

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner

IPR2023-00401
U.S. Patent No. 7,545,740

**PETITION FOR *INTER PARTES* REVIEW
UNDER 35 U.S.C. § 312 AND 37 C.F.R. § 42.104**

TABLE OF CONTENTS

PETITIONER’S EXHIBIT LIST	6
I. INTRODUCTION	8
II. GROUNDS FOR STANDING.....	8
III. NOTE.....	9
IV. TECHNOLOGY BACKGROUND.....	9
A. Link Aggregation	9
B. Multiplexing and Demultiplexing	10
V. SUMMARY OF THE ’740 PATENT	13
A. Overview of the ’740 Patent.....	13
B. Prosecution History	17
VI. LEVEL OF ORDINARY SKILL IN THE ART.....	18
VII. CLAIM CONSTRUCTION	18
A. “ <i>interface module</i> ”	19
B. “ <i>selecting, in a single computation</i> ”	19
VIII. RELIEF REQUESTED AND THE REASONS FOR THE REQUESTED RELIEF	20
IX. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE....	21
A. Challenged Claims and Statutory Grounds for Challenge	21
B. Grounds 1 & 2: Claims 1-31 are obvious under 35 U.S.C. § 103(a) over Bruckman alone or in view of Basso.	23
1. Summary of Bruckman	23

2.	Summary of Basso	27
3.	Reasons to Combine Bruckman and Basso	29
4.	Summary of Holdsworth.....	31
5.	Holdsworth shows background knowledge of a POSITA.....	33
6.	Claim 1	34
7.	Claim 2.....	47
8.	Claim 3.....	49
9.	Claim 4.....	49
10.	Claim 5.....	50
11.	Claim 6.....	51
12.	Claim 7.....	52
13.	Claim 8.....	53
14.	Claim 9.....	54
15.	Claim 10.....	55
16.	Claim 11	58
17.	Claim 12.....	63
18.	Claim 13.....	64
19.	Claim 14.....	66
20.	Claim 15.....	68
21.	Claim 16.....	69
22.	Claim 17.....	70

23.	Claim 18	72
24.	Claim 19	73
25.	Claim 20	73
26.	Claims 21-27	73
27.	Claim 28	74
28.	Claim 29	74
29.	Claim 30	75
30.	Claim 31	75
C.	Grounds 3 & 4: Claims 11 and 26 are obvious under 35 U.S.C. § 103(a) over Bruckman (alone or with Basso) in view of Holdsworth.	76
1.	Reasons to Combine Holdsworth with Bruckman.....	76
2.	Claims 11 and 26.....	79
X.	DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE.....	79
A.	Discretionary denial under the <i>Fintiv</i> factors is not appropriate.....	79
1.	No evidence regarding a stay	80
2.	Parallel proceeding trial date	80
3.	Investment in the parallel proceeding.....	81
4.	Overlapping issues with the parallel proceeding	82
5.	Identity of parties	82
6.	Other circumstances.....	82
B.	Discretionary denial under 35 U.S.C. § 325(d) is not appropriate	83

1.	<i>Becton, Dickinson</i> Factor (c).....	84
2.	<i>Becton, Dickinson</i> Factors (e) and (f):	85
C.	Discretionary denial under <i>General Plastic</i> is not appropriate.....	87
XI.	CONCLUSION.....	87
XII.	MANDATORY NOTICES	89
A.	Real Party-in-Interest	89
B.	Related Matters.....	89
C.	Lead and Back-up Counsel and Service Information	89
XIII.	CLAIMS APPENDIX	91
	CERTIFICATE OF WORD COUNT	105
	CERTIFICATE OF SERVICE	106

PETITIONER’S EXHIBIT LIST

Ex.1001	U.S. Patent No. 7,545,740 to Zelig et al.
Ex.1002	Prosecution History of U.S. 7,545,740
Ex.1003	Declaration of Dr. Henry Houh under 37 C.F.R. § 1.68
Ex.1004	<i>Curriculum Vitae</i> of Dr. Houh
Ex.1005	U.S. Patent Publication No. 2004/0228278 to Bruckman et al.
Ex.1006	U.S. Patent Publication No. 2003/0210688 to Basso et al.
Ex.1007	Digital Logic Design, by Holdsworth (“Holdsworth”) (2002)
Ex.1008	U.S. Patent No. 6,943,580 to Lewis et al. (“Lewis”)
Ex.1009	IEEE Standard Dictionary of Electrical and Electronics Terms, (1984)
Ex.1010	RFC 1042, A Standard for the Transmission of IP Datagrams over IEEE 802 Networks
Ex.1011	C++ Inside & Out, by Bruce Eckel (“Eckel”) (1993)
Ex.1012	Complaint, <i>Orckit Corporation v. Cisco Systems, Inc.</i> , Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Jul. 22, 2022)
Ex.1013	Federal Court Statistics
Ex.1014	Proposed Docket Control Order, <i>Orckit Corporation v. Cisco Systems, Inc.</i> , Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Dec. 2, 2022)
Ex.1015	Infringement Contentions, <i>Orckit Corporation v. Cisco Systems, Inc.</i> , Civil Action No. 2:22-cv-276-JRG-RSP (E.D. Tex, Nov. 3, 2022)
Ex.1016	Affidavit from the Internet Archive
Ex.1017	U.S. Patent No. 7,271,993 to Plojhar (“Plojhar”)

Ex.1018	RFC 793, Transmission Control Protocol
Ex.1019	RFC 2026, The Internet Standards Process – Revision 3
Ex.1020	U.S. Patent No. 5,960,428 to Lindsay et al. (“Lindsay”)
Ex.1021	U.S. Patent No. 5,999,538 to Haddock et al. (“Haddock”)
Ex.1022	Course Catalog for Electrical Engineering
Ex.1023	Copyright Record for Eckel (Ex.1011)
Ex.1024	U.S. Patent No. 6,473,424 to DeJager (“DeJager”)

I. INTRODUCTION

Pursuant to 35 U.S.C. §§ 311, 314(a), and 37 C.F.R. § 42.100, Cisco Systems, Inc. (“Petitioner”) respectfully requests that the Board review and cancel as unpatentable under (pre-AIA) 35 U.S.C. §103(a) claims 1-31 (the “Challenged Claims”) of U.S. Patent No. 7,545,740 (“’740 patent,” Ex.1001).

The ’740 patent relates to link aggregation, which involves joining a group of parallel physical links between two endpoints together into a single logical link. The ’740 patent, entitled “Two-way Link Aggregation,” was allowed after the Applicant amended the claims to recite that at least one of the links in a link aggregation group was “a bi-directional link operative to communicate in both an upstream direction and a downstream direction.” Ex.1002, 59.

As shown below and confirmed in the Declaration of Dr. Houh (Ex.1003), the concept of bi-directional, aggregated links was already known and would have been obvious to a POSITA. *See generally* Ex.1003. The references presented in this Petition render obvious the Challenged Claims, which should be canceled for unpatentability.

II. GROUNDS FOR STANDING

Petitioner certifies that the ’740 Patent is eligible for IPR, and that Petitioner is not barred or estopped from requesting IPR challenging the patent claims. 37 C.F.R. § 42.104(a).

III. NOTE

Petitioner cites to exhibits' original page numbers. Emphasis in quoted material has been added.

IV. TECHNOLOGY BACKGROUND

A. Link Aggregation

Link aggregation “is a technique by which a group of parallel physical links between two endpoints in a data network can be joined together into a single logical link.” Ex.1005, [0002]. “Link aggregation offers benefits of increased bandwidth, as well as increased availability, since the logical link can continue to function (possibly with reduced bandwidth) even when one of the physical links fails or is taken out of service.” Ex.1005, [0002]. Link aggregation for Ethernet networks is defined by industry standards. *See* Ex.1005, [0003].

Link aggregation includes a process for selecting which of the physical links in a Link Aggregation Group (LAG) will transmit a particular frame. This process employs “a distributor function, which distributes data frames submitted by MAC clients among the physical links in the group, and a collector function, which receives frames over the aggregated links and passes them to the appropriate MAC clients.” Ex.1005, [0003]. Industry standards describe “possible distribution algorithms that meet the requirements of the standard, while providing some measure of load balancing among the physical links in the aggregation group.”

Ex.1005, [0005]. These distribution algorithms commonly use a hash function. See Ex.1005, [0006]-[0012].

Hashing techniques used for distributing packets or data frames across link aggregation configurations were well-known before the '740 patent was filed.

Ex.1003, ¶¶14-16; *see also* Ex.1024, 5:28.

B. Multiplexing and Demultiplexing

Link aggregation uses a concept referred to as multiplexing. A multiplexer selects from a plurality of input lines for transmission over a single output line. An example 4-input multiplexer implementation is shown below:




Figure 5.1 (a) Block diagram of a 4-input multiplexer and (b) its gate implementation

Ex.1007, Fig. 5.1 (annotated); Ex.1003, ¶17.

A demultiplexer is the opposite of a multiplexer—it selects one of a plurality of output lines for transmission of data received over a single input line. An example 4-output demultiplexer implementation is shown below





Figure 5.10 Basic demultiplexer

Ex.1007, Fig. 5.10 (annotated); Ex.1003, ¶18.

Dr. Houh further explains the functionality of such multiplexers and demultiplexers. Ex.1003, ¶¶17-18.

To form larger de-multiplexers with more output lines, it was known to use two-level structures comprised of several demultiplexers. For example, as shown in the figure below, it was known to use several 1:4 demultiplexers in a two-level structure to create a 1:16 demultiplexer.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶19.

As can be seen from the figure above, the first level (on the right) includes four 1:4 demultiplexers that are controlled by lines C and D. The second level (on the left) includes a single 1:4 demultiplexer that is controlled by lines A and B. Through this two-level structure, a four-bit value (A, B, C, D) can select any of the

sixteen possible outputs (X0-X15) for a single input. It was also known to create large multiplexers using a similar two-level structure. *See* Ex.1008, 3:61-7:19; Fig. 8A; Ex.1003, ¶20.

Link aggregation uses demultiplexing to distribute data traffic from a single input among several available outputs. For example, a 1:16 demultiplexer such as the one shown above would be used to distribute data traffic across 16 different parallel output lines. A multiplexer is then used to combine incoming data from all sixteen lines to a single line. Ex.1003, ¶¶17-21.


V. SUMMARY OF THE '740 PATENT

A. Overview of the '740 Patent

The '740 patent is “directed to communication networks, and particularly to methods and systems for link aggregation in network elements.” Ex.1001, 1:5-7. The '740 patent describes a communication system 20 having a network element 32 that “interconnects a plurality of user ports 24 to a communication network 28.” Ex.1001, 4:7-8. “Network 28 may comprise a wide-area network (WAN), such as the Internet, a network internal to a particular organization (Intranet), or any other suitable communication network.” Ex.1001, 4:9-11. The user ports may connect “to a user node, such as a layer 2 or layer 3 switch.” Ex.1001, 6:38-39.

The network element 32 includes “one or more user interface modules (UIMs), such as line cards 40. Each line card is assigned to process data frames of

one or more user ports.” Ex.1001, 4:28-30. Network element 32 is shown below in Fig. 2.




Ex.1001, Fig. 2 (annotated); Ex.1003, ¶23.

The line cards 40 are connected to a user node through a set of ports 24. See Ex.1001, 4:5-11, Fig. 2. “User ports 24 forming port 64 are configured as an Ethernet LAG [link aggregation group] group, referred to as an external LAG group 68.” Ex.1001, 6:21-23. The ’740 patent claims these ports 24 as “a first group of first physical links arranged in parallel.” See e.g., Ex.1001, claim 1.

The interface modules are connected to the communication network 28 using backplane traces 56: “Backplane 52 comprises physical links, such as


backplane traces 56, typically in the form of printed circuit board (PCB) conductors.” Ex.1001, 4:34-37. The ’740 patent claims these traces 56 as “a second group of second physical links arranged in parallel.” See e.g., Ex.1001, claim 1.



Ex.1001, Fig. 2 (annotated); Ex.1003, ¶25.

The network element 32 further includes a control module 60 that, for each received packet, determines which trace and which port will transmit that packet. See Ex.1001, 3:45-53. Selection of both a port 24 from the first group and a trace 56 from the second group can be done in “single combined mapping operation that combines the two mapping operations described above.” Ex.1001, 6:63-65. The

“combined mapping comprises a single hashing operation that determines, for each such downstream frame, both the backplane trace 56 over which the frame is to be sent to one of line cards 40, and the user port 24 to be used within external LAG group 68.” Ex.1001, 6:66-7:3. This combined mapping operation includes several steps, as shown below in Fig. 4.



Ex.1001, Fig. 4.

At step 80, “[t]he method begins with control module 60 determining a hashing size parameter.” Ex.1001, 7:25-26. The hashing size parameter is based on the total number of traces multiplied by the total number of ports. *See* Ex.1001, 7:27-30, 42. Then, at step 82, a data frame is received. Ex.1001, 7:43-45. “For each downstream frame, control module 60 calculates a hashing key of the frame... by applying a suitable hashing function to the frame attributes of the downstream frame.” Ex.1001, 7:45-50. At step 86, “Module 60 divides the hashing key of the downstream frame by N_{bpow} [the hashing size parameter] and retains the modulo, or

the remainder of the division operation, as a mapping index.” Ex.1001, 7:51-53.

Accordingly, the steps shown in Fig. 4 are used to produce a single hash computation that is then used to determine both a port and a trace to transmit the packet. The ’740 patent provides an example in which “Module 60 partitions the binary representation of the mapping index into two parts having N1 and N2 bits.” Ex.1001, 7:60-61. The first set of bits N1 is used to select one of the ports: “Module 60 uses N1 bits as a user slot/port index, indicating over which user port 24 in external LAG group 68 the frame should be sent.” Ex.1001, 7:61-64. The second set of bits N2 is used to select one of the traces: “The remaining N2 bits are used as a backplane trace index, indicating over which of the backplane traces of the relevant line card the frame should be sent.” Ex.1001, 7:64-66.

As will be described in further detail below, the link aggregation concepts described and claimed in the ’740 patent are rendered obvious by the prior art. Ex.1003, ¶¶22-29.

B. Prosecution History

The ’740 patent was filed April 7, 2006. In a first non-final Office Action on October 2, 2008, the Office rejected many claims as anticipated by U.S. Patent No. 6,963,578 to Akahane. Ex.1002, 83. The Office indicated that dependent claims 4, 5, 13, 15, 19, 20, 28, and 30 contained allowable subject matter. Ex.1002, 85.

In response, the Applicant rewrote the dependent claims with allowable

subject matter in independent form. Ex.1002, 76; *see also* Ex.1002, 59-69.

Applicant also amended some claims to include limitations requiring the recited communication links to be “bi-directional.” Ex.1002, 59-69. The claims were then allowed with no statement of reasons for allowance. Ex.1002, 38.

VI. LEVEL OF ORDINARY SKILL IN THE ART

A Person of Ordinary Skill in The Art (“POSITA”) in April 2006 would have had a working knowledge of computer networking, including such commonly used technologies as Ethernet and the internet protocol / transmission control protocol (TCP/IP) suite. A POSITA would have had a bachelor’s degree in electrical engineering, or an equivalent, and two years of professional experience in the field of communication networks, including link aggregation in communication networks. Lack of professional experience can be remedied by additional education, and vice versa. Ex.1003, ¶¶7-9.

VII. CLAIM CONSTRUCTION

Claim terms in IPR are construed according to their “ordinary and customary meaning” to those of skill in the art. 37 C.F.R. § 42.100(b). *Phillips v. AWH Corp.*, 415 F.3d 1303 (Fed. Cir. 2005) (en banc). Petitioner submits that, for the purposes of this proceeding and the grounds presented herein, no claim term requires express construction. *Nidec Motor Corp. v. Zhongshan Broad Ocean Motor Co.*, 868 F.3d 1013, 1017 (Fed. Cir. 2017). For clarity, however, Petitioner notes below

example embodiments in the specification of certain terms. Ex.1003, ¶¶30-32.

A. “*interface module*”

Each independent claim of the ’740 patent recites “*one or more interface modules.*” The ’740 patent explicitly contemplates that a line card is within the scope of the term *interface module*: “Network element 32 comprises one or more user interface modules (UIMs), such as line cards 40.” Ex.1001, 4:28-29. A line card is thus an *interface module* as claimed. Ex.1003, ¶33.

B. “*selecting, in a single computation*”

As explained above at V.A, the ’740 patent describes a multi-step selecting process as shown in Fig. 4 and explained in accompanying text. This multi-step process includes determining a hash size (step 80), calculating a hash key (step 84), and determining a mapping (step 86). This multi-step process is recited in various dependent claims—including claims 8-10 which recite that the selecting process includes multiple steps, including (1) “*determining a hashing size responsively to a number of at least some of the first and second physical links,*” (2) “*applying the hashing function to the at least one of the frame attributes to produce a hashing key,*” and (3) “*calculating a modulo of a division operation of the hashing key by the hashing size.*”

It is apparent that the “*single computation*” does not refer to performing the *selecting* process in a single step, but rather to a single computational result of a

hashing function (e.g., single value) that *selects* the outcome. *See, e.g.*, Ex.1001, 5:29-31. Thus, the single value, or the final processing step that produces that single value, are examples of a “*single computation.*” Ex.1003, ¶¶34-36.

VIII. RELIEF REQUESTED AND THE REASONS FOR THE REQUESTED RELIEF

Petitioner asks that the Board institute a trial for *inter partes* review and cancel the Challenged Claims in view of the analysis below. Petitioner challenges all claims of the '740 patent because they are all asserted in co-pending litigation. A finding that the Challenged Claims are unpatentable in this proceeding will resolve the parties' dispute in the co-pending litigation and obviate any need for a trial regarding the '740 patent, substantially reducing the time and expense of litigation for all parties.

IX. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

A. Challenged Claims and Statutory Grounds for Challenge¹

Grounds	Claims	Basis
#1	1-31	35 U.S.C. § 103 (Pre-AIA) over Bruckman
#2	1-31	35 U.S.C. § 103 (Pre-AIA) over Bruckman and Basso
#3	11 and 26	35 U.S.C. § 103 (Pre-AIA) over Bruckman and Holdsworth
#4	11 and 26	35 U.S.C. § 103 (Pre-AIA) over Bruckman, Basso, and Holdsworth

U.S. Patent Publication No. 2004/0228278 to Bruckman et al. (“Bruckman”) published on November 18, 2004.

U.S. Patent Publication No. 2003/0210688 to Basso et al. (“Basso”) was published on November 13, 2003.

“Digital Logic Design,” by Brian Holdsworth and Clive Woods (“Holdsworth”) is a textbook with a copyright date of 2002 and is assigned ISBN Number 0-7506-45882. It was published in 2001 and available for sale as of at

¹ For each combination presented herein, Petitioner relies on the teachings, and not on a physical incorporation of elements. *See In re Mouttet*, 686 F.3d 1322, 1332 (Fed. Cir. 2012); *In re Etter*, 756 F.2d 852, 859 (Fed. Cir. 1985).

least August 12, 2001, as evidenced by the Internet Archive—a well-known archiving website. *See* Ex.1016. Holdsworth was also cited in a patent, showing that the book was known to POSITAs. *See* Ex.1017, 6:5-7 (“A useful discussion of logic circuits can be found in “Digital Logic Design”, B. Holdsworth and C. Woods, Newnes, 2002”). The totality of the evidence demonstrates that Holdsworth was publicly available to interested persons exercising reasonable diligence, and, therefore, was a printed publication as of 2005. *See Hulu, LLC v. Sound View Innovations, LLC*, IPR2018-01039, Paper 29, at 17-18 (P.T.A.B. Dec. 20, 2018) (precedential) (considering the “totality of the evidence” in assessing whether a reference is prior art). Ex.1003, ¶56.

Bruckman, Basso, and Holdsworth are each prior art under 35 U.S.C. § 102(b).

Petitioner also cites additional prior art as evidence of the background knowledge of a POSITA and to provide contemporaneous context to support Petitioners’ assertions regarding what a POSITA would have understood from the prior art in the grounds. *See Yeda Research v. Mylan Pharm. Inc.*, 906 F.3d 1031, 1041-1042 (Fed. Cir. 2018) (affirming the use of “supporting evidence relied upon to support the challenge”); 37 C.F.R. §42.104(b); *see also K/S HIMPP v. Hear-Wear Techs., LLC*, 751 F.3d 1362, 1365-66 (Fed. Cir. 2014); *Arendi S.A.R.L. v. Apple Inc.*, 832 F.3d 1355, 1363 (Fed. Cir. 2016). For instance, Dr. Houh and this

Petition cite to “Request for Comments” (RFC) documents published by the Internet Engineering Task Force (IETF). To the extent the Board determines that these IETF documents must qualify as prior art “printed publications” for the purposes for which they are cited, the documents do so qualify. *See, e.g.*, Ex.1019, 6, (“RFCs can be obtained from a number of Internet hosts...”), 8 (IETF documents are “readily available to a wide audience”), 26 (IETF participants “shall publicly announce...every activity” relating to the standardization process); Ex.1003, ¶38. The Board has repeatedly found IETF documents, including RFCs, to be “printed publications.” *See, e.g., Apple, Inc. v. VirnetX, Inc.*, IPR2017-00337, Paper 31, 46-47 (May 30, 2018) (RFCs are “precisely the type of documents whose main purpose is for public disclosure”); *Riot Games, Inc. v. Paltalk Holdings, Inc.*, IPR2018-00130, Paper 11, 30-33 (May 15, 2018) (RFCs are printed publications).


B. Grounds 1 & 2: Claims 1-31 are obvious under 35 U.S.C. § 103(a) over Bruckman alone or in view of Basso.

Analysis for Grounds 1 and 2 is very similar, and thus they are presented together. Ground 2 differs from Ground 1 by relying on additional obviousness teachings from Basso in limitation [1.6]. Ex.1003, ¶¶37-39.

1. Summary of Bruckman

Bruckman relates “to data communication systems, and specifically to methods and systems for link aggregation in a data communication network.”


Ex.1005, [0001]. Bruckman describes a communication system 20 having Equipment 22 that “is configured to convey packet data traffic between the customer nodes and a network (which may be a metro network, access network, or other type of core network, for example).” Ex.1005, [0047]. Equipment 22 further includes “a main switching card 32, which is connected to multiple line cards 34.” Ex.1005, [0047]. As shown in Fig. 1 below, the main switching card 32 and line cards 34 form a two-level multiplexer/demultiplexer.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶40.

Fig. 2 (shown below) provides more details of the communication system 20 shown in Fig. 1. The line cards 34 in Fig. 2 connect to the customer node via links

30. “Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30.” Ex.1005, [0056]. The line cards 34 connect to the switching core 40 and core network through backplane traces 52. Bruckman thus describes a two-level multiplexing/demultiplexing structure in which the first level includes the concentrators 50, each of which selects from one of four outputs (links 30) and the second level includes switching core 40, which selects one of a plurality of outputs (traces 52).



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶41.

When processing traffic from the customer nodes to the communication network, Bruckman’s two-level structure acts as a multiplexer: “The concentrators

multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40.” Ex.1005, [0056]. When processing traffic from the communication network to a customer node, Bruckman’s two-level structure acts as a demultiplexer by determining which trace and which port will transmit that data frame. The main card 32 includes a distributor 58, that “determines the link over which to send each frame based on information in the frame header.” Ex.1005, [0058]. To make this selection, the distributor “applies a predetermined hash function to the header information.” Ex.1005, [0058]. Bruckman provides an example of such a hash function shown in Table 1 below.

```
A unsigned short hash(unsigned char *hdr, short lagSize)
    {
        short i;
        unsigned short hash=178; // initialization value
        for (i=0; i<4; i++)        // hdr is 4 bytes length
B         hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
C     return (hash % lagSize)
    }
```

Ex.1005, Table 1 (annotated); Ex.1003, ¶42.

Like the selection process in the '740 patent, Bruckman’s hashing function includes multiple steps that produce a single hash computation. First, as shown in line A, the hashing function receives the lagSize parameter, which “is the number of active ports (available links 30) in link aggregation group.” Ex.1005, [0064]. In


line B, an operation is applied to header information (in variable *hdr*²) to create a hashing key stored in the “hash” variable. Then, in line C, a modulo of the division of the hash key and the LAG size is produced. This final, single modulo value is then used to select which link 30 will transmit the data frame. Ex.1005, [0058].

Accordingly, Bruckman shows that using a set of parallel traces and a set of parallel links/ports to transmit packets was known before the filing of the ’740 patent. Bruckman additionally shows that it was known to apply a hashing function to select one of the links. Ex.1003, ¶¶40-44.

2. Summary of Basso

Basso relates to link aggregation, particularly, “the field of packet switching networks, and more particularly to logically grouping physical ports of a network device.” Ex.1006, [0001]. Basso describes a network device, such as a router, that passes packets between the Internet and a client device. Like Bruckman, Basso describes a two-level structure, as shown in Fig. 1 below.

² “Here *hdr* is the header of the frame to be distributed.” Ex.1005, [0064].



Ex.1006, Fig. 4 (annotated); Ex.1003, ¶45.

Basso's network device 104a includes "a switch fabric coupled to a plurality of blades where each blade may comprise one or more network processors coupled to one or more ports." Ex.1006, [0009]. As shown in Fig. 4, Basso's two-level structure includes a first group of physical links—each associated with one of the ports 404—that connect the network device to a network node (in this case, Router 104b). Additionally, a second set of links connects each of the blades to the switching fabric: "[A] switch fabric 401 configured to direct the incoming packets of data to particular blades 402A-C coupled to switch fabric 401." Ex.1006, [0030].

Upon receipt of a packet, "a hash function may be performed on the source

and destination address in the packet header to generate a hash value.” Ex.1006, [0011]. Using that single hash value, “an appropriate **blade/port combination** may be identified to transmit the received packet of data.” Basso, [0012]. Basso thus uses a single hash computation to select both a first-level line (one of the ports) and second-level line (one of the blades)—a blade/port combination. Ex.1003, ¶¶45-47.

3. Reasons to Combine Bruckman and Basso

A POSITA would have found it obvious that when using a two-level demultiplexer structure such as the one described by Bruckman, a single hash computation would be used to control *both* a first-level and a second-level demultiplexer. Basso provides explicit evidence of doing so. Ex.1003, ¶48.

As an initial matter, both Bruckman and Basso are analogous art to the ’740 patent. Bruckman and Basso are directed to the same field of endeavor as the ’740 patent—network communications. *Compare* Ex.1001, 1:5-7 (“The present invention relates generally to communication networks, and particularly to methods and systems for link aggregation in network elements”) *with* Ex.1005, [0001] (“The present invention relates generally to data communication systems, and specifically to methods and systems for link aggregation in a data communication network.”) *and* Ex.1006, [0001] (“The present invention relates to the field of packet switching networks, and more particularly to logically grouping physical ports of a network device, e.g., router, into logical interfaces.”); *see also*

Ex.1003, ¶48.

Bruckman’s hashing technique produces a single hash computation that is used to determine which link 30 receives a data frame: The “distributor 58, which is responsible for distributing data frames arriving from the network **among links 30** in aggregation group 36 ... determines **the link** over which to send each frame based on information in the frame header.” Ex.1005, [0058].

Bruckman provides limited description of how the hash computation is used to select a trace. Basso provides further details regarding such two-level selections and includes explicit evidence of using a single hash computation to make a selection in both levels of a two-level multiplexer/demultiplexer structure. As explained above, Basso’s network device uses one hash value to select “an appropriate **blade/port combination** ... to transmit the received packet of data.” Ex.1006, [0012]. Basso thus uses a single hash computation to select a blade/port combination—both a first-level line (one of the ports) and second-level line (one of the blades). Accordingly, to the extent a POSITA implementing the device of Bruckman needed more information about how a single hash computation would be used to control both levels of a two-level multiplexer/demultiplexer structure, such teaching is provided by Basso.

A POSITA would have had a reasonable expectation of success because Basso describes a substantially similar two-level multiplexer/demultiplexer


structure as Bruckman. Ex.1003, ¶52. Thus, using a single hash computation to make selections in both levels of the two-level multiplexer/demultiplexer structure would have worked as well in Bruckman. A POSITA would have further found the combination beneficial as using a single hash computation (rather than multiple) would have allowed for improved operational efficiency due to the simplified design. Ex.1003, ¶52. Calculating a single hash value as opposed to multiple hash values reduces the computational load and thus allows for faster switching, which is consistent with Bruckman's goal to "ensure that sufficient bandwidth will be available on the links in the group in order to meet service guarantees." Ex.1005, [0015].

Thus, the combination of Bruckman and Basso represents the application of a known technique (Basso's single hash computation to select both levels in a two-level multiplexer/demultiplexer structure) to Bruckman's method (using a single hash computation for a two-level multiplexer/demultiplexer structure) to yield predictable results (fast and computationally efficient switching). Ex.1003, ¶¶48-53.

4. Summary of Holdsworth

Holdsworth is a textbook entitled "Digital Logic Design." Holdsworth describes a basic multiplexer and a basic demultiplexer. *See* Ex.1007, Fig. 5.1 and 5.10. Holdsworth further describes how 1:4 demultiplexers may be used in a two-

level fashion to create a 1:16 multiplexer, as shown in Fig. 5.17(a) below.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶54.

Holdsworth thus shows that when using two levels of 1:4 demultiplexers to create a 1:16 demultiplexer, a designer would use a first two-bit signal for selection at the first level and a second two-bit signal for selection at the second level. In

other words, Holdsworth shows that it was a well-known technique to use a first subset of bits for selection at the first level of a two-level demultiplexer, and a second subset of bits for selection at the second level of the two-level demultiplexer.

Holdsworth is analogous art to the '740 patent, as its engineering principles are applicable to the same field of endeavor as the '740 patent—communication networks. As explained above at IV.B, link aggregation in communication networks utilizes multiplexers, which Holdsworth describes in detail. *See generally*, Ex.1007. Because Holdsworth provides implementation details for multiplexers and demultiplexers, Holdsworth is reasonably pertinent to the technical problem of load balancing and link aggregation to which the '740 patent relates. Ex.1003, ¶¶54-56.

5. Holdsworth shows background knowledge of a POSITA

The Holdsworth textbook is “intended to cover all the material that is needed in a typical undergraduate or master’s course on Digital Logic Systems, and also to act as a reference text for graduates working in this field.” Ex.1007, preface.

Holdsworth represents the background knowledge of a POSITA because Holdsworth’s subject matter would have been part of the undergraduate education of a POSITA. Ex.1003, ¶57. For example, digital logic systems was a commonly required course in undergraduate electrical engineering degree plans. *See* Ex.1022.

6. Claim 1


[1.0] *A method for communication, comprising:*

Bruckman “relates generally to data communication systems, and specifically to **methods** and systems for link aggregation in a **data communication** network.” Ex.1005, [0001]; abstract.

Thus, because Bruckman describes methods for link aggregation in a data communication network, Bruckman renders obvious a “*method for communication*” as claimed. Ex.1003, ¶¶58-59.

[1.1] *coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel,*

Bruckman describes a communication system that includes a set of networked customer nodes 24, 26, 28 (any of which is a “*network node*”) that are connected to multiple line cards 34 (“*one or more interface modules*”) over parallel physical links 30 (“*a first group of physical links arranged in parallel*”). “In this example, central office equipment 22 communicates with customer nodes 24, 26, 28,... over physical links 30... For this purpose, equipment 22 comprises a main switching card 32, which is connected to multiple line cards 34 that serve links 30.” Ex.1005, [0047].



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶¶60.

The links 30 shown in Fig. 1 form an “aggregation group 36.” Bruckman, [0057]. “Link aggregation is a technique by which a group of **parallel physical links** between two endpoints in a data network can be joined together into a single logical link.” Ex.1005, [0002].

Thus, because Bruckman describes connecting a customer node to a set of line cards over a set of parallel links, Bruckman renders obvious “*coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel*” as claimed. Ex.1003, ¶¶60-62.

[1.2] at least one of said first physical links being a bi-directional link operative

to communicate in both an upstream direction and a downstream direction;

Bruckman's Equipment 22 transmits data over the physical links 30 in both directions. "Links 30 typically comprise **full-duplex** Ethernet links." Ex.1005, [0047]. The term full-duplex "[r]efers to a communication system or equipment capable of transmission simultaneously in two directions." Ex.1009, 376; *see also* Ex.1021, 8:11-13. Accordingly, the links 30 transmit data in both directions.


For data traffic in the downstream direction (from communication network to customer node), Bruckman's equipment 22 uses a distributor 58: "Aggregator 54 comprises **a distributor 58, which is responsible for distributing data frames arriving from the network among links 30** in aggregation group 36." Ex.1005, [0058]. In the upstream direction (from customer node to communication network), Bruckman utilizes a collector 56: "Aggregator 54 further comprises a collector 56, **which collects data frames that were received over different links 30** in group 36, and arranges the frames back into a single traffic stream." Ex.1005 [0065].

Thus, because Bruckman describes traffic flowing in both directions across full-duplex links, Bruckman renders obvious "*at least one of said first physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction*" as claimed. Ex.1003, ¶¶63-66.

[1.3] coupling each of the one or more interface modules to a communication

network using a second group of second physical links arranged in parallel,

In Bruckman's system, the line cards 34 (*"one or more interface modules"*) are coupled to the core network (*"communication network"*) with a set of parallel traces 52 (*"a second group of second physical links arranged in parallel"*). As shown in Fig. 2 below, "Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30. The concentrators multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40." Ex.1005, [0056].




Ex.1005, Fig. 2 (annotated); Ex.1003, ¶67.

Each of the traces 52 connects the switching core 40 to one of the

concentrators 50 and are arranged in parallel in Fig. 2. Traces 52 are part of a link aggregation group because Bruckman describes how “aggregation group extends over a number of concentrators 50,” and the traces 52 couple the concentrators to the switching core 40. Ex.1005, [0066]; *see also id.* [0057]. Since the traces 52 are part of a link aggregation group, Bruckman further renders obvious that the traces are “a group of parallel physical links.” Ex.1005, [0002].

As shown in Fig. 1 below, the main card 32 connects to the core network. Thus, by connecting the line cards 34 (“*interface modules*”) to the switching core 40 of the main card 32, the traces 52 connect the line cards 34 to a “*communication network*” as claimed.



Ex.1005, Fig. 1 (annotated); Ex.1003, ¶69.

Thus, because Bruckman’s line cards are connected to a core network using a set of parallel traces, Bruckman renders obvious “coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel” as claimed. Ex.1003, ¶¶67-70.

[1.4] at least one of said second physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;

As discussed at [1.2], Bruckman’s system handles traffic in both directions between the core network and the customer nodes. An aggregator distributes traffic not only to full-duplex links 30 (see [1.2]), “but also to traces 52 that connect to

multiplexers 50 that serve these links.” Ex.1005, [0057]. Since the aggregator includes both a traffic “distributor” and a traffic “collector,” it would have been obvious that traffic flows in both directions over traces 52. Ex.1005, [0058], [0065]; Ex.1003, ¶71.

Accordingly, at least one of the traces 52 (“*at least one of said second physical links*”) is “*a bi-directional link operative to communicate in both an upstream direction and a downstream direction*” as claimed. Ex.1003, ¶72.

[1.5] receiving a data frame having frame attributes sent between the communication network and the network node;

Bruckman’s main card 32 includes an aggregator 54 that receives data frames that are sent between the core network (“*communication network*”) and links 30 connected to customer nodes (a “*network node*”). The data frames include various attributes, including frame headers (“*frame attributes*”): “Typically, distributor 58 determines the link over which to send each frame based on information in the **frame header**, as described in the Background of the Invention. Ex.1005, [0058].

Thus, because Bruckman’s main card 32 receives data frames from the core network for transmission to the links 30 which connect to customer nodes, and those data frames have headers, Bruckman renders obvious “*receiving a data frame having frame attributes sent between the communication network and the*


network node” as claimed. Ex.1003, ¶¶73-74.

[1.6] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

First, Bruckman’s aggregator 54 includes a distributor 58 that “determines the link over which to send each frame.” Ex.1005, [0058]. “Preferably, distributor 58 applies a predetermined hash function to the header information [*frame attributes*].” Ex.1005, [0058].

Bruckman gives an example of a hash function to select a link in Table 1, below:

```
unsigned short hash(unsigned char *hdr, short lagSize)
{
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)      // hdr is 4 bytes length
        hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
    return (hash % lagSize)
}
```

 “single computation”

Ex.1005, Table 1, [0063] (annotated); Ex.1003, ¶76.

Bruckman’s hashing function thus selects which link 30 (“*first physical link out of the first group*”) will transmit the data frame from the core network to the customer node based on the value calculated from the “hash % lagSize” operation (“*in a single computation*”). A POSITA would have understood that the final line


of code in this hashing function returns the value of the operation “hash % lagSize.” Ex.1003, ¶77. A POSITA would have recognized that the code snippet provided in Table 1 is written using the syntax of the C or C++ programming languages. In both C and C++ languages, the % operator refers to a mathematical modulus operation, which returns the remainder from integer division. *See* Ex.1011³, 85; Ex.1020, 7:44-45; Ex.1003, ¶77. Either the modulus operation “hash % lagSize” itself or its result render obvious the claimed “*single computation.*”

Second, a POSITA would have found it obvious that Bruckman’s hashing function selects both which link 30 (“*first physical link out of the first group*”) **and** which trace 52 (“*second physical link out of the second group*”) will transmit the data frame from the core network to the customer node. Ex.1003, ¶78. Bruckman explicitly states that the hash function determines which link 30 receives the data frame: The “distributor 58, which is responsible for distributing data frames arriving from the network **among links 30** in aggregation group 36 ... determines **the link** over which to send each frame based on information in the frame header.” Bruckman, [0058].

Bruckman provides limited description of how the hash computation is used

³ Ex.1011 has a publication date of 12-18-1992 according to the copyright record. Ex.1023.

to select a trace. A POSITA would have recognized, however, that Bruckman's selection of a link 30 also determines which trace 52 will transmit the data frame. Ex.1003, ¶79. For example, referring to Fig. 2 below, if Link A is selected by the hash function, then Trace A is also selected since Trace A is the only trace that connects to link A to the switching core. Similarly, if Link B is selected by the hash function, then Trace B is also selected, since Trace B is the only trace that connects link B to the switching core. Accordingly, selection of a specific link 30 also determines a specific trace 52, e.g., whether to select Trace A or Trace B.




Ex.1005, Fig. 2 (annotated); Ex.1003, ¶79.

A POSITA would have found it obvious that selection of a link 30 also

determines which trace 52 will transmit the data frame because each link 30 is associated with only a single trace 52. Ex.1003, ¶80.

Ground 2: Alternatively, it would have been obvious to use the result of a single hash computation to select links from two groups from the combination of Bruckman and Basso. Basso describes a two-level demultiplexer that is analogous to Bruckman's two-level demultiplexer, as shown in Fig. 4 below.



Ex.1006, Fig. 4 (annotated); Ex.1003, ¶82.

Like Bruckman's links 30, Basso describes a set of physical links between blades of a router 104a to another network node (router 104B in the example of Fig. 4). Like Bruckman's traces 52, each of Basso's blades is connected to the switch fabric, which is connected to a communication network such as the Internet.

See Ex.1006, Fig. 1.

Basso describes using a single hash computation to select a blade/port combination. Upon receipt of a packet, “a hash function may be performed on the source and destination address in the packet header to generate a hash value.” Ex.1006, [0011]. Using that single hash value, Basso’s network device selects “an appropriate **blade/port combination** ... identified to transmit the received packet of data.” Ex.1006, [0012]. Basso thus uses a single hash computation to select both a first-level line (one of the ports) and second-level line (one of the blades)—a blade/port combination.

Selection of Bruckman’s blade/port combination is analogous to selection of a trace 30 and link 52 in Bruckman. For the reasons explained above at IX.B.3, a POSITA would have found it obvious for Bruckman’s hash computation to select a trace and link combination, as evidenced by Basso’s hash computation that similarly selects a blade/port combination.

Thus, Bruckman alone or together with Basso renders obvious “*selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group*” as claimed. Ex.1003, ¶¶75-84.

[1.7] *sending the data frame over the selected first and second physical links,*

Bruckman explains that its distributor “is responsible for **distributing data**

frames arriving from the network among links 30 in aggregation group 36.”

Ex.1005, [0058].

Consistent with the example above in [1.6], it would have been obvious that sending a frame over a given link (such as Link B) involves sending the data frame over the single trace (such as Trace B) that leads to that link as well. *See supra*, [1.6]; Ex.1003, ¶86.

Thus, because Bruckman describes transmitting frames over a specific link based on the result of the hash function, Bruckman renders obvious “*sending the data frame over the selected first and second physical links*” as claimed. Ex.1003, ¶¶85-87.

[1.8] *said sending comprising communicating along at least one of said bi-directional links.*

As explained above at [1.2] and [1.4], Bruckman’s device transmits data frames in both directions, and thus it would have been obvious for links 30 and traces 52 to be bi-directional. A POSITA would have thus understood that traffic passed from the Bruckman’s core network to the customer node is communicated across bi-directional links. Thus, Bruckman renders obvious “*said sending comprising communicating along at least one of said bi-directional links*” as claimed. Ex.1003, ¶88.

7. Claim 2

[2.1] *The method according to claim 1, wherein the network node comprises a user node, and*

Bruckman's customer nodes 24, 26, and 28 are “*user node[s]*” as claimed. Indeed, Bruckman uses the term customer and user interchangeably. *Compare* Ex.1005, [0013] (“**Service level agreements** between network service providers and **customers** commonly specify a certain committed bandwidth”) *with* Ex.1005, [0072] (“The connections on links 30, including any link aggregation groups, then compete for the remaining available bandwidth (typically in a weighted manner, based on the amount of excess bandwidth contracted for in the **users’ service level agreements**, as is known in the art)”).

Thus, because Bruckman's nodes 24, 26, and 28 are referred to as customer nodes, and Bruckman uses customer and user interchangeably, Bruckman renders obvious “*wherein the network node comprises a user node*” as claimed. Ex.1003, ¶¶89-90.

[2.2] *wherein sending the data frame comprises establishing a communication service between the user node and the communication network.*

First, as explained above at [1.7], Bruckman's device transmits data frames to customer nodes over the links selected from the link aggregation group.

Second, Bruckman's disclosure is directed to “**establishing** a connection with a guaranteed bandwidth for transmitting data over a logical link that includes

a plurality of parallel physical links between first and second endpoints.” Ex.1005, abstract; *see also* Ex.1005, claims 1, 11, 19, 32. Bruckman explains that establishment of the aggregated links “provides the total bandwidth guaranteed by the customer’s service level agreement” and provides “for bandwidth allocation in a link aggregation system to ensure that sufficient bandwidth will be available on the links in the group in order to meet service guarantees.” Ex.1005, [0014]-[0015].

A POSITA would have understood that the presence of an agreement means that there would have been a service established before the frames are actually sent and that sending the frames is in furtherance of that established service. Ex.1003, ¶93. Alternatively, a POSITA would have found it obvious for the data traffic sent via Bruckman’s system to include packets used to establish a common network communication service, a transmission control protocol (TCP) session. Ex.1003, ¶93 (citing Ex.1018). As of the priority date, TCP communications were one of the most common type of network communication service in use. Ex.1003, ¶93.

Thus, Bruckman’s communication of data frames across the aggregated links provides a service between the communication network and customer nodes, which renders obvious “*wherein sending the data frame comprises establishing a communication service between the user node and the communication network*” as claimed. Ex.1003, ¶¶91-94.

8. Claim 3

[3.1] *The method according to claim 1, wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled.*

First, Bruckman’s traces 52 (“*second physical links*”) are part of a “printed circuit back plane” and are thus “*backplane traces formed on a backplane*” as claimed. “Typically, main card 32 and line cards 34 are arranged in a card rack and plug into a printed circuit back plane, (not shown) which comprises traces 52.” Ex.1005, [0056].

Second, because the line cards 34 (“*one or more interface modules*”) are “arranged in a card rack and plug into [the] printed circuit back plane,” the line cards are “*coupled*” to the printed circuit back plane. Ex.1005, [0056].

Thus, Bruckman renders obvious “*wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled*” as claimed. Ex.1003, ¶¶95-97.

9. Claim 4

[4.0]-[4.5] *A method for ... second physical links,*

See [1.0]-[1.7]. Ex.1003, ¶¶98-103.

[4.6] *at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.*

First, Bruckman explains that the links 30 (“*first group of physical links*”) may be organized into link aggregation groups. “For example, an aggregation

group 36 of four physical links is defined between equipment 22 and node 24.”

Ex.1005, [0048]. “In the example shown in FIG. 2, aggregation group 36 comprises links L1 and L2, which are connected to LC1, and links L3 and L4, which are connected to LC2.” Ex.1005, [0057].

Second, Bruckman explains that “the embodiments described herein refer specifically to link aggregation in Ethernet (IEEE 802.3) networks.” Ex.1005, [0017]; *see also* [0003] (“For Ethernet networks, link aggregation is defined by Clause 43 of IEEE Standard 802.3”).

Thus, because Bruckman’s set of links 30 are Ethernet links that form a link aggregation group, Bruckman renders obvious “*at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group*” as claimed. Ex.1003, ¶¶104-06.

10. Claim 5

[5.0]-[5.5] *A method for ... and second physical links,*

See [1.0]-[1.7]. Ex.1003, ¶¶107-112.

[5.6] *coupling the network node to the one or more interface modules comprises aggregating two or more of the first physical links into an external Ethernet link aggregation (LAG) group*

See [4.6]. Ex.1003, ¶113.

[5.7] *so as to increase a data bandwidth provided to the network node.*

Bruckman explains that “[l]ink aggregation offers benefits of increased

bandwidth.” Ex.1005, [0002].

Thus, because link aggregation is used to increase bandwidth, Bruckman renders obvious using a LAG “*so as to increase a data bandwidth provided to the network node*” as claimed. Ex.1003, ¶¶114-15.

11. Claim 6

[6.1] *The method according to claim 1, wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network, and*

Bruckman’s line cards 34 (“*interface modules*”) include concentrators that “multiplex data traffic between links 30 and traces 52.” Ex.1005, [0056].

Bruckman also uses the term concentrators and multiplexers interchangeably. *See* Ex.1005, [0056] (referring to “concentrators 50”); Ex.1005, [0057] (referring to “multiplexers 50”).

Thus, because Bruckman’s concentrators/multiplexers 50 multiplex traffic from the customer nodes to the communication network, Bruckman renders obvious “*wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network*” as claimed. Ex.1003, ¶¶116-18.

[6.2] *[at least one of...] demultiplexing downstream data frames sent from the communication network to the network node.*

A POSITA would have understood that demultiplexing is the opposite of multiplexing. *See* Ex.1003, ¶119. As explained above at [1.2] and [1.4], Bruckman’s links and traces are bi-directional. *See* Ex.1005, [0058]. A POSITA would have understood that when the line card 34 receives a signal on a particular trace, and then selects one of several links with which to transmit the signal, the line card 34 is performing a demultiplexing function. Ex.1003, ¶119. In other words, while traffic in the upstream direction (customer node to communication network) is multiplexed by the line cards 34, downstream traffic (communication network to customer node) is demultiplexed by the line cards 34.

Thus, because Bruckman’s links are bi-directional and traffic in the downstream direction is demultiplexed by the line cards, Bruckman renders obvious “*demultiplexing downstream data frames sent from the communication network to the network node*” as claimed. Ex.1003, ¶120.

12. Claim 7

[7.1] *The method according to claim 1, wherein selecting the first and second physical links comprises balancing a frame data rate among at least some of the first and second physical links.*

Bruckman’s hashing technique for selecting a link “distributes traffic in an approximately uniform manner across the entire set of possible hash values.” Ex.1005, [0061]. By distributing the load equally among possible hash values (which correspond to specific links), Bruckman’s hashing technique is designed to

balance data among the links. Moreover, Bruckman incorporates by reference known algorithms for load balancing data traffic among links. *See* Ex.1005, [0005].

Thus, because Bruckman’s hashing technique seeks to distribute the load across the links in a uniform manner, Bruckman renders obvious “*wherein selecting the first and second physical links comprises balancing a frame data rate among at least some of the first and second physical links*” as claimed. Ex.1003, ¶¶121-22.

13. Claim 8

[8.1] *The method according to claim 1, wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes.*

As explained above at [1.6], Bruckman’s selection process includes applying a hashing function to the header. Bruckman’s hashing function is shown in Table 1 below:

hashing
function
“mapping
function”



DISTRIBUTOR HASH FUNCTION

```
unsigned short hash(unsigned char *hdr, short lagSize)
{
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)        // hdr is 4 bytes length
        hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
    return (hash % lagSize)
}
```

Ex.1005, Table 1 (annotated); Ex.1003, ¶123.

Bruckman’s hashing function is a mapping function because it maps a particular data frame to a specific trace 52 and link 30 through the system. Further, claim 9 explicitly recites that a hashing function is within the scope of a “*mapping function*.”

Thus, because Bruckman’s selecting process involves a hashing function, Bruckman renders obvious “*wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes*” as claimed. Ex.1003, ¶¶123-25.

14. Claim 9

[9.1] *The method according to claim 8, wherein applying the mapping function comprises applying a hashing function.*

As explained above at [8.1], Bruckman’s hash function is a mapping function. Thus, Bruckman renders obvious “*wherein applying the mapping function comprises applying a hashing function*” as claimed. Ex.1003, ¶¶126-27.

15. Claim 10

[10.1] *The method according to claim 9, wherein applying the hashing function comprises determining a hashing size responsively to a number of at least some of the first and second physical links,*

As shown in line A, Bruckman's hash function receives two parameters: `hdr` and `lagSize`. Bruckman explains that “lagsize is the number of active ports (available links 30) in link aggregation group.” Ex.1005, [0064].

DISTRIBUTOR HASH FUNCTION

```
A unsigned short hash(unsigned char *hdr, short lagSize)
{
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)      // hdr is 4 bytes length
        hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
    return (hash % lagSize)
}
```

“determining a hashing size”

Ex.1005, Table 1 (annotated); Ex.1003, ¶128.

Accordingly, by receiving the `lagSize` parameter, which is the number of active ports (available links 30) in the link aggregation group, Bruckman's hashing technique determines “a hashing size responsively to a number of at least some of the first and second physical links.” Thus, Bruckman renders this limitation obvious. Ex.1003, ¶129.

[10.2] *applying the hashing function to the at least one of the frame attributes to produce a hashing key,*

As shown on line A of the figure below, Bruckman’s hashing function receives the “hdr” parameter (“*frame attributes*”). “Here hdr is the header of the frame to be distributed.” Bruckman, [0064]. Then, in Line B, the hdr value is used to produce a “*hashing key*” represented by the variable “hash.”

DISTRIBUTOR HASH FUNCTION

```
A unsigned short hash(unsigned char *hdr, short lagSize)
{
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)        // hdr is 4 bytes length
B     hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
    return (hash % lagSize)
}
```

variable
representing
“*frame
attributes*”

“*hashing key*”

Ex.1005, Table 1 (annotated); Ex.1003, ¶130.

The “hash” variable represents a *hashing key* because it is used in a modulus operation (%) to produce the final hash result. *See supra* [1.6] & *infra* [10.3]. A POSITA would have been familiar with programming terminology and recognized that in many programming languages, including the code shown in Table 1, the % symbol represents the modulus operation. *See* Ex.1011, 85; Ex.1003, ¶131.

Thus, because Bruckman’s hashing function uses the header information to produce a hashing key, Bruckman renders obvious “*applying the hashing function to the at least one of the frame attributes to produce a hashing key*” as claimed.

Ex..1003, ¶132.

[10.3] *calculating a modulo of a division operation of the hashing key by the hashing size, and*

As shown in the annotated Bruckman hashing algorithm below, Line C includes the operation “hash % lagSize.” As known to POSITAs, the % operator represents the modulus (remainder) operation used to calculate the remainder from integer division. *See* Ex.1011, 85; Ex.1020, 7:44-45; *supra* [1.6]. A POSITA would have thus understood that the operation in Line C calculates the remainder (“*modulo*”) that results from dividing the value of “hash” (“*hashing key*”) by the value of “lagSize” (“*hashing size*”). Ex.1003, ¶133.

DISTRIBUTOR HASH FUNCTION

```
A unsigned short hash(unsigned char *hdr, short lagSize)
{
    short i;
    unsigned short hash=178; // initialization value
    for (i=0; i<4; i++)      // hdr is 4 bytes length
B     hash = (hash<<2) + hash + (*hdr>>(i*8) & 0xFF);
C     return (hash % lagSize)
}
```

“hashing key” *“hashing size”*

Bruckman, Table 1.

Bruckman’s hashing function thus calculates and returns a “*modulo of a division operation.*” Where there are 16 different links (as in Bruckman’s example), the value of the lagSize parameter (“*hashing size*”) would be 16. Ex.1003, ¶134.

Thus, Bruckman's hashing function, which returns "hash % lagSize," renders obvious "*calculating a modulo of a division operation of the hashing key by the hashing size*" as claimed. Ex.1003, ¶135.

[10.4] *selecting the first and second physical links responsively to the modulo.*


For the reasons explained above at [1.6], Bruckman's application of a hashing function teaches selecting first and second physical links responsive to the single hash result. Bruckman's hashing function returns a modulo, as shown in [10.3] by the line "return (hash % lagSize)." The modulus operation, or its resulting value (the modulo), is the single computation that is used to select a trace and a link as described above at [1.6]. Thus, Bruckman renders obvious "*selecting the first and second physical links responsively to the modulo*" as claimed. Ex.1003, ¶136.

16. Claim 11

[11.1] *The method according to claim 10, wherein selecting the first and second physical links responsively to the modulo comprises selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.*

This claim limitation merely recites an obvious aspect of digital logic design. As explained above at IX.B.1, Bruckman's equipment 22 is a two-level structure that resembles a two-level demultiplexer. Bruckman's concentrators 50 act as first-level demultiplexers as each concentrator 50 selects one of four links 30


switching core (second-level demultiplexer) shown below in green. Equipment 22 also includes control lines from the controller 42 to each of the concentrators (first-level demultiplexers) shown below in blue.



Bruckman, Fig. 2 (annotated); Ex.1003, ¶138.

Bruckman's description of these control lines is limited. Because each of the demultiplexing devices (switching core 40 and concentrators 50) in Fig.2 has one input and four outputs, a POSITA would have found it obvious to control each with a two-bit command signal. Ex.1003, ¶139. A two-bit command signal would efficiently be able to specify one of four operating states, each corresponding to the

four possible outputs. Ex.1003, ¶139. Thus, in Bruckman's example, a two-bit command signal would be used to control the switching core (second-level demultiplexer) to select a trace 52. A different two-bit signal would be used to control the concentrators (first-level demultiplexers) to select a link 30, consistent with known techniques for controlling two-level demultiplexers. Such techniques, illustrated below, would have been in the background knowledge of a POSITA. *See supra* IX.B.5; Ex.1003, ¶139.



Ex.1007, Fig. 5.17(a) (annotated); Ex.1003, ¶139.

As explained above at [1.6], Bruckman's hashing function produces a single hash value that is used to select both a line from the first level and a line from the second level. In Bruckman's example where there are 16 possible links, the lagSize parameter is 16. The operation "hash % 16" would produce a value ("modulo")

within the range of 0 and 15. Ex.1003, ¶140. The sixteen different values between 0 and 15 are represented in binary as a four-bit value (“*binary representation of the modulo*”).

Consistent with a POSITA’s background knowledge of using two two-bit control signals to control a two-level demultiplexer (formed from 1:4 demultiplexers), two of the four hash bits would be used to control the first level and the remaining two of the four hash bits would be used to control the second level. *See* Ex.1007, Fig. 5.17(a); Ex.1003, ¶141. Thus, a POSITA would have found it obvious for Bruckman’s control signals (represented by dotted lines in Fig. 2) to be two-bit signals formed from different bits of the hash value. Ex.1003, ¶141.

Thus, because Bruckman illustrates control signals to both levels of a two-level demultiplexer, and a POSITA would have known the technique for controlling such a structure involves subsets of bits for control of different levels, Bruckman renders obvious “*selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo*” as claimed. Ex.1003, ¶¶137-42.

17. Claim 12

[12.1] *The method according to claim 1, wherein the at least one of the frame attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP*

address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

As explained above at [1.5], Bruckman's equipment 22 receives a data frame with frame attributes. As noted in Bruckman's Background section, it was known to use "information carried in each Ethernet frame in order to make [a] decision as to the physical port to which the frame should be sent.... The information used to assign conversations to ports could thus include one or more of the following pieces of information: **a) Source MAC address b) Destination MAC address.**" Ex.1005, [0005]-[0011]. Moreover, POSITAs would have recognized that the list of options recited in claim 12 are well-known data fields and that are commonly part of a layer 2 data frame. *See e.g.*, Ex.1010; Ex.1003, ¶143.

Thus, because Bruckman's equipment receives a data frame having at least a source MAC address and a destination MAC address, Bruckman renders this limitation obvious. Ex.1003, ¶¶143-45.

18. Claim 13

[13.0]-[13.5] *A method for ... and second physical links,*

See [1.0]-[1.7]. Ex.1003, ¶¶146-51.

[13.6] *coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network comprising specifying bandwidth requirements comprising at least one of a committed information rate (CIR), a peak information rate (PIR) and an excess information rate (EIR) of a communication service provided by the*

communication network to the network node, and

First, as explained above at [1.1] and [1.3], Bruckman renders obvious “*coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network.*”

Second, Bruckman explains that service level agreements between network service providers (including a “*communication network*”) and customers (including a “*customer node*”) specify bandwidth in various terms. “In general, the bandwidth guaranteed by a service provider, referred to as the peak information rate (PIR), may include either CIR, or EIR, or both CIR and EIR (in which case $PIR=CIR+EIR$).” Ex.1005, [0013].

Thus, because Bruckman describes that providing guaranteed bandwidth in terms of either CIR, or EIR, or both CIR and EIR (in which case $PIR=CIR+EIR$), Bruckman renders this limitation obvious. Ex.1003, ¶¶152-54.

[13.7] allocating a bandwidth for the communication service over the first and second physical links responsively to the bandwidth requirements.

Bruckman describes allocating bandwidth in a manner so as to exceed the guaranteed bandwidth (“*responsively to the bandwidth requirements*”). See Ex.1005, [0030]. Bruckman describes “a method for establishing a connection with a guaranteed bandwidth for transmitting data between first and second endpoints,” Ex.1005, [0027]; *see also* [0016], [0030].

Accordingly, establishing the communication sessions in Bruckman are for the purpose of providing the guaranteed bandwidth and is thus “*responsively to the bandwidth requirements.*”

Thus, because Bruckman allocates bandwidth to exceed the guaranteed bandwidth, Bruckman renders this limitation obvious. Ex.1003, ¶¶155-57.

19. Claim 14

[14.0] *A method for connecting user ports to a communication network, comprising:*

As explained above at [1.0], [1.1], and [1.3], Bruckman describes equipment 22 that connects a plurality of links 30 (which may correspond to the claimed “*ports*”) to a communication network. Alternatively, Bruckman’s links are associated with ports. Ex.1005, [0056] (“multiple ports that serve respective links 30”). The ports associated with links 30 correspond to “*user ports*” because they connect to user nodes, as explained in [2.2]. Thus, Bruckman renders this limitation obvious. Ex.1003, ¶158.


[14.1] *coupling the user ports to one or more user interface modules;*

Consistent with the discussion at [1.1] and [14.0], Bruckman renders obvious “*coupling the user ports to one or more user interface modules*” as claimed. Ex.1003, ¶159.

[14.2] *coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel,*

First, consistent with the discussion at [1.3], Bruckman renders obvious “coupling each user interface module to the communication network via a backplane” as claimed. *See also* [3.2] (explaining that Bruckman’s traces 52 are backplane traces).

Second, as shown in Fig. 2 below, each of Bruckman’s two line cards 34 (“interface modules”) includes at least two traces 52 that connect that line card to the communication network.



Bruckman, Fig. 2 (annotated); Ex.1003, ¶161.

Thus, Bruckman renders obvious “coupling each user interface module to

the communication network via a backplane using two or more backplane traces arranged in parallel” as claimed. Ex.1003, ¶¶160-62.

[14.3] *at least one of said backplane traces being bi-directional and operative to communicate in both an upstream direction and a downstream direction;*

Consistent with the discussion above at [1.4], Bruckman renders this limitation obvious. Ex.1003, ¶163.

[14.4] *receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;*

Consistent with the discussion above at [1.6], Bruckman renders this limitation obvious. Ex.1003, ¶164.

[14.5] *for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces; and*

Consistent with the discussion above at [1.6], Bruckman renders this limitation obvious. *See also* [3.1] (explaining that Bruckman’s traces 52 are backplane traces). Ex.1003, ¶165.

[14.6] *sending the data frame over the selected backplane trace; said sending comprising communicating along said at least one of said backplane traces.*

Consistent with the discussion above at [1.7], Bruckman renders this limitation obvious. *See also* [3.1] (explaining that Bruckman’s traces 52 are backplane traces). Ex.1003, ¶166.

20. Claim 15

[15.0]-[15.4] *A method for ... backplane traces; and*

See [14.0]-[14.5]. Ex.1003, ¶¶167-71.

[15.5] *sending the data frame over the selected backplane trace, at least some of the backplane traces being aggregated into an Ethernet link aggregation (LAG) group.*

First, consistent with the discussion above at [1.7], Bruckman renders obvious “*sending the data frame over the selected backplane trace*” as claimed. See also [3.1] (explaining that Bruckman’s traces 52 are backplane traces).

Second, Bruckman explains that the traces 52 are part of a link aggregation group: “As a result of spreading group 36 over two (or more) line cards, **the link aggregation function applies** not only to links 30 in group 36 but also **to traces 52 that connect to multiplexers 50** that serve these links.” Ex.1005, [0057]. As also explained above at [4.6], Bruckman’s link aggregation groups may be Ethernet link aggregation groups.

Thus, because Bruckman’s equipment 22 sends data over selected traces 52 and links 30, and the traces 52 are part of an Ethernet link aggregation group, Bruckman renders this limitation obvious. Ex.1003, ¶¶172-74.

21. Claim 16

[16.1] *The method according to claim 14, wherein selecting the backplane trace comprises applying a hashing function to the at least one of the frame attributes.*

See claims 8 and 9. Ex.1003, ¶175.

22. Claim 17

[17.0] *Apparatus for connecting a network node with a communication network, comprising:*

As explained above at [1.1] and [1.3], Bruckman describes Equipment 22 (“*apparatus*”) that connects customer nodes (a “*network node*”) with a communication network. Thus, Bruckman renders this limitation obvious. Ex.1003, ¶176.

[17.1] *one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network,*

First, Bruckman’s Equipment 22 includes a plurality of line cards 34 (“*one or more interface modules*”). See Ex.1005, [0049].

Second, Bruckman’s line cards process data frames between a customer node (“*network node*”) and the communication network. “Each line card 34 comprises one or more concentrators 50, which comprise multiple ports that serve respective links 30. The concentrators multiplex data traffic between links 30 and traces 52, which connect the concentrators to switching core 40.” Ex.1005, [0056].

Third, consistent with the discussion at [1.5], Bruckman’s data frames include frame attributes.

Thus, because Bruckman’s equipment includes a plurality of line cards that multiplex traffic from a customer node to a communication network, and the traffic

includes data frames with frame attributes, Bruckman renders obvious “*one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network*” as claimed. Ex.1003, ¶¶177-80.

[17.2] *at least one of said interface modules being operative to communicate in both an upstream direction and a downstream direction;*

Consistent with the discussion above at [1.2] and [1.4], Bruckman’s Equipment processes traffic in both upstream and downstream directions. Accordingly, the line cards (“*interface modules*”) are “*operative to communicate in both an upstream direction and a downstream direction.*” Thus, Bruckman renders this limitation obvious. Ex.1003, ¶181.

[17.3] *a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;*

Consistent with the discussion above at [1.1], Bruckman renders this limitation obvious. Ex.1003, ¶182.

[17.4] *a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and*

Consistent with the discussion above at [1.3], Bruckman renders this limitation obvious. Ex.1003, ¶183.

[17.5] *a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which*

to send the data frame;

First, Bruckman's Equipment 22 includes a controller 42 ("*control module*"): "The operation of switch 40 is managed by a controller 42, typically an embedded microprocessor with suitable software for carrying out the functions described herein." Ex.1005, [0049]. Bruckman further describes additional software components running on controller 42, including "Connection Admission Control entity (CAC) 44," "traffic manager 46," and "Aggregator 54." Ex.1005, [0050], [0056]-[0057].

Second, consistent with the discussion above at [1.6], Bruckman's controller 42 (including aggregator 54) is "*arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame*" as claimed. Thus, Bruckman (alone or in combination with Basso) renders this limitation obvious. Ex.1003, ¶¶184-85.

[17.6] ...being bi-directional links....

See [1.2] and [1.4]. Ex.1003, ¶186.

23. Claim 18

[18.1] *The apparatus according to claim 17, and comprising a backplane to which the one or more interface modules are coupled, wherein the second physical links comprise backplane traces formed on the backplane.*

Consistent with the discussion above at [3.1], Bruckman renders this limitation obvious. Ex.1003, ¶187.

24. Claim 19

[19.0]-[19.4] *Apparatus for ... send the data frame,*

See [17.0]-[17.5]. Ex.1003, ¶¶188-92.

[19.5] *at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.*

See [4.6]. Ex.1003, ¶193.

25. Claim 20

[20.0]-[20.4] *Apparatus for ... send the data frame,*

See [17.0]-[17.5]. Ex.1003, ¶¶194-98.

[20.5] *...links being aggregated into an external Ethernet link aggregation (LAG) group....*

See [5.6] and [5.7]. Ex.1003, ¶199.

26. Claims 21-27

Claims 21-27 depend (directly or indirectly) from apparatus claim 17 and recite in apparatus form substantially the same subject matter limitations as claims 6-12 (respectively), which depend from method claim 1. Accordingly, claims 21-27 are obvious for the same reasons discussed previously for claims 6-12, respectively. Ex.1003, ¶¶200-10.

27. Claim 28

[28.0]-[28.4] *Apparatus for ... send the data frame,*

See [17.0]-[17.5]. Ex.1003, ¶¶211-15.

[28.5] *the communication network being arranged to provide a communication service...*

See [13.6]. Ex.1003, ¶216

[28.6] *the first and second groups of physical links being dimensioned to provide an allocated bandwidth for the communication service responsively to the bandwidth requirements.*

Consistent with the discussion above at [13.7], Bruckman renders this limitation obvious. A POSITA would have understood that if the links can provide the allocated bandwidth, they are thus “*dimensioned*” to provide such bandwidth. The links would otherwise not be able to provide the required bandwidth. Ex.1003, ¶217.

28. Claim 29

[29.0] *Apparatus for ... modules being bi-directional...;*

See [17.0]-[17.2]. Ex.1003, ¶¶218-20.

[29.3] *a backplane having the one or more user interface modules coupled thereto and comprising a plurality of backplane traces arranged in parallel so as to transfer the data frames between the one or more user interface modules and the communication network,*

Consistent with the discussion above at [3.1] and [18.1], Bruckman renders this limitation obvious. Ex.1003, ¶221.

[29.4] *at least one of said backplane traces being bi-directional...; and*

Consistent with the discussion above at [1.2] and [1.4], Bruckman renders this limitation obvious. Ex.1003, ¶222.

[29.5] *a control module,....*

See [19.4]. Ex.1003, ¶223.

29. Claim 30

[30.0] *Apparatus for ... comprising:*

See [17.0] and [14.0]. Ex.1003, ¶224.

[30.1] *one or more ...;*

Consistent with the discussion above at [17.1], Bruckman renders this limitation obvious. Ex.1003, ¶225.

[30.2] *a backplane having the one or more user interface modules...;*

See [29.3]. Ex.1003, ¶226.

[30.4] *a control module,...;*

See [19.4]. Ex.1003, ¶227.

[30.5] *at least some of the backplane traces are aggregated into an Ethernet link aggregation (LAG) group.*

See [15.5]. Ex.1003, ¶228.

30. Claim 31

[31.1] *The apparatus ... backplane trace.*

See claims 8 and 9. Ex.1003, ¶229.

C. Grounds 3 & 4: Claims 11 and 26 are obvious under 35 U.S.C. § 103(a) over Bruckman (alone or with Basso) in view of Holdsworth.


Grounds 1 and 2 show how Bruckman renders obvious limitations related to using different subsets of bits to select different levels in a multi-level demultiplexer. Grounds 1 and 2 further rely on the Holdsworth textbook as evidence of the background knowledge of a POSITA. To the extent Patent Owner argues that such digital logic design would not have been within the background knowledge of a POSITA, Grounds 3 and 4 are presented to show that combining the disclosures of Holdsworth with the other prior art would have been obvious to a POSITA. Ex.1003, ¶230.

1. Reasons to Combine Holdsworth with Bruckman

A POSITA would have found it obvious to use basic digital logic design to implement the control signals for Bruckman's two-level demultiplexer. For example, a POSITA would have found it obvious to use a first subset of bits for a first-level demultiplexer, and a second subset of bits for the second-level demultiplexers, as evidenced by Holdsworth. Ex.1003, ¶231.

As explained above, Bruckman's Equipment 22 uses a two-level demultiplexer structure for distributing incoming packets from the communication network among the links 30. Bruckman's depiction of two-level demultiplexer includes dotted lines from the controller to the first-level and second-level

demultiplexers, as shown in Fig. 2 below.



Ex.1005, Fig. 2 (annotated); Ex.1003, ¶232.

A POSITA would have recognized that these dotted lines represent command signals for controlling each of the demultiplexers (switching core 40 and concentrators 50). Ex.1003, ¶233. Bruckman illustrates these control lines but provides few implementation details about the command signals used on the control lines, indicating that such implementation would have been within the knowledge and skill of a POSITA. Ex.1003, ¶233. Accordingly, a POSITA looking at the teachings of Bruckman would have looked to known implementation

techniques for the command signals of Bruckman's two-level demultiplexer.

Ex.1003, ¶233. Bruckman explains that "at least some of the functions of the aggregator may be carried out by hard-wired logic or by a programmable logic component, such as a gate array." Ex.1005, [0057]. Accordingly, a POSITA would have looked to known logic design techniques. Ex.1003, ¶233.

Holdsworth represents basic logic design that is applicable to Bruckman's two-level demultiplexer structure. As stated in the preface, Holdsworth is intended to "act as a reference text for graduates working in this field." Holdsworth, preface. A POSITA would have had a reasonable expectation of success when using Holdsworth's technique because Holdsworth provides an example that is similar to that of Bruckman's example—a plurality of 1:4 demultiplexers that form a 1:16 demultiplexer. A POSITA would have therefore expected that using a two-bit command signal to control a first-level demultiplexer, and a second two-bit command signal to control second-level demultiplexers would have worked as well in Bruckman's two-level demultiplexer. Ex.1003, ¶234.

Given that Bruckman's hash function would produce a hash value in the range of zero to fifteen (and thus be readily represented as a four-bit value), a POSITA would have found it obvious to use two of the four bits to control the second-level demultiplexer (i.e., the switching core) and the other two of the four bits to control the first-level demultiplexers (i.e., the concentrators). Ex.1003, ¶235.

Regarding Ground 4, application of Holdsworth technique is also consistent with Basso's technique of using the single hash computation to select both a second-level line (one of the traces 52) and a first-level line (one of the links 30). This would have yielded the predictable result of effectively and efficiently controlling both levels of the two-level demultiplexer using the four-bit hash computation described by Bruckman.

Thus, the combination of Holdsworth with Bruckman represents application of a known technique (Holdsworth's different bit subsets for different levels of a two-level demultiplexer) to a known method (Bruckman's two-level demultiplexer) to yield predictable results (controlling both levels with a single hash computation). Ex.1003, ¶¶231-36.

2. Claims 11 and 26

The claim analysis for this ground is substantially the same as the analysis provided in Grounds 1 and 2. *See supra* IX.B.16, IX.B.26; Ex.1003, ¶237.

X. DISCRETIONARY DENIAL WOULD BE INAPPROPRIATE

A. Discretionary denial under the *Fintiv* factors is not appropriate

The six factors considered for § 314 denial strongly favor institution. *See Apple Inc. v. Fintiv, Inc.*, IPR2020-00019, Paper 11 (PTAB Mar. 20, 2020) (precedential).

1. No evidence regarding a stay

No motion to stay has been filed, so the Board should not infer the outcome of such a motion. *Sand Revolution II LLC v. Continental Intermodal Group – Trucking LLC*, IPR2019-01393, Paper 24 at 7 (PTAB June 16, 2020) (informative); *see also Dish Network L.L.C. v. Broadband iTV, Inc.*, IPR2020-01359, Paper 15 (Feb. 12, 2021) (“It would be improper to speculate, at this stage, what the Texas court might do regarding a motion to stay...”). Thus, this factor is neutral on discretionary denial.

2. Parallel proceeding trial date

This factor weighs strongly against discretionary denial because the projected trial date—based on median time-to-trial statistics—is in August 2024, after the Board’s Final Written Decision is expected in July 2024.⁴ While trial is currently proposed for March 4, 2024 (Ex.1014), the Board recognizes “that scheduled trial dates are unreliable and often change.” *See* Director’s June 21, 2022 Memorandum on Discretionary Denials (“Memo”), 8. Accordingly, the Board now uses median time-to-trial statistics in the relevant venue to determine a projected trial date for Fintiv purposes. Memo, 9.

⁴ July 2024 is 18 months after January 2023, when Petitioner expects a notice of accorded filing date for this petition.

The co-pending district court case was filed in the Eastern District of Texas on July 22, 2022. *See* Ex.1012. The current median time-to-trial in the Eastern District of Texas is 24.5 months. Ex.1013, 5. Accordingly, the projected trial date for *Fintiv* purposes is August of 2024—approximately 24 months after July 2022, and after the Board’s Final Written Decision is expected in July of 2024. Because the projected trial date is “around the same time or after” the Board’s expected final written decision, this factor weighs in favor of institution.

3. Investment in the parallel proceeding

The co-pending litigation is in its very early stages, and the investment in it has been minimal. The parties have not exchanged preliminary positions on claim construction or invalidity, expert discovery has not begun, and the parties have not exchanged their first set of discovery requests. *See PEAG LLC v. Varta Microbattery GmbH*, IPR2020-01214, Paper 8, 17 (Jan. 6, 2021). Further, the Markman hearing is not scheduled until September of 2023, two months after an expected institution decision by the Board. Ex.1014, 3.

Moreover, Petitioner only learned which claims were being asserted on November 3, 2022. *See* Ex.1015. Under *Fintiv*, Petitioner’s prompt filing “weigh[s] against exercising the authority to deny institution.” *Fintiv*, Paper 11 at 11 (“If the evidence shows that the petitioner filed the petition expeditiously, such as promptly after becoming aware of the claims being asserted, this fact has

weighed against exercising the authority to deny institution under NHK”). This factor favors institution.

4. Overlapping issues with the parallel proceeding

There is no present overlap of prior art issues due to the early stage of district court litigation. For example, Petitioner has not served its preliminary invalidity contentions in the district court proceeding. Consequently, this factor favors institution.

5. Identity of parties

Petitioner is a defendant in the litigation. That is true of most Petitioners in IPR proceedings. Accordingly, this factor should not be a basis for denying institution.

6. Other circumstances

The prior art presented in this Petition renders the Challenged Claims unpatentable as obvious. “[T]he PTAB will not deny institution of an IPR or PGR under Fintiv (i) when a petition presents compelling evidence of unpatentability.” Memo, 2. “Compelling, meritorious challenges are those in which the evidence, if unrebutted in trial, would plainly lead to a conclusion that one or more claims are unpatentable by a preponderance of the evidence.” Memo, 4. Here, the petition plainly shows that the ’740 patent claims no more than known concepts of bi-

directional link aggregation. The evidence of unpatentability is compelling, and thus the PTAB should not deny institution under *Fintiv*.

As such, because the *Fintiv* factors are either neutral or weigh against discretionary denial, and institution should not be denied on discretionary factors.

B. Discretionary denial under 35 U.S.C. § 325(d) is not appropriate

Denial under § 325(d) is not warranted because the challenges presented in this petition are neither cumulative nor redundant to the prosecution of the '740 Patent. To begin with, this petition relies in part on Basso and Holdsworth, neither of which was considered in prosecution. And while Bruckman was disclosed to the examiner after most claims had already been allowed, the Examiner erred by overlooking Bruckman's highly relevant teachings. *See* Ex.1002, 44. This was "material" error because the overlooked teachings of Bruckman render obvious the limitations that the Applicant had added to gain allowance. Discretionary denial is therefore not appropriate. *Advanced Bionics, LLC v. MED-EL Elektromedizinische Geräte GmbH*, IPR2019-01469, Paper 6 (PTAB Feb. 13, 2020) (precedential) ("*Advanced Bionics*").

Under the first prong of the *Advanced Bionics* framework, Petitioner acknowledges that Bruckman was "previously presented to the Office." *Advanced Bionics* at 7-8. Under the second prong, however, the evidence shows that the Examiner "erred in a manner material to the patentability of challenged claims."

Id. The second prong is guided by *Becton, Dickinson* factors (c), (e), and (f):

- (c) the extent to which the asserted art was evaluated during examination, including whether the prior art was the basis for rejection;
- (e) whether petitioner has pointed out sufficiently how the examiner erred in its evaluation of the asserted prior art; and
- (f) the extent to which additional evidence and facts presented in the petition warrant reconsideration of the prior art or arguments.

Id. at 9-10, n. 10 (citing *Becton, Dickinson & Co. v. B. Braun Melsungen AG*, IPR2017-01586, Paper 8 at 17-18 (Dec. 15, 2017)). These factors weigh against exercising discretion.

1. *Becton, Dickinson* Factor (c)

Bruckman was never cited in a rejection. Bruckman was introduced to the record after the Examiner had allowed several dependent claims, and after the Applicant had amended various independent claims to gain allowance. Just days after the examiner signed an IDS listing Bruckman, the Office mailed a Notice of Allowance for all claims with not statement of reasons for allowance. Ex.1002, 40, 44. Thus, the examiner made a “material error” by overlooking the disclosure of Bruckman that teaches the very concepts related to the bi-directional nature of the links that were added to the claims to gain allowance. Bruckman also teaches the very concepts that were in the dependent claims that examiner initially allowed. The sparse references to Bruckman in the file history indicate that Bruckman

received—at most—only insufficient, cursory attention during prosecution. A reference merely being of record is not sufficient reason to exercise discretion. *See Navistar, Inc. v. Fatigue Fracture Tech., LLC*, IPR2018-00853, Paper 13 at 17 (Sept. 12, 2018) (“Under [Becton factors] (c), (d), and (f) . . . the fact that [references] were of record, but not applied in any rejection by the Examiner . . . provides little impetus for us to exercise our discretion to deny institution under § 325(d).”)

Accordingly, because the Office overlooked the teachings of Bruckman when allowing the ’740 patent, and because Basso and Holdsworth have never been previously presented by the Office, factor (c) favors institution.

2. *Becton, Dickinson* Factors (e) and (f):

Advanced Bionics explains that “if the record of the Office’s previous consideration of the art is not well developed or silent, then a petitioner may show the Office erred by overlooking something persuasive under factors (e) and (f).”

Advanced Bionics at 10. Here, the Office’s underdeveloped consideration of Bruckman was “material error” because it overlooked Bruckman’s description of how data is both passed from the communication network to the customer nodes (*See* Ex.1005, [0058]) and passed from the customer nodes to the communication network (Ex.1005, [0065]). Bruckman also explains that the links may be “full-duplex” Ethernet links, which a POSITA would have understood to mean that they

are capable of transmitting data in both directions simultaneously. Ex.1005, [0047]; Ex.1009. By allowing claims based on limitations that were taught in the cited art of record—including the specific limitations the examiner had previously deemed allowable—the Office overlooked teachings that had a significant impact on patentability. *Advanced Bionics* at 8 n.9 (“material error may include misapprehending or overlooking specific teachings of the relevant prior art where those teachings impact patentability of the challenged claims”). This petition’s grounds of unpatentability are not merely “a disagreement with a specific finding of record by the Office.” *Advanced Bionics* at 10-11. Rather, this petition relies upon teachings in Bruckman that were simply overlooked during prosecution.

The Board has consistently declined to discretionarily deny institution under similar facts. *See, e.g., DISH Network LLC v. Sound View Innovations, LLC*, IPR2020-01041, Paper 13, 22 (PTAB Jan 19, 2021) (granting institution where “Petitioner has sufficiently shown Examiner error by pointing out specific teachings of the [previously-asserted] prior art that the Examiner overlooked”); *Trans Ova Genetics, LC v. XY, LLC*, IPR2018-00250, Paper 9, 18-19 (PTAB June 27, 2018) (granting institution where “the Examiner manifestly failed to appreciate [disclosures]...from [the previously-cited art]”).

With respect to factor (f), the Petition is supported by an expert declaration (Ex.1003) by Dr. Houh explaining how a POSITA would have understood

Bruckman and the obviousness of combining its teachings with those of Basso and Holdsworth. *See generally*, Ex.1003; *see also* Ex.1004. This new evidence also weighs against discretionary denial. *See, e.g., Puma N. Am., Inc. v. Nike, Inc.*, IPR2019-01058, Paper 10, 19 (PTAB Oct. 31, 2019) (instituting where petition presented “new non-cumulative evidence.... probative to issues of patentability and helpful to our consideration of a prior art combination that was not before the Examiner”).

Accordingly, the Board should not exercise its discretion to deny this petition under § 325(d).

C. Discretionary denial under *General Plastic* is not appropriate

The ’740 patent has not been challenged in any prior IPR petition, so none of *General Plastic* discretionary institution factors apply to this Petition. *See General Plastic Indus. Co., Ltd. v. Canon Kabushiki Kaisha*, IPR2016-01357, Paper 19 at 16 (PTAB Sept. 6, 2016) (Section II.B.4.i. precedential).

XI. CONCLUSION

Accordingly, Petitioner has established a reasonable likelihood that the Challenged Claims are unpatentable.

Respectfully submitted,

Dated: January 9, 2023
HAYNES AND BOONE, LLP
2323 Victory Avenue, Suite 700
Dallas, Texas 75219
Customer No. 27683

/Theodore M. Foster/
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456

XII. MANDATORY NOTICES

A. Real Party-in-Interest

Pursuant to 37 C.F.R. § 42.8(b)(1), Petitioner certifies that the real party-in-interest is Cisco Systems, Inc.

B. Related Matters

Pursuant to 37 C.F.R. § 42.8(b)(2), to the best knowledge of the Petitioner, the '740 patent is involved in the following case:

Case Heading	Number	Court	Date
<i>Orckit Corporation v. Cisco Systems, Inc.</i>	2:22-cv-276-JRG-RSP	EDTX	July 22, 2022

C. Lead and Back-up Counsel and Service Information

Lead Counsel

Theodore M. Foster
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (303) 382-6205
Fax: (214) 200-0853
ipr.theo.foster@haynesboone.com
USPTO Reg. No. 57,456

Back-up Counsel

David L. McCombs
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (214) 651-5533
Fax: (214) 200-0853
david.mccombs.ipr@haynesboone.com
USPTO Reg. No. 32,271

Gregory P. Huh
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (972) 739-6939
Fax: (214) 200-0853
gregory.huh.ipr@haynesboone.com
USPTO Reg. No. 70,480

Calmann J. Clements
HAYNES AND BOONE, LLP
2323 Victory Ave. Suite 700
Dallas, TX 75219

Phone: (972) 739-8638
Fax: (214) 200-0853
calmann.clements.ipr@haynesboone.com
USPTO Reg. No. 66,910

Please address all correspondence to lead and back-up counsel. Petitioner consents to service in this proceeding by email at the addresses above.

XIII. CLAIMS APPENDIX

- [1.0] 1. A method for communication, comprising:
- [1.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel,
- [1.2] at least one of said first physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;
- [1.3] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel,
- [1.4] at least one of said second physical links being a bi-directional link operative to communicate in both an upstream direction and a downstream direction;
- [1.5] receiving a data frame having frame attributes sent between the communication network and the network node;
- [1.6] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and
- [1.7] sending the data frame over the selected first and second physical links, said sending comprising communicating along at least one of said bi-directional links.
- [2.1] 2. The method according to claim 1, wherein the network node comprises a user node, and

- [2.2] wherein sending the data frame comprises establishing a communication service between the user node and the communication network.
- [3.1] 3. The method according to claim 1, wherein the second physical links comprise backplane traces formed on a backplane to which the one or more interface modules are coupled.
- [4.0] 4. A method for communication, comprising:
- [4.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel;
- [4.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;
- [4.3] receiving a data frame having frame attributes sent between the communication network and the network node;
- [4.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and
- [4.5] sending the data frame over the selected first and second physical links,
- [4.6] at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.
- [5.0] 5. A method for communication, comprising:
- [5.1] coupling a network node to one or more interface modules using a first

- group of first physical links arranged in parallel;
- [5.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;
- [5.3] receiving a data frame having frame attributes sent between the communication network and the network node;
- [5.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and
- [5.5] sending the data frame over the selected first and second physical links,
- [5.6] coupling the network node to the one or more interface modules comprises aggregating two or more of the first physical links into an external Ethernet link aggregation (LAG) group
- [5.7] so as to increase a data bandwidth provided to the network node.
- [6.1] 6. The method according to claim 1, wherein coupling each of the one or more interface modules to the communication network comprises at least one of multiplexing upstream data frames sent from the network node to the communication network, and
- [6.2] demultiplexing downstream data frames sent from the communication network to the network node.
- [7.1] 7. The method according to claim 1, wherein selecting the first and second

physical links comprises balancing a frame data rate among at least some of the first and second physical links.

[8.1] 8. The method according to claim 1, wherein selecting the first and second physical links comprises applying a mapping function to the at least one of the frame attributes.

[9.1] 9. The method according to claim 8, wherein applying the mapping function comprises applying a hashing function.

[10.1] 10. The method according to claim 9, wherein applying the hashing function comprises determining a hashing size responsively to a number of at least some of the first and second physical links,

[10.2] applying the hashing function to the at least one of the frame attributes to produce a hashing key,

[10.3] calculating a modulo of a division operation of the hashing key by the hashing size, and

[10.4] selecting the first and second physical links responsively to the modulo.

[11.1] 11. The method according to claim 10, wherein selecting the first and second physical links responsively to the modulo comprises selecting the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.

[12.1] 12. The method according to claim 1, wherein the at least one of the frame

attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

[13.0] 13. A method for communication, comprising:

[13.1] coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel;

[13.2] coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel;

[13.3] receiving a data frame having frame attributes sent between the communication network and the network node;

[13.4] selecting, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group; and

[13.5] sending the data frame over the selected first and second physical links,

[13.6] coupling the network node to the one or more interface modules and coupling each of the one or more interface modules to the communication network comprising specifying bandwidth requirements comprising at least one of a committed information rate (CIR), a peak information rate (PIR)

and an excess information rate (EIR) of a communication service provided by the communication network to the network node, and

[13.7] allocating a bandwidth for the communication service over the first and second physical links responsively to the bandwidth requirements.

[14.0] 14. A method for connecting user ports to a communication network, comprising:

[14.1] coupling the user ports to one or more user interface modules;

[14.2] coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel,

[14.3] at least one of said backplane traces being bi-directional and operative to communicate in both an upstream direction and a downstream direction;

[14.4] receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;

[14.5] for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces; and

[14.6] sending the data frame over the selected backplane trace; said sending comprising communicating along said at least one of said backplane traces.

[15.0] 15. A method for connecting user ports to a communication network, comprising:

- [15.1] coupling the user ports to one or more user interface modules;
- [15.2] coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel;
- [15.3] receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;
- [15.4] for each data frame, selecting responsively to at least one of the respective frame attributes a backplane trace from the two or more backplane traces;
and
- [15.5] sending the data frame over the selected backplane trace, at least some of the backplane traces being aggregated into an Ethernet link aggregation (LAG) group.
- [16.1] 16. The method according to claim 14, wherein selecting the backplane trace comprises applying a hashing function to the at least one of the frame attributes.
- [17.0] 17. Apparatus for connecting a network node with a communication network, comprising:
 - [17.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network,
 - [17.2] at least one of said interface modules being operative to communicate in

both an upstream direction and a downstream direction;

[17.3] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[17.4] a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and

[17.5] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame;

[17.6] at least one of said first physical links and at least one of said second links being bi-directional links operative to communicate in both said upstream direction and said downstream direction.

[18.1] 18. The apparatus according to claim 17, and comprising a backplane to which the one or more interface modules are coupled, wherein the second physical links comprise backplane traces formed on the backplane.

[19.0] 19. Apparatus for connecting a network node with a communication network, comprising:

[19.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the

communication network;

[19.2] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[19.3] a second group of second physical links arranged in parallel so as to couple the one or more interface modules to the communication network; and

[19.4] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame,

[19.5] at least one of the first and second groups of physical links comprising an Ethernet link aggregation (LAG) group.

[20.0] 20. Apparatus for connecting a network node with a communication network, comprising:

[20.1] one or more interface modules, which are arranged to process data frames having frame attributes sent between the network node and the communication network;

[20.2] a first group of first physical links arranged in parallel so as to couple the network node to the one or more interface modules;

[20.3] a second group of second physical links arranged in parallel so as to couple

- the one or more interface modules to the communication network; and
- [20.4] a control module, which is arranged to select for each data frame sent between the communication network and the network node, in a single computation based on at least one of the frame attributes, a first physical link out of the first group and a second physical link out of the second group over which to send the data frame,
- [20.5] two or more of the first physical links being aggregated into an external Ethernet link aggregation (LAG) group so as to increase a data bandwidth provided to the network node.
- [21.1] 21. The apparatus according to claim 17, and comprising a multiplexer, which is arranged to perform at least one of multiplexing upstream data frames sent from the network node to the communication network, and
- [21.2] demultiplexing downstream data frames sent from the communication network to the network node.
- [22.1] 22. The apparatus according to claim 17, wherein the control module is arranged to balance a frame data rate among at least some of the first and second physical links.
- [23.1] 23. The apparatus according to claim 17, wherein the control module is arranged to apply a mapping function to the at least one of the frame attributes so as to select the first and second physical links.

[24.1] 24. The apparatus according to claim 23, wherein the mapping function comprises a hashing function.

[25.1] 25. The apparatus according to claim 24, wherein the control module is arranged to determine a hashing size responsively to a number of at least some of the first and second physical links,

[25.2] to apply the hashing function to the at least one of the frame attributes to produce a hashing key,

[25.3] to calculate a modulo of a division operation of the hashing key by the hashing size, and

[25.4] to select the first and second physical links responsively to the modulo.

[26.1] 26. The apparatus according to claim 25, wherein the control module is arranged to select the first and second physical links responsively to respective first and second subsets of bits in a binary representation of the modulo.

[27.1] 27. The apparatus according to claim 17, wherein the at least one of the frame attributes comprises at least one of a layer 2 header field, a layer 3 header field, a layer 4 header field, a source Internet Protocol (IP) address, a destination IP address, a source medium access control (MAC) address, a destination MAC address, a source Transmission Control Protocol (TCP) port and a destination TCP port.

[28.0] 28. Apparatus for connecting a network node with a communication

network, comprising:

[28.1] one or more interface modules, which are arranged to process data frames

having frame attributes sent between the network node and the

communication network;

[28.2] a first group of first physical links arranged in parallel so as to couple the

network node to the one or more interface modules;

[28.3] a second group of second physical links arranged in parallel so as to couple

the one or more interface modules to the communication network; and

[28.4] a control module, which is arranged to select for each data frame sent

between the communication network and the network node, in a single

computation based on at least one of the frame attributes, a first physical link

out of the first group and a second physical link out of the second group over

which to send the data frame,

[28.5] the communication network being arranged to provide a communication

service to the network node, the service having specified bandwidth

requirements comprising at least one of a committed information rate (CR),

a peak information rate (PIR) and an excess information rate (EIR), and

[28.6] the first and second groups of physical links being dimensioned to provide

an allocated bandwidth for the communication service responsively to the

bandwidth requirements.

[29.0] 29. Apparatus for connecting user ports to a communication network,
comprising:

[29.1] one or more user interface modules coupled to the user ports, which are
arranged to process data frames having frame attributes sent between the
user ports and the communication network,

[29.2] at least one of said user interface modules being bi-directional and operative
to communicate in both an upstream direction and a downstream direction;

[29.3] a backplane having the one or more user interface modules coupled thereto
and comprising a plurality of backplane traces arranged in parallel so as to
transfer the data frames between the one or more user interface modules and
the communication network,

[29.4] at least one of said backplane traces being bi-directional and operative to
communicate in both said upstream direction and said downstream direction;
and

[29.5] a control module, which is arranged to select, for each data frame,
responsively to at least one of the frame attributes, a backplane trace from
the plurality of backplane traces over which to send the data frame.

[30.0] 30. Apparatus for connecting user ports to a communication network,
comprising:

- [30.1] one or more user interface modules coupled to the user ports,
- [30.2] which are arranged to process data frames having frame attributes sent between the user ports and the communication network;
- [30.3] a backplane having the one or more user interface modules coupled thereto and comprising a plurality of backplane traces arranged in parallel so as to transfer the data frames between the one or more user interface modules and the communication network;
- [30.4] a control module, which is arranged to select, for each data frame, responsively to at least one of the frame attributes, a backplane trace from the plurality of backplane traces over which to send the data frame;
- [30.5] at least some of the backplane traces are aggregated into an Ethernet link aggregation (LAG) group.
- [31.1] 31. The apparatus according to claim 29, wherein the control module is arranged to apply a hashing function to the at least one of the frame attributes so as to select the backplane trace.

CERTIFICATE OF WORD COUNT

Pursuant to 37 C.F.R. § 42.24(d), Petitioner hereby certifies, in accordance with and in reliance on the word count provided by the word-processing system used to prepare this Petition, that the number of words in this paper is 13,987.

Pursuant to 37 C.F.R. § 42.24(d), this word count excludes the table of contents, table of authorities, mandatory notices under § 42.8, certificate of service, certificate of word count, appendix of exhibits, and any claim listing.

Dated: January 9, 2023

/Theodore M. Foster/
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456

CERTIFICATE OF SERVICE

The undersigned certifies that, in accordance with 37 C.F.R. § 42.6(e) and 37 C.F.R. § 42.105, service was made on Patent Owner as detailed below.

Date of service January 9, 2023

Manner of service PRIORITY EXPRESS MAIL

Documents served Petition for *Inter Partes* Review Under 35 U.S.C. § 312 and 37 C.F.R. § 42.104 of U.S. 7,545,740; Petitioner's Exhibit List; All Exhibits; Petitioner's Power of Attorney.

Persons served May Patents Ltd.
c/o Dorit Shem-Tov
P.O.B. 7230
Ramat-Gan, 5217102
Israel

/Theodore M. Foster/
Theodore M. Foster
Lead Counsel for Petitioner
Registration No. 57,456