

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

JUNIPER NETWORKS, INC.,
Petitioner,

v.

SMART PATH CONNECTIONS, LLC,
Patent Owner.

IPR2021-01356
Patent 7,697,525 B2

Before JAMESON LEE, DAVID C. McKONE, and
STEVEN M. AMUNDSON, *Administrative Patent Judges*.

LEE, *Administrative Patent Judge*.

JUDGMENT
Final Written Decision
Determining No Challenged Claim Unpatentable
35 U.S.C. § 318(a)

I. INTRODUCTION

We instituted an *inter partes* review of claims 1–24 of U.S. Patent No. 7,697,525 B2 (Ex. 1001, “the ’525 patent”) owned by Smart Path Connections, LLC (“Patent Owner”). Paper 10 (“Decision to Institute” or “Inst. Dec.”).

We have authority to conduct this *inter partes* review under 35 U.S.C. § 6. This Final Written Decision is issued pursuant to 35 U.S.C. § 318(a) and 37 C.F.R. § 42.73. We determine that Juniper Networks, Inc. (“Petitioner”) has not proved by a preponderance of the evidence that any of claims 1–24 of the ’525 patent is unpatentable.

A. Background

Petitioner filed a Petition (Paper 1 (“Pet.”)) requesting *inter partes* review of claims 1–24 of the ’525 patent. Patent Owner filed a Preliminary Response. Paper 6. Petitioner filed a Preliminary Reply. Paper 8. Patent Owner filed a Preliminary Sur-Reply. Paper 9. The Decision to Institute was entered on February 1, 2022. Paper 10. Thereafter, Patent Owner filed a Patent Owner Response. Paper 14 (“PO Resp.”). Petitioner filed a Reply. Paper 17 (“Reply”). Patent Owner filed a Sur-Reply. Paper 19 (“Sur-Reply”). Oral hearing was held on November 1, 2022. A copy of the hearing transcript has been entered into the record as Paper 26 (“Tr.”).

B. Real Parties in Interest

Petitioner states that its real party-in-interest is Juniper Networks, Inc. Pet. xi. Patent Owner states that its real-party-in-interest is Smart Path Connections, LLC. Paper 4, 1.

C. Related Matters

Petitioner identifies the following as related litigation involving the ’525 patent: *Smart Path Connections, LLC v. Juniper Networks, Inc.*, No. 6:21-cv-327 (W.D. Tex.). Pet. xi. Patent Owner identifies the following as related litigation involving the ’525 patent: *Smart Path Connections, LLC v. Juniper Networks, Inc.*, No. 6:21-cv-328 (W.D. Tex.). Paper 4, 1. It appears that the two referenced district court cases are the same case and the

correct case number is: 6:21-cv-00328 (W.D. Tex.). Both parties identify IPR2021-01170 as an *inter partes* review involving a related patent. Pet. xi; Paper 4, 1.

D. The '525 Patent (Ex. 1001)

The '525 patent relates to methods and systems for forwarding multicast traffic in communication networks. Ex. 1001, 1:6–8. The '525 patent discloses use of a link aggregation (“LAG”) group. *Id.* at code (57).

A LAG is defined as follows:

Link aggregation (LAG) is a technique by which a group of parallel links between two endpoints in a data network can be joined together into a single logical link (referred to as a “LAG group”). Traffic transmitted between the endpoints is distributed among the physical links in a manner that is transparent to the clients that send and receive the traffic.

Id. at 1:25–30. The '525 patent describes a network node including a plurality of ports at least a subset of which is grouped in a LAG group.

Id. at 1:40–43. The network node also includes packet processing logic, coupled to receive data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of the ports in the LAG group and to forward only a single copy of each of the data packets via the output ports in the LAG group, while distributing the copies among the ports in the LAG group to balance the traffic load in the LAG group. *Id.* at 1:44–52. The network node includes line cards that are connected to the ports and a switching fabric that interconnects the line cards. *Id.* at 1:54–56.

The packet processing logic allocates to each of the received packets a fabric multicast identification (“FMID”) value selected from a range of possible FMID values. *Id.* at 2:7–10. Each FMID value is associated with

one of the ports in the LAG, and the packet processing logic forwards a single copy of the data packet to the port associated with the allocated FMID value for that packet. *Id.* at 2:7–12. In one embodiment, a line card allocates a line card FMID (“LC-FMID”) to each received packet, and then the switching fabric maps the LC-FMID values to respective central FMID (C-FMID) values selected from a second range of possible C-FMID values that is smaller than the first range of values for the LC-FMID value. *Id.* at 2:41–47. Thereafter, the switching fabric forwards the data packets responsively to the C-FMID values. *Id.* at 2:48–49.

Figure 1 of the ’525 patent (reproduced below) shows a network node:

FIG. 1




Figure 1 is a block diagram illustrating a network node. *Id.* at 3:7–9.

As shown in Figure 1, node 20 accepts multicast packets (having multiple destination addresses) from multicast source 24 and forwards the packets to two CTPs (connection termination ports) 32A and 32B.

Id. at 2:44–46. The ’525 patent describes:

CTP 32A comprises three physical ports 28A . . . 28C, which are configured as a LAG group. CTP 32B comprises two physical

ports 28D and 28E, which are configured as another LAG group. The terms CTP and LAG group are sometimes used interchangeably below. Node 20 may comprise additional CTPs (not shown in the figure), which comprise a single physical port. Such single ports are often regarded and managed by the node as LAG groups having a single member.

Id. at 2:47–54 (emphasis added).

“Node 20 comprises line cards, which process the data packets entering and leaving the node.” *Id.* at 2:64–65. Referring to Figure 1, the ’525 patent describes:

[N]ode 20 comprises five line cards denoted 36A . . . 36E. Line card 36A, which is referred to hereinbelow as the ingress card, accepts multicast packets from multicast source 24. Line cards 36B . . . 36E, which are referred to hereinbelow as the egress cards, are connected to the different output ports. Note that output ports belonging to a particular LAG group may be served by different line cards.

Id. at 5:1–8. Switching fabric 40 accepts packets from the ingress cards and routes them to the appropriate egress cards. *Id.* at 5:9–11. Node 20 forwards only a single copy of a multicast packet to each LAG group, and distributes the forwarding approximately evenly among the different ports within a LAG group. *Id.* at 5:37–50.

Each multicast source 24 is assigned a range of N virtual addresses, e.g., a range of FMIDs, and the ingress line card assigns one of the N FMIDs to each multicast packet arriving from the multicast source while attempting to distribute the packets approximately evenly among the different FMIDs in the range. *Id.* at 5:63–6:4. “The multicast packets are provided to fabric 40 along with their respective FMIDs. Both fabric 40 and the egress cards are pre-configured to forward each multicast packet to the appropriate output port based on its FMID.” *Id.* at 6:19–22. Fabric 40 may include a forwarding table, which maps FMIDs to line cards, and each line card

includes another forwarding table, which maps FMIDs to output ports. *Id.* at 6:22–27. “Thus, each multicast packet is assigned an FMID by the ingress card, and then forwarded to a particular output port by the fabric and egress card, based on its FMID.” *Id.* at 6:27–30.

For some embodiments, the ’525 patent describes:

In some embodiments, the total number of allocable FMIDs is lower in the switching fabric than in the line cards. In other words, fabric 40 may have an FMID space having a smaller number of bits than the FMID space of the line cards. For example, the line cards may support 13-bit FMIDs, referred to hereinbelow as line card FMIDs (LC-FMIDs), while the fabric may only support 11-bit FMIDs, referred to as central FMIDs (C-FMIDs). In these cases, multiple LC-FMIDs correspond to a particular C-FMID. Depending on the fabric definition, the C-FMID may comprise, for example, the least significant bits (LSB) or the most significant bits (MSB) of the LC-FMID.

The fabric forwards copies of the multicast packet to all egress cards having at least one output port that corresponds to the C-FMID. Because the association of multiple LC-FMIDs to a single C-FMID, superfluous multicast packets may be forwarded to egress cards unnecessarily, which causes unnecessary packet duplication and adds an unnecessary load to the interface between the fabric and the egress cards. The superfluous copies are typically discarded by the egress cards.

Id. at 8:35–55.

Of challenged claims 1–24, claims 1 and 12 are independent.

Claims 1 and 12 are illustrative and are reproduced below, with bracketed internal section numbering added to correspond to the labeling used by Petitioner:

1. [1.Pre] A network node in a communication network, comprising:

[1.1] a plurality of ports, at least a subset of which is grouped in a link aggregation (LAG) group; and packet processing logic, which is coupled to receive data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of the ports and

[1.2] to process the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group;

[1.3] wherein:

the packet processing logic is arranged to allocate to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, each FMID being associated with one of the ports in the subset, and to forward the single copy to the port associated with the allocated FMID value;

[1.4] the packet processing logic comprises multiple line cards connected to the ports and a switching fabric interconnecting the line cards, wherein, for each of the received data packets, a first line card connected to a first port via which the data packet is received is arranged to allocate the FMID value to the packet and to forward the packet to the switching fabric, and

[1.5] wherein the switching fabric and a second line card connected to a second port to which the data packet is to be sent are configured to forward the data packet responsively to the FMID value; and

[1.6] the first line card is arranged to assign to the data packets line card FMID (LC-FMID) values selected from a first range of possible LC-FMID values, and wherein the switching fabric is arranged to map the LC-FMID values to respective central FMID (C-FMID) values selected from a second range of possible C-FMID values that is smaller

than the first range and to forward the data packets responsively to the C-FMID values.

Id. at 10:15–55.¹

12. [12.Pre] A method of communication, comprising:

[12.1] receiving data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of a plurality of ports, at least a subset of which is grouped in a link aggregation (LAG) group; and

[12.2] processing the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group;

[12.3] wherein:

processing the data packets comprises allocating to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, the FMID value being associated with one of the ports in the subset, and forwarding the single copy to the port associated with the allocated FMID value;

[12.4] for each of the received data packets, processing the data packets comprises allocating the FMID value to the data packet by a first line card connected to a first port via which the data packet is received, and

[12.5] configuring a second line card connected to a second port to which the data packet is to be sent and a switching fabric interconnecting the first and second line cards to

¹ We read “(LC-FMID” in the claim as having the same meaning as “(LC-FMID)” and “(C-FMID” in the claim as having the same meaning as “(C-FMID),” as Petitioner suggests (Pet. vi n.2), because in both instances a close parenthesis character appears to be inadvertently left out of the claim. Patent Owner does not oppose Petitioner’s suggestion. It is reasonable to read the terms the same as though the close parenthesis is present. We recognize no other meaning. We took this approach in the Decision to Institute. Inst. Dec. 8 n.1. Post institution, neither party has objected.

forward the data packet responsively to the FMID value;
and

[12.6] allocating the FMID value comprises assigning to the data packets line card FMID (LC-FMID values selected from a first range of possible LC-FMID values, and wherein configuring the switching fabric comprises mapping the LC-FMID values to respective central FMID (C-FMID) values selected from a second range of possible C-FMID values that is smaller than the first range and forwarding the data packets responsively to the C-FMID values.

Id. at 11:34–12:3.²

*E. References and Other Evidence*³

The Petition relies on the following evidence:

References		Date	Exhibit
Devi	U.S. Pub. App. 2005/0238018 A1	Published Oct. 27, 2005	Ex. 1005
Mazzola	U.S. Patent 5,740,171	Issued Apr. 14, 1998	Ex. 1007
Borgione	U.S. Patent 8,155,125 B1	Issued Apr. 10, 2012, filed Sept. 17, 2004	Ex. 1006

The Petition also relies on a declaration of Tal Lavian, Ph.D.
(Ex. 1003).

² We read “(LC-FMID” in the claim as having the same meaning as “(LC-FMID),” as Petitioner suggests (Pet. viii n.5), because a close parenthesis character appears to be inadvertently left out of the claim. Patent Owner does not oppose Petitioner’s suggestion. It is reasonable to read the terms the same as though the close parenthesis is present. We recognize no other meaning. We took this approach in the Decision to Institute. Inst. Dec. 9 n.2. Post institution, neither party has objected.

³ The ’525 patent issued from Application 11/644,773, filed Dec. 21, 2006. Ex. 1001, codes (21), (22). The ’525 patent does not claim priority to an earlier-filed application.

The Patent Owner Response relies on a declaration of Robert Akl, D.Sc. (Ex. 2001).

F. Asserted Grounds of Unpatentability

Petitioner asserts that the challenged claims are unpatentable on the following grounds:

Claim(s) Challenged	35 U.S.C. §⁴	Reference(s)/Basis
1–24	103(a)	Devi, Mazzola
1–24	103(a)	Borgione, Mazzola

II. ANALYSIS

A. The Burden of Proof on Challenged Patent Claims

Petitioner must demonstrate by a preponderance of the evidence that the challenged claims are unpatentable. 35 U.S.C. § 316(e) (2018); 37 C.F.R. § 42.1(d) (2022). That burden never shifts to the patentee. *Dynamic Drinkware, LLC v. Nat’l Graphics, Inc.*, 800 F.3d 1375, 1378 (Fed. Cir. 2015).

B. Level of Ordinary Skill in the Art

Petitioner contends that a person of ordinary skill in the art at the time of the alleged invention “would have had a bachelor’s degree in electrical engineering, computer engineering, computer science, or a closely related field, and three to four years of experience in the design and development of telecommunication networks, systems, equipment, and components.” Pet. 7.

⁴ The Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (“AIA”), included revisions to 35 U.S.C. § 103 that became effective after the filing of the application that led to the ’525 patent. Therefore, we apply the pre-AIA version of 35 U.S.C. § 103.

Alternatively, Petitioner contends that a person of ordinary skill in the art “would have [had] a Master’s degree or similar post-graduate work in electrical engineering, computer engineering, computer science, or a closely related field, and fewer years of design and development experience.” *Id.*

In its Preliminary Response, Patent Owner took no position with respect to the level of ordinary skill in the art. In the Decision to Institute, we adopted Petitioner’s articulation of the level of ordinary skill in the art. Inst. Dec. 11. Post institution, Patent Owner does not dispute our adopted level of ordinary skill in the art.

Petitioner’s assertion is supported by the testimony of Dr. Lavian. Ex. 1003 ¶¶ 91–100. Petitioner’s assertion also appears consistent with the content of the applied prior art references. *Cf. Okajima v. Bourdeau*, 261 F.3d 1350, 1354–55 (Fed. Cir. 2001) (the applied prior art may reflect an appropriate level of skill).

Accordingly, in the absence of a contrary contention by Patent Owner, we continue to adopt Petitioner’s proposed level of ordinary skill in the art, as stated in two alternatives noted above.

C. Claim Construction

1. General Principles

We use the same claim construction standard that would be used to construe a claim in a civil action under 35 U.S.C. § 282(b), including construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent. 37 C.F.R. § 42.100(b) (2021). The claim construction standard set forth in *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc) is applicable.

Claim terms are generally given their ordinary and customary meaning as would be understood by one with ordinary skill in the art in the context of the specification, the prosecution history, other claims, and extrinsic evidence including expert and inventor testimony, dictionaries, and learned treatises, although extrinsic evidence is less significant than the intrinsic record. *Phillips*, 415 F.3d at 1312–17. Usually, the specification is dispositive, and it is the single best guide to the meaning of a disputed term. *Id.* at 1315.

Only those claim terms that are in controversy need to be construed, and only to the extent necessary to resolve the controversy. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017); *Wellman, Inc. v. Eastman Chem. Co.*, 642 F.3d 1355, 1361 (Fed. Cir. 2011).

Petitioner has not proposed a specific construction for any claim term, and states that no express construction is needed. Pet. 7. Further, Petitioner states that it understands each claim term according to its plain and ordinary meaning. *Id.* Patent Owner also has not proposed a specific construction for any claim term. PO Resp. 4–5. Patent Owner states that “for all claim terms, Patent Owner requests that the Board adopt the ordinary and customary meaning of the claim terms as understood by a POSITA [person with ordinary skill in the art].” *Id.* at 5.

2. *Limitation [1.2] and Limitation [12.2]*

Petitioner, however, introduces a claim construction position for the first time in this proceeding in its Reply:

If [Patent Owner] asserts these claim elements require forwarding only a single copy of each packet via only one port in each LAG, the '525 claims show this is an overly limiting interpretation. Dependent claims 3 and 14 add further limitations

to the independent claims by requiring the ports be grouped into multiple LAGs and to forward a copy of each packet to only one line card serving each LAG. These further limitations demonstrate the independent claims cover distributing forwarded copies of the same packet to different ports that may serve the same LAG.

Reply 7.

Limitation [1.2] reads as follow:

[1.2] to process the data packets so as to *forward only a single copy* of each of the data packets *via the output ports in the subset*, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group;

Limitation [12.2] reads as follows:

[12.2] processing the data packets so as to *forward only a single copy* of each of the data packets *via the output ports in the subset* while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group;

Ex. 1001, 10:22–27.

Because the above-quoted limitations of claims 1 and 12 plainly require that **only a single copy** of each data packet be forwarded to the output ports forming a subset (which according to limitations [1.1] and [12.1] forms a LAG), any position of Petitioner asserting the contrary must be contained in the Petition. 37 C.F.R. § 42.204(b)(3) (“the petition must set forth . . . [h]ow the challenged claim is to be construed.”). If not, Patent Owner is deprived of an adequate opportunity to respond, for example, by relying on expert testimony in a declaration to support its Response.

Petitioner’s articulated claim construction position regarding forwarding only a single copy is belated, and no explanation has been provided to justify omission of this claim construction position in the

Petition and to explain how the Petition complies with 37 C.F.R.

§ 42.204(b)(3). Thus, the argument is not entitled to consideration.

Even if we do consider Petitioner's belated claim construction position, it is nonetheless without merit and we are unpersuaded. Petitioner's contention is based on the doctrine of claim differentiation, which disfavors reading a limitation expressly stated in a dependent claim into a claim from which it depends. *See Liebel-Flarsheim Co. v. Medrad, Inc.*, 358 F.3d 898, 910 (Fed. Cir. 2004). Specifically, the presence of a dependent claim that adds a particular limitation raises a presumption that the limitation in question is not found in the independent claim. *Id.*

Here, dependent claims 3 and 14 respectively narrow independent claims 1 and 12 in other ways, i.e., requiring "two or more LAG groups." Therefore, even under the doctrine of claim differentiation it does not follow that the feature of forwarding only a single copy of each data packet via the output ports of a LAG is presumably not in independent claims 1 and 12. Furthermore, the doctrine of claim differentiation is merely a guide and not a rigid rule. *Laitram Corp. v. Rexnord, Inc.*, 939 F.2d 1533, 1538 (Fed. Cir. 1991). It still must yield to clear language to the contrary in the independent claims. *E.g., Wi-LAN USA, Inc. v. Apple Inc.*, 830 F.3d 1374, 1391 (Fed. Cir. 2016) (claim differentiation cannot overcome a contrary construction dictated by the claim language, the written description, or prosecution history). As noted above, claims 1 and 12 clearly and expressly require forwarding only a single copy of each data packet via the output ports of a LAG.

Other than the foregoing, we find that no express claim construction of any claim term is necessary for this Decision.

D. Alleged Obviousness of Claims 1–24 over Borgione and Mazzola

1. Principles of Obviousness

A patent claim is unpatentable under 35 U.S.C. § 103 if the differences between the claimed invention and the prior art are such that the claimed invention, as a whole, would have been obvious to a person having ordinary skill in the art to which the claimed invention pertains. *KSR Int’l Co. v. Teleflex Inc.*, 550 U.S. 398, 406 (2007). The question of obviousness is resolved on the basis of underlying factual determinations including:

(1) the scope and content of the prior art; (2) any differences between the claimed subject matter and the prior art; (3) the level of ordinary skill in the art; and (4) when in evidence, objective evidence of nonobviousness.⁵

Graham v. John Deere Co., 383 U.S. 1, 17–18 (1966). There must be reasoning with rational underpinning to support combining teachings from multiple references. *In re Kahn*, 441 F.3d 977, 988 (Fed. Cir. 2006).

2. Overview of Borgione

Borgione relates to transmitting multicast data packets across an aggregate of network connections between two network nodes. Ex. 1006, 1:7–10. It describes that pre-existing methods create a load imbalance for multicast packets among the plurality of network links. *Id.* at 3:37–49. Borgione states that its disclosed invention balances the transmission of multicast packets among an aggregate of network links that provide a logical channel or link between network nodes. *Id.* at 5:28–30.

Borgione discloses a method, system, and apparatus for transmitting multicast data packets over a plurality of physical network links that have been combined into one logical channel, in a manner that the packets are

⁵ Patent Owner has not presented objective evidence of nonobviousness in this proceeding.

distributed over more than one network link. *Id.* at code (57), 3:60–64.

Figure 1 is reproduced below:




Figure 1

Figure 1 is a block diagram illustrating the network topology. *Id.* at 5:11–12.

Link node 110 may be a switch, bridge, router, or a hub. *Id.* at 1:30–31. Link node 110 is connected to link node 120 through a plurality of network links 150, and link 120 also may be a switch, bridge, router, or a hub. *Id.* at 1:37–40. The plurality of network links 150 between link nodes 110 and 120 are aggregated as a single logical link over which all traffic between link nodes 110 and 120 is distributed. *Id.* at 1:57–60, 1:66–67. Routing protocols see the aggregated links as a single interface with a common IP address. *Id.* at 2:5–7. Link nodes 110 and 120 are in separate locations and each has its own associated local area network (“LAN”). *Id.* at 1:55–67.

Borgione describes a first network node selecting a destination interface identifier for an outgoing multicast packet, selecting one of the plurality of network links using the destination interface identifier, and transmitting the outgoing multicast packet to the second network node over the selected network link. *Id.* at 4:16–21. Borgione explains:

An added tag header can be used [for each multicast packet] to include a destination interface identifier. For example, in a VLAN network environment, such a tag header is included in packets per IEEE Std. 802.1Q. A portion of an IEEE Std. 802.1Q tag header is a VLAN identifier (VID), which is unique to a particular VLAN. A destination interface identifier within a tag header can be used to select which network link in a logical link is to be used to transmit a replicated multicast packet. Since the destination interface identifier often varies from replicated multicast packet to replicated multicast packet, use of the destination interface identifier to select a network link will lead to a more even distribution of multicast packet transmission across the logical link. Such a distribution can reduce the likelihood of a load imbalance in the logical link.

Id. at 5:37–50.⁶ Figure 6 is reproduced below:

⁶ VLAN means virtual local area network. Ex. 1007, 3:5–6.




Figure 6

Figure 6 is a flow chart of actions taken in accordance with Borgione's method and system. *Id.* at 5:21–22.

After a multicast packet is received by a link node, at step 630 the link node determines the VLAN location of receivers of the multicast packet using a table of multicast receivers, i.e., “multicast expansion table,”

corresponding to the multicast group ID of the multicast packet. *Id.* at 6:40–44. For each such destination VLAN, the link node at step 660 inserts a tag header that contains the respective destination interface identifier, i.e., a VID. *Id.* at 6:47–50. Then, at step 670, the link node calculates a network link identifier corresponding to a specific network link within the logical link. *Id.* at 6:58–61. To increase distribution among the physical links within a logical link, the network link identifier is generated by analyzing the VID portion of the tag header as well as the source and destination addresses. *Id.* at 6:61–64. A hash algorithm can be used to generate the network link identifier. *Id.* at 7:1–5. At step 680, the multicast packet is transmitted on the specific network link identified by the network link identifier. *Id.* at 7:6–8.

3. *Overview of Mazzola*

Mazzola generally relates to network switches and specifically to an address-translation mechanism for use in a network switch. Ex. 1007, 1:23–25. A network switch of a data communications network provides a switching function for transferring information such as data frames among network entities. *Id.* at 1:28–30. The switch includes a collection of cards interconnected by a backplane. *Id.* at 1:31–33. Each card includes a limited number of ports that couple the switch to other network entities over various types of media such as Ethernet. *Id.* at 1:33–35.

Specifically, Mazzola describes an address-translation mechanism that renders forwarding decisions for data frames transported among the ports of a switch. *Id.* at 3:3–6. According to Mazzola, each port is associated with a unique index value and a VLAN (“color”) identifier to facilitate multicast data transfer within the switch. *Id.* at 3:10–14. Additionally, the address-translation mechanism comprises a plurality of forwarding tables, each of

which contains entries having unique index values that translate to selection signals for ports destined to receive the data frames. *Id.* at 3:7–10.

A media access control (“MAC”) address of a data frame is combined with the “color” identifier of the data frame to produce a base line numerical quantity for searching the forwarding tables. *Id.* at 3:15–18. Each table entry, however, is directly accessed by a key comprising a hash transformation of the MAC/color quantity. *Id.* at 3:18–20. Mazzola explains:

Operationally, if the compared items [(the hashed MAC/color quantity and a lookup key for the table)] match, the index value stored in the table is provided to a target logic circuit for translation to a signal that selects a port or [a] group of ports for receiving the index frame. If the items do not match, the color identifier is passed to the target logic circuit with the result that all ports having that color identifier receive the frame in accordance with a multicast transfer.

Id. at 3:27–33.

Figure 1 is reproduced below:




FIG. 1

Figure 1 is a schematic block diagram of a network switch for interconnecting ports, implemented on various cards of the switch, that couple the switch to network entities of a data communications network *Id.* at 3:64–4:1. Network switch 100 includes supervisor card 120 and a plurality of line cards 130. *Id.* at 3:67–4:2. Each line card may have as many as 64 ports thereon. *Id.* at 4:2–3. Communication among the ports/cards occurs synchronously over switching bus 110. *Id.* at 4:4–6. Mazzola describes:

The architecture of the switching bus is preferably implemented as forwarding engine circuitry 125 and port interface circuitry 135 that cooperate to efficiently transmit data to, and receive data from, the switching bus 110. The forwarding engine 125 is located on the supervisor card 120 and comprises an encoded address recognition logic (EARL) circuit 400 coupled to a local target logic (LTL) circuit 450 over a result bus 475. Functionally, the EARL 400 executes all forwarding decisions for the switch 100, while the LTL 450 implements those forwarding decisions by selecting ports as destinations for receiving data frames transferred over the switching bus 110. As can be seen in FIG. 1, the LTL may be distributed among the cards of the switch via the result bus 475.

Id. at 4:8–21. Port interface circuitry 135 includes SAINT [synchronous advanced interface network termination (*id.* at 4:23–24)] circuit 200 which “is primarily responsible for (i) receiving data frames from the network and passing them to the switching bus, and (ii) receiving frames from the switching bus and passing them to the network.” *Id.* at 4:59–62. SAINT circuit 200 includes MAC circuit 210 coupled to DMA engine 220, and the DMA engine “modifies each frame by appending header information to it prior to arbitrating for access to the switching bus.” *Id.* at 5:8–9.

Figure 3 is reproduced below:




FIG. 3

Figure 3 is a block diagram illustrating the format of a data frame modified by port interface circuitry 135. *Id.* at 3:49–50. Two headers 310 are appended to data frame 300. *Id.* at 5:15–16. Mazzola describes:

When received at the port, the frame contains a 48-bit destination address in field 304 that identifies the destination network address of the frame and a 48-bit source address in field 306 identifying the source network address of the frame; the DMA engine 220 of SAINT 200 then appends two 48-bit headers 310 to the frame 300. Of interest to the present invention are a unique index value and a virtual local area network (VLAN) “color” identifier that are loaded in to a 10-bit index field 312 and a 10-bit color field 314, respectively, of the frame 300 by the engine.

As described herein, these values represent the source index of the port transferring the frame over the switching bus 110 and the VLAN association of the port.

Id. at 5:11–23.

“Each port or group of ports within switch 100 is assigned a unique index value.” *Id.* at 6:20–22. Mazzola further states elsewhere that each port has a unique index number. *Id.* at 5:48–49, 5:65–66, 6:38–39, 7:5. Additionally, a VLAN identification, i.e., “color,” is associated with each port of the switch. *Id.* at 6:27–30. Thus, each port has associated therewith a VLAN identification “color” and a unique index value. *Id.* at 6:38–39. EARL 400 uses address translation to map a data frame’s destination address and color to a destination port using forwarding tables. *Id.* at 5:45–49.

When a data frame appears on the bus, EARL 400 extracts certain information to produce a numerical input quantity for searching the forwarding tables. *Id.* at 7:56–61. Mazzola describes that either of the 48-bit destination address DA or 48-bit source address SA is combined with the 10-bit color identifier to produce the base line numerical input quantity for searching a table. *Id.* The combined 58-bit numerical quantity is hashed to produce a 15-bit key by a hashing function implemented by a conventional hashing algorithm. *Id.* at 8:17–20. The resulting key maps to an entry location in each of four tables. *Id.* at 8:25–27, 8:46–48. When accessing the tables, it is necessary to identify which of the four tables holds the correct index value for the MAC/color quantity of the current data frame. *Id.* at 8:49–52. “To do this, the MAC address and color are stored in the MAC/color field 514, and the index value is stored in the index field 516, when the entry is created.” *Id.* at 8:52–54. “After accessing a table with the hashed MAC/color quantity, the MAC address and color from that table is

compared with the MAC/color of the current frame to verify whether the index is correct.” *Id.* at 8:54–58.

Regarding determining the appropriate destination port for a data frame, Mazzola further describes:

Operationally, the comparator 532 compares the MAC/color quantity from data frame 300 with the MAC/color value stored in field 514 of a selected mapped entry to ensure that the entry contains the correct index value. If the compared items match, the index value stored in the proper entry is selected by multiplexer 536 and provided to the input of multiplexer 538. This latter mux is enabled by the matched output signal of comparator 532 to provide the index over the result bus 475 and to the LTL 450. In response, the LTL translates the received index to a signal that selects a port or a group of ports for receiving the data frame. If the MAC/color values do not match (i.e., a “miss”), the color identifier is selected by the multiplexer 538 and passed to the LTL 450 with the result that all ports having that color identifier receive the frame in accordance with a multicast transfer.

Id. at 9:20–35.

4. *Independent Claim 1*

a) *Recitation [1.Pre]*


Claim 1 recites in the preamble [1.Pre]: “A network node in a communication network, comprising:” Ex. 1001, 10:15–16. Petitioner asserts that Borgione teaches a network node in a communication network and a method for communication. Pet. 55 (citing Ex. 1006, 1:6–10, 3:53–56, 7:6–8). Referring to Figure 1 of Borgione, Petitioner identifies link node 110 as the network node. *Id.* Patent Owner does not present argument to the contrary. We are persuaded that Borgione discloses “[a] network node in a communication network.”⁷

⁷ We need not decide whether this preamble recitation is limiting.

b) *Limitation [1.1]*

Claim 1 further recites limitation [1.1]: “a plurality of ports, at least a subset of which is grouped in a link aggregation (LAG) group; and packet processing logic, which is coupled to receive data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of the ports.” Ex. 1001, 10:17–22.

Petitioner presents the following annotated partial drawing of Borgione’s Figure 1 to show satisfaction of limitation [1.1] by Borgione:



The figure above is Borgione’s Figure 1 with annotations in gold for Network Node, in purple for Egress Links, and in green for Ingress Links added by Petitioner (Pet. 56). Petitioner cites the follow text of Borgione:

[T]he plurality of network links between 110 and 120 *can be aggregated as a single logical link*. In this manner, each link node 110 and 120 sees the plurality of network links between

them as one logical interface. . . . Routing protocols treat *the aggregated links as a single, routed interface* with a common IP address.

Id. (alterations by Petitioner) (citing Ex. 1006, 1:17–23, 1:37–39, 1:57–2:7, 3:60–64; Ex. 1003 ¶¶ 216–217).

Although Borgione’s disclosure does not refer to the presence of ports, Dr. Lavian testifies:

Based on Borgione’s description of the structure and function of its links, a POSITA would have understood that Borgione discloses its links include a plurality of ports for receiving and forwarding packets. For instance, Borgione explains that its links include network link identifiers which function to identify ports:

The *network link identifier* identifies which of the plurality of network links is to be used for sending the packet between link nodes 110 and 120.

Once a *network link identifier* has been calculated, the replicated multicast packet can then be transmitted on the identified network link (680).

Ex. 1003 ¶ 219 (citing Ex. 1006, 3:34–36, 3:53–56, 4:45–54, 6:35–37, 7:6–8). We understand Dr. Lavian as testifying that each identified physical network link represents an output port on the link node sending the packet and an input port on the link node receiving the packet, and that the ports provide access of a data packet to the link. The assertion is logical and supported by the disclosure of Borgione. Patent Owner presents no contrary argument disputing the assertion.

We are persuaded by Petitioner that each of Borgione’s physical egress links identifies a specific output port on link node 110 sending the packet and a specific input port on link node 120 receiving the packet.

Furthermore, we are persuaded by Petitioner that the plurality of purple colored “Egress Links” in the above annotated Figure 1 of Borgione

constitutes a LAG group (labeled as such by Petitioner) in which at least a subset of a plurality of ports, specifically all of the output ports on link node 110, is grouped. Patent Owner presents no contrary argument.

Regarding processing logic, Petitioner asserts:

Further, Borgione's network node includes packet processing logic for receiving multicast traffic from VLAN nodes 130(1)–(5) wherein each packet has a respective destination address that specifies forwarding the packet to a group of multiple recipients via egress links. And Borgione's network node includes a processor processing instructions for forwarding packets by analyzing multicast destination addresses and packet headers:

[T]he present invention provides an apparatus comprising a processor, network interfaces coupled to a plurality of network links configured to form a logical channel, and a memory for storing instructions that upon execution cause the processor to transmit a replicated multicast packet over one of the plurality of network links, and to select one of the plurality of network links by analyzing a destination ethernet address of the replicated multicast packet and a non-ethernet address component of a header of the replicated multicast packet.

Pet. 57 (emphases omitted) (alteration by Petitioner) (quoting Ex. 1006, 4:45–54). Petitioner's assertions are reasonable and supported by the cited evidence. Patent Owner presents no contrary argument in that regard. We are persuaded that Borgione discloses processing logic as recited in limitation [1.1].

For the foregoing reasons, we are persuaded that Borgione discloses limitation [1.1].

c) Limitation [1.2]

Claim 1 further recites limitation [1.2]: “[and] to process the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group.” Ex. 1001, 10:22–27.

Petitioner asserts that Borgione discloses this limitation. Pet. 59. Alternatively, Petitioner contends that Borgione either alone or in combination with Mazzola suggests this limitation. *Id.* Although obviousness is alleged, we find no obviousness reasoning or analysis in the Petition with regard to the requirement of forwarding only a single copy of each of the data packets via the output ports in the subset. Instead, the obviousness assertion in the Petition for limitation [1.2] is to account for the presence of output ports in a network node, and that assertion merely makes reference to Petitioner’s accounting for limitation [1.1] discussed above. *Id.* at 60. Accordingly, we focus on whether Borgione discloses forwarding only a single copy of each of the data packets via the output ports in the LAG, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG.

Petitioner specifically cites to the following disclosures in Borgione:

Upon receipt of a packet including a multicast destination address, a linking node [e.g., node 110] can identify the network nodes subscribing to the multicast group and replicate the multicast packets. The replicated multicast packets can then be forwarded to the identified network nodes. Each replicated multicast packet will have the same source and destination addresses in the MAC header.

Ex. 1006, 6:14–20 (cited at Pet. 59).

[T]he present invention provides a method for transmitting a replicated multicast packet over one of a plurality of network links that form one logical channel. . . . [¶] The first network node selects a destination interface identifier for an outgoing multicast packet, selects one of the plurality of network links using the destination interface identifier, and transmits the outgoing multicast packet to the second network node over the selected network link.

Id. at 4:3–21 (cited at Pet. 59).

For several reasons, we are not persuaded by Petitioner that Borgione forwards only a single copy of a packet via the output ports of Borgione’s LAG as identified by Petitioner.

Nowhere in the above-cited portions of Borgione is an indication that no other copy or replicated packet will be forwarded via the output ports of Borgione’s LAG. The above-cited description discusses one forwarding instance but does not represent that there is no other forwarding of a copied or replicated packet. To the contrary, the cited text indicates that other copies or replicated versions would still be forwarded via the output ports of the same LAG.

Borgione describes that upon receipt of a packet with a multicast destination address, i.e., a multicast packet, a linking node such as node 110 identifies the network nodes subscribing to the multicast group and replicates the packet. Ex. 1006, 6:14–17. Borgione further states: “The replicated multicast packets can then be forwarded to the identified network nodes [such as node 120].” *Id.* at 6:17–18.

With regard to distributing forwarded copies to balance the traffic load within the LAG, Petitioner cites to the following disclosure of Borgione:

The present invention balances the transmission of replicated multicast packets among an aggregate of network links that provide a logical channel or link between network nodes An added tag header can be used to include a destination interface identifier A destination interface identifier within a tag header can be used to select which network link in a logical link is to be used to transmit a replicated multicast packet. Since the destination interface identifier often varies from replicated multicast packet to replicated multicast packet, use of the destination interface identifier to select a network link will lead to a more even distribution of multicast packet transmission across the logical link. Such a distribution can reduce the likelihood of a load imbalance in the logical link.

Pet. 59–60 (alterations by Petitioner) (citing Ex. 1006, 5:27–50; Ex. 1003 ¶ 225).

This description about balancing the transmission of replicated multicast packets among a plurality of physical network links in a logical channel forming a LAG also indicates that not just a single packet is forwarded across the logical channel to another node. Instead, it indicates that all replicated packets, as many as there are eventual destinations, are forwarded across the logical channel or link, but that the choice of the physical link within the LAG is varied for each replicated packet to avoid an imbalance caused by sending all replicated packets across the logical channel on the same physical network link.

As discussed above, the disclosures of Borgione cited by Petitioner do not support Petitioner's assertions. Rather, they indicate a packet with multiple destinations is replicated multiple times and each replicated packet is transmitted across the logical channel including a number of physical network links. They also indicate the particular network link used for transmission of a replicated packet is selected to avoid load imbalance created by using the same network link each time.

Patent Owner cites to the following description in the Summary of the Invention section of Borgione, which even more clearly indicates that a multicast packet in Borgione, when received by one node for forwarding to another node, is replicated into multiple copies and that all replicated copies are transmitted over a plurality of physical network links that are combined into one logical channel (LAG): “The present invention presents a method, system, and apparatus to transmit replicated multicast packets into a plurality of physical network links that are combined into one logical channel or link *so that the replicated multicast packets are distributed over more than one network link.*” Ex. 1006, 3:60–64 (cited at PO Resp. 29–30) (emphasis added).⁸

Patent Owner further explains, with persuasion, how Borgione accomplishes the more balanced distribution among the plurality of network links, by adding to each replicated multicast packet a tag header that includes a VLAN (virtual local area network) identifier (“VID”) that is unique to a VLAN, such that packets destined to different VLANs will have a different VID. PO Resp. 30 (citing Ex. 1006, 3:60–42, 5:33–36, 5:38–40, 6:25–27; Ex. 2002, 29:3–6 (Lavian deposition)). Patent Owner further explains, with persuasion, that “[i]f the ‘non-ethernet address components’ of the replicated multicast packets differ, then Borgione will send the replicated multicast packets (i.e., copies of the received data packet) over more than one of the plurality of network links.” *Id.* at 28 (citing Ex. 2001


⁸ Borgione describes that preexisting methods and systems would select the same physical network link to send each replicated multicast packet and therefore would send all of the replicated packets over the same physical link, which creates “an undesirable load imbalance among the plurality of network links.” Ex. 1006, 3:45–49.

¶ 66). Thus, replicated packets going to different VLANs at the receiving node would be transmitted over different network links within the same logical channel or LAG. This description conveys that Borgione does not forward only a single multicast packet from the output ports of a LAG as is required by limitation [1.2].

In its Reply, Petitioner argues that according to the '525 patent, each packet is duplicated “when servicing multiple link aggregation groups (‘LAGs’).” Reply 3. According to Petitioner, that is what Borgione’s node does, when modified to include a switching fabric. *Id.* at 5. Petitioner explains:

Independent claims 1 and 12 recite grouping a subset of ports into a LAG. And those claims recite forwarding a single copy of a data packet via the ports in the LAG subset. But the claims don’t exclude the forwarded packet being a duplicate of another packet. Rather, the '525 Patent establishes the opposite—namely, the node will often forward a packet’s duplicate to an output port.

The '525 Patent states that “[w]hen the packet needs to be sent to multiple LAG groups, the switching fabric duplicates the packet.” As shown below, the patent’s Figure 1 [color annotations and lower half of the illustration added by Petitioner] demonstrates forwarding a duplicated data packet to an output port of a LAG subset. And Figure 1 shows forwarding multiple duplicated data packets to output ports of multiple LAGs. For example, as Dr. Akl admitted at deposition, Figure 1’s switching fabric 40 forwards Packet 1 to port 28A of LAG 1 and forwards Duplicate Packet 1 to port 28D of LAG 2.



[Annotated Figure 1 of Borgione illustrates
a network node (Ex. 1001, 3:7–9)]

Further, '525 claim 3 recites “the switching fabric is configured to duplicate the data packet and to forward a copy of the data packet.” PO and Dr. Akl don’t address that the claimed invention requires the ability to duplicate packets before forwarding them on one of the LAGs.

Id. at 3–5 (first alteration by Petitioner).⁹

⁹ During oral hearing, Patent Owner argued that annotated Figure 1 was not presented in either the Petition or Dr. Lavian’s declaration (Ex. 1003) and therefore is a new argument that should not be considered. Tr. 51:21–52:3. This figure appears on page 4 of the Reply. First, Patent Owner did not make this objection in its Sur-Reply or seek to file a motion to strike. Second, we do not see this figure as presenting a new argument that should have, if at all, been included in the Petition. By this figure, Petitioner appropriately responds to Patent Owner’s argument distinguishing Borgione from the claimed invention by relying on a more vivid illustration of Petitioner’s position. The Petition already asserts that Borgione discloses limitation [1.2]. Pet. 59.

Petitioner's argument is misplaced. The issue is not about whether claim 1 excludes the making of copies or duplicates of a data packet or whether a packet sent to a LAG is a copy or duplicate of another packet sent elsewhere. Rather, the issue is whether Borgione discloses forwarding only a single copy of a multicast packet via the output ports of a LAG, as is required by limitation [1.2].

Petitioner's annotated version of Figure 1 of the '525 patent is consistent with and reflects what is claimed. For instance, it is illustrated that Packet 1 is sent to LAG 32A only once, and sent to LAG 32B only once. Similarly, it is illustrated that Packet 2 is sent to LAG 32A only once, and sent to LAG 32B only once. Borgione, however, as discussed above, operates differently by sending multiple copies or duplicates of the same packet from the output ports of one LAG.

Petitioner asserts that Patent Owner's description about how Borgione sends packet copies with different non-ethernet addresses on more than one of the plurality of the network links forming a LAG does not distinguish from the claim. Reply 5–6. But it does. We determined in Section II.C.2 above that claim 1 requires forwarding only a single copy of each of the data packets via the output ports of a LAG. Borgione discloses one LAG, which comprises multiple network links 150 between link node 110 and link node 120 (Ex. 1006, Fig. 1). Borgione's sending multiple copies of a data packet on multiple network links 150 does not satisfy limitation [1.2].

Petitioner further asserts:

PO's hypothetical [when packets have different non-ethernet addresses, Borgione sends them over different network links in the same LAG] confirms that, when packets have the same non-ethernet address components, Borgione will forward a single packet over one network link in the LAG. And Borgione

describes forwarding packets over separate links of a “logical channel” (i.e., LAG), showing that Borgione discloses or suggests forwarding packets over one LAG and over a link of another LAG. Thus, even under PO’s hypothetical, Borgione discloses or renders obvious claim elements 1.2 and 12.2.

Reply 7. There is little support, if any, for Petitioner’s assertion, and we are unpersuaded.

The reasonable inference is that if Borgione is not selecting different network links to use based on different non-ethernet addresses, such as when the non-ethernet addresses are the same, then Borgione would use the same network link to transmit the plurality of duplicated packets.¹⁰ Similarly, the description of forwarding multiple packets over a plurality of links within a LAG would not have reasonably suggested forwarding only a single data packet among all network links in the LAG. The idea of forwarding only a single packet via the output ports in the LAG simply is absent, not disclosed or reasonably suggested by Borgione.

Further, Figure 6 of Borgione, reproduced above in Section II.D.2, shows that when a multicast packet is received, the multiple destinations are expanded and placed in a table (MET), and that each table entry is processed by replicating the packet at step 650, and by transmitting the packet on a selected network link at step 680 if the destination is across the logical link, and then the process repeats for the next destination entry in the table.

Ex. 1006, Fig. 6. The concept of forwarding only a single copy of a multicast packet via the output ports of a LAG is also absent from Figure 6.

¹⁰ Petitioner does not explain why packets that have the same ethernet address portion as well as the same non-ethernet address portion are not the same packet going to the same destination, in which case there would not be a multicast packet as is required by claim 1.

Because we find that Borgione does not teach or suggest the feature of forwarding only a single copy of a multicast packet via the output ports of a LAG, as required by limitation [1.2], we determine that Petitioner accounts for this feature through reliance on improper hindsight in light of the disclosure of the '525 patent.

d) Conclusion

Petitioner has not sufficiently accounted for limitation [1.2]. Hence, Petitioner has not proved by a preponderance of the evidence that claim 1 would have been obvious over Borgione and Mazzola.

5. Independent Claim 12

Limitation [12.2], like limitation [1.1], also requires forwarding only a single copy of each data packet via the output ports of a LAG. The parties' arguments and supporting evidence for limitation [12.2] are essentially the same as those for limitation [1.2]. For the same reasons discussed above regarding the deficiency of Petitioner's arguments and evidence for limitation [1.2], Petitioner's arguments and evidence for limitation [12.2] are equally deficient. Accordingly, Petitioner has not proved by a preponderance of the evidence that claim 12 would have been obvious over Borgione and Mazzola.

6. Dependent Claims 2–11 and 13–24

Claims 2–11 each depend, directly or indirectly, from independent claim 1. Claims 13–24 each depend, directly or indirectly, from independent claim 12. Thus, claims 2–11 each incorporate all of the limitations of claim 1, including limitation [1.2], and claims 13–24 each incorporate all of the limitations of claim 12, including limitation [12.2].

All of the deficiencies of Petitioner's assertions discussed above in the context of limitation [1.2] equally apply to claims 2–11, and all of the

deficiencies of Petitioner's assertions regarding limitation [12.2] as noted above equally apply to claims 13–24. Accordingly, Petitioner has not proved by a preponderance of the evidence that any of claims 2–11 and 13–24 would have been obvious over Borgione and Mazzola.

E. Alleged Obviousness of Claims 1–24 over Devi and Mazzola

1. Overview of Devi

Devi discloses a method performed by an ingress card including receiving a packet with a destination identifier and determining an aggregated link associated with the destination identifier that includes multiple egress ports. Ex. 1005 ¶ 5. The method further includes determining a destination port from the egress ports of the aggregated link, and communicating the packet to the destination port. *Id.*

Figure 1 of Devi shows a system coupled to a network and is reproduced below:




FIG. 1

Figure 1 illustrates system 100 including switch 104 coupled to network 102. *Id.* ¶ 15. “Using network 102 and switch 104, packets are communicated from sources 103 to destinations 112. Components of switch 104 are referred to as ‘ingress’ components when receiving packets from sources 103 and ‘egress’ components when sending packets to destinations 112.” *Id.* Switch 104 includes cards 106 coupled to physical links 108 and a backplane 114 that represents hardware and/or software allowing cards 106 to exchange information with one another. *Id.* ¶ 18. “Each card 106 has one or more ports 110 coupled to physical links 108.” *Id.* ¶ 19. Links 108 couple to cards 106 through ports 110. *Id.* ¶ 20.

Multiple physical links 108 to a particular destination 112 may be aggregated to form an aggregated link 116. *Id.* ¶ 23. It is treated as a single link between switch 104 and the particular destination 112. *Id.* “Because aggregated link 116 includes multiple physical links 108 that go to the same destination 112, a packet sent to destination 112 need only be sent to one port 110 corresponding to one physical link 108 of aggregated link 116.” *Id.*

Devi refers to destination identifier 215, destination information 216, and port information. *Id.* ¶¶ 29–32. “Destination identifier 215 contemplates any information determinable from a packet that allows card 106 to identify one or more destinations 112 for the packet.” *Id.* ¶ 30. “Destination information 216 associates destination identifiers 215 in packets with egress links.” *Id.* ¶ 31. Card 106 uses destination information 216 to determine an egress link or links using destination identifier 215 from a packet. *Id.* Port information 218 refers to any form of data that identifies ports 110 associated with a particular link. *Id.* ¶ 32. “Port information 218 for an aggregated link 116 identifies the multiple ports 110 coupled to the physical links 108 of aggregated link 116, which may include ports 110 on a

single card 106 or multiple cards 106.” *Id.* Destination information 216 and port information 218 are stored within card 106. *Id.* ¶ 29.

Figure 3 of Devi shows a flow chart for forwarding packets and is reproduced below:




Figure 3 is a flow chart of a method in Devi for forwarding data packets. *Id.* ¶ 13.

Card 106 receives a packet at step 302. *Id.* ¶ 37. At step 304, card 106 determines destination identifier 215 for the packet and uses destination information 216 to determine if card 106 has previously learned the destination identifier 215. *Id.* If destination identifier 215 has been learned, card 106 next determines if destination identifier 215 identifies a multicast or unicast destination at step 308. *Id.* If the packet is a multicast packet, card 106 selects one of the destinations for the packet using destination information 216 at step 310. *Id.* Card 106 identifies an egress link associated with the destination using destination information 216 at step 312. *Id.* ¶ 39. “If the link is an aggregated link 116, card 106 determines a destination port 110 of aggregated link 116 using the information in port information 218 at step 320.” *Id.* ¶ 41. The selection of a particular destination port 110 from the multiple ports 110 of the aggregated link may be by random selection, load balancing, pseudo-random selection or hashing using part of the packet information, historical tracking, round robin, or any other method of selection. *Id.* “Card 106 then communicates the packet to the destination port 110 at step 322.” *Id.*

2. *Independent Claim 1*

a) *Recitation [1.Pre]*

Claim 1 recites in the preamble [1.Pre]: “A network node in a communication network, comprising:” Ex. 1001, 10:15–16. Referring to Figure 1 of *Devi*, Petitioner asserts that *Devi* discloses a network node in a communication network and a method for communication. Pet. 17. Petitioner states: “*Devi* couples a network node (switch 104) to a communication network 102.” *Id.* (citing Ex. 1005 ¶¶ 18–19, 23; Ex. 1003 ¶¶ 108–110). Patent Owner does not present argument to the contrary. We

are persuaded that Devi discloses “[a] network node in a communication network.” We need not decide whether this preamble recitation is limiting.

b) Limitation [1.1]

Claim 1 further recites limitation [1.1]: “a plurality of ports, at least a subset of which is grouped in a link aggregation (LAG) group; and packet processing logic, which is coupled to receive data packets having respective destination addresses that specify forwarding the packets to groups of multiple recipients through at least one of the ports.” Ex. 1001, 10:17–22. Referring to Figure 1 of Devi, Petitioner asserts: “Devi’s network node 104 has multiple ports 110 and packet processing logic that includes ingress and egress line cards 106 and a switching fabric 114 interconnecting the line cards. And Devi describes grouping one or more ports 110 into a link aggregation group (‘LAG’) 116.” Pet. 19 (citing Ex. 1005 ¶¶ 18–21, 23–25, 32, 34, 41; Ex. 1003 ¶¶ 111–112).

Petitioner further asserts: “Devi’s logic receives packets having respective destination addresses that specify forwarding the packets to multicast groups through one or more ports 110 of LAG 116. For example, Devi’s ingress line card stores destination addresses of received multicast packets for forwarding to output ports of LAGs.” *Id.* at 19–20 (citing Ex. 1005 ¶¶ 8, 26, 29–32, 34, 41; Ex. 1003 ¶¶ 113–115).

Patent Owner does not present argument to the contrary. Petitioner’s assertions are supported by the cited evidence. We are persuaded that Devi discloses limitation [1.1].

c) Limitation [1.2]

Claim 1 further recites limitation [1.2]: “[and] to process the data packets so as to forward only a single copy of each of the data packets via the output ports in the subset, while distributing forwarded copies of the data packets among the output ports in the subset so as to balance a traffic load within the LAG group.” Ex. 1001, 10:22–27. Petitioner explains:

Devi processes received packets to forward only one copy of each packet via output ports in a respective LAG. Devi explains that its “method allows card 106 to communicate the packet to all of its destinations 112 without sending multiple copies of the packet to the same destination.” To do so, Devi forwards one copy of each packet to a particular output port of an associated LAG.

Pet. 20–21 (citing Ex. 1005 ¶¶ 2, 23, 34, 36, 42, 44, 46; Ex. 1003 ¶¶ 116–117). Referring to Devi’s Figure 4, Petitioner further explains: “Devi determines the destination port within the LAG 116 (step 416) and forwards a copy of the packet to the destination port in the LAG (step 418).” *Id.* at 22 (citing Ex. 1003 ¶ 118).

Additionally, Petitioner explains that Devi forwards copies among the output ports 110 to balance a traffic load within LAG 116. *Id.* at 23 (citing Ex. 1005 ¶ 23; Ex. 1003 ¶ 119). Referring to Figure 3 of Devi, Petitioner states: “Devi determines the output port for load balancing within the LAG 116 (step 320) and then forwards the packet copy to the port (step 322).” *Id.* (citing Ex. 1005 ¶ 41; Ex. 1003 ¶¶ 120–122).

Patent Owner does not present argument to the contrary. Petitioner’s assertions are supported by the cited evidence. We are persuaded that Devi discloses limitation [1.2].

d) Limitation [1.3]

Claim 1 further recites limitation [1.3]:

wherein:

the packet processing logic is arranged to allocate to each of the received data packets a fabric multicast identification (FMID) value selected from a range of possible FMID values, each FMID being associated with one of the ports in the subset, and to forward the single copy to the port associated with the allocated FMID value;

Ex. 1001, 10:28–35.

Petitioner makes two assertions regarding limitation [1.3]. First, Petitioner asserts that *Devi* alone renders obvious limitation [1.3]. Pet. 25–28. Alternatively, Petitioner asserts that *Devi* in combination with *Mazzola* teaches limitation [1.3]. *Id.* at 28–29. For reasons discussed below, we are not sufficiently persuaded by either assertion. We first discuss the contention based on *Devi* alone, and then the contention based on a combination of *Devi* and *Mazzola*.

As noted above, *Devi* discloses that “[d]estination identifier 215 contemplates any information determinable from a packet that allows card 106 to identify one or more destinations 112 for the packet.” Ex. 1005 ¶ 30. *Devi* describes: “Destination identifier 215 represents information identifying an intermediate or final destination 112 for the packet. Destination identifier 215 may include a network address for destination 112, a logical address for an egress link, or any other information useful for identifying the destination of a packet.” *Id.*

Petitioner identifies *Devi*’s destination identifier 215 as an FMID value and asserts that for each received packet, *Devi*’s line card allocates destination identifier 215, a multicast value, to the packet before forwarding the packet to switching fabric 114. Pet. 25 (citing Ex. 1005 ¶¶ 30–34).

Although the destination identifier is something already contained within each packet when it is received by the line card, nothing requires the allocation be entirely done from scratch. Mere recognizing the destination identifier within each packet and taking note of it in connection with that received packet, such as by storing it, reasonably constitutes allocating the destination identifier to that packet.

Petitioner further asserts that Devi's destination identifier 215, as FMID value, is associated with one egress port in a subset of ports which have been grouped in a LAG. *Id.* at 26–27. Petitioner has not made a sufficient showing for that contention. For reasons discussed below, we find that Devi's destination identifier 215 is associated with an aggregated link, not with any port. We also agree with Patent Owner that Devi uses port information 218, which is not associated with any packet, to determine a port. Ex. 1005 ¶ 34.

Devi describes that destination information 216 associates destination identifiers 215 in packets with egress links, and that card 106 uses destination information 216 to determine an egress link or links using destination identifier 215 from a packet. *Id.* ¶ 31. Regarding a port to which a packet would be forwarded, Devi describes that after the associated egress link is determined from the destination information, card 106 uses port information 218 to determine a destination egress port 110 of aggregated egress link 116 and forwards the packet to destination egress port 110. *Id.* ¶ 34. Devi describes that port information 218 for an aggregated link 116 identifies multiple ports 110 coupled to the physical links 108 of aggregated link 116. *Id.* ¶ 32. Specifically, Devi states:

If the link is an aggregated link 116, card 106 determines a destination port 110 of aggregated link 116 using the information in port information 218 at step 320. Card 106 may select a particular destination port 110 from the multiple ports 110 of the aggregated link using random selection, load balancing, pseudo-random selection or hashing using part of the packet information, historical tracking, round robin, or any other method of selection. Card 106 then communicates the packet to the destination port 110 at step 322.

Id. ¶ 41. These manners of selecting a port once an aggregated link 116 is identified indicate that destination identifier 215 as an FMID value is not associated with any specific port.

Petitioner asserts that Devi “suggests allocating such an FMID value associated with a port to each received packet.” Pet. 27. The Petition states:

Devi discloses that each aggregated link includes a virtual address associated with the physical links of the egress ports. And Lasserre, which the ’525 Patent incorporates-by-reference, illustrates that it was generally known to allocate identifiers (e.g., VLAN ID tags) to each multicast packet. Thus, it would have been obvious to allocate Devi’s multicast ID values as FMID values and to associate each FMID with one port in the subset. A POSITA would have been motivated to do so because assigning virtual addresses based on FMID values would allow for dynamic address assignment to provide flexibility when adding or removing line cards if one card fails.

Id. at 27–28 (citing Ex. 1005 ¶¶ 27, 32, 44, claims 3, 17; Ex. 1009, 11–13, 21–22; Ex. 1003 ¶¶ 129–130).

Petitioner’s arguments are not persuasive. That each aggregated link has a virtual address associated with all of the physical links of the egress ports in the aggregated link means the virtual address is associated with all of the ports in the aggregated link collectively and not to any one port. Thus, Devi’s disclosure cited by Petitioner does not reasonably suggest the virtual address is associated with any specific port. Claims 3 and 17 of

Devi, to which Petitioner cites, each read: “wherein the aggregated link comprises a virtual address that specifies a plurality of physical links coupled to the egress ports.” Ex. 1005, 10–11. That language does not suggest separate identification or addressing of each port, and assigning the virtual address of the aggregated link as an FMID value does not result in or reasonably suggest an FMID that is associated with a specific port.

Similarly, Petitioner does not point to any disclosure in Lasserre that any information in a multicast packet is associated with a specific port. Lasserre makes no mention of associating ports with packets using VLAN ID tags. Even if each individual port has its own virtual address, which it does not, Petitioner has not adequately explained why associating a packet to a specific port would provide more flexibility than what Devi already provides by associating a packet to an aggregated egress link. Petitioner leaves much in that regard insufficiently explained.

Alternatively, Petitioner argues that Devi and Mazzola together teach allocating an FMID value to each received packet, with each “FMID value associated with a port.” Pet. 28–29. Petitioner asserts:

Alternatively, a POSITA would have found it obvious to allocate an FMID value associated with a port to each received packet. Mazzola teaches its ingress line card assigns a 58-bit value to each multicast packet based on the packet’s destination address and a VLAN “color” tag. Mazzola explains that VLAN color tag “facilitate[s] multicast data transfers within the switch at accelerated speeds and addressing capabilities.”

Based on Mazzola, *it would have been obvious to allocate one of Devi’s multicast ID values to each received packet as an FMID value and to associate that FMID with one port in the subset. A POSITA would have been motivated to do so because Mazzola teaches that assigning virtual addresses based on multicast ID values provides a quicker, more efficient address-*

to-port translation mechanism particularly suited for VLAN arrangements.

Id. (alteration by Petitioner) (citing Ex. 1007, 2:64–67, 3:13–15, 5:15–23, 6:28–31, 6:39–40, 7:56–61, 8:14–20, 10:46–50, code (57); Ex. 1003 ¶¶ 131–133). We have two difficulties with Petitioner’s proposed combination.

First, the articulated motivation to combine is overly general and is disconnected with the circumstances and requirements of Devi. Mazzola discloses an address-translation mechanism using table look-up, as discussed above in Section II.D.3. It makes no mention of and is not concerned with aggregated links, sending a packet to just one port in an aggregated link, or load distribution among the ports in an aggregated link. It makes no distinction between one port and a plurality of ports, as it discloses that the index value retrieved from a table may be associated with either one port or a plurality of ports. Ex. 1007, 9:28–31 (“In response, the LTL translates the received index to a signal that selects a port or a group of ports for receiving the data frame.”). Although Mazzola describes that each port is associated with a unique index value and a VLAN “color” identifier (Ex. 1007, 3:12–14), Mazzola also describes that “[t]he destination of any frame may be one or more ports” (*id.* at 6:16–17) and that a port can be a multicast port including a group of ports (*id.* at 6:20–24).

The “efficiency” Mazzola mentions appears to be using a look-up table for address conversion versus not using a look-up table, and has little, if anything, to do with associating an FMID with a port versus associating an FMID with a group of ports or a LAG. Further, with Devi’s disclosure of sending only one packet to a LAG and load balancing among the ports in a LAG (limitation [1.2]), it is uncertain whether table look-up as described in Mazzola would even work in that environment. Petitioner does not address

whether the table entries can be predetermined or how frequently the table entries have to be revised. The statement of increased efficiency is speculative in the context of Devi.

Second, Petitioner has over-simplified the issue with limitation [1.3] and has not adequately explained how Mazzola's address-translation mechanism would be incorporated into Devi without affecting those portions of Devi relied on by Petitioner to account for limitation [1.2]. It is not adequately explained how Devi would be modified to incorporate teachings from Mazzola. Petitioner does not adequately address whether the proposed combination would impact Devi's components and functionality that are also relied on by Petitioner. It is unclear why the purported combination would have been successful.

As discussed above, Devi does not disclose allocating to each of the received data packets an FMID value, where each FMID is associated with one of the ports among the plurality of grouped ports in a LAG group. It also does not disclose any circuitry or step to process such FMID values.

What Petitioner has proposed for limitation [1.3], as best we can understand, is to incorporate Mazzola's address-translation mechanism into Devi, and use it, rather than the circuitry and steps disclosed in Devi, to process the incoming data packets. That means in the proposed combination, there would not be, as disclosed in Devi, a mapping of each received data packet to a particular associated egress link, and then a determination of a particular port based on other criteria such as round robin distribution among the ports in a LAG group. But such disclosures are what Petitioner has relied on to account for limitation [1.2] as discussed above. The Petition does not explain how limitation [1.2] would still be met when Mazzola's address-translation mechanism is incorporated into Devi. It is

unclear how Devi's components are modified without disturbing Petitioner's accounting for limitation [1.2].

In its Reply, Petitioner argues that Devi's destination identifier 215, as an FMID, is mapped to destination information 216 and port information 218, and "Devi's port information 218 is associated with a LAG's port." Reply 22–23. On that basis, and without further explanation, Petitioner concludes that "Devi discloses or renders obvious" limitation [1.3]. *Id.* at 23.

The assertion that Devi itself discloses limitation [1.3] was raised by Petitioner for the first time in the Reply, and thus is improper. *See* 37 C.F.R. § 42.23(b) ("A reply may only respond to arguments raised in the corresponding opposition, patent owner preliminary response or patent owner response."); *Intelligent Bio-Sys., Inc. v. Illumina Cambridge Ltd.*, 821 F.3d 1359, 1369 (Fed. Cir. 2016); *Ariosa Diagnostics v. Verinata Health, Inc.*, 805 F.3d 1359, 1366–67 (Fed. Cir. 2015). This new contention is not entitled to consideration. Even if we do consider it, it is without persuasion.

For reasons already discussed above, Devi does not disclose and would not have rendered obvious limitation [1.3]. Petitioner's characterization in its Reply that Devi maps destination identifier 215 to destination information 216 and port information 218, and that port information 218 is associated with a LAG port is vague and inaccurate.

In Devi, the destination identifier is not mapped to information, either destination information or port information. Rather, Devi describes that destination information 216 associates destination identifier 215 with an egress link, and that card 106 uses destination information 216 to determine an associated egress link for the destination identifier. Ex. 1005 ¶ 31. Devi

also describes that after the associated egress link is determined, card 106 uses port information 218 to determine a port of the egress link. *Id.* ¶ 34. That means port information is used to determine a port, based on Devi's sending just one packet to a LAG and Devi's load balancing among ports in a LAG. Petitioner's argument regarding use of information to make a determination related directly or indirectly to the destination identifier as mapping the destination identifier to that information is misplaced and inaccurate. Additionally, contrary to Petitioner's assertion, as discussed above, Devi's port information is not associated with any one port. For these reasons, the separate determination of a corresponding port, based on load balancing for a LAG and sending only one packet to a LAG, subsequent to having determined a corresponding aggregated egress link, does not disclose and reasonably would not have suggested limitation [1.3].

Petitioner's articulated motivation to combine the teachings of Devi and Mazzola for limitation [1.3] also lacks rational underpinning, and the proposed combination of Devi and Mazzola lacks sufficient specificity to convey the combined structure and operation needed to adequately account for limitation [1.3].

In summary, for the foregoing reasons, Petitioner's accounting for limitation [1.3], either based on Devi alone or based on Petitioner's combination of Devi and Mazzola, is unpersuasive.

e) Conclusion

Petitioner has not proved by a preponderance of the evidence that claim 1 would have been obvious over Devi and Mazzola.

3. Independent Claim 12

Each of limitations [12.1] through [12.6] substantially corresponds to limitations [1.1] through [1.6], respectively. Petitioner does not

substantively argue claims 1 and 12 separately. Neither does Patent Owner. Our discussions above for claim 1 and the deficiencies of Petitioner's evidence and arguments for claim 1 equally apply to claim 12. Petitioner has not proved by a preponderance of the evidence that claim 12 would have been obvious over Devi and Mazzola.

4. *Dependent Claims 2–11 and 13–24*

Claims 2–11 each depend, directly or indirectly, from independent claim 1. Claims 13–24 each depend, directly or indirectly, from independent claim 12. Thus, claims 2–11 each incorporate all of the limitations of claim 1, including limitation [1.3], and claims 13–24 each incorporate all of the limitations of claim 12, including limitation [12.3].

All of the deficiencies of Petitioner's assertions discussed above in the context of claim 1 equally apply to claims 2–11, and all of the deficiencies of Petitioner's assertions regarding claim 12 as noted above equally apply to claims 13–24. Accordingly, Petitioner has not proved by a preponderance of the evidence that any of claims 2–11 and 13–24 would have been obvious over Devi and Mazzola.

III. CONCLUSION

Petitioner has not proved by a preponderance of the evidence that claims 1–24 of the '525 patent are unpatentable. The outcome for the challenged claims of the '525 patent is set forth in the table below. In summary:

Claims	35 U.S.C. §	Reference(s)/Basis	Claims Shown Unpatent- able	Claims Not shown Unpatent- able
1–24	103	Borgione, Mazzola		1–24
1–24	103	Devi, Mazzola		1–24
Overall Outcome				1–24

IV. ORDER

In consideration of the foregoing, it is hereby:

ORDERED that Petitioner has not shown by a preponderance of the evidence that claims 1–24 of the '525 patent are unpatentable;

FURTHER ORDERED that, because this is a Final Written Decision, parties to the proceeding seeking judicial review of the decision must comply with the notice and service requirements of 37 C.F.R. § 90.2.

IPR2021-01356
Patent 7,697,525 B2

FOR PETITIONER:

Joseph F. Edell
Kyle K. Tsui
FISCH SIGLER LLP
joe.edell.IPR@fischllp.com
kyle.tsui@fischllp.com

FOR PATENT OWNER:

Brad Liddle
Seth A. Lindner
Scott Breedlove
Michael Pomeroy
CARTER ARNETT PLLC
bliddle@carterarnett.com
slindner@carterarnett.com
sbreedlove@carterarnett.com
mpomeroy@carterarnett.com