

United States Patent and Trademark Office
Before the Patent Trial and Appeal Board

JUNIPER NETWORKS, INC.

(Petitioner)

V.

PORTSMOUTH NETWORK CORP.

(Patent Owner)

Inter Partes Review

Nos. IPR2024-00913 and IPR2024-00989 | **U.S. Patent No. 8,014,394**

Petitioner's Demonstrative Exhibits


Juniper Exhibit 1032
Juniper v. Portsmouth
IPR2024-00913

Petition Grounds

GROUND	REFERENCE(S)	CLAIM(S)	BASIS
1.1	Blease in view of Weyman and Hu	1-2, 4, 6-7, 11-12, 14, 17	§103
1.2	Blease in view of Weyman, Hu, and Deng	16	§103
2.1	Blease in view of Weyman	8, 18	§103
2.2	Blease in view of Weyman and Hu	9	§103
2.3	Blease in view of Weyman, Hu, and Deng	10	§103
2.4	Blease in view of Weyman and Rao	8, 18	§103
2.5	Blease in view of Weyman, Rao and Hu	9	§103
2.6	Blease in view of Weyman, Rao and Deng	10	§103

'394 Patent — Distributed Routing Architecture

1. A method for communication, comprising:
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 aggregating two or more of the ports to form a high-capacity port, and
wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.



Blease Teaches a Distributed Routing Architecture Embodiment

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.




FIG. 4

In one embodiment, for example, control unit 44 and its corresponding functionality may be distributed within IFCs 54.

Weyman — Distributed Routing Architecture

(19) **United States**

(12) **Patent Application Publication**
Weyman et al.

(10) **Pub. No.: US 2005/0041665 A1**

(43) **Pub. Date: Feb. 24, 2005**

(54) **SYSTEM AND METHOD FOR DISTRIBUTED
MULTICAST ROUTING**

(21) **Appl. No.: 10/644,155**

(22) **Filed: Aug. 20, 2003**

(75) **Inventors: Raphael J. Weyman, Hemel**
Hempstead (GB); **Francisco Aquino,**
Marlborough, MA (US); **Maurice A.**
Goodfellow, Dunstable (GB)

Publication Classification

(51) **Int. Cl.⁷ H04L 12/56**


(52) **U.S. Cl. 370/390**

Correspondence Address:
MICHAELSON AND WALLACE
PARKWAY 109 OFFICE CENTER
328 NEWMAN SPRINGS RD
P O BOX 8489
RED BANK, NJ 07701

ABSTRACT

A distributed router composed of individual units connected by links in a routing fabric which appears to be a single network entity routes multicast packets by bridging to the final unit in the fabric. At the final unit the packet is replicated as necessary and routed. The scheme avoids multiple transmission of the same packet across the fabric.

(73) **Assignee: 3Com Corporation**



Distributed Routing is One of Three Basic Architectures

a. Centralized, Decentralized, and Distributed Router Architectures

32. There are three basic router architectures: (1) centralized, (2) decentralized, and (3) distributed.

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.

March 12, 2025

Date

James Olivier Ph.D.
James Olivier, Ph.D.

Declaration of James Olivier, Ph.D. in 90/019,441 (EX1027), ¶¶32, 34, 88

Nowhere in Blease is it described how this functionality may be distributed or what is meant by “distributed.”

May 9th, 2025

Date

James Olivier, Ph.D.

Declaration of James Olivier, Ph.D. (EX2005), ¶¶50, 117

'394 Patent — Setting an Aggregation Limit

1. A method for communication, comprising:
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 aggregating two or more of the ports to form a high-capacity port, and wherein aggregating the ports comprises setting an upper limit on a number of the multicast packet streams that are permitted for forwarding over the high-capacity port, wherein the request indicates that the given multicast packet stream is to be forwarded over the high-capacity port, and wherein updating the list comprises modifying the list to indicate that the given multicast packet stream is permitted for forwarding over the high-capacity port only when the number of the multicast packet streams that are permitted for forwarding over the high-capacity port, including the given multicast packet stream, does not exceed the upper limit.

For example, ports can be aggregated to form a Link Aggregation Group (LAG), as specified in the IEEE 802.1AX standard, entitled "IEEE Standard for Local and Metropolitan Area Networks—Link Aggregation," Nov. 3, 2008,

'394 patent (EX1001), 6:64-67

1017	SysKconnect GmbH, Link Aggregation according to IEEE Standard 802.3ad, White Paper, v1.10 (Oct. 10, 2002)
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	Jeffrey, "Link Aggregation Control Protocol," Presentation to the Link Aggregation Study Group (Apr. 1998)
1023	Jeffrey, 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)

Declaration of Tal Lavian (EX1003), ¶¶34, 38, 40-41

Blease and Hu are directed to similar subject matter, processing multicast stream requests, and Hu's link aggregation addresses the problem of potentially overloading a port when doing so.

Institution Decision (Paper 8) at 37

Hu Teaches LAG Limits

Hu teaches aggregating a plurality of ports into a single logical port, as a link aggregation group (LAG), to form a high-bandwidth port. Hu, 1:30-39, 3:10-32. High-bandwidth ports increase the utility of the ports because they are able to carry more streams and/or higher bitrate streams. Additionally, through the aggregation of ports, packets may be shared or balanced among the ports, “increas[ing] overall utilization efficiency of network components and quality of service to end-users.” Hu, 1:13-15, 1:30-35; *see also* Chen, 5:55-58. Lavian, ¶93.

Hu further teaches using pre-defined thresholds to limit the number of streams using a particular LAG. For example, if video traffic loads at all of the ports of a LAG would exceed a pre-defined threshold, then a request for a video stream may be rejected. Hu, 5:13-22, 6:39-50. Hu teaches that limits on the number of streams on a LAG are important for video streams because if “insufficient bandwidth [is allocated] to video streams, the quality of service degrades.” Hu, 1:20-29. Moreover, fluctuations from variable video bitrates exacerbate the issue, further highlighting the need for limits. Hu, 1:16-19, 1:58-62. Thus, by ensuring that a limit to the number of connections to a LAG are not exceeded, and by rejecting subscriber requests if such limits would be exceeded, quality of service is maintained. Lavian, ¶¶93-94.

8. A method for communication,

...

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.

Weyman Adds New Line Cards

Furthermore, Weyman describes adding new cards to the distributed router, stating “when a unit with a default configuration joins the fabric, it needs to update its configuration to that of the distributed router.” -989 Ex. 1006 ¶ 69. Weyman continues, “[i]f a unit with a non-default configuration joins the fabric then there should be a selection mechanism which determines the distributed router’s configuration.” *Id.* As Petitioner points out, Weyman does not describe in detail how this configuration would be accomplished, so a POSITA would look to Rao, which does. -989 Pet. 64–65; *see also* -989 Ex. 1003 ¶¶ 169, 175.

Institution Decision (Paper 8) at 50-51

³ A POSITA would understand that the term ‘non-default’ was a typo and should have been ‘default’ given the prior teaching of Weyman [0069].

Declaration of James Olivier, Ph.D. (EX2005), ¶60, fn. 3

Q. So then is the use of non-default in paragraph 156 also a typo?

A. Yes.

Deposition of James Olivier, Ph.D. (Jul. 17, 2025) (EX1026), 26:11-13

Furthermore, Weyman describes adding new cards to the distributed router, stating “when a unit with a default configuration joins the fabric, it needs to update its configuration to that of the distributed router.” -989 Ex. 1006 ¶ 69. Weyman continues, “[i]f a unit with a non-default configuration joins the fabric then there should be a selection mechanism which determines the distributed router’s configuration.” *Id.* As Petitioner points out, Weyman does not describe in detail how this configuration would be accomplished, so a POSITA would look to Rao, which does. -989 Pet. 64–65; *see also* -989 Ex. 1003 ¶¶ 169, 175.

Institution Decision (Paper 8) at 50-51