

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,

Petitioner,

v.

PORTSMOUTH NETWORK CORPORATION,

Patent Owner.

In re *Inter Partes* Review of:
U.S. Patent No. 8,014,394

**DECLARATION OF DR. TAL LAVIAN
IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW**

Juniper Exhibit 1003 Juniper v. Portsmouth

TABLE OF CONTENTS

I.	INTRODUCTION	1
	A. Qualifications	1
	B. Documents and Information Considered	6
	C. Summary of Opinions	6
II.	LEGAL STANDARDS TO BE APPLIED	7
III.	STATE OF THE ART	9
	A. Line Cards and Switching Fabrics	10
	B. Link Aggregation Groups.....	12
	1. IEEE Standards Relating to Link Aggregation.....	13
IV.	SUMMARY OF THE '394 PATENT	18
	A. '394 Patent Overview.....	18
	B. Prosecution History Overview	22
V.	SUMMARY OF THE PRIOR ART.....	25
	A. Blease Overview.....	25
	B. Weyman Overview	31
	C. Hu Overview	35
	D. Deng Overview.....	37
	E. Rao Overview.....	39
VI.	A PERSON OF ORDINARY SKILL IN THE ART	41
VII.	CLAIM CONSTRUCTION	46
VIII.	GROUND 1 – CLAIMS 8 AND 18 ARE UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN.....	48

A.	Motivation to Combine the Teachings of Blease and Weyman	49
B.	Independent Claims 8 and 18	53
1.	Preambles of Claims 8 and 18	56
2.	Claim Elements 8[a.i], 18[a], and 18[b.i]	58
3.	Claim Elements 8[a.ii] and 18[b.ii]-18[b.iv]	61
4.	Claim Elements 8[b] and 18[b.v]	64
5.	Claim Elements 8[c] and 18[b.vi]	66
6.	Claim Elements 8[d] and 18[c.i]	67
7.	Claim Elements 8[e] and 18[c.ii]	71
8.	Claim Elements 8[f.i] and 18[d.i]	72
9.	Claim Elements 8[f.ii] and 18[d.ii]	73
IX.	GROUND 2 – CLAIM 9 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND HU	75
A.	Motivation to Combine the Teachings of Blease, Weyman, and Hu76	
B.	Dependent Claim 9	77
X.	GROUND 3 – CLAIM 10 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND DENG	82
A.	Motivation to Combine the Teachings of Blease, Weyman, and Deng	82
B.	Dependent Claim 10	84
XI.	GROUND 4 – CLAIMS 8 AND 18 ARE UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND RAO	86
A.	Motivation to Combine the Teachings of Blease, Weyman, and Rao	87
B.	Independent Claims 8 and 18	88

1. Claim Elements 8[f.ii] and 18[d.ii] 89

XII. GROUND 5 – CLAIM 9 IS UNPATENTABLE OVER BLEASE IN VIEW
OF WEYMAN, RAO, AND HU.....93

XIII. GROUND 6 – CLAIM 10 IS UNPATENTABLE OVER BLEASE IN
VIEW OF WEYMAN, RAO, AND DENG 94

XIV. CONCLUSION.....95

TABLE OF DOCUMENTS CONSIDERED

Exhibit	Description
1001	U.S. Patent No. 8,014,394 (“394”)
1002	File History of U.S. Patent Application No. 12/369,011
1005	U.S. Patent 7,983,205 (“ Blease ”)
1006	U.S. Patent Application Publication No. 2005/0041665 (“ Weyman ”)
1007	U.S. Patent 8,068,416 (“ Hu ”)
1008	U.S. Patent Application Publication No. 2009/0290585 (“ Deng ”)
1009	U.S. Patent 8,238,344 (“ Chen ”)
1012	First Amended Complaint, <i>Portsmouth Network Corp. v. Juniper Networks, Inc.</i> , 1-23-cv-12201, Dkt. No. 61 (D. Mass. Dec. 22, 2023)
1013	First Amended Complaint – Exhibit 4, <i>Portsmouth Network Corp. v. Juniper Networks, Inc.</i> , 1-23-cv-12201, Dkt. No. 61-4 (D. Mass. Dec. 22, 2023)
1014	Kolias and Tomkos, “Switch Fabrics,” IEEE Circuits & Devices Magazine, pages 12-17 (Sept./Oct. 2005)
1015	Elhanany et al., “Switch Fabric Interfaces,” Computer, Vol. 36, Issue 9, pages 106-108 (Sept. 2003)
1016	Robert Wood, Next-Generation Network Services, Cisco Press, pages 125-131, 275-280 (2006)
1017	SysKonnect GmbH, Link Aggregation according to IEEE Standard 802.3ad, White Paper, v1.10 (Oct. 10, 2002)
1018	IEEE Standard 802.3, Amendment to Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications – Aggregation of Multiple Link Segments (2000 Edition) (“802.3ad-Std”)
1021	Congdon and Subramaniam, “Why we don’t need to specify the load balancing algorithm in the Standard,” Load Balancing Algorithm, Aggregation Study Group (Feb. 4, 1998)
1022	Jeffree, “Link Aggregation Control Protocol,” Presentation to the Link Aggregation Study Group (Apr. 1998)
1023	Jeffree, Contribution to the November ‘98 Link Aggregation Meeting, Strawman/D2½: Changes/Additions to 802.3 required in order to specify Link Aggregation (Nov. 9, 1998)
1024	U.S. Patent No. 8,102,848 (“ Rao ”)

Exhibit	Description
1025	U.S. Patent Application Publication No. 2007/0140155 (“ Yu ”)

I. INTRODUCTION

1. My name is Dr. Tal Lavian. I am submitting a declaration in connection with *inter partes* review (“IPR”) proceedings before the United States Patent and Trademark Office for U.S. Patent No. 8,014,394 (“the ’394 Patent”).

2. I have been retained on behalf of Juniper Networks, Inc. to offer technical opinions with respect to the ’394 Patent and the prior art cited in this IPR.

A. Qualifications

3. Based on my qualifications, education, knowledge, expertise, experience, and work background, I believe I am qualified to offer opinions relating to the technology described in the ’394 Patent. A copy of my *curriculum vitae*, detailing my education and experience, is also being submitted in this IPR proceeding as EX1004. Additionally, the following overview of my background pertains to my qualifications for providing expert testimony in this matter.

4. I have worked as a scientist, a technologist, an educator, an engineer, and an entrepreneur in network communications, telecommunications, and computer science for over 35 years.

5. I am currently an inventor of over 100 issued patents and patent applications, most of which are in the related network communications fields. A list of issued patents and patent applications is included in my curriculum vitae. Several of these patents and patent applications are related to the subject matter of the ’394 Patent.

6. Based on my education, knowledge, expertise, technical experience, and training outlined below, I believe I am qualified to offer technical opinions regarding the '394 Patent.

7. Specifically, I received a Ph.D. in Computer Science, specializing in networking and communications, from the University of California at Berkeley in 2006. I also obtained a Master of Science ("M.Sc.") degree in Electrical Engineering from Tel Aviv University, Israel, in 1996. In 1987, I obtained a Bachelor of Science ("B.Sc.") in Mathematics and Computer Science, also from Tel Aviv University.

8. I was with the University of California at Berkeley for almost 20 years, serving as Berkeley Industry Fellow, Lecturer, Skydeck member, Visiting Scientist, Ph.D. Candidate, and Nortel's Scientist Liaison. I was employed by the University of California at Berkeley and was appointed as a lecturer and Industry Fellow in the Center of Entrepreneurship and Technology ("CET") as part of UC Berkeley College of Engineering. I have taught several classes on wireless devices and smartphones. Some positions and projects were held concurrently, while others were held sequentially. My research has been focused on network communications, network services, communication infrastructure, and telecommunication systems.

- 9.** From 2007 to the present, I served (and continue to serve) as a Principal Scientist at my company, TelecommNet Engineering Inc., where I develop network communication technologies and provide research and consulting in advanced technologies, mainly in computer networking and Internet technologies. Further, I am currently the CEO and CTO of Aybell (previously VisuMenu, Inc.). I founded VisuMenu, Inc. in 2010. At VisuMenu, Inc., I led the software design and development of a visual interactive voice response system for smartphones and mobile devices based on the innovative use of wireless and network communications technologies. In 2016, VisuMenu, Inc. was rebranded as Aybell. At Aybell, I have facilitated the design, architectural development, and implementation of a cloud data center for connecting any smartphone user to any company and service by digitizing interactive voice systems and exposing them through cloud-service application programming interfaces to other applications.
- 10.** In 2008, I served as a Communications Consultant at Ixia. I researched and developed advanced network communications testing technologies and ensured wireless service quality, including developing testing for IP routing and switching devices and broadband access equipment. This includes also providing traffic generation and emulation for the full range of protocols, such as routing, Multi-Protocol Label Switching (“MPLS”), Layer 2 and Layer 3 Virtual Private Networks (“VPN”), carrier Ethernet, broadband access, and data center bridging.

11. For eleven years, from 1996 to 2007, I worked for Bay Networks and Nortel Networks. Bay Networks was in the business of making and selling computer network hardware and software. Nortel Networks acquired Bay Networks in 1998, and I continued to work at Nortel after the acquisition. During my tenure at Bay and Nortel, I held positions as Principal Scientist, Principal Architect, Principal Engineer, and Senior Software Engineer. I led the development and research involving several networking technologies, including active networks, DWDM-RAM, and networking computation projects.

12. Before that, from 1994 to 1995, I worked as a software engineer and team leader for Aptel Communications, designing and developing wireless technologies, mobile wireless devices, and network software products.

13. From 1990 to 1993, I worked as a software engineer and team leader at Scitex Ltd., developing system and network communications tools (mainly in C and C++).

14. From 1987 to 1990, I worked as a software engineer for Shalev Inc., designing and developing real-time software and algorithms.

15. I have extensive experience in communications technologies, including wireless technologies, routing and switching architectures, and protocols, including MPLS Networks, Layer 2 and Layer 3 VPN, and Pseudowire technologies. Much of my work for Nortel Networks (mentioned above) involved researching and

developing these technologies. For example, I wrote software for Bay Networks and Nortel Networks switches and routers, developed network technologies for the Accelar 8600 family of switches and routers, the OPTera 3500 SONET switches, the OPTera 5000 DWDM family, and the Alteon L4-7 switching product family. I wrote software for Java-based device management, including a software interface for device management and network management in the Accelar routing switch family's network management system. I have also worked on enterprise Wi-Fi solutions, wireless mobility management, and wireless infrastructure.

16. I am a member of several professional organizations, including the Association of Computing Machinery ("ACM") and the Institute of Electrical and Electronics Engineers ("IEEE") (senior member). I am also certified under the IEEE WCET (Wireless Communications Engineering Technologies) 2012 Program, specifically designed by the IEEE Communications Society (ComSoc) to address the worldwide wireless industry's growing and ever-evolving need for qualified communications professionals.

17. Based on my education and experience, I believe I am qualified to render the opinions outlined in this declaration.

18. My *curriculum vitae* contains further details on my education, experience, publications, and other qualifications to render an expert opinion. I am being compensated at my standard consulting rate for my work on this declaration. My

compensation is not dependent on the outcome of this case, and I have no financial interest in the outcome.

B. Documents and Information Considered

19. I have reviewed the '394 Patent, including the patent claims in view of the specification. In addition, I have reviewed the '394 Patent's prosecution history, the prior art discussed herein, and the remaining exhibits listed herein.

C. Summary of Opinions

20. In my opinion, claims 8-10 and 18 of the '394 Patent are rendered obvious by the prior art. In the remainder of this declaration, I demonstrate that:

- The combined teachings of Blease¹ and Weyman² renders obvious claims 8 and 18 of the '394 Patent;
- The combined teachings of Blease, Weyman, and Hu³ renders obvious claim 9 of the '394 Patent;
- The combined teachings of Blease, Weyman, and Deng⁴ renders obvious claim 10 of the '394 Patent;
- The combined teachings of Blease, Weyman, and Rao⁵ renders obvious claims 8 and 18 of the '394 Patent;

¹ U.S. Patent 7,983,205 (EX1005, "Blease").

² U.S. Patent Application Publication No. 2005/0041665 (EX1006, "Weyman").

³ U.S. Patent 8,068,416 (EX1007, "Hu").

⁴ U.S. Patent Application Publication No. 2009/0290585 (EX1008, "Deng").

- The combined teachings of Blease, Weyman, Rao, and Hu renders obvious claim 9 of the '394 Patent; and
- The combined teachings of Blease, Weyman, Rao, and Deng, renders obvious claim 10 of the '394 Patent.

'394 claims	Obviousness in light of
8, 18	Blease and Weyman
9	Blease, Weyman, and Hu
10	Blease, Weyman, and Deng
8, 18	Blease, Weyman, and Rao
9	Blease, Weyman, Rao, and Hu
10	Blease, Weyman, Rao, and Deng

II. LEGAL STANDARDS TO BE APPLIED

21. Counsel has advised me of the following legal standards that are relevant to IPR proceedings, which I have applied in forming my opinions.

22. I understand that there are two ways in which prior art may render a patent claim unpatentable. First, the prior art can “anticipate” the claim. Second, the prior art can render the claim “obvious” to a person of ordinary skill in the art. I understand that for an invention claimed in a patent to be patentable, it must not be anticipated and must not be obvious based on what was known before the invention was made.

⁵ U.S. Patent 8,102,848 (EX1025, “Rao”).

23. I understand that a patent claim may be invalid as anticipated under 35 U.S.C. § 102 if a prior art reference discloses each element of the claimed subject matter.

24. I understand that a patent claim may be invalid as obvious under 35 U.S.C. § 103 if the differences between the subject matter claimed and the prior art are such that the claimed subject matter as a whole would have been obvious to a person of ordinary skill in the art at the time the invention was made. I also understand that several factual inquiries underlie a determination of obviousness. These inquiries include (1) the scope and content of the prior art, (2) the level of ordinary skill in the field of the invention, (3) the differences between the claimed invention and the prior art, and (4) any objective evidence of non-obviousness (i.e., “secondary considerations”).

25. I also understand that when a party alleges obviousness based on a combination of references, that party should identify a reason why a person of ordinary skill in the art would have been motivated to combine the asserted references in the manner recited in the claims and to explain why a person of ordinary skill in the art would have had a reasonable expectation of success in making such combinations.

26. I understand that an obviousness analysis permits the application of “common sense” in examining whether a claimed invention would have been

obvious to a person of ordinary skill in the art. For example, I understand that combining familiar elements according to known methods and in a predictable way may suggest obviousness when such a combination would yield nothing more than predictable results. And I understand a party asserting obviousness should still provide a specific motivation to combine the references as recited in the claims and explain why one would have reasonably expected to succeed in doing so.

27. I am not aware of any evidence of secondary considerations that would support a determination of non-obviousness of the claimed subject matter in the '394 Patent.

28. I have been informed that, in IPR proceedings, the party challenging the patent bears the burden of proving unpatentability by a preponderance of the evidence. I understand that a preponderance of the evidence means "more likely than not."

29. I understand that, in IPR proceedings, patent claim terms are given their ordinary and customary meaning in view of the specification and prosecution history, unless those sources show an intent to depart from such meaning.

III. STATE OF THE ART

30. Prior to the filing date of the '394 Patent, network communication was a well-developed field of technology. It was well known that networks include a set of network nodes interconnected by communication link segments. Devices, such


as servers and user computers, are the sources and destinations of data traffic flowing through the network. Typically, data traffic flows from sources to destinations through intermediate network nodes. These network nodes use addresses contained in headers of data packets to forward packets to other nodes closer to the packets' intended destination.

A. Line Cards and Switching Fabrics

31. By the '394 Patent's filing date, most, if not all, network nodes included line cards and switching fabrics for transmitting data packets. A line card is a circuitry that provides an interface with a network for data transmission. Line cards may be fixed or replaceable within network nodes. As shown in the figure below, a line card often includes a network processor, a traffic manager, and a queue manager to examine and transmit data packets.⁶ Typically, a network node includes one or more line cards for receiving data packets via one or more ingress ports. And a network node includes one or more line cards for sending data packets via one or more egress ports.⁷ A line card may function bidirectionally as both an ingress line card and an egress line card.

⁶ EX1014 (Kolias), 2.

⁷ EX1014 (Kolias), 2; EX1015 (Elhanany), 2-3.



32. Switching fabric is a networking topology that allows for transferring data between one or more ingress ports and one or more egress ports. For example, the article “Switching Fabrics” by Christos Kolias explains that a switching fabric functions “to provide connectivity for transferring data between ingress (incoming) and egress (outgoing) ports and to satisfy and support (and enforce) any quality-of-service (QoS) demands.”⁸ As another, the article “Switch Fabric Interfaces” by Itamar Elhanany explains that a “switch fabric moves incoming data from a set of ingress ports to a single egress port, in the case of unicast devices, or multiple egress ports, in the case of multicast devices.”⁹ A switching fabric also may be called a switched fabric, a fabric card, a backplane, or a switch backplane.¹⁰

⁸ EX1014 (Kolias), 1.

⁹ EX1015 (Elhanany), 1.

¹⁰ EX1014 (Kolias), 2; EX1016 (Next Generations Network Services), 7.

33. Switching fabrics may take various forms. Two popular switching fabric architectures are single-stage and multi-stage.¹¹ For example, a crossbar fabric is a single-stage architecture having a series of crosspoints forming an interconnected grid of pathways between ingress and egress ports. As another example, a Clos switch is a multi-stage architecture having three stages of disjointed paths between ingress and egress ports.¹²

B. Link Aggregation Groups

34. For packet transmission, networks typically connect devices to one another via “links.” Link aggregation is a method in which a network uses multiple links in parallel to increase throughput beyond what a single link could provide. These parallel links also provide redundancies in case one or more links fail. A link aggregation group, or LAG, is a collection of physical links bundled together to form a single logical link. Link aggregation may be referred to as Ethernet port trunking, network interface card (“NIC”) teaming or bonding, or port teaming.¹³

35. Before the filing date of the '394 Patent, LAGs were well known and widely used in the field of network communications. By March 2000, the IEEE had

¹¹ EX1014 (Kolias), 2-4; EX1016 (Next Generations Network Services), 9.

¹² EX1014 (Kolias), 2-4.

¹³ EX1016 (Next-Generation Network Services), 14; Ex. 1017 (SysKonnnect White Paper), 3, 5, 9.

standardized the protocol for link aggregation as 802.3ad.¹⁴ The '394 Patent incorporates-by-reference IEEE 802.1AX, which followed, and is substantially similar to, 802.3ad.¹⁵

36. IEEE 802.3ad provides the protocol for transmitting data through link aggregate groups, the Link Aggregation Control Protocol (LACP).¹⁶ Typically, network systems employing LAGs use a distribution and load balancing algorithm to distribute traffic across the multiple links of a LAG. For example, a network could distribute traffic via a round-robin scheduling method in which packets are sent through links in equal portions, in a circular order among the links, and without priority considerations. As another example, a network could distribute traffic through a load balancing process in which links are chosen to transmit packets with the goal of ensuring no single link or queue is overloaded.¹⁷

1. IEEE Standards Relating to Link Aggregation

37. The IEEE published the 802.3ad Standard in March 2000 and focused on “aggregating multiple physical links ... into one logical link” as an amendment to the 802.3 Standard, “CSMA/CD Access Method and Physical Layer

¹⁴ EX1018 (802.3ad Standard), 6.

¹⁵ '394 Patent, 6:64-7 :4.

¹⁶ EX1018 (802.3ad Standard), 126-152.

¹⁷ EX1018 (802.3ad Standard), 109-111, 170.

Specifications.”¹⁸ During development of the IEEE 802.3ad Standard, many vendors presented their concepts, ideas, proposals, and drafts. And multiple vendors released to the market products supporting such standards published the products’ documentations. Further, multiple scientific publications published on the proposed standard in the development process, as illustrated in the following examples.

a. Hewlett-Packard’s 1990s Work on Load Balancing Algorithms

38. As part of the IEEE 802.3ad “Aggregation Study Group,” representatives from HP presented their proposal for “Load Balancing Algorithms” (February 4th, 1998).¹⁹ The title of their presentation was “Why we don’t need to specify the load-balancing algorithm in the Standard.”²⁰


39. The slide below is one of the slides presented by HP representatives (a similar diagram ended up as part of the final published standard). Those slides illustrated multiple concepts for link aggregation and specified the reasons why the

¹⁸ EX1018 (802.3ad standard), 1, 11.

¹⁹ EX1021 (HP Aggregation Proposal).

²⁰ EX1021 (HP Aggregation Proposal), 1, 8. This presentation provides examples for load balancing algorithms, some considerations for specific mechanisms, and analysis of different mechanisms. The presenter argues: “Don’t Specify the Algorithm, Only i[t]s basic requirements,” and provides examples on “Why we shouldn’t standardize the algorithm.” He concludes with: “We don’t need to standardize the algorithm, only its requirements.” The HP presenter discussed different options for specific algorithms, and concluded that the 802.3 Standard shouldn’t include specific algorithms and only requirements.

standard should not “specify the load balancing algorithm in the Standard” and let each vendor implement their internal mechanisms for load balancing.²¹



b. Jeffree's 1998 IEEE Study Group Presentation for LACP

40. In an April 1998 Study Group presentation, Mr. Tony Jeffree presented some concepts, ideas, and challenges to address the Link Aggregation Control Protocol (“LACP”). In his proposal, Mr. Jeffree addressed some solutions for

²¹ EX1021 (HP Aggregation Proposal), 1, 7.

identifying LAGs, detection aggregation possibility, and preventing duplications.

This included proposed protocol solutions and operations.²²

c. Juniper Networks' 1998 Link Aggregation Draft Standard

41. While working for Juniper Networks, Mr. Jeffree also presented a “Strawman” design draft standard for Link Aggregation on November 9, 1998. As shown below, this draft discloses and depicts the physical port.²³


92.4.8.1 Selection Logic

Each physical MAC has both a physical Port and an Aggregator associated with it (i.e., there are the same number of Aggregators as physical Ports); both the physical Port and the Aggregator are allocated the same Key and Port number.

Aggregation is represented by a physical Port *selecting* an appropriate Aggregator, and then *attaching* to that Aggregator. When there are multiple physical Ports in an aggregation, the Aggregator that the set of Ports selects is the Aggregator with the same Port number as the lowest numbered physical Port. The lowest numbered physical Port may not be in a state that allows data to be transferred on its physical link; however, it has selected that Aggregator. This is illustrated in Figure 92-5.

²² EX1022 (Jeffree 1998 Proposal), 1, 4-15.

²³ EX1023 (Jeffree Strawman Proposal) 34-35.



42. Mr. Jeffree’s proposed draft discusses the changes of identifiers (IDs) and keys for the physical ports and the logical aggregators.

1.2.7 Key

A parameter of each physical Port of a system identifying those Ports that can be aggregated together.

NOTE—Ports in a system that share the same value of this parameter are potentially able to aggregate together.

1.2.8 Link Aggregation Group

A grouping of Link Segments, of the same medium type and speed, that are treated as if they are all part of a single Link Segment. The MDIs associated with each Link Segment in a Link Aggregation Group are associated with the same pair of devices. For the purposes

of this definition, a device is a MAC Bridge, an end station or a repeater.²⁴

IV. SUMMARY OF THE '394 PATENT

A. '394 Patent Overview

43. The '394 Patent is titled “High-Speed Processing of Multicast Content Requests.” As the cover shows, the application for the '394 Patent was filed on February 11, 2009.


44. The '394 Patent discloses a method for communication using a network element with multiple ports and multiple interconnected processing units, each processing unit including a list of permitted multicast packet streams and the ports over which the permitted streams should be forwarded.²⁵ The '394 Patent alleges that its method is “particularly effective in adapting rapidly to requests to change [a] multicast forwarding configuration” because it utilizes a distributed network element including a plurality of processing units, each processing unit having assigned ports and processing packets according to a locally stored filtering database.²⁶

45. An example of a **network element 32** according to the '394 Patent is illustrated in Figure 1 below.

²⁴ Ex. 1032 (Jeffree Strawman Proposal), 8.

²⁵ '394 Patent, Abstract, 1:51-2:3.

²⁶ '394 Patent, 4:60-5:34.



46. **Network element 32** includes, for example, a plurality of **line cards 36** connected by **interconnection fabric 40** and including a plurality of **ports 44**. The **line cards 36** further include **Network Processors (NPs) 48**, which process packets communicated over **ports 44** based in part on **filtering databases 52**.²⁷

47. In one example of operating **network element 32**, the '394 Patent describes:²⁸

- Receiving, at a **port** of a **line card**, a multicast content request to join a certain multicast service;

²⁷ '394 Patent, 6:23-46.

²⁸ '394 Patent, 7:27-8:35.

- Updating **databases 52**, according to the request, to permit forwarding of the requested multicast stream on a specific port; and
- Forwarding incoming multicast streams in accordance with updated **databases 52**.

48. The '394 Patent also describes some features that are included in some embodiments. For example, the '394 describes a process of inserting a new line card into a network element and disabling the local ports for a time while the local filtering database for the line card is updated via a multicast status update request or IGMP general query to all multicast content clients.²⁹ In other embodiments, **ports 44** are aggregated to form LAGs, to which the '394 Patent references IEEE 802.1AX.³⁰ In other embodiments, the number of streams permitted for forwarding over a port or a LAG is limited, so as to not exceed their maximum bandwidths.³¹ And, in yet other embodiments, **network element 32** runs an aging process that removes inactive multicast stream group members.³²

49. Representative independent claim 8 recites:

A method for communication, comprising:

²⁹ '394 Patent, 9:41-53.

³⁰ '394 Patent, 6:61-7:6.

³¹ '394 Patent, 8:36-45, 9:54-67.

³² '394 Patent, 9:24-27.

operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports, wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and 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;

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;

receiving from a client a request relating to reception of a given multicast packet stream;

distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request; and

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

wherein operating the network element comprises initializing a new processing unit added to the network element by:

deactivating the ports to which the new processing unit is assigned;

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

50. Claims 8 and 18 are independent claims. Claims 9 and 10 depend from claim 8.

B. Prosecution History Overview

51. The application issued as the '394 Patent was filed on February 11, 2009.³³ During prosecution, the Examiner rejected the claims based on prior art, such as Naik (U.S. 2006/0146857) and Liu (U.S. 2002/0191628), that shows it was known in the art to:

- Operate a network element comprising multiple interconnected processing units and multiple ports, wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and 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;
- Storing in each processing unit a respective list, each list indicating all of one or more multicast packet streams that are permitted for forwarding

³³ EX1002 (“’394 History”), 1.

by the network element, and further indicating which of the one or more ports each multicast packet stream is permitted to be forwarded to;

- Receiving from a client a request relating to reception of a given multicast packet stream;
- Distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request; and
- Selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists.³⁴

52. During prosecution, the Examiner indicated originally-filed dependent claims 7, 8, 25, 26, and 28 included allowable subject matter.³⁵ Originally-filed claim 10, which depends from independent claim 1, recites, in pertinent part, “wherein operating the network element comprises initializing a new processing unit added to the network element by: deactivating the ports to which the new processing unit is assigned; while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request, receiving multicast

³⁴ ’394 History, 186-215 (04/14/2011 Office Action providing prior art anticipation and obviousness arguments).

³⁵ ’394 History, 106-131 (10/29/2010 Office Action).

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 after constructing the list, activating the ports to which the new processing unit is assigned.”³⁶

53. In response, the applicant cancelled independent claim 1 and rewrote dependent claim 10 to be in independent form (eventually issuing as claim 8), including the limitations from originally-filed claim 1.³⁷ The Examiner subsequently allowed the claims, identifying the limitations of originally-filed dependent claims 10 and 28 as reasons for allowance for claims 8 and 18.³⁸ That is, the Examiner allowed the claims because the dependent claims introduced limitations related to initializing a processing unit.

54. It is my understanding that the USPTO did not consider the Blease, Weyman, Rao, and Hu references during the original prosecution of the '394 Patent. Further, it is my understanding that the USPTO did not consider the obviousness combinations of Blease in view of Weyman, or further in view of Rao, Hu, and/or Deng, during the original prosecution of the '394 Patent. As detailed

³⁶ '394 History, 32 (02/11/2009 originally-filed claims).

³⁷ '394 History, 218-228 (04/22/2011 Applicant Amendment/Remarks making similar amendments in originally-filed dependent claim 28, issued as claim 18).

³⁸ '394 History, 236-240 (05/12/2011 Notice of Allowance).

below, these references show it was known or obvious to initialize processing units by sending multicast status requests.

V. SUMMARY OF THE PRIOR ART

55. In Section III (above), I described the state of the pertinent art prior to 2009. The following prior art references show that every element of the Challenged Claims was known in the art by February 11, 2009.

A. Blease Overview

56. U.S. Patent No. 7,983,205 (EX1005; “Blease”) is titled “Outgoing Interface Mapping For Multicast Traffic.” As reflected on its cover, Blease was filed on February 28, 2005, and issued on July 19, 2011. I understand that Blease qualifies as prior art under at least 35 U.S.C. § 102(e) based on its filing date.

57. Blease is directed to “techniques for transmitting multicast traffic from a network device to subscriber devices within a computer network.”³⁹ As shown below in Figure 1, Blease describes that the system provides subscriber devices 8 with connectivity to network 10 through network switches 6 and **network device 4**.⁴⁰ As Blease explains, “Network device 4 is capable of providing subscriber devices 8 with multicast traffic, e.g., Internet Protocol television (IPTV), from network 10. Subscriber devices 8 request multicast streams, e.g.,

³⁹ Blease, 2:10-12.

⁴⁰ Blease, 3:34-50.

IPTV channels, by sending a membership request for a specific multicast group to network device 4.”⁴¹ When network device 4 receives a membership request for a multicast group, it maps the request to an outgoing interface through which the requested multicast stream will be forwarded.⁴²




FIG. 1

58. Blease describes an exemplary **network device 4** as **router 42**, shown below in Figure 4. Blease describes that **router 42** includes **interface cards**

⁴¹ Blease, 3:53-58.

⁴² Blease, 3:65-4:17.

(“IFCs”) 54A-54N, inbound links 55A-55N, outbound links 56A-56N, and a control unit 44.⁴³




FIG. 4

59. As shown, control unit 44 includes routing information 46, multicast forwarding information 48, OIF (outgoing interface) mapper 50, OIF maps 52, and OIF reference counts 53.⁴⁴ Multicast forwarding information 48 “controls replication and forwarding of multicast packets to subscriber devices.”⁴⁵

⁴³ Blease, 8:39-67.

⁴⁴ Blease, 9:43-45.

⁴⁵ Blease, 10:49-61.

60. Blease offers an example of processing multicast traffic in Figures 6A-6B, reproduced below. The process includes receiving a membership request for a multicast group (step 60), performing various mapping-related functions to update multicast forwarding information 48 (steps 64-80), and finally sending a multicast stream on a mapped OIF (step 82).⁴⁶

⁴⁶ Blease, 13:64-14:57.

[receiving a membership request for a multicast group]

[various mapping-related functions to update multicast forwarding information 48]




FIG. 6A




FIG. 6B

61. Please states that in some instances, a membership request may be received on an interface that is not configured to use OIF maps 52, in which case OIF mapper 50 “adds an unmapped interface to multicast forwarding information 48”

(NO at step 64 → step 65).⁴⁷ OIF mapper 50 may choose to map the request to the interface that received the membership request.⁴⁸ That mapped interface in multicast forwarding information 48 will be used to send the multicast stream, as discussed above.

62. Furthermore, while Blease appears to show control unit 44 as a centralized control unit in at least Figure 4 shown above, Blease also explains that “The architecture of **router 42** illustrated in FIG. 4 is shown for exemplary purposes only. The invention is not limited to this architecture. In other embodiments, **router 42** may be configured in a variety of ways. In one embodiment, for example, **control unit 44** and its corresponding functionality may be distributed within **IFCs 54**.”⁴⁹ As I have explained above, network communication was a well-developed field of technology prior to the '394 Patent, and so even though Blease does not provide extensive details of the “variety of ways,” a person of ordinary skill in the art would have recognized a number of ways to configure multicast routers with distributed functionality.

B. Weyman Overview

63. U.S. Patent Application Publication No. 2005/0041665 (EX1006; “Weyman”) is titled “System And Method For Distributed Multicast Routing.” As

⁴⁷ Blease, 14:10-17.

⁴⁸ Blease, 6:62-7:3, 8:29-38, 14:17-20.

⁴⁹ Blease, 11:52-57.

reflected on its cover, Weyman was filed on August 20, 2003, and was published on February 24, 2005. I understand that Weyman qualifies as prior art under at least 35 U.S.C. §§ 102(a), (b), and (e) based on its filing date and publication date. I understand that the USPTO did not consider Weyman during the prosecution of the '394 Patent.

64. Weyman is directed to a “fabric of units organized as a single virtual multicast router and units intended for or capable of use in such a fabric” that “allows a symmetrical operation with the units of [the] fabric participating similarly in a multicast routing functionality.”⁵⁰ An example of a **multicast router** according to Weyman is shown below in annotated Figure 10.

⁵⁰ Weyman, [0001], [0006].




FIG.10

65. As shown, the **multicast router** includes two **units 10 and 20** and a **fabric link 30** and is configured to forward a multicast stream from a source S to multicast clients C1 through C5 over **ports 15**.⁵¹ **Units 10 and 20** include Vbridges

⁵¹ Weyman, [0122].

with tables that identify, for example, multicast groups and interfaces on which to accept and forward streams.⁵²

66. Further, Weyman teaches that Internet Group Management Protocol (IGMP) reports, used by clients to register for reception of multicast streams, for example, may be processed in at least two different ways.⁵³ In one way, an IGMP report received at one unit will “process the IGMP report and ... send updates to the remainder of the units in the fabric,” and, alternatively, “the IGMP message could be delivered to every unit in the fabric for processing.”⁵⁴ client units generate IGMP reports to register or change their membership in multicast groups. An IGMP report received at one unit is either “delivered to every unit in the fabric for processing” or processed at one unit that “send[s] updates to the remainder of the units in the fabric.” Weyman thus teaches updating distribution and routing information across all units equally because changes “must be kept synchronized.”⁵⁵ In furtherance of the teaching to keep units synchronized, Weyman teaches that when a unit joins a fabric with a default configuration, that

⁵² Weyman, [0106]-[0116]. *See also* [0028] describing a Vbridge as a bridge that will forward data within one VLAN. In networks without VLANs, “Vbridges” are the same as “bridges” well-known in the art.

⁵³ Weyman, [0160]-[0162].

⁵⁴ Weyman, [0162].

⁵⁵ Weyman, [0162].

unit's configuration must be updated to that of the router through an update process.⁵⁶

C. Hu Overview

67. U.S. Patent No. 8,238,344 (EX1007; "Hu") is titled "System And Method Of Communicating A Media Stream." As reflected on its cover, Hu was filed on September 20, 2007 and issued on November 29, 2011. I understand that Hu qualifies as prior art under at least 35 U.S.C. § 102(e) based on its filing date. It is my understanding that the USPTO did not consider Hu during the prosecution of the '394 Patent.

68. Hu describes a system and method for "allocating media streams include[ing] measuring real-time port usage data related to each physical port associated with a link aggregation group of a network device."⁵⁷ An example of a **content switch 102** according to Hu is shown below in Figure 1. **Content switch 102** is, for example, "a network content switch/router system or other edge system of an Internet Protocol Television (IPTV) network or other packet-based video distribution network" that communicates with a plurality of content sources 132 to send multicast streams through a plurality of **ports 122 and 126**.⁵⁸, such as a network content switch/router system or other edge system of an Internet Protocol

⁵⁶ Weyman, [0069], [0156].

⁵⁷ Hu, 2:13-16.

⁵⁸ Hu, 2:57-67, 4:3-8.

Television (IPTV) network or other packet-based video distribution network. Hu,
Fig. 1, Abstract, 2:56-61.




FIG. 1

69. Also shown in Figure 1, Hu describes that multiple **physical ports 122 and 126** are logically grouped into **link aggregation groups (LAGs) 124 and 128**, respectively. Hu states: “link aggregation refers to a grouping of multiple network links (i.e., communication channels or circuits) into one logical high bandwidth link, which may be used to route audio data, text data, video data, or any combination thereof. A LAG link aggregation group refers to a set or group of physical ports, such as the first plurality of physical ports 122, which are logically grouped to form a single logical port to provide the high-bandwidth link.”⁵⁹

70. Further with respect to the LAGs, Hu teaches managing requests based on pre-defined traffic thresholds. For example, when a request is received through a LAG, the router system measures the traffic load through the ports associated with the LAG, both individually and in combination, and determines whether the request should be granted or rejected in response.⁶⁰

D. Deng Overview

71. U.S. Patent Application Publication No. 2009/0290585 (EX1008; “Deng”) is titled “Data Link Layer Switch [With] Multicast Capability.” As reflected on its cover, Deng was filed on August 4, 2009, claims priority to several applications including U.S. Patent Application No. 09/291,705, filed on April 14, 1999, and

⁵⁹ Hu, 3:7-32.

⁶⁰ Hu, 6:39-7:8.

was published on November 26, 2009. I understand that Deng qualifies as prior art under at least 35 U.S.C. § 102(e) based on its filing date. It is my understanding that the USPTO considered Deng during the prosecution of the '394 Patent but did not consider Deng in view of the teachings of Blease, Weyman, or Hu.

72. Deng describes a system and method for providing multicasting via a switching mechanism coupled to a plurality of port interface controllers, multicast address tables, multicast matchers, and timers.⁶¹ Deng seeks to “reduce[] network traffic load and allow[] the Ethernet switch to be used for a large number of multicast streams.”⁶²

73. In one aspect of its system, Deng discloses utilizing a timer to age out addresses in a multicast address table. For example, when receiving a request to join a multicast group, a multicast address—identifying attached hosts—is stored in multicast address table 320.⁶³ Deng continues: “In addition to storing the multicast address, the matcher 310 starts or restarts the timer 330. The timer 310 is used to determine the length of time that an address remains in the multicast

⁶¹ Deng, [0017]-[0018].

⁶² Deng, [0026].

⁶³ Deng, [0018], [0036].

address table 320. An address is removed from the table upon the expiration of the timer.”⁶⁴


E. Rao Overview

74. U.S. Patent No. 8,102,848 (EX1025; “Rao”) is titled “Multicast High Availability Enhancements For Faster Convergence.” As reflected on its cover, Rao was filed on November 19, 2008 and issued on January 24, 2012. I understand that Rao qualifies as prior art under at least 35 U.S.C. § 102(e) based on its filing date.

75. Rao describes multicast data packet forwarding and distributed routers and switches configured for multicast data packet forwarding, including their management of failover behavior—for example “when a route processor module ... fails and route processing duties failover to a standby route processor module.”⁶⁵ Rao shows one example of router behavior upon failover to a standby route processor in Figure 4 shown below.

⁶⁴ Deng, [0036].

⁶⁵ Rao, abstract, 3:24-67.



76. For example, when failover causes a standby processor to become a master processor, an announcement is made of the new master processor, and the new master processor generates Host Membership Queries “to solicit membership information from attached IGMP hosts.” At that time, a delay timer is also started.⁶⁶ Replies from neighbors in the multicast network and from connected IGMP hosts are analyzed to add, remove, or confirm routing table entries as

⁶⁶ Rao, 9:3-17.

needed.⁶⁷ Finally, when end markers are received (for example an “End of Tree Information Base” (EOTIB) marker as shown in Figure 4) or the delay timer expires, the current multicast state is then established and normal operation begins.⁶⁸

VI. A PERSON OF ORDINARY SKILL IN THE ART

77. It is my understanding that I must analyze and apply the prior art cited above from the perspective of a person of ordinary skill in the art as of the time of the invention, which I understand is February 11, 2009, for purposes of IPR proceedings. I have also been advised that to determine the appropriate level of a person having ordinary skill in the art. The following factors may be considered: (1) the types of problems encountered by those working in the field and prior art solutions thereto; (2) the sophistication of the technology in question, and the rapidity with which innovations occur in the field; (3) the educational level of active workers in the field; and (4) the educational level of the inventor. When forming my opinions herein, I analyzed and applied the prior art and claim constructions from the perspective of a person of ordinary skill in the art as of the time of the invention.

⁶⁷ Rao, 9:18-46.

⁶⁸ Rao, 9:47-65.

78. The '394 Patent relates to a communication network and method for transmitting multicast data packets, wherein a network element includes multiple ports and processing units:⁶⁹

A method for communication includes operating a network element, which includes multiple ports and multiple interconnected processing units, such that each processing unit is assigned to process packets that are communicated over a respective set of one or more of the ports. Each processing unit stores a respective list, which indicates one or more multicast packet streams that are permitted for forwarding by the network element, and further indicates a respective subset of the ports over which each of the multicast packet streams is permitted for forwarding. A request relating is disturbed [*sic*] among the multiple processing units, and the respective list in each processing unit is updated responsively to the request. Multicast packets associated with the given multicast packet stream are forwarded selectively by the processing units in accordance with the respective updated lists.

79. Accordingly, an embodiment of the '394 "invention provides a method for communication" including operating a network element, storing lists that indicate multicast packet streams permitted for forwarding, receiving a request from a client, distributing the request among multiple processing units, updating lists responsively to the request, and selectively forwarding multicast packets in accordance with updated lists.⁷⁰ Furthermore, "[in] some embodiments, operating

⁶⁹ '394 Patent, abstract.

⁷⁰ '394 Patent, 1:51-2:3.

the network element includes aggregating two or more of the ports to form a high-capacity port.”⁷¹

80. Additionally, the '394 Patent references and utilizes various known standards, industry knowledge, and content delivery systems existing at the time, including Internet Protocol Television (IPTV), the Internet Group Management Protocol (IGMP), the multicast Listener Discovery (MDL) protocol, and IEEE 802.1AX.⁷²

BACKGROUND OF THE INVENTION

Various communication applications deliver content to multiple recipients over packet networks by transmitting multicast traffic. Applications of this sort include, for example, Internet Protocol Television (IPTV), as well as different applications that provide video and audio content. In order to avoid unnecessary congestion and processing load, it is often desirable not to flood content indiscriminately over the network, but to selectively forward multicast content to the appropriate recipients. A number of communication protocols can be used to introduce such selectivity. Two example protocols that can be used for this purpose are the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocols.

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,”

⁷¹ '394 Patent, 2:24-26.

⁷² '394 Patent, 1:10-47, 6:64-7:1.

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.

Additional aspects of IGMP and MLD are described by Christensen in “Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches,” RFC 4541, May, 2006, by Holbrook et al. in “Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast,” RFC 4604, August, 2006, and by Fenner et al. in “Internet Group Management Protocol (IGMP)/Multicast Listener Discovery (MLD)-Based Multicast Forwarding (‘IGMP/MLD Proxying’),” RFC 4605, August, 2006, which are incorporated herein by reference.

[¶¶]

[P]orts 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.

81. It is my opinion that a person of ordinary skill in the art would have possessed at least a bachelor’s degree in electrical engineering, computer engineering, computer science, or a closely related field, and approximately two or more years of experience in computer and networking technology. Alternatively, a person of ordinary skill in the art would have had additional education (for

example a Master's degree or similar post-graduate work) and fewer years of experience, or vice versa.

82. The person of ordinary skill in the art also would have possessed the following areas of knowledge, which were part of the general knowledge in the art of the '394 patent as of February 2009. The '394 patent deals with the core concepts and architecture of network communication services. The ideas contained therein were mainstream at the time of the invention, including products, markets, business, networking vendors, networking service providers, publications, books, academic educations, and industry standards. Indeed, the concepts of switch architecture, line cards, and switching fabric were well-known and well-published, as described in the State of the Art Section above.⁷³ These concepts were essential components of packet-switching systems, switching architecture, and protocol design.

83. My opinions regarding the level of ordinary skill in the art are based on, among other things, my over 30 years of experience in computer science and network communications, my understanding of the essential qualifications that would be relevant to an engineer or scientist tasked with investigating methods and systems in the relevant area, and my familiarity with the backgrounds of colleagues, co-workers, and employees, both past and present.

⁷³ Section III above.

84. I am well aware of the qualifications of such a skilled artisan because I have worked with, supervised, and hired engineers with similar capabilities. By 2009, I had over 22 years of practical experience in the industry. And by 2009, I had received a Ph.D. in computer science specializing in network communications from the University of California at Berkeley. Further, by 2009, I was working with individuals who met the above criteria of persons of ordinary skill in the art. Many of my colleagues possessed at least a B.S. in the relevant fields and had several years of experience designing and developing network communication systems.

85. Thus, I am familiar with the understanding and knowledge of persons of ordinary skill in the art as of 2009. I was at least as qualified as the person of ordinary skill in the art that I have identified above. Thus, I understand the perspective of a person of ordinary skill in the art, which I have applied in my analysis.

VII. CLAIM CONSTRUCTION

86. For IPR petitions filed after November 13, 2018, I understand the Patent Trial and Appeal Board (“Board”) construes claim terms as having their ordinary and customary meaning to a person of ordinary skill in the art at the time of invention. I have applied this standard in rendering my opinions on claim

construction. In addition, I have been asked to assume that the claims are not indefinite.

87. At this time, I do not propose any terms for construction because the ordinary and customary meaning of the terms is appropriate. However, I do not waive, and expressly reserve, the right to respond to any constructions proposed by the Patent Owner or the Board.

VIII. GROUND 1 – CLAIMS 8 AND 18 ARE UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN

88. The Challenged Claims include independent claims 8 and 18. As I detail below, a person of ordinary skill in the art would have recognized that each and every limitation of claims 8 and 18 is disclosed or rendered obvious over Blease in view of Weyman.

89. In summary, Blease discloses a system and method including the same distributed network element architecture as the '394 Patent, including pluralities of processing units and associated ports, and where each processing unit includes and updates a list of forwarding information for use in forwarding multicast packet streams. And, to the extent Blease does not provide explicit details about a particular configuration, component, or operating step, for example in claims 8 and 18—which a person of ordinary skill in the art would not have found necessary to disclose because distributed architectures for network components were well-known at the time of the '394 Patent. In fact, such distributed architecture was the mainstream fundamental core architecture used by most of the industry and publications. For example, Weyman teaches such examples for distributed multicast routing, including a plurality of processing units that are interconnected by a fabric link, that process multicast membership requests in a synchronized fashion, and that further accommodates new processing units being added to the fabric. The teachings of Blease and Weyman, and a person of ordinary skill in the

art's motivation for combining their teachings, summarized above, are described in detail below.

A. Motivation to Combine the Teachings of Blease and Weyman

90. Like the '394 Patent, Blease and Weyman describe systems and methods relating to communication networks, multicast traffic and are in the same field of endeavor and contain overlapping disclosures with similar purposes.⁷⁴

91. As discussed above, Blease discloses a router including a plurality of interface cards (IFCs) 54 and a control unit 44 including, for example, multicast forwarding information 48 and an outgoing interface (OIF) mapper 50. Also discussed above, Blease describes that its exemplary configuration is not absolute, and "control unit 44 and its corresponding functionality may be distributed within IFCs 54."⁷⁵ Blease does not dictate the exact manner in which control unit 44 and its functionality should be distributed. Blease leaves that implementation decision to the persons of ordinary skill in the art implementing its disclosed system and method, and such persons of ordinary skill in the art would have been motivated to look to other prior art systems in implementing distributed systems for processing multicast traffic.

⁷⁴ '394 Patent, 1:6-8; Blease, 1:10-12; Weyman, [0001].

⁷⁵ Blease, 11:54-57.

92. At the time of the '394 Patent's filing date, various router configurations were known to persons of ordinary skill in the art, including those with control functionality distributed amongst line cards (or interface cards)—such as Weyman's. As discussed above in Section V.B, Weyman teaches a “System And Method For Distributed Multicast Routing” that “allows a symmetrical operation with the units of a fabric participating similarly in a multicast routing functionality,” including a plurality of units connected by a fabric, each unit having essentially the same components but servicing different ports of a router.⁷⁶ And, Weyman further teaches that functions such as processing updates in response to IGMP reports are performed symmetrically across units, such that the system's multicast distribution tree is kept synchronized.⁷⁷

93. Similarly, Weyman explains that new processing units added to a router should be synchronized with the rest of the router by updating its configuration from that of a default configuration to that of the distributed router so it can “operate as a single distributed router.”⁷⁸ While Weyman does not describe an exact method that must be used to update a new unit, Weyman explains that IGMP queries are transmitted to “determine the group membership of clients on an

⁷⁶ See my discussion in Section V.B above; *see also* Weyman, Title, Figure 10, [0001]-[0006], [0024], [0122].

⁷⁷ Weyman, [0159]-[0162].

⁷⁸ Weyman, [0069], [0156].

attached sub-network” and update a multicast distribution tree based on responses received.⁷⁹ Weyman therefore teaches a process of adding a processing unit to a router by first updating the processing unit for use, wherein one taught method of updating includes sending IGMP membership queries. A person of ordinary skill in the art would have recognized this as a known and practiced procedure to obtain the current status of a multicast network and its clients.⁸⁰

94. In other words, a person of ordinary skill in the art would not expect a processing unit being put into service to be immediately and accurately configured to operate in the current state of any particular multicast network. The processing unit may have a configuration from being installed in a different or prior-in-time network, or perhaps may never have been put into service and instead have a “default” configuration, as described by Weyman above. A processing unit having an incorrect or unknown configuration would thus unpredictably process packets or drop packets completely. Therefore, a person of ordinary skill would ensure that the processing unit is deactivated prior to being updated for use in the current multicast network—such that it is neither sending nor receiving packets on any of its ports—and then activate the processing unit after updating—so that it may

⁷⁹ Weyman, [0161]-[0162].

⁸⁰ See, for example, Yu (EX2025), [0053], [0059]. Up-to-date multicast membership information is obtained from the multicast clients because membership changes may have been made prior to normal processing.

begin the normal operation of forwarding packets to the correct ports according to the current state of the multicast network.

95. As a person of ordinary skill in the art implementing Blease's system would have been motivated to look for implementations like Weyman's, it would have been obvious to implement Blease's distributed control by configuring each of Blease's IFCs with control unit functionality, to connect the IFCs through a fabric, to synchronize their multicast distribution information in response to requests, and to accommodate the addition of an IFC by disabling the IFC, updating its multicast forwarding table to a current state by sending a multicast membership query, and then activating the IFC. Weyman explains that benefits of configuring units in a fabric include appearing as a single entity to others, sharing workloads between units, increasing forwarding performance, and minimizing dependence on any single unit or connection in the fabric.⁸¹ And, redundancies are created such that if any single unit or connection fails, the system may recover quickly with the remaining units.⁸² Built-in redundancies and failover concepts are fundamental design principles of network architecture and design, and this was well-known and the norm in the industry. Furthermore, the addition of processing units also allows for the replacement of such failed units, as a sufficient or threshold number of units

⁸¹ Weyman, [0024].

⁸² Weyman, [0052], [0072]-[0075], [0086], [0093], [0162]; Blease, 11:34-45.

and/or backup units may be required for normal operation and for handling a desired amount of network traffic.

96. A person of ordinary skill in the art would have had a reasonable expectation of success configuring Blease with a distributed control unit and functionality, as taught by Weyman. Indeed, Blease explicitly suggests to distribute the control unit, and switching fabrics were well-known in the art at the time of the '394 Patent.⁸³

Furthermore, Blease explicitly discloses that interfaces may be deactivated and activated, administratively or operationally, or for any other reason.⁸⁴

Implementation of known components in known configurations would have been well within the skill of a person of ordinary skill in the art.

B. Independent Claims 8 and 18

97. It is my opinion that Blease in view of Weyman renders obvious each and every element of independent claims 8 and 18. A brief summary of the claim elements and applicable art are provided below:

'394 Patent Claim	'394 Patent Claim language	Applied Prior Art ⁸⁵
-------------------	----------------------------	---------------------------------

⁸³ Blease, 11:53-57; *see also* my discussion of fabrics in Section III.A.

⁸⁴ Blease, 11:34-45.

⁸⁵ The identification of prior art references in the table is only provided for convenience and is only intended to be a broad indicator of which reference(s) are applied to the claim language in the discussions that follow. It is not intended to be limiting to the analysis in any way. As claim limitations build upon and/or refer to earlier claim limitations, each prior art reference should be considered to be applicable to the combined teachings as a whole. In concluding that the claims of

'394 Patent Claim	'394 Patent Claim language	Applied Prior Art⁸⁵
Claim 8		
8[pre]	A method for communication, comprising:	Blease
8[a.i]	operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports,	Blease, Weyman
8[a.ii]	wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and 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;	Blease, Weyman
8[b]	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;	Blease, Weyman
8[c]	receiving from a client a request relating to reception of a given multicast packet stream;	Blease
8[d]	distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request; and	Blease, Weyman
8[e]	selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists,	Blease
8[f.i]	wherein operating the network element comprises initializing a new processing unit added to the network element by:	Weyman

the '394 Patent are obvious over Blease in view of Weyman and Hu (in Section VIII), or further in view of Deng (in Section IX), I have considered the combined teachings of the prior art as a whole from the perspective of a person of ordinary skill in the art.

'394 Patent Claim	'394 Patent Claim language	Applied Prior Art⁸⁵
8[f.ii]	deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned.	Weyman
Claim 18		
18[pre]	A network element, comprising;	Blease
18[a]	multiple ports; and	Blease, Weyman
18[b.i]	multiple interconnected processing units,	Blease, Weyman
18[b.ii]	wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and	Blease, Weyman
18[b.iii]	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,	Blease, Weyman
18[b.iv]	wherein the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,	Blease, Weyman
18[b.v]	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,	Blease, Weyman
18[b.vi]	to receive from a client a request relating to reception of a given multicast packet stream, and	Blease

'394 Patent Claim	'394 Patent Claim language	Applied Prior Art⁸⁵
18[c.i]	to distribute the request among the processing units within the network element so as to update the lists responsively to the distributed request and	Blease, Weyman
18[c.ii]	to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists,	Blease
18[d.i]	wherein a new processing unit added to the network element is configured to initialize by:	Weyman
18[d.ii]	deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned.	Weyman

1. Preambles of Claims 8 and 18

98. Claim 8's preamble recites: "A method for communication, comprising," and claim 18's preamble recites: "A network element, comprising." Blease discloses all elements of these preambles to the extent either limits the scope of the claims.

99. Blease discloses a method for providing subscriber devices 8 with traffic from a network 10 through a network device 4.⁸⁶ See Figure 1:




FIG. 1

100. For example, Blease explains that “Network device 4 is capable of providing subscriber devices 8 with multicast traffic, e.g., Internet Protocol television (IPTV), from network 10. Subscriber devices 8 request multicast streams, e.g., IPTV channels, by sending a membership request for a specific multicast group to network device 4.”⁸⁷

⁸⁶ Blease, 3:34-58.

⁸⁷ Blease, 3:53-55.

101. In Figure 4 below, network device 4 is shown as router 42.⁸⁸




FIG. 4

102. Thus, Blease discloses a method for communication and a network element.

2. Claim Elements 8[a.i], 18[a], and 18[b.i]

103. Claim elements 8[a.i], 18[a], and 18[b.i] are set forth in the below table, and recite substantively similar limitations. Blease discloses, or in view of Weyman renders obvious, these claim elements.

⁸⁸ Blease, 8:39-52.

Claim Element 8[a.i]	Claim Elements 18[a] and 18[b.i]
operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports	multiple ports; and multiple interconnected processing units,

104. As I explained above with respect to at least the preamble of claim 1, Blease discloses that router 42 is a network device (*i.e.*, a network element) operated to provide subscriber devices with multicast traffic.

105. In Figure 4, annotated below, network device 4 is shown as **router 42**.⁸⁹




FIG. 4

106. **Router 42** includes a plurality of **interface cards (IFCs) 54A-54N**, each having an associated **inbound link 55A-55N** and **outbound link 56A-56N**. The

⁸⁹ Blease, 8:39-52.

inbound and outbound links are the interfaces through which **router 42** communicates with other entities and to, for example, receive multicast stream packets, forward multicast stream packets, and receive membership requests.⁹⁰ Blease uses the terms “link” and “interface” to refer to the connections between network entities and points on which those connections are made—which may be physical or virtual / logical elements.⁹¹ From these disclosures, a person of ordinary skill in the art would have recognized that **router 42** includes a plurality of “ports” as interfaces used to form links with other network entities—that is, identifiable endpoints where connections start or end.⁹² Thus, Blease discloses that **router 42** includes multiple ports.

107. A person of ordinary skill in the art would have also recognized that each **IFC 54** is a “processing unit” because, at the least, receiving, forwarding, and transmitting packet information requires a processor to interpret and format information as necessary for interpretation by other devices, such as when

⁹⁰ Blease, 1:28-33, 3:65-4:9, 8:39-52, 10:6-13.

⁹¹ For example, *see* Blease, 1:28-33, 3:65-4:9.

⁹² For example, a person of ordinary skill in the art would have been aware of prior art multicast traffic being sent over “links” and being received and forwarded through “ports.” Chen (EX1009), 1:5-27, 5:47-62. *See also* Hu, 3:19-26 (ports used to provide links).

transmitting packets to subscribers, network switches, or **control units**.⁹³ Thus, Blease discloses that **router 42** includes multiple processing units.

108. Additionally, Blease discloses that **control unit 44** and its functionality is distributed within **IFCs 54**.⁹⁴ As discussed above in Section VII.A.1, it would have been obvious, in view of Weyman, to configure each of Blease's **IFCs** with **control unit** functionality, to connect the **IFCs** through a fabric, and to synchronize their multicast distribution information. Thus, Blease further discloses that **router 42** includes **multiple processing units**, which are interconnected.

109. Thus, Blease in view of Weyman renders obvious operating a network element, wherein the network element further comprises multiple interconnected processing units and multiple ports.

3. Claim Elements 8[a.ii] and 18[b.ii]-18[b.iv]


110. Claim elements 8[a.ii] and 18[b.ii]-18[b.iv] are set forth in the below table, and recite substantively similar limitations. Blease discloses, or in view of Weyman renders obvious, these claim elements.

⁹³ For example, a person of ordinary skill in the art would have been aware of processors performing touring functions and processing packet transfers to and from ports and switch fabrics) Chen, 4:40-56. *See also* my discussion of line cards that, similarly to interface cards, include circuitry that provide an interface with a network for data transmission, in Section III.A.

⁹⁴ Blease, 11:54-57.

Claim Element 8[a.ii]	Claim Elements 18[b.ii]-[b.iv]
wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and 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	wherein each processing unit is assigned to a different subset of the multiple ports of the network element, and 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,

111. As I discussed with respect to claim element 8[a.i] above, Blease discloses that each IFC 54 of router 42 includes inbound links 55 and outbound links 56 representative of connections to ports of router 42. As shown in the excerpted portion of Figure 4 below, each IFC 54A-54N has its own respective links 55A-55N and 56A-56N for receiving and sending packets. Thus, Blease discloses that each processing unit is assigned to a different subset of the multiple ports of the network element.



112. As I further discussed with respect to claim element 8[a.i], Blease, alone or in view of Weyman, further discloses that each IFC 54 is a processing unit. For example, each IFC includes control unit 44, which includes and utilizes multicast

forwarding information 48 to forward packets that it receives.⁹⁵ Blease explains that control unit may be implemented as any combination of software and hardware, such as one or more processors executing stored instructions.⁹⁶ Thus, Blease, alone or in view of Weyman, discloses that 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.

4. Claim Elements 8[b] and 18[b.v]

113. Claim elements 8[b] and 18[b.v] are set forth in the below table, and recite substantively similar limitations. Blease discloses, or in view of Weyman renders obvious, these claim elements.

Claim Element 8[b]	Claim Element 18[b.v]
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	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

114. As I discussed with respect to at least claim element 8[a.i] and in Section VIII.A above, Blease, alone or in view of Weyman, discloses that control unit 44

⁹⁵ Blease, 8:53-67.

⁹⁶ Blease, 11:63-12:3.

and its functionality is included in each IFC 54. Control unit 44 includes components such as OIF mapper 50 and multicast forwarding information 48.⁹⁷

115. Blease discloses that multicast forwarding information 48 “controls replication and forwarding of multicast packets to subscriber devices” and is updated “based on messages received from the subscriber devices that indicate a desire to join or leave multicast groups, i.e., to receive or stop receiving multicast stream.”⁹⁸ For example, outgoing interface (OIF) information is mapped to multicast forwarding information 48 to indicate which OIF⁹⁹ is used to stream the multicast stream.¹⁰⁰ When a multicast stream no longer has any subscribers, or if an OIF becomes disabled, multicast forwarding information 48 is updated to remove the associated mappings.¹⁰¹ While Blease does not illustrate an example of multicast forwarding information 48, a person of ordinary skill in the art would have recognized that multicast forwarding information 48 is a “list” because it

⁹⁷ Blease, Fig. 4, 9:43-44, 10:49-61.

⁹⁸ Blease, 10:50-55.

⁹⁹ As I discussed in with respect to claim element 1[a.i], Blease’s interfaces represent “ports” of router 42.

¹⁰⁰ Blease, 8:64-67, 10:49-52, 14:50-57.

¹⁰¹ Blease, 11:5-17, 11:34-51.

includes a structure in which entries are added, removed, and processed as individual items, as discussed above.¹⁰²

116. Thus, Blease discloses 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.

5. Claim Elements 8[c] and 18[b.vi]

117. Claim elements 8[c] and 18[b.vi] are set forth in the below table, and recite substantively similar limitations. Blease discloses these claim elements.

Claim Element 8[c]	Claim Element 18[b.vi]
receiving from a client a request relating to reception of a given multicast packet stream	to receive from a client a request relating to reception of a given multicast packet stream, and

118. As I discussed with respect to claim 1’s preamble, Blease explains that “Network device 4 is capable of providing subscriber devices 8 with multicast traffic, e.g., Internet Protocol television (IPTV), from network 10. Subscriber devices 8 request multicast streams, e.g., IPTV channels, by sending a membership request for a specific multicast group to network device 4.”¹⁰³ Blease further states that “router 42 receives a membership request for a multicast group from a

¹⁰² Also, Blease describes that similar mapped information are “listed.” Blease, 6:62-63, 8:29-30.

¹⁰³ Blease, 3:53-55.

subscriber device via a network switch on an interface of one of IFCs 54,” for example using IGMP messages.¹⁰⁴ Multicast streams are communicated as “packets.”¹⁰⁵

119. Thus, Blease discloses receiving from a client a request relating to reception of a given multicast packet stream.

6. Claim Elements 8[d] and 18[c.i]

120. Claim elements 8[d] and 18[c.i] are set forth in the below table, and recite substantively similar limitations. Blease, in view of Weyman, discloses these claim elements.

Claim Element 8[d]	Claim Element 18[c.i]
distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request; and	to distribute the request among the processing units within the network element so as to update the lists responsively to the distributed request and

121. As I discussed with respect to claim element 8[c] above, Blease discloses receiving a request relating to reception of a multicast packet stream. Blease discloses that such requests are forwarded to control unit 44 for processing.¹⁰⁶ And, as I discussed with respect to at least claim element 8[a.i], Blease discloses, or in view of Weyman renders obvious, that control unit 44 is distributed within IFCs

¹⁰⁴ Blease, 10:6-12, 14:3-9.

¹⁰⁵ Blease, 8:53-67.

¹⁰⁶ Blease, 10:12-13.

54, and so each of router 42's IFCs 54 includes control unit 44. Weyman further teaches that in a distributed system, multicast distribution information must be kept synchronized: "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."¹⁰⁷

122. Weyman therefore teaches that Blease's control units 44 (in each IFC 54) may process IGMP requests in at least two ways: (1) the request is distributed to each control unit 44 in each IFC 54 for processing, or (2) the request is processed by the control unit 44 of the IFC 54 that received the request, and result are distributed accordingly. In both situations, Blease in view of Weyman renders obvious distributing the request among the multiple processing units within the network element and updating the respective list in each processing unit responsively to the request.

¹⁰⁷ Weyman, [0162]

a. (1) The Request is Distributed to Each Control Unit for Processing

123. In this implementation, the request is received at a first control unit (for example, at an IFC 54A) and then distributed to the other control units (for example, to IFCs 54B-54N), and therefore the request is distributed among the multiple processing units within the network element as claimed.

124. Each control unit 44 of each IFC 54 includes OIF mapper 50 and multicast forwarding information 48. If the request was received on an interface that is not configured to use an OIF map—and instead operate similar to conventional transmission techniques—then OIF mapper 50 will map the request by updating multicast forwarding information 48 to indicate that the multicast stream will be forwarded from the same interface that received the request.¹⁰⁸ Therefore, the respective list in each processing unit is updated responsively to the request as claimed.

125. Thus, Blease in view of Weyman renders obvious distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request.

¹⁰⁸ Blease, 6:62-7:3, 14:10-20.

b. (2) The Request is Processed by One Control Unit and the Results are Distributed to Each Other IFC

126. In this implementation, the request is received at a first control unit (for example, at an IFC 54A, similar to the first implementation above), processed at the first control unit (for example, determining to update its own multicast forwarding information 48 to indicate that the multicast stream will be forwarded from the interface that received the request, similar to the first implementation above), and then that determination will be distributed to the other control units (for example, to IFCs 54B-54N) so that they may each update their own multicast forwarding information 48, thereby keeping distribution information synchronized, as taught by Weyman. This configuration also renders obvious distributing the request and updating the respective list in each processing unit, as claimed.

127. In particular, dependent claims 4 and 14 of the '394 Patent explain that the claimed distributing and updating are satisfied by “distributing among the multiple processing units an identity of an ingress port over which the request was received at the network element” and “causing the list to indicate that the given multicast packet stream is permitted for forwarding over the ingress port.”

128. As I explained in this implementation, the control unit 44 of the IFC 54 that receives the request instructs the other IFCs 54 (each including their own control units 44) to update their own multicast forwarding information 48 to map the request to the interface that received the request. The transmission of update

information from the first control unit 44 to each other control unit 44 that causes the other control units 44 to update their multicast forwarding information 48 is therefore “distributing” and “updating” according to dependent claims 4 and 14 of the ’394 Patent.

129. Thus, Blease in view of Weyman further renders obvious, in another way, distributing the request among the multiple processing units within the network element, and updating the respective list in each processing unit responsively to the request.

7. Claim Elements 8[e] and 18[c.ii]

130. Claim elements 8[e] and 18[c.ii] are set forth in the below table, and recite substantively similar limitations. Blease discloses these claim elements.

Claim Element 8[e]	Claim Element 18[c.ii]
selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists	to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists,

131. Blease discloses that multicast packets of requested multicast packet streams are forwarded through particular interfaces of IFCs 54 by first identifying them in multicast forwarding information 48. For example, Blease states that “control unit

44 forwards the packet on a set of outbound links 56 in accordance with multicast forwarding information 48.”¹⁰⁹

132. Thus, Blease discloses selectively forwarding multicast packets associated with the given multicast packet stream by the processing units in accordance with the respective updated lists.

8. Claim Elements 8[f.i] and 18[d.i]

133. Claim elements 8[f.i] and 18[d.i] are set forth in the below table, and recite substantively similar limitations. Blease in view of Weyman renders obvious these claim elements.

Claim Element 8[f.i]	Claim Element 18[d.i]
wherein operating the network element comprises initializing a new processing unit added to the network element by:	wherein a new processing unit added to the network element is configured to initialize by:

134. Weyman teaches that a “system can accommodate the addition or removal of units” and such a unit “needs to update its configuration to that of the distributed router.”¹¹⁰ As I discussed in Section VIII.A, a person of ordinary skill in the art would have been motivated to modify Blease’s router to accommodate the addition of a new processing unit and to update that new processing unit, as taught by Weyman.

¹⁰⁹ Blease, 8:64-67; *see also* Blease, 10:49-52, 14:50-57.

¹¹⁰ Weyman, [0069]; *see also* Weyman, [0156].

135. Thus, Blease in view of Weyman renders obvious operating the network element comprises initializing a new processing unit added to the network element.

9. Claim Elements 8[f.ii] and 18[d.ii]

136. Claim elements 8[f.ii] and 18[d.ii] are set forth in the below table, and recite substantively similar limitations. Blease in view of Weyman renders obvious these claim elements.

Claim Element 8[f.ii]	Claim Element 18[d.ii]
deactivating the ports to which the new processing unit is assigned;	deactivating the ports to which the new processing unit is assigned;
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	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
after constructing the list, activating the ports to which the new processing unit is assigned.	after constructing the list, activating the ports to which the new processing unit is assigned.

137. As I discussed with respect to claim element 8[f.i] above, Weyman teaches initializing a new processing unit added to the network element by updating the new processing unit's configuration to that of the distributed router.

138. Weyman further teaches that a process for determining current multicast group memberships includes sending query messages and updating a multicast

distribution tree in response. For example, “[q]uery messages are sent by IP multicast routers to determine the group membership of clients,” and “[c]hanges in the multicast membership reported may affect the multicast distribution tree that the router maintains” and that are kept synchronized across units.¹¹¹

139. As I discussed in Section VIII.A, a person of ordinary skill in the art would have been motivated to modify Blease’s router to accommodate the addition of a new processing unit, such as an IFC, and to update that new processing unit by sending a multicast status request, updating the multicast forwarding table in response, and activating the IFC for processing, as taught by Weyman.

140. Thus, Blease in view of Weyman renders obvious deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned..

¹¹¹ Weyman, [0161]-[0162].

IX. GROUND 2 – CLAIM 9 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND HU

141. Dependent claim 9 depends from claim 8, which is obvious over Blease in view of Weyman. As I detail below, it is my opinion that Blease in view of Weyman and Hu renders obvious dependent claim 9.

142. In summary, although Blease and Weyman teach a distributed multicast router that grants requests to join multicast groups by updating multicast forwarding information (as discussed for example in Section VIII above), Blease and Weyman do not explicitly disclose criteria for granting requests. Specifically, Blease and Weyman do not appear to explicitly disclose that “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,” as recited in claim 9 of the ’394 Patent. Hu teaches setting pre-defined thresholds that place limits on the number of streams permitted over router links, which would have been obvious to use to increase the quality of transmissions and to avoid exceeding the available bandwidth on any link (as discussed in Section IX.A below). The teachings of Blease, Weyman, and Hu, and a person of ordinary skill in the art’s motivation for combining their teachings, summarized above, are described in detail below.

A. Motivation to Combine the Teachings of Blease, Weyman, and Hu

143. Hu is in the same field of endeavor and contains overlapping disclosures with similar purposes as the '394 Patent, Blease, and Weyman.¹¹²

144. Blease discloses network device 4, such as router 42, and its various functions including forwarding multicast streams according to updated routing information described in multicast forwarding information 48, as I have discussed in Section VIII above.¹¹³ Blease also describes tracking the number of subscriber devices receiving multicast streams for each outgoing interface (OIF). Blease describes the tracking with respect to a pruning function—removing a multicast group from multicast forwarding information 48 when a group no longer has any subscribers.¹¹⁴ With respect to tracking, Hu further teaches to improve the efficiency of links by, for example, introducing pre-defined thresholds to limit the number of streams using a particular link aggregated group (LAG) of ports and denying requests to join multicast groups that will exceed those thresholds, as I have discussed in Section V.C.

145. Therefore, a person of ordinary skill in the art would have been motivated to place threshold limits on Blease's links, as taught by Hu. A person of ordinary skill

¹¹² See my discussion in Section VIII.A and Hu, 1:6-7, 6:10-22.

¹¹³ See, for example, Blease, 6:62-7:2, 8:64-67, 9:14-24, 10:6-13, 10:49-52, 14:3-5, 14:10-20, 14:50-57.

¹¹⁴ Blease, 11:5-17.

in the art would have recognized that placing limits on connections (and even on non-LAG connections) would have increased the quality of those connections.

When attempting to transmit more bandwidth than is available on a link, the quality of the transmission suffers.¹¹⁵ For multicast streams like IPTV that involves video, decreased quality streams may manifest as delays, artifacts, dropped frames, dropped connections, and the like. And, for videos with variable bitrates, sufficient overhead is even more desirable.¹¹⁶

146. A person of ordinary skill in the art would have had a reasonable expectation of success utilizing pre-defined thresholds to limit streams on connections, as taught by Hu. The application of threshold-based limits to manage connections—by comparing measured values—would have been well within the skill of a person of ordinary skill in the art.

B. Dependent Claim 9

147. It is my opinion that Blease in view of Weyman and Hu renders obvious each and every element of dependent claim 9.

148. Claim 9 depends from claim 8 and includes the limitations set forth in the below table. Blease in view of Weyman and Hu renders obvious the additional limitations of this claim.

¹¹⁵ Hu, 1:20-29.

¹¹⁶ Hu, 1:16-19, 1:58-62.

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.

149. Hu teaches setting and referencing pre-defined thresholds for LAGs when deciding whether to grant a request to join a multicast group.¹¹⁷ For example, Hu teaches processing multicast communications including the following exemplary steps:

- Receiving an IGMP request to join a multicast group and receive a video stream;¹¹⁸
- Determining a traffic load at each port associated with the LAG that received the request, taking into account recent join and leave requests and changing bitrates, for example;¹¹⁹
- Determining whether a total traffic load, taking into account the request, would exceed a pre-defined threshold;¹²⁰

¹¹⁷ Hu, 3:10-32, 5:13-22.

¹¹⁸ Hu, 6:10-17.

¹¹⁹ Hu, 6:23-38.

¹²⁰ Hu, 6:51-60.

- If the total traffic load is lower than the pre-defined threshold, granting the request and assigning the stream to a port in the LAG having the lowest load;¹²¹
- If the total traffic load is higher than the pre-defined threshold, rejecting the request unless an existing data stream can be dropped to accommodate the request.¹²²

Hu therefore teaches setting an upper limit on the number of streams permitted for forwarding over a link such as a LAG.¹²³

150. Furthermore, I have been informed that Patent Owner, in its district court litigation, may be asserting that this limitation may be met and infringed broadly by load-balancing techniques.¹²⁴

151. The technique of setting an upper limit on the number of streams based on measured loads across the LAG and assigning a stream to a lowest-load port *is* a load-balancing algorithm that is designed to prevent issues that arise when faced with insufficient bandwidth, as I’ve discussed in at least Section IX.A above. In

¹²¹ Hu, 6:51-52.

¹²² Hu, 6:62-7:8.

¹²³ *See also* Hu, 5:13-22, 5:30-6:3, 13:8-10.

¹²⁴ EX1013, 17-18. The terms used in claim 9 are not identical to those identified in the exhibit from claim 1, but they appear to both relate to not-updating a list when a “maximum value” (claim 9) or “upper limit” (claim 1) related to a number of multicast packets streams is exceeded.

fact, Hu explains that its teachings seek to improve on “static load-balancing algorithm[s]” by “measuring real-time port usage related to each physical port associated with a link aggregation group.”¹²⁵ Hu therefore teaches this limitation to the extent “load-balancing” is required.

152. Indeed, a person of ordinary skill in the art would have recognized that the load-balancing was typically implemented for systems having multiple ports, links, or connections, such as LAGs, as I have discussed in Section III.B, and so a person of ordinary skill in the art would have been motivated to implement any of a number of known algorithms, such as Hu’s teaching using real-time port usage, or other traditional load-balancing algorithms that limit traffic based on “stream count.”¹²⁶ As I discussed in Section IX.A, a person of ordinary skill in the art would have been motivated to modify Blease’s links to have limits, as taught by Hu. The well-known traditional load-balancing techniques discussed above provide the same benefits with respect to network usage and congestion and would have been obvious to implement for the same reasons. Therefore, where a request to join a multicast group would normally result in an updated multicast forwarding information 48, as I have discussed with respect to at least claim element 8[d]

¹²⁵ Hu, 1:43-45, 2:13-16.

¹²⁶ Hu, 1:46-50, 1:59-61.

above, it follows from Hu's teachings that multicast forwarding information 48 is only updated when the limit is not exceeded.

153. Thus, Blease in view of Weyman and Hu renders obvious that 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.

X. GROUND 3 – CLAIM 10 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND DENG

154. Dependent claim 10 depends from claim 8, which is obvious over Blease in view of Weyman. As I detail below, it is my opinion that Blease in view of Weyman and Deng renders obvious dependent claim 10.

155. In summary, although Blease and Weyman do not appear to explicitly disclose criteria for dropping streams including “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,” as recited in claim 10 of the ’394 Patent. Deng teaches these tracking and removing functions with the use of a timer, which would have been obvious to use to drop old streams in the system and method of Blease in view of Weyman (as discussed in Section X.A below). The teachings of Blease, Weyman, Deng, and a person of ordinary skill in the art’s motivation for combining their teachings, summarized above, are described in detail below.

A. Motivation to Combine the Teachings of Blease, Weyman, and Deng

156. Deng is in the same field of endeavor and contains overlapping disclosures with similar purposes as the ’394 Patent, Blease, and Weyman.¹²⁷

¹²⁷ See my discussion in Section VIII.A and Deng, [0002].

157. Blease discloses network device 4, such as router 42, and its various functions including forwarding multicast streams according to updated routing information described in multicast forwarding information 48, as I have discussed in Section VIII above.¹²⁸ As I have also discussed in Section IX above, Blease describes a pruning function of tracking the number of subscriber devices receiving multicast streams for each outgoing interface (OIF) and removing a multicast group from multicast forwarding information 48 when a group no longer has any subscribers.¹²⁹ Deng also teaches a pruning procedure that tracks the amount of time after a request to join a multicast group has been granted, and removes the corresponding address from a multicast address table once an amount of time has expired.¹³⁰

158. A person of ordinary skill in the art would have been motivated to modify Blease and Weyman to additionally prune aged permissions from multicast forwarding information 48 and drop those streams upon expiration of a timer, as taught by Deng. The addition of a time-based function to drop streams provides for

¹²⁸ See, for example, Blease, 6:62-7:2, 8:64-67, 9:14-24, 10:6-13, 10:49-52, 14:3-5, 14:10-20, 14:50-57.

¹²⁹ Blease, 11:5-17.

¹³⁰ Deng, [0018], [0031], [0035]-[0036]. See also Section V above.

increased bandwidth sufficient to grant requests to join multicast groups or otherwise maintain a desired quality of existing streams.¹³¹

159. Additionally, a person of ordinary skill in the art would have had a reasonable expectation of success configuring Blease with timer-based functionality. Indeed, timers and clocks were well-known processing components, and Deng explains that only a “conventional” counter or counting mechanism is used.

B. Dependent Claim 10

160. It is my opinion that Blease in view of Weyman and Deng renders obvious each and every element of dependent claim 10.

161. Claim 10 depends from claim 8 and includes the limitations set forth in the below table. Blease in view of Weyman and Deng renders obvious the additional limitations of this claim.

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.

¹³¹ As I discussed in Section IX.A, transmission quality and video quality suffers when bandwidth is unavailable.

162. As I discussed in Section X.A above, Blease in view of Weyman and Deng renders obvious 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 by, for example, tracking an amount of time that a host has been included on a multicast table since being added in response to a request.¹³² And, Blease in view of Weyman and Deng further renders obvious removing the permission from the respective lists responsively to determining that the permission is aged by, for example, removing an aged permission through updating multicast forwarding table 48, which as I have discussed in at least Section VIII.A, is kept synchronized at IFCs 54.

163. Thus, Blease in view of Weyman and Deng renders obvious that 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.

¹³² Deng, [0035]-[0036].

XI. GROUND 4 – CLAIMS 8 AND 18 ARE UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN AND RAO

164. The Challenged Claims include independent claims 8 and 18. As I detail below, a person of ordinary skill in the art would have recognized that each and every limitation of claims 8 and 18 is disclosed or rendered obvious over Blease in view of Weyman and Rao.

165. In summary, as I have discussed in Section VIII above, Blease in view of Weyman renders obvious claims 8 and 18 of the '394 Patent, including the elements 8[f.i] and 8[f.ii] related to initializing a new processing unit by “deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned,” as recited in claim 8. For example, as I discussed, a person of ordinary skill in the art would have been motivated to modify Blease’s router to accommodate the addition of a new processing unit, such as an IFC, and to update that new processing unit by deactivating the processing unit’s ports, sending a multicast status request, updating the multicast forwarding table in response, and activating the IFC ports for processing, as taught by Weyman.

166. However, to the extent Weyman does not explicitly spell out certain steps for configuring a new processing unit that would have been obvious to person of ordinary skill in the art, Rao teaches these procedures so that the processing unit is configured for the current multicast state. For example, Rao teaches to only activate a new processing unit after all multicast membership information is received or a period for reply has elapsed. The teachings of Blease, Weyman, and Rao, and a person of ordinary skill in the art's motivation for combining their teachings, summarized above, are described in detail below.

A. Motivation to Combine the Teachings of Blease, Weyman, and Rao

167. Rao is in the same field of endeavor and contains overlapping disclosures with similar purposes as the '394 Patent, Blease, and Weyman.¹³³

168. As discussed above, Blease in view of Weyman renders obvious the initialization of a new processing unit such as an IFC, and to update that new processing unit by deactivating the processing unit's ports, sending a multicast status request, updating the multicast forwarding table in response, and activating the IFC ports for processing. Rao teaches these same steps for initializing a processing unit by "generat[ing] Host Membership Queries," starting a "delay

¹³³ See my discussion in Section VIII.A and Rao, 1:6-7, 6:10-22.

timer,” updating entries in routing tables in response to query replies, and finally activating multicast forwarding information for data path forwarding.¹³⁴

169. To the extent that Blease and Weyman do not explicitly instruct a person of ordinary skill in the art with steps to initialize a processing unit, as recited in claims 8 and 18 of the '394 Patent, a person of ordinary skill in the art would have been thus motivated to look to known implementations, as in Rao, for details. Rao explains that by learning of the multicast state prior to activation, “deleterious effects are minimized” from any inconsistencies caused by different timings in the updating process.¹³⁵

170. A person of ordinary skill in the art would have had a reasonable expectation of success configuring Blease with the initialization processes taught by both Weyman and Rao. Indeed, router functions of disabling ports, updating multicast forwarding information, and activating ports were well-known in the art and well within the skill of a person of ordinary skill in the art.¹³⁶

B. Independent Claims 8 and 18

171. It is my opinion that Blease in view of Weyman and Rao renders obvious each and every element of independent claims 8 and 18. My analysis for claim elements 8[a] through 8[f.i] and 18[a] through 18[d.i], based on the teachings of

¹³⁴ Rao, 11:20-22, 11:29-32, 11:33-44.

¹³⁵ Rao, 9:66-10:7.

¹³⁶ See Blease, 11:34-45, 11:53-57.

Blease and Weyman, remain the same for my analysis based on the combined teachings of Blease, Weyman, and Rao below. Blease in view of Weyman and Rao renders obvious the additional limitations of claims 8 and 18—8[f.ii] and 18[d.ii].

1. Claim Elements 8[f.ii] and 18[d.ii]

172. Claim elements 8[f.ii] and 18[d.ii] are set forth in the below table, and recite substantively similar limitations. Blease in view of Weyman and Rao renders obvious these claim elements.

Claim Element 8[f.ii]	Claim Element 18[d.ii]
deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned.	deactivating the ports to which the new processing unit is assigned; 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 after constructing the list, activating the ports to which the new processing unit is assigned.

173. As I discussed with respect to claim element 8[f.i] above, Weyman teaches initializing a new processing unit added to the network element by updating the new processing unit’s configuration to that of the distributed router.

174. Rao further teaches to initialize a new processing unit by sending a multicast status request. For example, Rao states to “generate Host Membership Queries on its connected networks to solicit membership information from attached IGMP hosts” and to start a “delay timer.”¹³⁷ Then, based on received responses to the Host Membership Queries, “current [routing] tables” are updated.¹³⁸ Once the current multicast state is determined, for example indicated by the receipt of EOTIB markers or the end of the delay timer, normal operation may begin.¹³⁹ Rao teaches that this includes “activating the internal router path IP multicast forwarding information for data path forwarding.”¹⁴⁰

175. Although Rao does not describe the processing unit and its data path forwarding as first being deactivated prior to the updating described above, a person of ordinary skill in the art would have recognized based on Rao’s teachings that the processing unit is deactivated. In particular, in Rao’s specific example, the processing unit is a “standby” processor.¹⁴¹ A standby processor would not be expected to actively analyze and forward multicast stream packets, and so a person of ordinary skill in the art would have recognized that the processor is placed in a

¹³⁷ Rao, 9:3-17.

¹³⁸ Rao, 9:18-46.

¹³⁹ Rao, 9:47-65,

¹⁴⁰ Rao, 11:20-22, 11:29-32, 11:33-44.

¹⁴¹ Rao, 9:3-5.

deactivate state prior to initialization. Additionally, Rao teaches that, after updating, “the internal router path IP multicast forwarding information for data path forwarding” is activated.¹⁴² A person of ordinary skill in the art would have recognized this as further evidence that the processor was previously deactivated. Furthermore, as I have discussed with respect to Weyman in Section VIII.A above, a new processing unit would not be expected to have current forwarding information, and a person of ordinary skill in the art would not be motivated to put a processing unit into operation with incorrect forwarding information. Accordingly, the reasonable course of action would be to prevent the processing unit from forwarding packets until it is updated to the current multicast state, by deactivating and activating the processing unit.

176. As I discussed in Section XI.A, a person of ordinary skill in the art would have been motivated to modify Blease’s router to accommodate the addition of a new processing unit, such as an IFC, and to update that new processing unit by sending a multicast status request, updating the multicast forwarding table in response, and activating the IFC for processing, as taught by Weyman and Rao.

177. Thus, Blease in view of Weyman and Rao renders obvious deactivating the ports to which the new processing unit is assigned; while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request,

¹⁴² Rao, 11:20-22, 11:29-32, 11:33-44.

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 after constructing the list, activating the ports to which the new processing unit is assigned.”

XII. GROUND 5 – CLAIM 9 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN, RAO, AND HU

178. The Challenged Claims include dependent claim 9. As I explained in Section IX above, Blease in view of Weyman and Hu renders obvious claim 9 of the '394 Patent. For the same reasons that a person of ordinary skill in the art would have been motivated to combine the teachings of Blease, Weyman, and Hu, a person of ordinary skill in the art would have been motivated to combine the teachings of Blease, Weyman, Rao, and Hu.

179. Therefore, it is my opinion that Blease in view of Weyman, Rao, and Hu renders obvious each and every element of dependent claim 9.

XIII. GROUND 6 – CLAIM 10 IS UNPATENTABLE OVER BLEASE IN VIEW OF WEYMAN, RAO, AND DENG

180. The Challenged Claims include dependent claim 10. As I explained in Section X above, Blease in view of Weyman and Deng renders obvious claim 10 of the '394 Patent. For the same reasons that a person of ordinary skill in the art would have been motivated to combine the teachings of Blease, Weyman, and Deng, a person of ordinary skill in the art would have been motivated to combine the teachings of Blease, Weyman, Rao, and Deng.

181. Therefore, it is my opinion that Blease in view of Weyman, Rao, and Hu renders obvious each and every element of dependent claim 9.

XIV. CONCLUSION

182. For all the reasons explained above, it is my opinion that Blease in view of Weyman, or further in view of Hu or Deng, renders obvious claims 8-10 and 18 of the '394 Patent. Further, it is my opinion that Blease in view of Weyman and Rao, or further in view of Hu or Deng, renders obvious claims 8-10 and 18 of the '394 Patent.

183. 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 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 1001 of Title 18 of the United States Code.

5/31/2024 | 1:54:08 PM PDT

Dated: May 31, 2024

By: 
Dr. Tal Lavian