

IN THE UNITED STATES PATENT & TRADEMARK OFFICE

Reexamination Control No: 90/019,441

Filed: March 5, 2024

Patent No. 8,014,394

Title: HIGH-SPEED PROCESSING OF
MULTICAST CONTENT REQUEST

Examiner: Hotaling, John M.

Group Art Unit: 3992

Mail Stop "Ex Parte Reexam"
Attn: Central Reexamination Unit
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

DECLARATION OF JAMES OLIVIER, Ph.D.

I, Dr. James Olivier, declare under penalty of perjury as follows:

I. INTRODUCTION

1. My name is James L. Olivier, Ph.D. I have been retained on behalf of Patent Owner, Portsmouth Network Corporation, as an independent expert and have been asked to provide my expert opinions regarding U.S. Patent No. 8,014,394. Ex. 1001 (“the ’394 Patent”).

II. PROFESSIONAL BACKGROUND

2. My qualifications for forming the opinions given in this declaration are summarized here and are addressed more fully in my curriculum vitae, which is submitted as Exhibit 2002. My curriculum vitae also includes a list of my publications and technical consulting experience.

3. I am a recent Fulbright award winner in the areas of computer science and blockchain technologies, and I currently serve as an adjunct professor at Southern Methodist University.

4. I possess the knowledge, skills, experience, training, and education to form an expert opinion and testimony in this case. I am an expert in the field of networking technologies and have been an expert in this field since before 2002. I have designed and developed access systems such as those described in the ’394 Patent while working at Marconi Communications and teach MAC bridging and Virtual Private LAN Services like those systems described in the ’394 Patent in my Multiprotocol Label Switching class.

5. I have a Ph.D. from The Ohio State University in Electrical Engineering with minors in Discrete Mathematics, Computer Science, and Microelectronics. My specialty was fault tolerant computing and my Ph.D. dissertation was based on coding theory and was titled “Concurrent Error Detection in Arithmetic Processors using GAN Codes,” in which I developed new codes for use in Arithmetic Processors such as microprocessors. I have published papers on

fault tolerant computing and coding theory and continued my work while at the General Motor Research Laboratory.

6. I have extensive experience in network systems design and development and have specialized in networking product development. I have developed and designed equipment for networks since my time at AT&T Bell Laboratories in the 1990s, where I worked on AT&T 5ESS, Network Control Point, and Autoplex Series base stations as a Member of the Technical Staff. These systems made extensive use of a counter-rotating ring network topology known as the Common Network Interface, much like the network described in the '394 Patent. It was there where I first began my work with telecommunication standards bodies as a contributing member of the Asynchronous Transfer Mode ("ATM") ATM Forum. Later at DSC Communications, I was the Senior Manager of the ATM systems engineering group developing ATM packet switches for a new generation base station for Motorola's use in their Centralized Base Station Controller. While at DSC, I was also their corporate representative to the ATM Forum, participating in bi-monthly standards body development. I was one of the first contributors to the wireless standards at the ATM Forum.

7. At Samsung Telecommunications America, I was a Principal Engineer for wireless broadband services over Universal Mobile Telecommunications System ("UMTS"). UMTS is a third-generation ("3G") broadband standard developed by the 3rd Generation Partnership Project ("3GPP"). At Samsung, I worked on designing their next-generation cellular switch, a UMTS mobile switching center ("MSC"). While at Samsung, I was also their corporate representative to the International Telecommunication Union ("ITU"), the United Nations agency responsible for standardizing information and communication technologies. The ITU, under its International Mobile Telecommunications-2000 effort, was part of the 3GPP standards body. It was there that

I participated in the development of standards for advanced wireless networks. While at Samsung, I was responsible for the “services portion” of the UMTS switch. This subsystem provided Internet service, email, and the like to the subscribers.

8. At Marconi, I worked on several systems for the access market, such as Digital Subscriber Line (“DSL”) modems and Digital Subscriber Line Access Multiplexers (“DSLAMs”), along with the design of point-to-point wireless systems. At Navini Networks, I was responsible for layer 2 and layer 3 network protocols for their Wideband Code Division Multiplexed Access (“WCDMA”) wireless base stations and broadband modems. These layers were responsible for packet transmissions for various services over the WCDMA air interface. I was also responsible for the development of a Single Sign On (“SSO”) system which allowed users to gain information relating to their wireless subscription and delete or add services.

9. Currently I am an Adjunct Professor in the Network Engineering Graduate Program at Southern Methodist University’s School of Electrical Engineering, where I teach MPLS networked-enabled applications. Among the MPLS applications I teach is the use of MPLS in advanced network design, including content delivery networks, (CDNs) service provider network design and Intranet vs. Internet network designs. In my class I also teach DSLAMS, SONET/SDH rings, ADMs, Ethernet switch design and bandwidth allocation models. I also discuss multicast streams in MPLS and Virtual Private networks. I am also the owner of Olivier Consulting, where I provide consulting services for advanced network/product design along with Intellectual Property consulting. This year I received a three-year Fulbright specialist award in the area of computer science for my work at the Hunt Institute for Engineering and Humanity, which is part of the Lyle School of Engineering at Southern Methodist University in Dallas, Texas. There, I served as the Program Lead for Transformational Technology. Among the transformational technologies I

investigate at the Hunt Institute is the Internet of Things (“IoT”), which makes extensive use of Fourth-Generation and Fifth-Generation networks.

10. I am a co-inventor of U.S. patent No. 8,334,775 issued Dec. 18, 2012 and entitled “RFID-Based Asset Security and Tracking System, Apparatus and Method.” This invention relates to a Radio-Frequency Identification (RFID)-based Global Positioning System (GPS) tracking system to provide control and security of assets. This system integrates an RFID-based detection system with a conventional GPS tracking system. The GPS tracking system includes GPS receivers connected over a network to a centrally based GPS monitoring system.

11. Because of my background, training, and experience, I am qualified as an expert to explain the background of the technology encompassed by the '394 Patent as well as to offer opinions related to the validity of the claims of the '394 Patent in view of the prior art references discussed herein.

12. If called upon to do so, I could and would testify truthfully regarding the points stated below in this declaration.

III. COMPENSATION

13. I am being compensated for my independent analysis as an expert with respect to this *inter partes* review proceeding. My work on this matter is being billed at my customary rate of \$685.00 per hour. Also, I am being reimbursed for reasonable expenses I incur in relation to my services. My compensation is not contingent in any way on the content of my analysis, the opinions that I offer, or the outcome of this proceeding. I have no financial interest in Petitioner or Patent Owner, and I have no financial interest in the '394 Patent.

IV. OVERVIEW OF TASK & BASIS FOR OPINIONS

14. I have been asked to review the '394 Patent in light of the materials cited and discussed herein and my knowledge and experience in this field during the relevant time frame, which I understand to be February 11, 2009. I have been asked to provide opinions related to certain issues from the perspective of a person of ordinary skill, having knowledge of the relevant art, in the relevant time frame of February 11, 2009 and—except where otherwise noted—the opinions stated in this declaration are from that perspective. I describe the qualifications and abilities of such a person in the Section “Level of Ordinary Skill in the Art” below. I have also been asked to consider whether certain combinations of references have rendered obvious the inventions described in the claims of the '394 Patent.

15. My findings, as explained below, are based on my study, education, training, and experience as well as items that I reviewed to prepare my opinions, including the documents discussed in Mr. Wechselberger’s declaration and any other items I reference in my below analysis.

16. I understand that other issues may arise that require further explanation, and I will provide that explanation if appropriate. As a result, I respectfully reserve the right to update and supplement this declaration and the information and opinions provided herein.

V. MATERIALS CONSIDERED

17. In preparing this declaration and forming my opinions, I have considered: the claims, specification, and prosecution history of the '394 Patent; the Request for Re-Examination and the accompanying declaration by Mr. Wechselberger (including materials discussed in each); all materials cited and discussed in Mr. Wechselberger’s declaration; and the December 17, 2024 Office Action and all materials cited in the Office Action. As part of my analysis, I have also considered my own knowledge and experience, including my work and experience in the field.

VI. LEGAL CONSIDERATIONS

18. I am not an attorney. My understanding of the law is based on information provided by counsel for Patent Owner.

19. It is my understanding that assessing the validity/patentability of a U.S. patent claim based on a prior art analysis requires two essential steps. First, one must analyze the terms of the patent claims to understand what meaning one of ordinary skill in the art would give the terms. Second, one may then assess validity by comparing a patent claim to the “prior art.” I understand that the teaching of the prior art is viewed through the eyes of a person of ordinary skill in the art at the time the invention was made. My opinion as to what constitutes a relevant person of ordinary skill in the art is set forth below.

20. I understand that a claim is invalid as obvious under 35 U.S.C. § 103 if the differences between the subject matter sought to be patented and the prior art are such that the subject matter of the claim as a whole would have been obvious at the time of the invention to a person having ordinary skill in the art. I have been informed that the following factors are used to determine whether or not the claimed subject matter would have been obvious: (i) the scope and content of the prior art; (ii) the differences between the prior art and the claimed invention; (iii) the level of ordinary skill in the field of the invention; and (iv) any relevant objective considerations of non-obviousness.

21. I also understand that a party asserting obviousness based on a combination of prior art references must demonstrate that one of ordinary skill in the art would have been motivated to combine the teachings of those references to achieve the claimed invention and that the skilled artisan would have had a reasonable expectation of success in doing so. It is my understanding that it is not enough to show that one skilled in the art could combine elements of multiple

references, but instead there must be some reason that would have prompted a person of ordinary skill in the art to combine the elements in the way the claimed invention does. I understand that there must be some reasoned explanation as to why one of ordinary skill in the art would combine the references.

22. I understand that a combination of references would not have been obvious if the alleged modification(s) to be made to the reference(s) would have been inconsistent with the reference's stated goals or would have rendered the combination inoperable for its intended purpose. I further understand that for something to have been obvious, the party asserting obviousness must explain why a POSITA would have selected components for combination in the manner claimed.

23. I also understand that a party asserting obviousness based on a single prior art reference must nonetheless show a motivation to modify the reference.

24. I have been informed that a prior art reference must be considered in its entirety, i.e. as a whole, including portions that would lead away from the claimed invention.

25. It is my further understanding that obviousness cannot be based on a hindsight combination of components selected from prior art references. I also understand that an invention would not have been obvious simply because all of the elements of the invention may have been known separately in the prior art.

26. My opinions address what would have been logical to, and within the skill level of, a person of ordinary skill in the art ("POSITA") at the time of the invention, given the state of the relevant art, the knowledge and skill that a POSITA would have, the teachings of the references discussed below, and how a POSITA would have understood those teachings. My opinions also address whether a POSITA would have had a reasonable expectation of the modified systems

discussed below successfully yielding the results discussed below. I understand that a POSITA is a hypothetical person. I understand that a POSITA thinks along the lines of conventional wisdom. I understand that a POSITA is also a person of ordinary creativity, not an automaton (*see* Section “Level of Ordinary Skill in the Art,” below, for qualifications of a POSITA).

27. I understand that a claim term is given the meaning that the term would have to a POSITA at the time of the invention, which generally is the ordinary and customary meaning of the term. I further understand that the ordinary and customary meaning of a term may be evidenced by a variety of sources, including the words of the claims themselves, the specification, the prosecution history, and extrinsic evidence concerning relevant scientific principles, the meaning of technical terms, and the state of the art. I also understand that a patentee can act as a lexicographer by clearly expressing an intent to define a claim term.

VII. LEVEL OF ORDINARY SKILL IN THE ART

28. Counsel for Patent Owner has informed me that a person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention.

29. For the purpose of this Declaration, and based on my opinions herein, the time of invention for the '394 Patent is February 11, 2009. However, I understand that the conception dates of the inventions may have been earlier than that date. Any earlier or later conception date around February 11, 2009 would not alter my opinions herein.

30. Mr. Wechselberger opined that a POSITA “would have had a B.S. degree in electrical engineering, computer engineering, or equivalent training or job experience, with two years of experience in telecommunications systems, or a Master's degree in electrical engineering, computer engineering, or other equivalent degree. Additional education could compensate for less


practical experience, and vice versa.” Ex. 1003 (Wechselberger Decl.) at ¶ 19. In my opinion, a POSITA would have had a bachelor’s degree in computer science, computer engineering, electrical engineering, or an equivalent field, or approximately two years of experience working in the field of information technology and networking.

31. For the purposes of this reexamination, there is no appreciable difference between Mr. Wechselberger’s POSITA and my own. As a result, the opinions expressed in this declaration remain same under Mr. Wechselberger’s POSITA standard and my own. I have applied both POSITA standards in forming my opinions in this declaration.

VIII. BACKGROUND

a. Centralized, Decentralized, and Distributed Router Architectures


32. There are three basic router architectures: (1) centralized, (2) decentralized, and (3) distributed. As the name implies, a *centralized router architecture* centralizes decision-making regarding handling the multicast packets. The entire router has a single (central) processor and a single (central) database. There are no processors in the line card, only interface logic for the ports. Every client request the router receives is sent to the central database, which updates its list of instructions for how to handle incoming multimedia packets. When the router receives a multimedia packet, the central processor decides where to send the packet based on the instructions stored in the central database. The Hu prior art reference, cited in Grounds 1 and 2 of the Office Action, is an example of a centralized router architecture:



Ex. 1009 at FIG. 1. In the figure above, the centralized Processing Logic 104 uses the centralized Memory 106 to make decisions on how to handle any packets the router receives via Physical Ports 1 and Physical Ports 2. The advantage of the centralized router architecture is that there is very little risk of inconsistent treatment of incoming packets. Because there is only one centralized processor and database, every packet will be handled consistently and there is minimal risk the router will send the packet to the wrong destination.

33. While centralized router architectures maximize consistency, they run the risk of overburdening the central processor and database as the number of interfaces and ports increases. The *distributed router architecture* minimizes the risk of overburdening by replacing the central processor and database with processors and databases distributed to each line card. As the number of line cards increases, so does the processing power. As it utilizes a distributed router architecture,

the '394 patent does not have a central processor or database, relying instead on databases 52 and network processors 48 distributed to each line card 36 in the router (network element 32).



Ex. 1001 at FIG. 1. The '394 patent refers to the line cards 36 as “processing units” since they do the processing to determine how to handle incoming multicast packets. For example, when the router receives a multicast packet it arrives on one of the line cards 36 via a port 44. The network processor 48 of the receiving line card queries the database 52 of the same line card for instructions on how to handle the packet. Once the receiving line card’s processor determines how to handle the packet, it executes the instruction and sends the packet to the appropriate destination.

34. While the distributed router architecture solves the overburdening problem, it creates a new risk of inconsistent treatment of packets. A distributed router architecture must have some mechanism to make sure that each line card processing unit has the same instructions as all the other line card processing units in the router. Otherwise, the line cards will behave inconsistently and multicast packets will be delivered to the wrong destinations. The *decentralized*

router architecture attempts to address both the overburdening and the inconsistency problems by creating a hybrid between the two architectures. Specifically, the decentralized router architecture maintains centralized decision-making while distributing some of the processing burden to line card processing units. For example, Naik's "route processors 210 handle routing and control of the data traffic forwarding." Ex. 1005 at [0046]. Thus, Naik's Route Processors (RP) 210 centralize decision-making for how to **route** packets to their destinations. Naik's "line cards 260 perform **forwarding** of the data traffic." *Id.* The processors on Line Cards 260 (FP Rx 280 and FP Tx 290) follow the instructions provided by the centralized Route Processors (RP) 210 to forward multicast packets. In this way Naik's centralized processor maintains consistent treatment without getting overburdened.




FIG. 2

Ex. 1005 at FIG. 2.


35. One important consequence of differing router architectures is that features required for one type of architecture are not necessary for, and even counterproductive to, others.

For the distributed router architecture of the '394 patent to function properly, each line card processing unit must have a "list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element." In other words, every line card processing unit must have the same packet handling instructions as all other line card processing units or packets will be mishandled. But centralized or decentralized router architectures, like that of Hu or Naik, do not need this list in each line card processing unit because they have a centralized processor ensuring consistent routing of packets. In fact, the presence of such lists in the line cards of centralized or decentralized routers would be counterproductive because they would use up precious memory space and create the possibility that the lists' instructions conflict with those of the centralized processor.

36. Another feature required for one architecture but unnecessary or counterproductive for other architectures is the distribution of client requests among the line card processing units. The distributed router architecture of the '394 patent requires each line card processing unit must distribute "the request among the multiple processing units within the network element" so that the other line card processing units can update their list of multicast packet streams permitted for forwarding. Without distributing every client request to every other line card processing unit, the line card processing units cannot update their lists and ensure consistent treatment of incoming packets. But for a centralized or decentralized architecture, like that of Hu or Naik, there is no need to distribute client requests to all the other line card. Instead, each line card sends the received client requests to the centralized processor, which ensures each line card has the instructions it needs. As a result, many of the '394 patent's requirements for line card processing units are not necessary, and counterproductive, for centralized and decentralized router architectures like those of Naik and Hu.

b. Overview of the '394 Patent


37. The '394 patent implements a distributed router architecture that lacks any centralized processor and database. Instead, the '394 patent distributes decision-making to multiple processing units (line cards 36) that each have their own processor (NP 48) and database (DB 52). When the router (network element 32) receives a multicast packet, it arrives at one of the line cards 36. The processor on this "ingress" line card queries the database on the same line card and decides where to send the multicast packet.



Ex. 1001 at FIG. 1.

38. But in order for the '394 patent to function properly, every line card 36 must handle packets consistently. The router will mishandle multicast packets if one line card follows a set of instructions and another follows different instructions. As a result, the '394 patent must have a mechanism to update instructions for handling multicast packets so that every line card processing unit follows the same rules. The '394 patent describes this mechanism in its discussion of Figure 2.

FIG. 2



Id. at FIG. 2.

39. “The [updating] method begins with network element 32 receiving a multicast content request message, at a request reception stage 60.” *Id.* at 7:30-32. The specification provides more detail about how the request is received by the network element:

The request is sent from a given multicast client 28, is addressed to a certain multicast source and arrives at one or [sic: of] ports 44 of the network element. The line card at which the request is received is referred to herein as the ingress line card, and the specific port in this line card over which the request arrives is referred to herein as the ingress port.

Id. at 7:34-39.

40. The ingress line card **duplicates this request** during duplication step 64, and distributes the copies as follows:

The NP in the ingress line card duplicates the request, at a duplication step 64. The first copy of the request is forwarded to the appropriate multicast sources, at a copy forwarding step 68. In addition, *the second copy of the request is forwarded to all the NPs in the network element (including the originating NP), via fabric 40, at a copy distribution step 72.*

Id. at 7:53-54 (emphasis added).

41. The independent claims of the '394 patent incorporate this “distribution step 72” through the “distributing the request among the multiple processing units within the network element” limitation. The '394 patent clarifies why the request must be distributed to every line card processing unit in the network element:

The NPs in the line cards of network element 32 receive the distributed copy of the request via fabric 40. Each NP receiving the distributed copy updates its respective database 52 accordingly, at an updating step 76.

...
The network element selectively forwards incoming multicast streams in accordance with the updated filtering data bases, at a selective forwarding step 80. Selective forwarding is carried out in a distributed manner by the different NPs in line cards 36. When a given NP 48 in a given line card 36 receives packets that are associated with a given multicast stream, the NP looks for a matching entry in its filtering database. If the database indicates that the multicast stream in question is permitted for forwarding to one or more ports, the NP forwards the packets associated with this stream to the permitted ports. Otherwise, the NP discards the packets.

Id. at 7:62-8:35.

42. Figure 2 and its accompanying description make clear that, for the '394 patent to function, each line card processing unit must distribute **the client requests** it receives to all other line cards in the router so that those other line cards can update their databases with the current instructions for how to handle incoming multicast packets.


43. If a line card processing unit receives multicast packets that relates to a client request that did not pass through the line card's ports, the only way it will know what to do with the packet is if the line card processing unit that received the client request relating to that packet (which the '394 patent calls the “ingress line card”) distributed that request to every line card in the network element. Without that distribution, the line card processing unit that received the multicast packets, but did not receive the client request, simply “discards the packets.” *Id.* at 8:35.

44. Without the requirement that each line card processing unit have a list of all multicast packet streams permitted for forwarding in the router, and the requirement that each line

card processing unit distribute any client request it receives to all the other line card processing units, the '394 patent line card processing units would not be synchronized and may have different instructions for handling multicast packets, resulting in the router processing multicast packets inconsistently.

c. Overview of Naik

45. Naik relates to an Edge Router R1 (blue box), which acts as a boundary, or edge, between the end users on a Local Area Network (LAN) and the internet. Ex. 1005 at [0004].



Id. at FIG. 1 (color added).

46. According to Naik, Edge Router R1 has three main components: (1) “a plurality of [line] cards 260” (green boxes below); (2) “a plurality of loosely coupled route processors 210” (red boxes below); and (3) “interconnection fabric 250.” *Id.* at [0040].





FIG. 2

Id. at FIG. 2 (color added). Naik also explains the role Line Cards 260 and Route Processors (RP) 210 play in Edge Router R1's handling of multicast packets.

47. Naik's Line Cards 260 "perform forwarding of the data traffic," and do not control, or determine the **routing** of, data traffic. *Id.* at [0046]. Each Line Card 260 has an Interface 270 that "has a plurality of ports 256." *Id.* "The interface 270 is coupled to a receive forwarding processor (FP Rx) and a transmit forwarding processor (FP Tx). The FP Rx renders a forwarding decision for each *packet* received at a port and determines where to forward the *packet*. To that end, the FP Rx makes use of a multicast forwarding information base (MFIB) 480." *Id.* at [0046] (emphasis added).



Id. at FIG. 2 (color added).

48. “In the event that the packet is to be forwarded to one of the router’s two route processors, the FP Rx [280] makes use of an internal forwarding information base (IFIB) to determine to which route processor the packet should be forwarded.” *Id.* at [0046]. The ACM 500 in Route Processor (RP) 210 monitors “incoming request[s] issued by the end nodes to join, leave or modify membership in multicast groups designated to receive data. The ACM accepts or denies these requests based on a configurable service policy.” *Id.* at [0053].




FIG. 2

Id. at FIG. 2 (color and annotation added).

49. If ACM 500 accepts the request, MRIB 238 is updated. *Id.* at [0045]. MRIB 238 distributes a derivative of the original client request by populating MFIB 480 and MFIB Forwarding Tables 400 of the line cards. *Id.* “The MFIB 480 on the Line Cards 260 is derived from the MRIB [238].” *Id.*




FIG. 2

Id. at FIG. 2 (color and annotation added).

50. Each Line Card 260 has its own distinct MFIB 480. For example, Naik describes a “plurality” of Line Cards 260 and “Multicast Forwarding Information Bases (MFIB’s) 1-N 480 executing on the line cards.” *Id.* at [0040] (“plurality”), [0045] (“1-N”). By numbering the MFIB 480 on each line card “1-N,” Naik makes clear that the MFIB 480 on each line card is different from the MFIBs 480 on other line cards. Likewise, Naik notes, “The MFIB 480 *is derived from the MRIB.*” *Id.* at [0045] (emphasis added). Thus, MFIBs 480 are not copies of MRIB 238 and have only a subset of the information contained in MRIB 238.

IX. THE CITED PRIOR ART DO NOT INVALIDATE THE CLAIMS UNDER REEXAMINATION

A. Naik’s Route Processors (RP) 210 are not “processing units”

51. The ’394 patent claim language sets out several requirements for the claimed “processing units.” For example, independent claim 11 of the ’394 patent states:

11. A network element, comprising:

multiple ports; and
multiple interconnected processing units, 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 the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,
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,
to receive from a client a request relating to reception of a given multicast packet stream, 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 to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists,
wherein two or more of the ports are aggregated to form a high-capacity port, and wherein an upper limit is set on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and wherein the processing units are configured to update the lists to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

Ex. 1001 at 13:48-14:16. Claim 11 identifies at least eight requirements of “processing units”: (1) assigned to a different subset of the multiple ports of the network element; (2) contain a processor for processing packets; (3) store respective lists, each of which indicates all of one or more multicast packet streams that are permitted for forwarding by the network element; (4) receive from a client a request relating to reception of a given multicast packet stream; (5) distribute the request among the processing units within the network element; (6) update the lists responsively to the distributed request and to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists; (7) selectively forward multicast packets in accordance with the updated lists; and (8) update the lists to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the

number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

52. Line Cards 36 are the only structures in the '394 patent that satisfy all these requirements. *See, e.g.*, FIG. 1, 6:26-28 (“each line card 36 is assigned to process the packets communicated over a given subset of the ports.”); 6:32-33 (“Each line card 36 comprises a processor, such as a Network Processor (NP) 48 ... which processes the packets communicated over the ports assigned to the line card.”); 6:35-38 (“Each line card 36 further comprises a filtering database 52, which holds a list of multicast streams that are permitted for forwarding by the network element”); 7:34-39 (“The request is sent from a given multicast client 28, is addressed to a certain multicast source and arrives at one or 35 ports 44 of the network element. The line card at which the request is received is referred to herein as the ingress line card, and the specific port in this line card over which the request arrives is referred to herein as the ingress port.”); 7:53-61 (“The NP in the ingress line card duplicates the request, at a duplication step 64. The first copy of the request is forwarded to the appropriate multicast sources, at a copy forwarding step 68. In addition, the second copy of the request is forwarded to all the NPs in the network element (including the originating NP), via fabric 40, at a copy distribution step 72. The second copy can be distributed, for example, by sending it in a multicast message addressed to all the NPs of the network element.”); 7:62-65 (“The NPs in the line cards of network element 32 receive the distributed copy of the request via fabric 40. Each NP receiving the distributed copy updates its respective database 52 accordingly, at an updating step 76.”); 8:60-9:4 (“In some embodiments, when the network element receives a request for a new multicast stream, the NP of the ingress line card checks whether the ingress port can support this additional stream. For example, the NP may query the port counter of the ingress port, which is stored in database 52 of the ingress card, and

check whether the current number of multicast streams permitted for forwarding over the ingress port has reached the maximum allowed number for that port. If the maximum number of streams has already been exhausted, the NP of the ingress card does not duplicate the request and does not distribute it among the NPs. As a result, the new multicast stream will be denied.”). Furthermore, the ’394 patent explicitly states that Line Cards 36 are “processing units.” 6:23-24 (“Network element 32 comprises multiple processing units, such as line cards 36.”).

53. Consistent with the ’394 patent’s claim language and written description, the Office Action cites Naik’s Line Cards 260 to satisfy the initial “processing units” limitation. Office Action at pp.7-9. However, Naik’s Line Cards 260 cannot satisfy all the “processing units” requirements of the ’394 patent claim language. The Office Action also cites Naik’s Route Processor (RP) 210 for the “processing units” limitations Naik’s Line Cards 260 cannot satisfy. The table below identifies when the Office Action relies on Line Cards 260 (cells shaded green) and when it relies on Route Processor (RP) 210 (cells shaded red) to satisfy the eight “processing unit” requirements of claim 11:

“Processing units” limitation	Naik evidence cited by Office Action
Assigned to a different subset of the multiple ports of the network element	Line Cards 260. Office Action at pp. 9 (claim 1), 75 (claim 11), 83 (claim 8), 99 (claim 18).
Contains a processor for processing packets	Line Cards 260 (specifically receive forwarding processor (FP Rx) and transmit forwarding processor (FP Tx)). Office Action at pp. 10-12 (claim 1), 75 (claim 11), 83 (claim 8), 99 (claim 18).
Store respective lists, each of which indicates all of one or more multicast packet streams that are permitted for forwarding by the network element	Line Cards 260 (specifically MFIB 480 and its MFIB Forwarding Table 400). Office Action at pp. 12-16 (claim 1), 75 (claim 11), 83 (claim 8), 99-100 (claim 18).
Receive from a client a request relating to reception of a given multicast packet stream	Route Processors (RP) 210 (specifically ACM 500). Office Action at pp. 18-19 (claim 1), 76 (claim 11), 84 (claim 8), 101 (claim 18).
Distribute the request among the processing units within the network element	Route Processors (RP) 210 (specifically MRIB 238). Office Action at pp. 19-27 (claim

	1), 75 (claim 11), 84 (claim 8), 101 (claim 18).
Update the lists responsively to the distributed request and to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists	Route Processors (RP) 210 (specifically MRIB 238). Office Action at pp. 27-28 (claim 1), 76 (claim 11), 84 (claim 8), 101 (claim 18).
selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists	Route Processors (RP) 210 (specifically ACM 500 and MRIB 238). Office Action at pp. 31-32 (claim 1), 76 (claim 11), 84 (claim 8), 101 (claim 18).
update the lists to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit	Route Processors (RP) 210 (specifically ACM 500 and MRIB 238). Office Action at pp. 51-58 (claim 1), 77 (claim 11).

54. However, as Figure 2 of Naik demonstrates, the Route Processors (RP) 210 (shaded in red) and the Line Cards 260 (shaded in green) are two distinct components.




FIG. 2

Ex. 1005 at FIG. 2 (color added). As Naik clarifies, “the router R1 comprises a plurality of loosely coupled route processors 210 connected to a plurality of interface cards 260 that input (Rx) and output (Tx) messages.” *Id.* at [0040]. Since Route Processors (RP) 210 are not assigned to any ports, much less different subsets of Naik’s ports, they cannot be characterized as “processing units” and cannot be relied upon to satisfy the ’394 patent’s “processing units” limitations.

B. The patentability of all reexamined claims should be confirmed because Naik does not disclose “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.”

55. Every claim of the ’394 patent requires “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.” Ex. 1001 at 12:3-6 (independent claim 1 and dependent claims), 13:6-9 (independent claim 8 and dependent claims), 13:59-61 (independent claim 11 and dependent claims), 14:65-67 (independent claim 18). Naik is the only prior art that the Office Action cites for this limitation. Office Action at pp. 12-16 (Ground 1 – claim 1 and dependent claims), 75 (Grounds 1 and 2 – claim 11 and dependent claims), 83 (Grounds 3 and 4 – claim 8 and dependent claims), 99-100 (Ground 3 – claim 18). Because Naik does not teach this limitation, the patentability of all reexamined claims should be confirmed.

56. The ’394 patent claims make clear that the list stored in each line card processing unit must be the same as the list stored in all other processing units in the network element insofar as every list must have all multicast packet streams that are permitted for forwarding by the network element. For example, each claim of the ’394 patent states that each line card must have a list that contains all multicast packet streams that can be forwarded by any line card in the network element. Ex. 1001 at 12:3-8 (claim 1 and dependents); 13:6-11 (claim 8 and dependents);

13:59-63 (claim 11 and dependents); 14:65-15:2 (claim 18). If the list in each respective line card must indicate all multicast packet streams permitted for forwarding by the network element, then all lists must identify the same set of multicast packet streams. These lists must include all multicast packet streams that are permitted for forwarding in order for the '394 patent to operate. As shown in FIG. 1 of the '394 patent, the system does not have a centralized processor to forward packets. Instead, each line card processing unit has its own network processor (NP) 48, which relies on the list stored in the database (DB) 52 of each line card processing unit to properly handle incoming packets. For the line card processing units to behave consistently, the list stored on each line card processing unit must be identical to the list stored on every other line card processing unit with respect to the identification of all multicast packet streams that can be forwarded by the network element.

57. Naik discloses the opposite of what is required by the '394 patent, namely that the list contained in each Line Card 260 is maintained differently, and that these lists are a subset of all multicast packet streams that can be forwarded by the network element. The closest Naik gets to disclosing the "list of all permitted multicast streams" stored on Line Cards 260 is MFIB 480, which contains a forwarding table 400.




FIG. 2

Ex. 1005 at FIG. 2 (color added).

58. Naik explains that Line Card 260 includes “receive forwarding processor (FP Rx) 280” and “transmit forwarding processor (FP Tx) 290,” which determine where to forward and transmit packets by looking up information in MFIB 480. *Id.* at [0046]. As a result, the MFIB 480 in each Line Card 260 identifies *some* multicast packet streams that are permitted for forwarding. However, Naik also makes clear that the MFIB 480 in each line card is different:

In the illustrative embodiment, the router operating system 232 implements multicast components as separate software process modules that are distributed over active route processors 210 and line cards 260 which, in turn, are individually hot-swappable and may have redundant (standby) backups. To that end, the multicast modules include *Multicast Forwarding Information Bases (MFIB's) 1-N 480* executing on the line cards and multicast components including the RIB 236, PIM 234, IGMPv3 235, ACM 500, and a MRIB 238 executing on the route processors 210.

Id. at [0045] (emphasis added). Thus, Line Card 1 has MFIB 480-1, Line Card 2 has MFIB 480-2, and so on until Line Card N, which has MFIB 480-N. If the MFIBs in each line card were the same, then it would not be necessary for Naik to separately number each MFIB.

59. Naik also makes clear that MRIB 238, which is shown as a component of Route Processor (RP) 210 in FIG. 2, collects routing information for multicast content.




FIG. 2

Ex. 1001 at FIG. 2 (color and annotation added).

60. Specifically, Naik explains:

Operationally, the router builds the MRIB 238 (and, more specifically, its multicast routing table) after exchanging routing information with the neighbors. For example, PIM provides MRIB with a chosen reverse path forwarding (RPF) address and interface for each route, which is used to accept packets for forwarding through the router. The primary role of the MRIB 238 is to facilitate communication between the various multicast components, i.e., coordinate distribution of state between the protocol applications (e.g., PIM, IGMPv3, etc.) and between these protocols and MFIB. As a result, MRIB provides the *routing information* and a multicast-capable path to each multicast destination. *The MFIB 480 is derived from the MRIB* and is embodied as one or more multicast forwarding tables 400 that are populated on the line cards 260 and whose contents describe how to *forward data packets* through the router.

Ex. 1005 at [0045] (emphasis added). If anything in Naik is a “list of all permitted multicast streams,” it is the MRIB 238. But this list is stored in Route Processor (RP) 210, not the Line Card 260. As a result, the MRIB 238 is not located in a “processing unit,” and, since this limitation

requires the processing unit store the list, MRIB 238 cannot teach the limitation. Because “the MFIB 480 is “derived from the MRIB” 238, it must have *different* information from MRIB 238. But if the MRIB 238 has all multicast packet streams that are permitted for forwarding, then the derived MFIBs 480 must list something less than all such streams.

61. The Office Action made four arguments that Naik satisfied this limitation. Office Action at pp. 12-16. First, the Office Action argued:

Each MFIB 480 via its forwarding table 400, shown in FIG. 4, indicates all of one or more multicast packets streams that are permitted for forwarding by the network element (edge router R1). Ex. 1005 ¶ [0045]; Ex. 1003 ¶ [0115].

Office Action at p. 13. Neither of the cited sources support this argument. Paragraph [0045] of Naik (Ex. 1005) discusses “MFIB 480” and “forwarding tables 400” but does not say that every MFIB 480 and forwarding table 400 must include *all* multicast packet streams that are permitted for forwarding by the network element. Likewise, Paragraph [0115] of the Wechselberger Declaration (Ex. 1003), does not say every MFIB 480 and forwarding table 400 must include *all* multicast packet streams that are permitted for forwarding by the network element.

62. Second, the Office Action noted that “route processors 210 essentially build MFIB 480” of the Line Cards 260 using the MRIB 238. Office Action at pp. 13-15. The Office Action also found:

Using information from the MRIB, the MFIBs are “populated” on each of line cards 260: “The MFIB 480 *is derived from the MRIB* and is embodied as one or more multicast forwarding tables 400 that are populated on the line cards 260 and whose contents describe how to forward data packets through the router.”

Id. at p. 15 (emphasis added). But this observation defeats the argument that Naik teaches this limitation. The MRIB 238 is the closest thing Naik gets to a list indicating all multicast packet streams that are permitted for forwarding by the network element. The MFIBs 480 are derivations of the MRIB 238. A derivation is not a copy, and, as a result, there is no teaching, implicit or


explicit, that the various MFIBs 480-1 through 480-N each include all the multicast packet streams that are permitted for forwarding by the network element. As a result, Naik cannot teach this limitation.

63. Third, Naik's identification of each MFIB by number teaches, if anything, that the MFIBs are different from each other. The Office Action, relying on the Request for Reexamination, argues that Naik FIG. 2 "indicates that each line card 260 contains an identical forwarding table 400 by indicating 'MFIB 480' as the same element in every line card 260, not for example as 'MFIB 480-1 ... MFIB 480-N.'" *Id.* But the request for reexamination implicitly concedes that if each MFIB 480 on each Line Card 260 was separately numbered (e.g., MFIB 480-1, MFIB 480-2, ..., MFIB 480-N), then that would prove the MFIBs were different. This is *exactly* what Naik does. As noted above, Naik's description of FIG. 2 clarifies, "Multicast Forwarding Information Bases (MFIB's) 1-N 480 executing on the line cards." Ex. 1005 at [0045] (emphasis added). Thus, Line Card 1 has MFIB 480-1, Line Card 2 has MFIB 480-2, and so on until Line Card N, which has MFIB 480-N. As a result, the corollary to the reasoning in the Office Action and Request for Reexamination is that Naik's MFIBs are not all identical, cannot all contain every multicast packet stream that is permitted for forwarding, and, therefore, cannot teach this limitation.

64. Finally, the Office Action argues that FIG. 4 of Naik teaches a "list indicating all of one or more multicast packet streams that are permitted for forwarding by the network element":

Specifically, each entry 410 in the table of FIG. 4 represents a data packet. *Id.* ¶ [0050]. Moreover, the fields within each entry 410 indicate information relevant to the forwarding of a packet by the network element (edge router R1). *Id.* For instance, the illustrated fields within MFIB Forwarding Table 400 include a multicast group address 412, an ingress interface 42 [sic: 420], and an egress interface 430, indicating permission for forwarding the packet. *Id.* ¶¶ [0050]-[0051]. The group address field 412 contains "a multicast group network (IP) address," and egress interface 430 field contains "a list of output (forwarding) interfaces over which the incoming packet should be forwarded." *Id.* ¶ [0051].

Office Action at pp. 15-16. The Office Action concludes that because the rows in FIG. 4 are “one or more multicast packet streams that are permitted for forwarding by the network element,” and because each row includes “all of” the information for each multicast packet stream, FIG. 4 satisfies the “all of one or more multicast packet streams that are permitted for forwarding by the network element” limitation. *Id.* at p. 16.



Id. In other words, the Office Action argues that if you have a list of “one or more packet streams that are permitted for forwarding by the network element,” and for each packet stream the list has “all of” the data fields for each packet streams, you satisfy the ’394 patent’s requirement that the processing unit store “all of one or more packet streams that are permitted for forwarding by the network element.” This is different from how the ’394 patent works, as reflected by its claims. The claim phrase “all of” requires that the list stored in each processing unit include every multicast packet stream that can be forwarded by the network element, which is necessary for a network element to function without a central processor to control packet forwarding. For Naik’s FIG. 4 to disclose this limitation, there would need to be some indication that the Entries 410 in Table 400

included all multicast packet streams that were permitted for forwarding in the Edge Router R1.

No such disclosure exists. As a result, FIG. 4 does not teach this limitation.

C. The patentability of reexamined claims 11-18 should be confirmed because Naik's line card processing units do not "distribute the request among the processing units within the network element."


65. Claims 11 through 18 of the '394 patent require line card processing units that "distribute the request among the processing units within the network element." Ex. 1001 at 13:66-67 (independent claim 11 and dependent claims), 15:5-6 (independent claim 18). Naik is the only prior art that the Office Action cites for this limitation. Office Action at pp. 19-27 and 76 (Ground 1 – claim 11 and dependent claims), 19-27 and 101 (Ground 3 – claim 18). Because Naik does not teach this limitation, the patentability of reexamined claims 11-18 should be confirmed.

66. As discussed in section IV.B ("Overview of the '394 Patent"), the '394 patent makes clear that the line card processing units must "distribut[e] the request among the multiple processing units within the network element." The '394 patent text accompanying Figure 2 states the line card that receives a client request (which it labels the "ingress line card") duplicates the request during duplication step 64, and distributes the copies as follows:

The [Network Processor] in the ingress line card duplicates the request, at a duplication step 64. The first copy of the request is forwarded to the appropriate multicast sources, at a copy forwarding step 68. In addition, *the second copy of the request is forwarded to all the [Network Processors] in the network element (including the originating [Network Processor]), via fabric 40, at a copy distribution step 72.*

Ex. 1001 at 7:53-54 (emphasis added).

FIG. 2



Id. at FIG. 2. This “copy distribution step 72,” where line card processing units forward a copy of the client request to all other line card processing units in the router so every line card can update the list in their respective databases, is covered by the “distribute the request among the multiple processing units within the network element” limitation in claims 11-18.

67. The '394 patent claim language also makes clear that the line card processing unit receiving the client request must distribute copies of the request to every other line card processing unit in the network element. For example, by indenting the “distribute the request” limitation so that it is organized under the “processing units” limitation, claim 11 makes clear that the line card processing unit will “distribute the request among the processing units within the network element.”

11. A network element, comprising:
 multiple ports; and
 multiple interconnected **processing units**, 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 the processing units are configured to process packets that are communicated over respective subsets of one or more of the ports,
 ...
 to **receive from a client a request** relating to reception of a given multicast packet stream, 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 to selectively forward multicast packets associated with the given multicast packet stream in accordance with the updated lists, ...

Id. at 13:46-14:4 (emphasis and annotation added); *see also id.* at 14:54-15:4.

68. Furthermore, the '394 patent claim language makes clear that the line card processing unit must distribute the request to all processing units within the network element. For example, claim 8 states:

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*

Id. at 13:14-17 (emphasis added); *see also* 12:11-14 (claim 1 requiring “updating the respective list in each processing unit responsively to the request”). It would not be possible to update “the respective list in each processing unit responsively to the request” unless the request has been distributed to every line card processing unit in the router.

69. Naik does not disclose line card processing units that distribute requests received from clients to all other line card processing units in the network element, and, as a result, cannot teach this limitation. The Office Action’s discussion of the “distribute the request among the processing units within the network element” limitation for claims 11 and 18 both refer back to the arguments the Office Action made for the corresponding limitation in claim 1. Office Action at pp. 76 (claim 11 analysis referring back to [1.4] and [1.4.1]), 101 (claim 18 analysis referring back to [1.4] and [1.4.1]). Sections [1.4] and [1.4.1] present three arguments why Naik discloses this limitation. None of these arguments are valid.

70. First, the Office Action argues “the phrase ‘distributing the request’ of element [1.4.1] at least includes a router forwarding a multicast request to the source of the multicast packet stream within the network.” Office Action at p. 19. The reasoning for this conclusion is that claim 3, which depends from claim 1, specifies “wherein distributing the request comprises ... forwarding [a] first copy [of the request] to a source of the given multicast packet stream.” *Id.* at pp. 19-21. The Office Action then cites *Alcon Research Ltd. v. Apotex Inc.* for the proposition that “a dependent claim cannot be broader than the claim from which it depends.” *Id.* at p. 21 (citing 687 F.3d 1362, 1367 (Fed. Cir. 2012)). Claim 3 actually states:

3. The method according to claim 1, wherein distributing the request comprises duplicating the request to produce first and second copies, forwarding the first copy to a source of the given multicast packet stream, and distributing the second copy among the multiple processing units.

Ex. 1001 at 12:36-40. If you tried to prove infringement of claim 3 but only established the accused process forwarded a first copy of a client request to a source of the given multicast packet stream, you would clearly fail. Likewise, attempting to invalidate claim 3 by only pointing to prior art disclosure of a first copy of a client request being forwarded to a source of the given multicast packet stream will also fail. This not only is true for claim 3, but also for claim 1, which, as the Office Action points out, must be broader than claim 3. In any event, the language of claims 1 and 3 make clear that “distributing the request” must include sending the request to all processing units in the router. Ex. 1001 at 12:11-14 (“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”), 12:36-40 (“The method according to claim 1, wherein distributing the request comprises duplicating the request to produce first and second copies, forwarding the first copy to a source of the given multicast packet stream, and distributing the second copy among the multiple processing units.”). As discussed in Section IV.B above, the distribution of the client request must

be to all line card processing units in the router because that is how the line card processing units update their lists of all multicast streams permitted for forwarding by the router, which, in turn, is necessary for the '394 patent to function. Sending a copy of the request to a source of the given multicast packet stream will not distribute the request to all processing units. At most, the request will pass through one processing unit on its way to the source of the given multicast packet stream.

71. Second, the Office Action argues that Naik discloses “distributing the request” through edge router R1 communicating the request “to its neighboring routers (neighbors) on the link toward the rendezvous point (for a shared tree) or source for a SPT.” Office Action at pp. 21-23. The Office Action goes on to argue that “to be forwarded from edge router R1 to ‘neighboring routers (neighbors) on the link,’ as taught by Naik, the forwarded multicast request must be distributed ‘among the multiple processing units,’ i.e., among the *one or more line cards 260*, in order to exit router R1.” *Id.* at 23 (emphasis added). In other words, the Office Action’s second argument again relies on Naik sending a first copy of the client request to the source of the given multicast packet stream. By the Office Action’s logic, sending the request to the source necessarily means the request must pass through one of edge router R1’s ports, which in turn means it must pass through, and update, one Line Card 260. But both the '394 patent claim language and written description require distributing the request to *all* line card processing units, not just one. This is so, because every line card processing unit must update its list with each client request in order for the line card processing units to behave consistently. In light of this, Naik’s disclosure of sending the client request through one port, and one Line Card 260, so that it can travel on the path to the source of the multicast packet stream, does not disclose this limitation.

72. Finally, the Office Action argues that Naik discloses “distributing the request” when “((i) multicast request arrives, (ii) MFIB executing on line cards 260 notifies route processors

210 of arrival, (iii) Admission Control Module (ACM) 500 executing on route processors 210 snoops information from the multicast request, and (iv) Multicast Route Information Base (MRIB) executing on route processors 210 populates the MFIB 480 and MFIB Forwarding Tables 400 based on the multicast request.)” Office Action at pp. 23-27. The Office Action’s final argument for this limitation relies on the Route Processor (RP) 210’s ACM 500 to accept the client request, the Route Processor (RP) 210 to update the MRIB 238, and the MRIB 238 to populate the MFIB 480 and MFIB Forwarding Table 400 on a Line Card 260. This argument fails for several reasons.

73. First, a line card processing unit must “distribute the request among the multiple processing units.” As explained in Section V.A, Naik’s Route Processor (RP) 210 (red boxes below) is separate from Line Cards 260 (green boxes below), and cannot be a line card processing unit because, among other things, it is not assigned to a subset of Naik’s Ports 256.




FIG. 2

Ex. 1005 at FIG. 2 (color added). As a result, Naik’s Route Processors (RP) 210 cannot be relied upon to satisfy this limitation.

74. Second, even if Naik's Route Processors (RP) 210 could be used to disclose this limitation, whatever the MRIB 238 passes back to the MFIB 480 and MFIB Forwarding Table 400 in Line Card 260 is not the original client request. The Route Processor (RB) 210 processes the client request to populate MRIB 238. That is one modification of the client request. Subsequently, "the MFIB 480 is *derived* from the MRIB." Ex. 1005 at [0045] (emphasis added). A derivation is not a copy, making the MFIB 480 contents a second modification of the original client request. Whatever is sent back to Line Card 260 from Route Processor (RP) 210 is twice removed from the original client request.

75. Third, even if Naik's Route Processors (RP) 210 could be used to disclose this limitation, and even if what they sent back to Line Cards 260 was the client request, they do not distribute the request to *all* line card processing units in Naik's edge router R1. In explaining the steps of Figure 7, Naik states, "The ACM snoops the IGMPv3 packets received and if a join message [multicast request] is detected at 710, the ACM ... checks the policy map at 712. If an action is found, it is automatically applied *to the requesting port 714.*" Ex. 1005 at [0073] (emphasis added).




FIG. 7

Id. at FIG. 7 (color and annotation added). As Step 712 in Figure 7 indicates, the ACM 500 response is only with respect to the “edge router port that received the join,” not all ports and all line card processing units. In other words, if the Router Processor (RP) 210 updates the MRIB 238, which in turn updates the MFIB 480, the update only goes to the Line Card 260 containing the port that received the client request. As a result, the request is not distributed to all line card processing units, as required by the '394 patent.

- D. The patentability of reexamined claims 8-10 and 18 should be confirmed because Naik's hot-swapping does not initialize new route processors or line cards.**

76. Claims 8-10 and 18 of the '394 patent require initializing a new processing unit in a running network element. Ex. 1001 at 13:21-33 (independent claim 8 and dependent claims), 15:10-16:10 (independent claim 18). Naik is the only prior art that the Office Action cites for this

limitation. Office Action at pp. 84-86 (Grounds 3 and 4 – claim 8 and dependent claims), 101 (Ground 3 – claim 18). Because Naik does not teach this limitation, the patentability of reexamined claims 8-10 and 18 should be confirmed.

77. The '394 patent makes clear that the claimed “initializing a new processing unit” is performed after the router is operating. For example, the claim language states that the first step in initializing a new processing unit is “deactivating the ports to which the new processing unit is assigned.” Ex. 1001 at 13:24-25 (claim 8), 16:1-2 (claim 18). If the router was not operational when the new line card processing unit was added, there would be no need to deactivate the ports of the new processing unit. Likewise, the '394 patent describes activating a new line card processing unit that is “inserted into an active network element”:

When a new line card is inserted into an active network element, the NP of the new line card runs a process that updates its filtering database. Typically, the NP of the new line card disables its local ports 44 for a certain period of time (e.g., by setting the ADMIN_STATUS of these ports to 45 DOWN), during which a multicast status update request (e.g., an IGMP general query) is transmitted through all the ports of the network element. The multicast content clients respond to the status update request, thus allowing the newly-added NP to update its local filtering database, port counters and aging 50 counters. When the database and counters are up-to-date (or after a certain time out), the NP activates the ports of the newly-inserted line card, and may start normal operation.

Ex. 1001 at 9:41-53 (emphasis added).

78. The only argument the Office Action makes that Naik discloses “initializing a new processing unit” is that Naik’s Route Processors (RP) 210 and Line Cards 260 can hot-swapped with “redundant (standby) backups.” Office Action at pp. 85-86. The Office Action argues, “Because hot-swapping or bringing online a redundant (standby) backup for one of route processors 210 or line cards 260 would require initializing that new processing unit into the active operating environment, Naik discloses or suggests that operating the network element may involve the claimed ‘initializing a new processing unit added to the network element.’” *Id.*

79. Hot-swapping an active Route Processor (RP) 210 or Line Card 260 with a “redundant (standby) backup” does not require initializing the redundant (standby) backup. In fact, the only way you can hot-swap a standby for an active element is if the standby was already initialized. Hot-swappable backups do not go through initialization, since they are already actively running software. As a result, Naik’s hot-swapping would not involve initializing a new line card processing unit. Second, the ’394 patent claim language makes clear what it means to initialize a new line card processing unit, namely deactivating the ports of the new line card processing unit, sending a multicast status request, receiving multicast membership reports, constructing a list of all multicast packet streams permitted for forwarding, and, after the list is constructed, activating the ports of the new line card processing unit. Ex. 1001 at 13:24-33, 16:1-10. In other words, initializing a new line card processing unit means giving it the instructions it needs to handle incoming multicast packets consistently with the pre-existing line card processing units in the network element. Naik disclosure does not describe doing any of these steps as part of hot-swapping active Route Processors (RP) 210 or Line Cards 260 with redundant or standby backups. In fact, Naik’s redundant or standby backup Route Processors (RP) 210 and Line Cards 260 would already have these lists before any hot-swapping occurred.

E. The patentability of reexamined claims 8-10 and 18 should be confirmed because Ramanathan does not deactivate ports during activation of a new line card processing unit.

80. Claims 8-10 and 18 of the ’394 patent require “deactivating the ports to which the new processing unit is assigned” while initializing the new line card processing unit. Ex. 1001 at 13:24-25 (independent claim 8 and dependent claims), 16:1-2 (independent claim 18). Ramanathan is the only prior art that the Office Action cited for this limitation. Office Action at pp. 86-89 (Grounds 3 and 4 – claim 8 and dependent claims), 101 (Ground 3 – claim 18). The Office Action

argues that during Ramanathan's core reset and reboot of a Linecard Processor (LP) 204, "when the LP application is loaded, it will populate its software structures while blocking the access to the HW components during LP initialization, module configuration, and port configuration." *Id.* at p. 87 (citing Ex. 1010 at [0081]). The Office Action concludes:

Ramanathan explains the blocking of hardware access, or deactivating the ports of an LP, during initializing from an LP upgrade or reset in other contexts as well. *See id.* ¶ [0090] ("Before the LP upgrade begins, the hardware MAC table information is stored into a protected portion of memory Until this process is completed, all hardware accesses may be disabled.") ¶ [0092] ("After an LP reset, the VLAN hardware accesses are disabled by the underlying system software. Once the LP application is up, hardware accesses are allowed."). For at least these reasons, Naik and Ramanathan discloses or suggest "deactivating the ports to which the new processing unit is assigned," as recited in element [8.6.1]. Ex. 1003 ¶ [0322].

Id. at pp. 87-88.

81. The Office Action's argument falsely equates blocking or disabling *access* to a port with deactivating it. While the Ramanathan paragraphs cited by the Office Action use words like "disabled," what is disabled is access to the hardware, e.g. the protected memory, not the hardware itself. *See, e.g.,* Ex. 1010 at [0090] ("Before the LP upgrade begins, the hardware MAC table information is stored into a protected portion of memory Until this process is completed, all hardware *accesses* may be disabled."); [0092] ("After an LP reset, the VLAN hardware *accesses* are disabled by the underlying system software. Once the LP application is up, hardware accesses are allowed.").

82. Blocking hardware access to a portion of memory is not the same as deactivating the port. For example, one can block access to a portion of memory in a port while the port remains active. And, in fact, that is exactly what Ramanathan does when it states, "The MP will *not* send card down and port down events when LPs perform a core reset." Ex. 1010 at [0082] (emphasis added). By not sending a port down event during a core reset, Ramanathan actually teaches the

opposite of what the '394 patent claims, namely that the ports of a new processing unit will remain active (although access is blocked to a portion of memory) during initializing. This is done to preserve the MAC table during initialization. Because Ramanathan does not teach deactivating ports during activation of a new line card processing unit, the patentability of reexamined claims 8-10 and 18 should be confirmed.

F. The patentability of reexamined claims 8-10 and 18 should be confirmed because there is no motivation to combine Naik with Ramanathan and Schulter for the “while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request” limitation.

83. As part of the initialization of a new line card processing unit, claims 8-10 and 18 of the '394 patent require “while the ports to which the new processing unit is assigned are deactivated, sending a multicast status request.” Ex. 1001 at 13:26-27 (independent claim 8 and dependent claims), 16:3-4 (independent claim 18). The Office Action relies on Naik in combination with Schulter and Ramanathan to cover this limitation. Office Action at pp. 89-93 (Grounds 3 and 4 – claim 8 and dependent claims), 101 (Ground 3 – claim 18).

84. As the '394 patent claim language makes clear, this limitation requires that during the activation of a new line card processing unit, while the ports are deactivated, the new line card processing unit sends a multicast status request. Just like the previous limitation (Section V.E of this response), the Office Action relies on Ramanathan's disclosure of blocking or disabling access to hardware to meet this limitation's requirement that the ports be deactivated. For the same reasons discussed in Section V.E, Ramanathan does not disclose this limitation. But beyond not disclosing this limitation, Ramanathan actually teaches away from it. For example, Ramanathan states:

Before the LP upgrade begins, the hardware MAC table information is stored into a protected portion of memory. After the LP core reset, the MP synchronizes the software MAC table (exactly in the same way that the synchronization takes place on a normal reset).

After this is completed, the MAC entry to hardware index association is recreated. *Until this process is completed, all hardware accesses may be disabled.*

Ex. 1010 at [0090] (emphasis added).

85. Ramanathan also teaches:

VLAN configuration is synchronized to the LP on card up event. After an LP reset, *the VLAN hardware accesses are disabled by the underlying system software*. Once the LP application is up, hardware accesses are allowed.

Id. at [0092] (emphasis added).

86. If “all hardware accesses may be disabled” and “VLAN hardware accesses are disabled” during LP upgrades and resets, then Ramanathan does not have components that can send a multicast status request during an LP upgrade or reset, as required by claims 8-10 and 18 of the '394 patent. As a result, Ramanathan teaches away from the '394 patent's requirement that a multicast status request is sent while the ports to the new processing unit are deactivated, and the Office Action's motivation to combine does not exist.

X. CONCLUSION

87. This declaration reflects my opinions as currently held. However, my analysis may continue such that I may consider additional information that may result in additional observations.

88. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

March 12, 2025

Date


James Olivier, Ph.D.