so forwarded.” EX1005, [0127].

129. In some examples, Weyman discloses that the multicast forwarding tables can maintain certain forwarding rules for the *entire* distributed router (in other words, for “all of one or more multicast packet streams [for] the network element”). For instance, Weyman discloses that the distributed router’s forwarding tables that reside in each unit are “synchronised.” EX1005, [0046] Indeed, Weyman states that “[t]he 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.*” EX1005, [0046].

130. Further, Weyman discloses that “the forwarding databases on the units are *synchronized.*” EX1005, [0065]. Weyman further discloses that “the multicast router will also have synchronised router interfaces [in the router forwarding table].” EX1005, [0100]; *see also* EX1005, [0162] (“Report messages can be generated by the client units in order to route query messages or a change in an applicant's multicast membership. Changes in the multicast membership reported may affect the multicast distribution tree that the router maintains. *Any changes to the multicast distribution tree must be kept synchronised for all the units in the fabric.* One possibility is to cause merely one unit to process the IGMP report and to send updates to the remainder of the units in the fabric. Alternatively the IGMP message could be delivered to every unit in the fabric for processing.”).

131. **Third**, Weyman/Li discloses “indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to.” For instance, 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] (“For each multicast stream arriving into the fabric on a particular LAN and then for each user port the unit must bridge the stream to the user port if the user port has membership of the same VLAN as the stream’s ingress VLAN ***and this VLAN’s Vbridge multicast forwarding table has this user port in the egress for the multicast stream’s destination group***. Further, for each user port, the unit must route the stream to the user port if the router’s multicast forwarding table indicates that the stream should be accepted from its interface to the stream’s ingress VLAN; ***and the router’s multicast forwarding table indicates that the stream should be routed to the router interface attached to the same Vbridge as this port***; and for the Vbridge forwarding to the VLAN of which the port is a member, ***the multicast forwarding table has this user port in the egress list for the multicast stream’s destination group***.”).

132. Figures 9 and 10 of Weyman depict two exemplary processes. Among them, Weyman discloses that when a Vbridge 11 of unit 10 receives a multicast stream, then Vbridge 11 may look up its Vbridge forwarding table to determine the

appropriate destination of the multicast stream (e.g., the multicast client). In one example, Weyman describes that the Vbridge forwarding table of Vbridge 11 indicates that the multicast packet stream should be delivered to three particular destinations, which include: (1) a local user port that is connected to a multicast client C1; (2) an router 31 that is internal to unit 10; and (3) a fabric port 16 that is connected to unit 20 which serves multiple destination multicast clients. EX1005, [0105], [0115], [0128], FIG. 10. Thus, Weyman discloses that unit 10's Vbridge 11 forwards the multicast packet stream to the local user port, the internal router 31, and to unit 20 via the fabric port 16. *See, e.g.*, EX1005, [0105], [0122]-[0124], FIG. 10. I have reproduced Figure 10 of Weyman below.




FIG.10

EX1005, FIG. 10.

133. As shown above, Weyman depicts that when router 31 receives the multicast packet stream (e.g., from the multicast source and via Vbridge 11), the router 31 may use its router forwarding table (e.g., RIB) to determine that Vbridge 12, which connected via a router interface, will subsequently distributed the multicast packet stream, and as a result, the router 31 sends the multicast packet stream to Vbridge 12. EX1005, [0112] (“The forwarding table used by the data

plane needs to include, for each source/group pair, the router interface on which to accept the stream and the router interfaces to which to forward the stream. It follows that the stream must not be routed if the stream is received on any other router interface and the stream will not be routed to any other router interfaces that are not specified.”), [0123]. Weyman’s Vbridge 12 then uses its Vbridge forwarding table (e.g., FDB) to route the multicast packet stream to the appropriate destination (e.g., to multicast client C2). EX1005, [0105], [0123].

134. Likewise, when Weyman’s unit 20 receives the multicast packet stream from fabric port 16, the Vbridges (e.g., Vbridges 21, 22) and router (e.g., router 31) in unit 20 can use their forwarding tables (e.g., FDB and RIB) to forward the multicast packet stream to the appropriate destination (e.g., via the requesting user ports). EX1005, [0124], [0131].

135. As I explained, each of Weyman’s units can use the Vbridge forwarding table(s) (e.g., FDB) and the router forwarding table (e.g., RIB) (collectively, the “list”) to “forward a multicast stream only to the ports that require it.” EX1005, [0127]. Thus, the “list” in Weyman indicates “which of the one or more ports each multicast packet stream is permitted to be forwarded to.”

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

136. Weyman/Li teaches claim limitation [1f].

137. For example, Weyman discloses 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]. In particular, Weyman discloses, with respect to Figure 11, that “[i]n order to join the destination group for the stream, the multicast client would need to send an IGMP host report.” *See* EX1005, [0137]. I have reproduced and annotated Figure 11 of Weyman below, which highlights the IGMP host report (outlined in green) that is sent from the multicast client to the Vbridge 22 of the unit 20 within Weyman’s distributed router.




FIG.11

EX1005, FIG. 11.

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

138. Weyman/Li teaches claim limitation [1g].

139. In one example, Weyman discloses that when Vbridge 22 of unit 20 receives an IGMP host report (e.g., “the request”), then “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.” EX1005, [0137]. I have reproduced and annotated Figure 11 of Weyman below. In this annotated version, the green boxes highlight that the IGMP host report bridged to the fabric link (e.g., from the multicast client C) is sent from Vbridge 22 across fabric link 30, and the IGMP host report carried on VLAN 2 is sent to router 31 via Vbridge 12.




FIG.11

EX1005, FIG. 11.

140. As shown above, Figure 11 of Weyman only shows the forwarding of the IGMP host report being distributed to a single unit 10 because the figure only two depicts units. However, Weyman expressly discloses that the distribution of the IGMP host report is sent to all other units within the distributed router, such that “[a]ny changes to the multicast distribution tree [is] kept synchronised for all the units in the fabric.” EX1005, [0162] (“One possibility is to cause merely one unit to process the IGMP report and to send updates to the remainder of the units in

the fabric. *Alternatively the IGMP message could be delivered to every unit in the fabric for processing.*”).

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

141. Weyman/Li teaches claim limitation [1h].

142. As Weyman explains, once the IGMP host report is distributed to each unit in the fabric, each of the distributed router’s units “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]. Weyman also discloses that “[IGMP] Report messages can be generated by the client units in order to route query messages or a *change in an applicant's multicast membership*. Changes in the multicast membership reported may affect the *multicast distribution tree* that the router maintains. Any changes to the multicast distribution tree must be *kept synchronised for all the units* in the fabric.” EX1005, [0162]. Therefore, Weyman explicitly teaches updating the forwarding table in response to the IGMP Report message, in addition to the construction of the forwarding table in response to the IGMP Report message.

143. Consequently, it is my conclusion that Weyman teaches that each of the units within the distributed router updates the “forwarding table of each

multicast router [in each unit]” (e.g., “the respective list”) in response to receiving the IGMP host report.

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

144. Weyman/Li teaches claim limitation [1i].

145. As I previously discussed, in Section IX.A.2.vi [claim limitation [1[e]], Weyman explains:

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]. Therefore, Weyman teaches “selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective [] lists.”

146. Further, as discussed in Section IX.A.2.xi[claim limitation 1[j]], Weyman discloses that the “router’s forwarding tables and the bridge forwarding tables” is updated in response to receiving the IGMP host report. Thus, Weyman

discloses that the combined “router’s forwarding tables and the bridge forwarding tables” is an “updated list.”

147. Therefore, Weyman/Li teaches that the unit in the distributed router can “selectively forward” the multicast packet stream in accordance with the router’s forwarding tables and the bridge forwarding tables (e.g., the “respective updated list”).

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

148. Weyman/Li teaches claim limitation [1j].

149. Weyman discloses that “operating the network element comprises aggregating two or more of the [physical connections] to form [an aggregated link].” For example, Weyman discloses the distributed router can forward a multicast packet stream to an aggregated link. EX1005, [0148] (“Each replica of a multicast stream should be forwarded to an *aggregated link* only once regardless of how many units over which the aggregated link is distributed.”). Weyman explains that “[i]t is possible for a host to be connected to a multiplicity of units in a fabric by means of an *aggregated link*, which comprises a *multiplicity of physical connections constituting a single logical connection*.” EX1005, [0050]; *see also* EX1005, [0050] (“For example, for the purpose of transmitting and receiving protocol data frames on an aggregated link where *the protocol considers*

the aggregated link to be a single logical link, each unit in the fabric needs to be coordinated so that protocol frames are handled consistently.”).

150. A POSITA would have understood that the physical connections aggregated in Weyman are the claimed “ports” and that the aggregated link in Weyman is a “high-capacity port.” With respect to Figure 6, Weyman describes that one physical port (e.g., link) in unit 10 and another physical port (e.g., link) in unit 20 are aggregated to form an aggregated link with the switch 91. EX1005, [0013], [0088]. I have reproduced Figure 6 of Weyman below.




FIG.6

EX1005, FIG. 6.

151. In this example, Weyman discloses that the aggregated link shown in Figure 6 can be formed using “a link aggregation control protocol” or “LACP.” EX1005, [0025], [0088]. Those skilled in the art would have readily recognized that link aggregation was understood to mean a bundle of physical ports that are logically grouped together. *See, e.g.,* EX1023, [0002] (“Network devices are interconnected by communication links for transmitting data packets through physical ports. In order to provide redundancy or bandwidth, some networks bundle multiple communication links between a pair of network devices. ***Link Aggregation Control Protocol (LACP) is part of an IEEE specification (802.3ad) that allows several physical ports to be bundled together to form a single logical channel.*** The logical channel has multiple links that connect the network infrastructure device and the network device.”). Thus, it is my conclusion that a POSITA would have understood that Weyman bundles physical ports across both of unit 10 and unit 20 to form an aggregated link. Also, the aggregated link formed by LACP was understood to be a link aggregation group (LAG). EX1024, [0028]-[0029]. The LAG (i.e., Weyman’s aggregated link) is exactly the “high-capacity port” that is disclosed in the ’394 Patent. *See, e.g.,* EX1001, 6:61-7:1 (“In some embodiments, ***two or more ports 44 are aggregated to form a high-capacity port 56,*** in order to provide high traffic capacity and redundancy to a certain content client or content source. For example, ***ports can be aggregated to form a***

Link Aggregation Group (LAG), as specified in the IEEE 802.1AX standard, entitled “IEEE Standard for Local and Metropolitan Area Networks—Link Aggregation,” Nov. 3, 2008, which is incorporated herein by reference.”); *see also* EX1025, 5:49-53.

152. Accordingly, Weyman/Li teaches “aggregating two or more of the ports to form a high-capacity port.”

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,

153. Weyman/Li teaches claim limitation [1k].

154. For instance, Li discloses that “each xDSL service port is configured with the *maximum number of accessed multicast channels allowable*.” EX1007, 28-23-29:1. To do so, Li’s channel joining processing unit is able to determine whether a maximum number of the channels has been exceeded. Specifically, in the context of Weyman/Li’s distributed router, when the distributed router receives a multicast join request, Weyman/Li’s channel joining processing unit determines if the number of channels that are currently being forwarded over the particular port or an aggregated link is below the maximum number of multicast streams that are permitted for forwarding over that particular port or aggregated link, and if not,

then Li's channel joining processing unit rejects the client request to join the multicast stream. EX1007, 31:4-19.

155. Therefore, Weyman/Li's distributed router teaches setting an upper limit (e.g., a maximum number) of the number of multicast streams that are permitted for forwarding over one of its ports and/or aggregated links (e.g., "high-capacity port"), as described in Section IX.A.1. [Weyman/Li].

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

156. Weyman/Li teaches claim limitation [11].

157. Claim limitation [11] is ambiguous as to whether "the request" refers to the "request" recited in claim limitation [1f], or the "request" recited in claim limitation [1g]. However, 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.

158. Consequently, a POSITA would have understood that the "request" includes a "distributed copy of the request." And Weyman discloses that the distributed copy of the IGMP membership report (e.g., the "request") satisfies claim limitation [11] because the data frame encapsulating the IGMP membership report sent via the fabric link.

159. Specifically, Weyman discloses that its units can distribute a copy of the IGMP membership report to other units in the distributed router, for example, by encapsulating the IGMP membership report in a data frame that also includes an identity of an ingress port that may also include an aggregated link (e.g., the “high-capacity port”). For instance, Weyman discloses using a “frame” to communicate a message, such as an IGMP membership/host report (EX1005, [0137], [0162]), over the fabric link. *See, e.g.,* EX1005, [0042] (“Some products, for example the units described in Donoghue et al. *supra*, have hardware support for the learning of MAC addresses directly against any user port within the fabric, regardless of the unit on which the user port is located. In these circumstances, it 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. For example, as described in Donoghue et al. *supra*, ***while a frame is within the fabric, it may include a special header which among other things identifies an ingress port and an egress port employing multiple bit words of which the less significant bits identify a port within a unit and the more significant bits identify a unit, so that the special header can identify uniquely within the fabric both a source port and unit and a destination port and unit.***”).

160. Weyman further describes that the fabric link includes a frame that has “a special header which among other things identifies ***an ingress port*** and an

egress port employing multiple bit words of which the less significant bits identify a port within a unit and the more significant bits identify a unit, 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] (“If the fabric header conforms to that just described it may comprise (as disclosed in the aforementioned application) *a source port identification*, a destination port identification, a field indicating whether a (single) destination port within the fabric is known for the packet, a source trunk identifier and other fields not particularly relevant to the invention.”).

161. Thus, when a unit within Weyman/Li’s distributed router distributes, via the fabric link, an IGMP host report that received from a multicast client over an aggregated link (e.g., the “high-capacity port”) (*see* Section IX.A.2.viii [claim limitation [1g]], EX1005, [0137], [0162]), a POSITA would have readily recognized that the data frame carrying the IGMP host report (e.g., the “distributed copy of the request”) would also have included an identity of the high-capacity port over which the IGMP host report was received at the distributed router. EX1005, [0042], [0043] (“If *the fabric header* conforms to that just described it *may comprise (as disclosed in the aforementioned application) a source port identification*, a destination port identification, a field indicating whether a (single)

destination port within the fabric is known for the packet, a source trunk identifier and other fields not particularly relevant to the invention.”), FIG. 7.

162. Weyman uses the such identification of “an ingress port” or “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]). Therefore, Weyman/Li teaches that the data frame carrying the IGMP host report (e.g., the “distributed copy of the request”)—specifically the identification of “an ingress port” or “source port identification”—indicates that the multicast stream requested by the client sending IGMP host report “is to be forwarded over the high-capacity port.”

- 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.**

163. Weyman/Li teaches claim limitation [1m].

164. As I explained in Sections IX.A.2.xii [1k], IX.A.2.xiii [1l], and IX.A.1 [Weyman/Li], when a unit in Weyman/Li’s distributed router receives the data frame, which includes the IGMP host report, Weyman/Li’s channel joining processing unit within that unit of Weyman/Li’s distributed router determines (e.g., using an identity of which aggregated link is in the header of the received data

frame) whether the number of multicast streams that are currently being forwarded over the aggregated link (e.g., the “high-capacity port”) is below the maximum number of multicast channels that are permitted for forwarding over that aggregated link. *See, e.g.*, EX1007, 31:4-12.

165. And if not, Weyman/Li’s channel joining processing unit rejects the client request to join the multicast packet stream based on the IGMP host report. EX1007, 31:4-31:24. However, if the number of streams currently being sent over a particular port is below the maximum number of multicast channels (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”), then Weyman/Li’s channel joining processing unit in the unit of Weyman/Li’s distributed router may permit client to join the multicast packet stream by accepting the IGMP host report (e.g., the multicast join request). *See, e.g.*, EX1007, 30:21-24.

166. Then, the unit within Weyman/Li’s distributed router processes the IGMP host report (e.g., along with the identity of the aggregate link in the header of the received frame) to modify the particular unit’s multicast forwarding tables (e.g., the “respective list”). *See* EX1005, [0137]-[0139], [0162], Section IX.A.2.ix [claim limitation [1h]]. In doing so, to the modified “list” indicates that the multicast stream requested by the IGMP host report 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*...”); IX.A.2.vi [claim limitation [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).

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**

167. Weyman/Li teaches the additional claim limitation of [4[a]].

168. As I explained, in Section IX.A.2.xiii [claim limitation [1[l]],

Weyman teaches distributing a copy of the IGMP membership report (e.g., the “request”) to other units by encapsulating the IGMP membership report in a frame. The frame also includes a header that also includes “an identity of an ingress port”—the port over which the IGMP membership report was received by the network element. *See also* EX1005, [0041]-[0043], [0137]-[0138], [0162].

- 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.**

169. Weyman/Li teaches claim limitation [4[b]].

170. Specifically, as I explained in Section IX.A.2.xiii [claim limitation [1[I]], when a unit within Weyman/Li's distributed router 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 can modify its multicast forwarding tables (e.g., the "respective list") (EX1005, [0137]-[0139], [0162]) indicating that the requested multicast stream is "permitted for forwarding over the ingress port" (EX1005, [0042], [0127], [0128]).

171. **First**, Weyman discloses that the multicast forwarding table in a unit of the distributed router may identify a particular user port (e.g., the "ingress port") where the message from the multicast client was received, regardless of which unit was assigned to the user port associated with the multicast request. EX1005, [0042]. Specifically, Weyman discloses that "it 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." EX1005, [0042].

172. **Second**, as discussed in Section IX.A.2.xiii[claim limitation [1]], Weyman discloses that the data frame carrying the IGMP host report would

include a unique identifier corresponding to the “ingress port.” EX1005, [0042]-[0043].

173. *Third*, Weyman discloses using the IGMP host report to update the multicast forwarding table in the unit. EX1005, [0138]-[0139], [0162]. As discussed in Sections IX.A.2.vi [claim limitation [1e]] and IX.A.2.xiv [claim limitation [1m]], the unique identifier for the “ingress port” in the data frame of the IGMP host report may be used to update the multicast forwarding table in the unit of the distributed router, for example, to indicate that the requested multicast packet stream that is indicated by the IGMP host report is permitted for forwarding over the ingress port, which satisfies claim limitation [4[b]].

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.

174. Weyman/Li teaches the additional claim limitation of claim 6.

175. Specifically, 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]. Since Weyman discloses that the IGMP host report (e.g., the “request”) is a request to receive a multicast stream by joining the destination group, Weyman/Li teaches the additional claim limitation of claim 6.

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.

176. Weyman/Li teaches the additional claim limitation of claim 7.

177. Indeed, Weyman discloses that “to join the destination group for the stream, the client would need to send an IGMP host report.” EX1005, [0137]. Further, Weyman explains that “internet group management protocol” may be abbreviated as IGMP. EX1005, [0025].


6. Independent Claim 11

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

178. To the extent the preamble may be interpreted to be a claim limitation, Weyman/Li teaches the preamble as I explained in Section IX.A.2.i [element [1[pre]]].


ii. [11a] multiple ports; and

179. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.iii [claim limitation [1[b]]].




iii. [11b] multiple interconnected processing units,

180. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.iii. [claim limitation [1b]].



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

181. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.iv. [claim limitation [1c]].



- 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,

182. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.v. [claim limitation [1d]].

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

183. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.v. [claim limitation [1d]].

- 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**

184. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.vi. [claim limitation 1e]].

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

185. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.vii. [claim limitation [1f]].

- 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,**

186. Weyman/Li teaches this claim limitation as I explained in Sections IX.A.2.viii. [claim limitation [1g]]-IX.A.2.x [claim limitation [1i]].

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

187. Weyman/Li teaches this claim limitation as I explained in Section IX.A.2.xi. [claim limitation [1j]].

- 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**

188. Weyman/Li teaches this claim limitation as I explained in Section

153. [claim limitation [1k]].

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

189. Weyman/Li teaches this claim limitation as I explained in Section

IX.A.2.xiii. [claim limitation [11]].

- 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.**

190. Weyman/Li teaches this claim limitation as I explained in Section

IX.A.2.xiv. [claim limitation [1m]].

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.

191. Weyman/Li teaches the additional claim limitation of claim 14 as I explained in Section IX.A.3. [Claim 4].

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.

192. Weyman/Li teaches the additional claim limitation of claim 17 as I explained in Section IX.A.5. [Claim 7].

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

1. Weyman/Li in view of Rooholamini

193. As I explained in Section IX.A.1 [Weyman/Li], the channel joining processing unit in each unit within Weyman/Li's distributed router uses the maximum number of multicast streams that are currently being transmitted over a particular port to determine whether a multicast packet stream is permitted for forwarding over the particular port. As a result, a POSITA would have been motivated to monitor and record the number of multicast streams that are currently being distributed over the particular port to be used in Weyman/Li's distributed router.

194. Indeed, recording a number of multicast streams that are being transmitted over a specific port was well known in the art. For instance,

Rooholamini provides support for packet-based multicasting across links using network components, including ports on a line card. *See, e.g.*, EX1008, [0003], [0042], [0047]-[0049]. Rooholamini further discloses that a switch fabric (e.g., a distributed multicast router) is able to manage delivery of multiple multicast streams by maintaining a “port path count table.” EX1008, [0029]-[0030]. According to Rooholamini, the port path count table tracks a number of the multicast streams that are passing through a particular port of a switch in the switch fabric/distributed multicast router. EX1008, [0029]-[0030]. In one example, Figure of Rooholamini depicts an illustrative port path count table. I have reproduced Figure 4 of Rooholamini below.

	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

EX1008, FIG. 4.

195. As shown above, Figure 4 of Rooholamini depicts that the rows in the port path count table correspond to switches $S_1, S_2, \dots, S_n$; and the columns correspond to ports $P_0, P_1, \dots, P_m$ that are disposed of within each of those switches. EX1008, [0030]. Figure 4 of Rooholamini further depicts that the numbers in the body of the port path count table correspond to a number of multicast streams that are currently being transmitted over the corresponding port of the particular switch within the switch fabric/distributed multicast router. EX1008, [0030]. Further, Rooholamini discloses updating the port path count table may be performed during operations of the distributed multicast router, for example, in a manner that maintains a count of the number of multicast streams passing through a particular port of the particular switch within the distributed multicast router in real time. *See, e.g.,* EX1008, [0029], [0035].

196. A POSITA would have found it obvious to combine Weyman/Li with the port path count table of Rooholamini in order to identify the number of multicast packet streams that is being transmitted over a particular port (including any aggregated link) within the distributed router. The combination of Weyman's distributed multicast router, Li's channel joining processing unit, and Rooholamini's port path count table is illustrated in the figure below.

Weyman in combination with Li and Rooholamini



By combining Weyman, Li, and Rooholamini as described above, a POSITA would have understood that Weyman/Li/Rooholamini's unit in the distributed router would control the number of multicast streams that is permitted for forwarding over a particular port (including any aggregated link), as taught by the channel joining processing unit in Li, by using the port path count that is provided locally in each unit, as taught by Rooholamini.

197. A POSITA would have been motivated to combine Weyman/Li with Rooholamini as shown above because a POSITA would have known that, in order to implement Li's channel joining processing unit, the POSITA would need to reliably track the number of streams being transmitted over a given 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 Rooholamini expressly contemplates that its port count table may be used to manage a number of multicast streams that is currently being served by the distributed router’s ports. *See, e.g.,* EX1008, [0029] (“By tracking the number of paths that go through the ingress and egress ports in the switches 212-1 to 212-3 using the port path count tables 234, 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.*”). It is my conclusion that a POSITA would have been motivated to use the port count table to monitor and keep track of the number of multicast streams that are concurrently being handled by any port in Weyman/Li’s system, including any aggregated link, because such information would have allowed Weyman/Li’s channel joining processing unit to manage the bandwidth of any port in Weyman/Li’s system, including any aggregated link.

198. Further, 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] (“Furthermore, a fabric containing a multicast router

needs to *minimise its use of the fabric link or links.*”); *see also* EX1005, [0076], [0098].

199. A POSITA would have had a reasonable expectation of success in combining Weyman/Li’s distributed multicast router with Rooholamini’s port path count table because Weyman, Li, and Rooholamini all relate to a distributed router (e.g., a switch fabric) that is capable of routing multicast packet streams over a network. EX1005, [0004]; EX1007, 48:5-18; EX1008, [0004]. Therefore, the combination of Weyman/Li and Rooholamini (e.g., Weyman/Li/Rooholamini) would have merely required a combination of known prior art elements (e.g., Weyman/Li’s multicast forwarding tables and Rooholamini’s port path count table) according to known methods (e.g., implementing various tables in a memory on the unit within the distributed router) to yield predictable results (e.g., maintaining count information in Weyman/Li’s multicast forwarding tables that allow Li’s channel-number comparison unit to use the port path table to determine whether to accept a multicast join request for a particular port.).


200. As explained below, the Weymay/Li/Rooholamini combination renders claims 5 and 15 obvious.

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

201. Weyman/Li/Rooholamini teaches the additional claim limitation of 5[a].

202. As I explained in Section IX.A.2.xi. [claim limitation 1[j]], Weyman/Li's distributed router may use an aggregated link (e.g., the "high-capacity port") that aggregates one port in a unit (e.g., unit 10) and another port that is disposed of in a different unit (e.g., unit 20). By way of an example, I have reproduced Figure 6 of Weyman below.



EX1005, FIG. 6.


203. Therefore, Weyman/Li/Rooholamini's teaches that "the ports forming the high-capacity port are assigned to at least two of the processing units."

[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.

204. Weyman/Li/Rooholamini teaches claim limitation 5[b].

205. **First**, Weyman/Li/Rooholamini discloses “holding respective local copies of the number of the multicast packet streams that are permitted for forwarding over the high-capacity port.”

206. As I explained in Section IX.B.1 [Weyman/Li/Rooholamini], each unit within Weyman/Li/Rooholamini's distributed router is capable of holding a port count table, which is a copy local to the unit as shown below.



207. As shown above, the port count table in the unit within Weyman/Li/Rooholamini's distributed router includes the current number of the multicast packet streams that are being forwarded and that are permitted for forwarding over the ports in the distributed router, including the aggregated link.

208. **Second**, Weyman/Li/Rooholamini discloses “updating the local copies based on the distributed request.” Specifically, Rooholamini discloses 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].

209. Thus, Weyman/Li/Rooholamini's port count table in each of the units are updated when there is a change to the number of multicast streams that are currently being transmitted over a particular port (such as the aggregated link) for that particular unit. Each change may be triggered by the reception of an IGMP host report (e.g., the “request”), for example, when it is distributed and processed by each of the units in the distributed router. *See, e.g.*, EX1005, [0137]-[0139], [0162] (“IGMP message could be delivered to every unit in the fabric for processing.”).

210. **Third**, Weyman/Li/Rooholamini discloses “selectively updating the list in each of the at least two processing units based on the respective local copies.” In one example, Weyman/Li/Rooholamini discloses that

Weyman/Li/Rooholamini's channel joining processing unit in each unit within the distributed router is capable of using an updated local port count table to determine whether a newly-requested multicast packet stream is permitted for forwarding over a particular port. *See, e.g.*, EX1007, 28:23-29:1, 31:4-12. Further, Weyman/Li/Rooholamini discloses that Weyman's unit has an ability to then selectively update its multicast forwarding table to note that the requested multicast stream is permitted forwarding. EX1005, [0138]-[0139], [0162].

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.

211. Weyman/Li teaches the additional limitation of claim 15 as I explained in Section IX.B.2. [Claim 5].

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

1. Weyman/Li In View Of Pitcher

212. As discussed in Sections VIII.A. [Weyman] and IX.A.1 [Weyman/Li], each unit in Weyman/Li's distributed router includes multicast forwarding tables.

Those multicast forwarding tables include the ports (including any aggregated link) over which the multicast stream should be forwarded. EX1005, [0112], [0115].

213. A POSITA would have understood that entries in the multicast forwarding tables may become stale (e.g., have not been used to route multicast streams for a predetermined period of time), and that stale entries are undesirable because they unnecessarily consume valuable network resources, such as space in a memory. EX1005, [0156] (The multicast forwarding table is stored “in non-volatile storage”); EX1026, 5:64-6:4. As a result, a POSITA would have been motivated to identify and remove stale entries in multicast forwarding tables.

214. A POSITA would have also known that stale entries in multicast forwarding tables may be removed periodically, for example, using a timer. Indeed, Pitcher discloses that each entry of a multicast group destination list (e.g., a multicast forwarding table) may be associated with a port report expiration timer. EX1009, 7:56-62. The POSITA would have recognized that when a switch (e.g., a unit in Weyman/Li’s distributed router) receives an IGMP membership report (e.g., a multicast request to join a multicast stream), then the switch would add the “receiving port” to the appropriate multicast forwarding table, for example, “if it is not already a current member” and would reset the corresponding port report timer. EX1009, 7:62-66. Further, the POSITA would have understood that “[i]f the port

report timer expires, the port is removed from the multicast group destination list [multicast forwarding table],” as disclosed by Pitcher. EX1009, 7:66-67.

215. 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, in order to remove stale entries in the unit’s multicast forwarding tables. A POSITA would have understood that stale entries create inefficient use of memory space in the unit, and would have been motivated to remove such stale entries. EX1026, 5:64-6:4.

216. In addition, a POSITA would have had a reasonable expectation of success when combining the unit within Weyman/Li’s distributed router and Pitcher’s timer-based entry removal technique because doing so would have merely required using a known technique (Pitcher’s a table entry removal technique using a timer) to improve a similar device (Weyman/Li’s unit in the distributed router, which performs substantially similar function as a switch) in substantially the same way (e.g., by removing stale entries from the unit’s multicast forwarding tables).

217. The combination of Weyman/Li’s unit within the distributed router and Pitcher’s timer-based table entry removal would have merely been an application of a known technique (Pitcher’s timer-based entry removal technique) to a known device (Weyman/Li’s unit) that was ready for improvement (e.g., the

stale entries in multicast forwarding tables unnecessarily consuming valuable space in memory), which would have yielded predictable results (e.g., removing stale entries from the unit's multicast forwarding tables and improving the overall efficiency of the storage in Weyman/Li/Pitcher's unit).

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**

218. Weyman/Li/Pitcher teaches the additional claim limitation of [16a].

As I explained in Section IX.C.1 [Weyman/Li/Pitcher], each of Weyman/Li/Pitcher's unit in the distributed router tracks the use of entries in the multicast forwarding tables using a timer. EX1009, 7:56-62. For instance, Weyman describes that entries in the multicast forwarding tables (e.g., the Vbridge forwarding table) may indicate "a permission to forward a certain multicast packet stream over a certain port." EX1005, [0112], [0115]. And Weyman/Li/Pitcher's timer is updated, for example, by monitoring IGMP membership reports (e.g., the "requests") relating to the certain multicast stream that are received at the Weyman/Li/Pitcher's unit, as disclosed by Pitcher. EX1009, 7:62-67.

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

219. Weyman/Li/Pitcher teaches claim limitation [16b].

220. Specifically, Pitcher discloses that “[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 multicast forwarding tables].” EX1009, 7:66-67. Since each unit has its own port report timer for each entry of the forwarding table as I explained in Section IX.C.1 [Weyman/Li/Pitcher], the “processing units are configured ... to remove the permission from the respective lists” as required by limitation [16b].

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

1. Weyman In View Of Jain And LeMaire

221. As I explained in Section VIII.A. [Weyman], Weyman discloses a distributed router with multiple units, each of which is capable of selectively forwarding multicast streams using a multicast forwarding table. Further, Weyman discloses that “the system [e.g., the distributed router] can accommodate the *addition* or removal of unit.” EX1005, [0069] (“For the routers to operate as a single distributed router every defined router configuration should be identical on all units in the fabric. Preferably the configuration is stored in non-volatile storage

in all the units. *If the system can accommodate the addition or removal of units, then when a unit with a default configuration joins the fabric, it needs to update its configuration to that of the distributed router.*”).

222. Indeed, 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]. However, Weyman does not explicitly disclose how the newly added unit would “update its configuration to that of the distributed multicast router,” including the multicast forwarding tables in the unit. As a result, a POSITA would have been motivated to identify one or more ways to perform an update to the initial configuration of the newly added unit.

223. As explained below, a POSITA would have known that certain features of an 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.

i. Disabling Ports During Initialization In View Of Jain

224. Jain discloses that adding a node in a network (e.g., a switch), for example, after the node has been down (e.g., “reset, power down, rebooting”), may involve executing a “controlled boot-up process” by the node and for its configuration. EX1010, [0019] (“When network node 102A is down for some reason, e.g., a catastrophic switch software and/or hardware failure, reset, power

down, rebooting, etc., its link port interfaces are deactivated or de-energized, including the device-side and network-side interfaces. Accordingly, upon sensing the link down condition on link 110-1, remote device 104A switches over to its hot standby port 108-2 for communicating with remote device 104B via network node 102B. *When network node 102A is powered up again (i.e., power-on or system reset condition), it engages in a controlled boot-up process thereafter, as will be set forth in detail below, such that its link ports do not 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.* As pointed out in the Background section of the present disclosure, such a situation gives rise to various deleterious effects with respect to the integrity of the traffic received at the network node.”).

225. Jain discloses, with respect to the flowchart in Figure 2, that the controlled boot-up process may involve issuing a “Link Down” command on at least some of the node’s ports handling data traffic to disable them. EX1010, [0020] (“Upon detecting the reset condition and in response thereto, 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).”). I have reproduced the flowchart shown in Figure 2 of Jain below.




FIG. 2

EX1010, FIG. 2.

226. According to Jain, “[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, upon 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.” EX1010, [0021].

227. A POSITA would have been motivated to combine the distributed router disclosed by Weyman with the controlled boot-up process, as disclosed in Jain, to ensure that, when a unit is newly added to Weyman’s distributed router, the user ports are deactivated or disabled to prevent the ports from erroneously providing multicast streams to the network during the initialization process. To be clear, a POSITA would not have disabled the fabric port on the newly added unit because Weyman 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 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). *See also* EX1005, [0156] (“For the routers to operate as a single distributed multicast router

each defined multicast configuration *needs to be identical* on all units in the fabric.”).

228. Indeed, a POSITA would have been motivated to implement Jain’s controlled boot-up process because when a 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 recognized that beginning network communications without proper initialization would have been undesirable because doing so “gives rise to various deleterious effects with respect to the integrity of the traffic received at the network node.” EX1010, [0019].

229. It is my conclusion that a POSITA would have had a reasonable expectation of success when combining Weyman’s distributed router with Jain’s controlled boot-up process because the combination would merely involve using a known technique (e.g., Jain’s disabling of certain user ports for a new network node during initialization) to improve a similar device (e.g., a newly added unit in Weyman’s distributed router) in the same way (e.g., by temporarily disabling certain user ports such that the newly added unit would not inadvertently transmit unnecessary/erroneous communications or cause potential network disruptions).

ii. Initialization of New Unit In View Of LeMaire

230. LeMaire discloses a bridge for handling multicast streams. EX1011, Abstract (“A bridge (62) capable of reducing unwanted WAN multicast packet traffic in a LAN and an associated method and an associated LAN is disclosed. The bridge examines the contents of WAN multicast query and report packets and WAN multicast router-to-router discovery packets and includes this information in its filtering database (88). The bridge designates ports on which query packets or MRRDPs are received as router ports. When there is no WAN router generating multicast query packets into the LAN, the bridge simulates the behavior of a WAN router and generates WAN multicast query packets so as to cause report packet generation by end systems. A timeout interval for undesignating router ports and a method for determining an appropriate timeout for a port is disclosed.”).

231. LeMaire describes that, in Figure 5, a bridge 62 includes multiple ports that are in communication with other devices in the network. EX1011, 7:64-8:2. LeMaire discloses that the bridge 62 also includes a “Bridge Filtering Table 88,” and that “[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. I have reproduced Figure 5 of LeMaire below.




FIG. 5

EX1011, FIG. 5.

232. LeMaire further discloses that when the bridge 62 is first added to the network (i.e., during the “power up”), the bridge 62 transmits IGMP queries that request a multicast membership status from each of the multicast clients coupled to the ports of bridge 62. EX1011, 9:12-19. “In response to the IGMP Queries, [multicast clients that are] attached to the bridge 62 that wish to receive [multicast streams] will transmit IGMP Reports” to the bridge 62. EX1011, 9:21-23. According to LeMaire, the bridge is then able to use the IGMP membership reports to populate the bridge filtering table 88. EX1011, 9:23-34.

233. A POSITA would have found it obvious to combine Weyman/Jain’s distributed router with LeMaire’s bridge initialization technique when a new unit is added to Weyman/Jain’s distributed router such that Weyman/Jain’s distributed router uses LeMaire’s initialization technique to update its multicast forwarding tables while the new unit’s user ports are disabled.

234. For example, when a new unit is added to Weyman/Jain’s distributed router, its user ports are disabled, as discussed in Section IX.D.1.i [Weyman/Jain], and Weyman/Jain’s distributed router would then be able to transmit IGMP queries to its multicast clients, as taught by LeMaire, using other units in the distributed router. *See, e.g.*, EX1011, 9:12-19 (“On power up or system reset or when none of its ports are designated as Router Ports, an improved bridge 62 according to the invention, transmits IGMP Query packets at regular intervals out of each of its

ports A-E. Initially, each of these ports are designated by bridge 62 as non-Router Ports. Bridge 62 transmits IGMP Queries with a WAN source address which is assigned to it according to the invention.”).

235. The other units in Weyman/Jain’s distributed router would then receive IGMP membership reports from the multicast clients in response to the IGMP queries, and those other units would forward those IGMP membership reports to the new unit via the fabric link and the new unit’s fabric port, as taught by LeMaire and Weyman/Jain. *See* EX1011, 9:21-23; EX1005, [0137]-[0139], [0162]. Weyman/Jain’s new unit would then update its multicast forwarding table in accordance with the received IGMP membership reports, as taught by both Weyman/Jain and LeMaire. EX1005, [0138]-[0139]; EX1011, 9:23-34. Further, the POSITA would have understood that once the initialization process is completed, the unit can activate the previously disabled user ports to initiate communication. EX1008, [0021], [0007].

236. A POSITA would have been motivated to combine Weyman/Jain’s distributed router with LeMaire’s bridge initialization technique as described above. Although Weyman discloses adding a new unit to the distributed router and that some initialization is needed (EX1005, [0069], [0156]), Weyman does not provide details of such initialization process. However, a POSITA would have

readily recognized that LeMaire provides the missing implementation details for the initialization in Weyman.

237. It is my conclusion that a POSITA would have had a reasonable expectation of success in combining Weyman/Jain's distributed router with bridge initialization technique, as disclosed by LeMaire, because both Weyman/Jain and LeMaire are directed to provisioning a network element that supports multicast packets streams, for example, using an IGMP multicast protocol to create multicast forwarding tables. *See, e.g.*, EX1005, [0104]-[0105]; EX1011, 4:42-5:4.

238. A POSITA would have found the combination obvious because it would merely require using a known technique (e.g., LeMaire's bridge initialization technique) to improve a similar device (e.g., a new unit in Weyman/Jain's distributed router, which performs substantially similar functions as the bridge disclosed in LeMaire that routes multicast streams, as discussed above) in the same way (e.g., the newly added unit can initialize the multicast forwarding table and function as a unit within Weyman/Jain's distributed router).

239. In addition, a POSITA would have found the combination obvious also because it would have merely been a combination of prior art elements (e.g., adding a new unit to 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 multicast forwarding table and functioning as a unit within Weyman/Jain's distributed router).

2. Independent Claim 8

240. I note that claim elements [8pre]-[8i] are identical to claim elements [1pre]-[1i].

i. [8pre] A method for communication, comprising:

241. To the extent the preamble may be interpreted to be a claim limitation, Weyman/Jain/LeMaire teaches the preamble as I explained above, in Section IX.A.2.i. [element [1pre]]. *See* EX1005, Abstract. EX1005, [0136]-[0138], FIG. 11.

ii. [8a] operating a network element,

242. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.ii. [claim limitation [1a]]. *See* EX1005, Abstract, [0006], [0093], [0107], [0136]-[0137], [0148], [0151]-[0153], FIG. 11.

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

243. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.iii. [claim limitation [1b]]. *See* EX1005, [0024]-[0025], [0029], [0060], [0062], [0064], [0122]-[0125], [0140], [0151], [0158], [0162], FIG. 11.

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

244. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.iv. [claim limitation [1c]]. *See* EX1005, [0024]-[0025], [0029], [0060], [0062], [0064], [0122]-[0125], [0140], [0151], [0158], [0162], FIGS. 8, 11.

- v. **[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;**

245. Weyman/Jain/LeMaire teaches or renders obvious this claim limitation as I explained in Section IX.A.2.v. [claim limitation [1d]]. *See* EX1005, [0025], [0093], [0097]-[0098], [0104], [0122]-[0125], [0127], FIGS. 8, 11.

- vi. **[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;**

246. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.vi. [claim limitation [1e]]. *See* EX1005, [0046], [0065], [0100], [0105], [0111]-[0112], [0115], [0122]-[0124], [0127]-[0128], [0131], [0162], Figures 2, 10.

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

247. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.vii. [claim limitation [1f]]. *See* EX1005, [0137], [0160], FIG. 11.

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

248. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.viii. [claim limitation [1g]]. *See* EX1005, [0137], [0160], [0162], FIG. 11.

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

249. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.A.2.ix. [claim limitation [1h]]. *See* EX1005, [0138], [0162].

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

250. Weyman/Jain/LeMaire teaches this claim limitation, as I explained in Section IX.A.2.x. [claim limitation [1i]]. *See, e.g.,* EX1005, [0127]; *see also* Sections IX.D.2.v.-IX.D.2.vi., IX.D.2.ix. [claim limitations [8[d], 8[e], [8h]].

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

251. Weyman/Jain/LeMaire teaches this claim limitation.

252. For instance, Weyman discloses that its 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]).

253. And, as I explained in Section IX.D.1.ii [Weyman/Jain/LeMaire], when adding a new unit, the new unit may be configured to be initialized by sending and receiving IGMP query messages and subsequently updating the multicast forwarding table. *See also* Sections IX.D.2.xii.-IX.D.2.xiv. [claim limitations [8k]-[8m]].

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

254. Weyman/Jain/LeMaire teaches this claim limitation.

255. As I explained in Section IX.D.1.ii [Weyman/Jain/LeMaire], adding a new unit to Weyman/Jain/LeMaire's distributed router may involve the distributed router disabling (“deactivat[e]”) the user ports of the new unit. *See* EX1010, [0020] (“FIG. 2 is a flowchart of a scheme for controlling boot-up process of an Ethernet node in accordance with one embodiment. When the Ethernet node (such as, e.g., network node 102A described above in reference to FIG. 1) is put into a system reset or power-up condition, the boot-up process involving initialization and configuration of both hardware and system software is commenced (blocks 202 and 204). Upon detecting the reset condition and in response thereto, 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).”).

xiii. [81] 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

256. Weyman/Jain/LeMaire teaches this claim limitation.

257. *First*, Weyman/Jain/LeMaire discloses initializing a newly added unit “while the ports to which the new processing unit is assigned are deactivated.”

258. As I explained in Section IX.D.1.ii [Weyman/Jain/LeMaire], when adding a new unit to Weyman/Jain/LeMaire’s distributed router, the distributed router can disable user ports. Section IX.D.1.i [Weyman/Jain] (*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.”).

259. *Second*, Weyman/Jain/LeMaire discloses that the initialization step includes “sending a multicast status request.”

260. Further, as I explained in Sections IX.D.1.i- IX.D.1.ii. [Weyman/Jain/LeMaire], while the ports are disabled, the distributed router transmits IGMP queries to its multicast clients, as taught by LeMaire, using other units in Weyman/Jain/LeMaire's distributed router. EX1011, 9:12-19.

261. A POSITA would have understood that a IGMP query is a “multicast status request,” which involves asking for a multicast “group membership” of the requesting multicast clients. EX1005, [0161]; *see also* EX1017, [0061].

262. **Third**, Weyman/Jain/LeMaire discloses that the initialization step includes “receiving multicast membership reports sent by multicast clients in response to the multicast status request.”

263. As I explained in Sections IX.D.1.i- IX.D.1.ii. [Weyman/Jain/LeMaire], while the ports are disabled, “[i]n response to IGMP Queries,” Weyman/Jain/LeMaire's other units can receive IGMP membership reports from the multicast clients and forward those IGMP membership 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].

264. **Fourth**, Weyman/Jain/LeMaire discloses the initialization step involves “constructing the list by the new processing unit based on the received multicast membership report.”

265. As I explained, in Sections IX.D.1.i- IX.D.1.ii.

[Weyman/Jain/LeMaire], while the ports of Weyman/Jain/LeMaire's new unit are disabled, the new unit can use the IGMP membership reports it receives to update the multicast forwarding table. *See, e.g.*, 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.

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

266. Weyman/Jain/LeMaire teaches this claim limitation.

267. As I explained, in Sections IX.D.1.i- IX.D.1.ii.

[Weyman/Jain/LeMaire], after the new unit completes the initialization process, which includes constructing the multicast forwarding table, Weyman/Jain/LeMaire's new unit can activate the previously disabled user ports to initiate network communications. 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.").

3. Independent Claim 18

268. I note that claim elements [18pre]-[18h] are identical to claim elements [11pre]-[11h].

i. [18pre] A network element, comprising;

269. To the extent the preamble may be interpreted to be a claim limitation, Weyman/Jain/LeMaire teaches the preamble as I explained in Section IX.D.2.i. [element [8pre]].

ii. [18a] multiple ports; and

270. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.iii. [claim limitation [8b]].

iii. [18b] multiple interconnected processing units,

271. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.iii. [claim limitation [8b]].

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

272. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.iv. [claim limitation [8c]].

v. [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,

273. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.v. [claim limitation [8d]].

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

274. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.v. [claim limitation [8d]].

- vii. **[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,**

275. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.vi. [claim limitation [8e]].

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

276. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.vii. [claim limitation [8f]].

- ix. **[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,**

277. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Sections IX.D.2.viii-IX.D.2.x. [claim limitations [8g]-8[i]].

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

278. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.xi. [claim limitation [8j]].

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

279. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.xii. [claim limitation [8k]].

- 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**

280. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.xiii. [claim limitation [8l]].

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

281. Weyman/Jain/LeMaire teaches this claim limitation as I explained in Section IX.D.2.xiv. [claim limitation [8m]].

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

1. Weyman/Jain/LeMaire In View Of Li

282. As I explained, in Section IX.A.1. [Weiman/Li], a POSITA would have been motivated to seek out implementation techniques that are designed to share a workload between units to control the amount of bandwidth that is concurrently being utilized across its ports (including aggregated link), as taught by Weyman/Jain/LeMaire's distributed router.

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

284. A POSITA would have found it obvious to combine Weyman/Jain/LeMaire's distributed multicast router with Li's channel joining processing unit to perform Li's channel joining processing technique in each of Weyman/Jain/LeMaire's units. As I described in Section IX.A.1. [Weiman/Li], the reasons that a POSITA would be motivated to combine Weyman and Li are also equally applicable here because the initialization process, as described by Jain and LeMaire, are independent of the port bandwidth control described in Li. Additionally, combining Weyman/Jain/LeMaire's unit with Li's channel joining processing unit would have yielded predictable results, for example, for the same reasons I provided in Section IX.A.1. [Weiman/Li].

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.

285. Weyman/Jain/LeMaire/Li teaches this claim limitation as I explained in Section IX.A.2.xiv. [claim limitation [1m]].

286. For example, as I explained in Section IX.A.2.xiv. [claim limitation [1m]] and in Sections IX.A.1. [Weyman/Li] and IX.E.1 [Weyman/Jain/LeMaire/Li], when the distributed router receives a request for a multicast packet stream (e.g., an IGMP membership report), the channel joining processing unit may determine whether the maximum number of multicast streams that are currently being forwarded over a particular port (e.g., a high-capacity port) is below the maximum number of multicast packet streams that are permitted to be forwarded over that particular port, and if not, the request to receive the multicast packet stream is rejected by the channel joining processing unit. EX1007, 31:4-24.

287. However, if number of multicast streams that are currently being forwarded over a particular port (e.g., a high-capacity port) is below the maximum number of multicast packet streams that are permitted to be forwarded over that particular port(i.e., “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”), then the channel joining processing unit may receive the multicast stream (e.g., by processing the multicast membership report / join request). EX1007, 30:21-24. Then the unit of the distributed router would process the multicast join request and modify the unit’s multicast forwarding tables (e.g., the “respective list”) (EX1005, [0137]-[0139], [0162]) to indicate that the requested multicast packet stream is “permitted for forwarding over a given port.”

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

1. Weyman/Jain/LeMaire In View Of Pitcher

288. As I explained in Section IX.C.1. [Weiman/Li/Pitcher], a POSITA would have understood that it is undesirable to have stale entries in multicast forwarding tables because keeping stale entries in the multicast forwarding tables would unnecessarily waste valuable resources (e.g., space in memory). EX1026, 5:64-6:4; EX1027, [0015]; EX1005, [0156]. Thus, a POSITA would have been motivated to efficiently identify and remove stale entries in the multicast forwarding tables of Weyman/Jain/LeMaire’s distributed multicast router and remove them.

289. As described in Section IX.C.1 [Weyman/Li/Pitcher], a POSITA would have understood that Pitcher discloses a timer-based method to remove stale entries from the multicast forwarding tables.

290. 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 multicast forwarding tables.

291. The reasons that I provided in Section IX.C.1 [Weiman/Li/Pitcher] for combining Weyman, Li, and Pitcher, are equally applicable here because Jain and LeMaire's initialization technique is independent of Pitcher's timer-based stale entry removal method and combining Weyman/Jain/LeMaire's unit with Pitcher's timer-based stale entry removal method would have yielded the predictable results I provided in Section IX.C.1. [Weiman/Li/Pitcher].

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.

292. Weyman/Jain/LeMaire/Pitcher teaches this claim limitation as I explained in Section IX.C.2. [Claim 16].

293. As I explained in Section IX.C.2. [Claim 16] and in Sections IX.C.1. [Weyman/Li/Pitcher] and IX.F.1 [Weyman/Jain/LeMaire/Pitcher], each unit within the distributed router is configured to track the entries that are being utilized in the multicast forwarding tables by deploying a timer. EX1009, 7:56-62. Further, the entries in each of the multicast forwarding tables (e.g., the Vbridge multicast forwarding table or FDB) can indicate “a permission to forward a certain multicast packet stream over a certain port.” EX1005, [0112], [0115].

294. Further, Pitcher discloses that the timer is updated, for example, by monitoring IGMP membership reports (e.g., the “requests”) that correspond to the particularl multicast stream that is received by Weyman/Jain/LeMaire/Pitcher’s distributed router. EX1009, 7:62-67. And, according to Pitcher, “[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 multicast forwarding tables].” EX1009, 7:66-67.

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.

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

296. In Weyman/Li, the steps of storing the multicast 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 IX.A.2.vi, ¶¶138, 144 above [claim limitations [1e], [1g], [1i]]. It was well known in the art that such functionalities may be embodied in dedicated hardware logic.

297. For example, Ding discloses a “method for forwarding multicast packets by a module” that may be implemented in hardware logic that includes ***“integrated circuitry, ... a Field Programmable Gate Array (FPGA).”*** EX1012, Claim 1, [0086] (“Alternative embodiments of the invention may be implemented using discrete components, integrated circuitry, programmable logic used in conjunction with a programmable logic device such as a Field Programmable Gate Array (FPGA) or microprocessor, or any other means including any combination thereof.”).

298. It would have been obvious to a POSITA to use dedicated hardware to implement the “storing,” “distributing,” and “selectively forwarding” functionalities of the unit, as described in Weyman/Li because that would have merely amounted to a combination of known prior art elements (Weyman/Li’s unit

and Ding's dedicated hardware logic) that are implemented according to a known method (e.g., integration of the hardware logic into a distributed router) to yield a predictable result of seamless operation by the distributed router and using Ding's dedicated hardware logic. A POSITA would have known that implementing Weyman/Li's unit with Ding's dedicated hardware logic would have merely been a routine design choice (e.g., a choice between using a generic processor and a dedicated hardware logic), which would have been obvious to do.

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.

299. Weyman/Li/Ding teaches the additional claim limitation of claim 12 as I explained above for claim 2. *See* Section IX.G.1 [Claim 2].

X. Conclusion

300. It is my conclusion that the challenged claims 1-2, 4-12 and 14-18 of the '394 patent are unpatentable.

XI. Availability for Cross-Examination

301. In signing this declaration, I recognize that the declaration will be filed as evidence in a case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross examination in the case and that cross examination will take place within the

United States. If cross examination is required of me, I will appear for cross examination within the United States during the time allotted for cross examination.

XII. Right to Supplement

302. I reserve the right to supplement my opinions in the future to respond to any arguments that the Patent Owner may raise and to take into account new information as it becomes available to me.

XIII. Jurat

303. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1101 of Title 18 of the United States Code.

DATE: March 28, 2024



Dr. Henry H. Houh