

Cisco Systems and Juniper Networks, Inc. v. Orckit Corporation

Patent Owner Demonstratives IPR2023-00401 U.S. Patent No. 7,545,740 June 12, 2024

James Carmichael
Steve McBride
Mitch Yang
Carmichael IP, PLLC

Michael Ng
George Stamatopoulos
Kobre & Kim LLP

Overview

1.The '740 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

The '740 Patent

1.The '740 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

The '740 Patent


US007545740B2

(12) **United States Patent**
Zelig et al.

(10) **Patent No.:** **US 7,545,740 B2**
(45) **Date of Patent:** **Jun. 9, 2009**

(54) **TWO-WAY LINK AGGREGATION**

(75) Inventors: **David Zelig**, Givatayim (IL.); **Ronen Solomon**, Givatayim (IL.); **Uzi Khilil**, Netanya (IL.)

(73) Assignee: **Corrigent Systems Ltd.**, Tel Aviv (IL.)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 461 days.

(21) Appl. No.: **11/279,045**

(22) Filed: **Apr. 7, 2006**

(65) **Prior Publication Data**
US 2007/0237172 A1 Oct. 11, 2007

(51) **Int. Cl.** **H04L 12/56** (2006.01)

(52) **U.S. Cl.** **370/230; 370/235; 370/445**

(58) **Field of Classification Search** None
See application file for complete search history.

(56) **References Cited**
U.S. PATENT DOCUMENTS
3,411,507 A 11/1968 Wingrove
4,535,785 A 8/1985 van den Honert
4,573,481 A 3/1986 Bullara
4,602,624 A 7/1986 Nuples et al.
4,608,985 A 9/1986 Crish
4,628,942 A 12/1986 Sweeney et al.
4,649,936 A 3/1987 Ungar
4,702,254 A 10/1987 Zabara
4,739,764 A 4/1988 Luc et al.
4,867,164 A 9/1989 Zabara
4,962,751 A 10/1990 Krauter
5,025,807 A 6/1991 Zabara
5,188,104 A 2/1993 Wernicke

(Continued)

FOREIGN PATENT DOCUMENTS
WO WO-01/10375 A2 2/2001


OTHER PUBLICATIONS
Mamalos, et al., "A Method for Transmitting PPP Over Ethernet (PPPoE)", in IETF RFC 2516, Feb. 1999, p. 1-15.

(Continued)

Primary Examiner—Chi H. Pham
Assistant Examiner—Seon D. Hyun
(74) *Attorney, Agent, or Firm*—Darby & Darby P.C.

(57) **ABSTRACT**
A method for communication includes coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel. Each of the one or more interface modules is coupled to a communication network using a second group of second physical links arranged in parallel. A data frame having frame attributes sent between the communication network and the network node is received. A first physical link out of the first group and a second physical link out of the second group are selected in a single computation based on at least one of the frame attributes. The data frame is sent over the selected first and second physical links. This method allows two or more link aggregation groups to be concatenated, using a single processing stage to determine port assignment for each frame in each of the link aggregation groups.

31 Claims, 3 Drawing Sheets



Ex. 1001
CISCO SYSTEMS, INC. / Page 1 of 13

(10) **Patent No.:** **US 7,545,740 B2**
(45) **Date of Patent:** **Jun. 9, 2009**

(57) **ABSTRACT**

A method for communication includes coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel. Each of the one or more interface modules is coupled to a communication network using a second group of second physical links arranged in parallel. A data frame having frame attributes sent between the communication network and the network node is received. A first physical link out of the first group and a second physical link out of the second group are selected in a single computation based on at least one of the frame attributes. The data frame is sent over the selected first and second physical links. This method allows two or more link aggregation groups to be concatenated, using a single processing stage to determine port assignment for each frame in each of the link aggregation groups.


Paper 21 [POR] at pp. 8-14; Ex. 1001 ['740 Patent], p. 1

The '740 Patent

SUMMARY OF THE INVENTION

In various communication applications, users connect to a communication network through a network element, such as an access concentrator or aggregator for obtaining different data services.


Embodiments of the present invention that are described hereinbelow provide improved methods and systems for connecting users to a communication network with increased capacity and quality of service. The network element comprises one or more user interface modules (UIMs), each serving one or more user ports. In some embodiments, each UIM is connected to the communication network using two or more physical links arranged in parallel, in order to provide sufficient bandwidth at a given quality-of-service (QoS) level.



Paper 21 at pp. 8-14; Ex. 1001, 1:34-47, Fig. 1

The '740 Patent


Upstream data frames sent from the user ports to the communication network and downstream data frames sent from the communication network to the user ports are distributed among the parallel physical links, so as to balance the traffic load among the links. The load balancing enables each UIM, and the network element as a whole, to deliver a higher bandwidth at a given QoS or to improve the QoS at a given bandwidth.



Paper 21 at pp. 8-14; Ex. 1001, 1:48-55, Fig. 1

The '740 Patent

In some embodiments, the UIMs are coupled to a backplane of the network element, and the parallel physical links comprise backplane traces. In some embodiments, the physical links are configured as an Ethernet link aggregation (LAG) group. Distribution of frames to individual physical links typically comprises applying a suitable mapping function, such as a hashing function. The mapping function typically uses frame attributes, such as various header fields of the frame, to determine a physical link over which to send each frame.




Paper 21 at pp. 8-14, 63-67;
Ex. 1001, 1:55-65, Fig. 4

The '740 Patent

FIG. 4 is a flow chart that schematically illustrates a method for single-stage downstream hashing, in accordance with an embodiment of the present invention. Although the following description demonstrates the method using the simplified system configuration of FIG. 3 above, the method can also be applied in any other system configuration comprising two stages of link aggregation, such as the configurations discussed in the descriptions of FIGS. 1 and 2 above.

The method begins with control module 60 determining a hashing size parameter denoted N_{bpow} , at a hash size definition step 80. N_{bpow} is defined as $N_{bpow} = N_{extp} \cdot N_{bpt}$, wherein N_{extp} denotes the number of user ports in external LAG group 68, and N_{bpt} denotes the number of backplane traces 56 in each LAG group 58. For simplicity, it is assumed that all line cards 40 having a user port in external LAG group 68 have the same number of backplane traces 56 connecting them to MUX 44. In the present example, $N_{extp} = 4$, $N_{bpt} = 4$, therefore $N_{bpow} = 16$.

In different embodiments of network element 32, the value of hashing size N_{bpow} can be hard-wired in the concentrator design. Alternatively, the hashing size can be provided to the network element as part of its configuration, or determined by control module 60 responsively to the configuration setting of network element 32, as detected by module 60. In general, different line cards 40 may comprise different numbers of user ports belonging to external LAG group 68.




Paper 21 at pp. 8-14, 29; Ex. 1001, 7:17-42.

The '740 Patent

During normal operation, network element 32 receives downstream frames via network connection 36, at a frame reception step 82. For each downstream frame, control module 60 calculates a hashing key of the frame, at a hash key calculation step 84. As explained above, the hashing key is typically calculated by applying a suitable hashing function to the frame attributes of the downstream frame.

Module 60 performs an integer modulo- N_{bpow} division of the hashing key, at a mapping calculation step 86. Module 60 divides the hashing key of the downstream frame by N_{bpow} and retains the modulo, or the remainder of the division operation, as a mapping index. The mapping index can be represented as a binary number having a fixed number of bits, depending on the value of N_{bpow} . In the present example in which $N_{bpow}=16$, the mapping index, being a remainder after division by 16, is an integer number in the range $0 \dots 15$, which can be represented using 4 bits.

Module 60 partitions the binary representation of the mapping index into two parts having N1 and N2 bits. 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. 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. In the present



Paper 21 at pp. 8-14, 29; Ex. 1001, 7:43-65, Fig. 4.

The '740 Patent

1. A method for communication, comprising:
coupling a network node to **one or more interface modules**
using a first group of first physical links arranged in parallel, 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;
coupling **each of the one or more interface modules** to a communication network using a **second group of second physical links arranged in parallel**, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Paper 21 at p. 16; Ex. 1001, Cl. 1.

The '740 Patent

15. A method for connecting user ports to a communication network, comprising:
coupling the user ports to one or more user interface modules;
coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel;
receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;
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
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.

Paper 21 at p. 16; Ex. 1001, Cl. 15.

Claim Construction

1.The '740 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

Claim Construction

1.The '740 Patent

2.Claim Construction

- “**selecting**” limitation (“selecting, in a single computation ... a first physical link out of a first group and a second physical link out of the second group”)
- “**arranged in parallel**”

3.The Challenged Claims Are Patentable

The “selecting” limitation

1. A method for communication, comprising:
coupling a network node to **one or more interface modules** **using a first group of first physical links** arranged in parallel, 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;
coupling **each of the one or more interface modules** to a communication network using a **second group of second physical links arranged in parallel**, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Paper 21 at p. 21-32; Paper [Sur-reply] 32 at pp. 2-8; Ex. 1001, Cl. 1

The “selecting” limitation

Further, there must be two selections because two links are “selected” from two different “groups”: a first group of links and a second group of links. Ex1001 [’740 Patent], claim 1. This means both selections must select from a group consisting of a plurality of links. *See, e.g.,* Ex2027 [Merriam-Webster] (“Select: chosen from a number or group by fitness or preference”; “Group: an assemblage of objects regarded as a unit”). If there is only one predetermined second physical link that must be used, there is neither a “selection” of “a second physical link” nor a “second group” as expressly required by the claims. And the claims require the “second group” to comprise of “second physical links,” thus the group must include a plurality of second physical links to choose from. Ex1001 [’740 Patent], claim 1.

Paper 21 at p. 24

The “selecting” limitation

- PO’s Construction: “selecting, in a single computation ... a first physical link out of a first group and a second physical link out of the second group, wherein the first physical link is not fixedly assigned to the second physical link”
- The structural limitation “wherein the first physical link is not fixedly assigned to the second physical link” is required by the claims
- Accordingly, “selecting, in a single computation ... a first physical link out of a first group and a second physical link out of the second group” should include “wherein the first physical link is not fixedly assigned to the second physical link”

Paper 21 at 21-32; Paper 32 at 2-8

The “selecting” limitation

In some embodiments, the UIMs are coupled to a backplane of the network element, and the parallel physical links comprise backplane traces. In some embodiments, the physical links are configured as an Ethernet link aggregation (LAG) group. Distribution of frames to individual physical links typically comprises applying a suitable mapping function, such as a hashing function. The mapping function typically uses frame attributes, such as various header fields of the frame, to determine a physical link over which to send each frame.

Unlike some known network element configurations, in which each user port is fixedly assigned to a specific backplane trace, the load balancing operation in embodiments of the present invention enables statistical multiplexing of the frames, in which there is no direct relationship or connection between user ports and backplane traces.

In some embodiments, two or more physical user ports are aggregated into a LAG group external to the network element, so as to form an aggregated user port having a higher bandwidth. The frames to and from the aggregated user port are distributed among the physical user ports of the external LAG group to balance the load and enable higher bandwidth and higher QoS. In the downstream direction, a combined mapping operation for frames addressed to the aggregated user port determines an individual backplane trace and a physical user port over which to send each frame in a single mapping computation.

Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 1:56-2:15

The “selecting” limitation

One way of balancing the load is to distribute the data rate of the Ethernet frames as uniformly as possible among the backplane traces. Alternatively, any other load balancing criterion can be used. In order to achieve load balancing, each group of backplane traces belonging to a particular line card is configured as an Ethernet link aggregation (LAG) group 58. Each LAG group 58 is considered by the relevant line card and by multiplexer 44 to be a single logical link having an aggregated bandwidth (i.e., capacity) equal to the sum of the bandwidths of the individual backplane traces in the group. As a result, there is no pre-assigned relationship or connection between any given user port 24 and a specific backplane trace 56, as in some conventional network elements. Ethernet frames are statistically multiplexed so as to balance the load among the backplane traces.

Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 5:6-22

The “selecting” limitation

In some embodiments, in the downstream direction, a control module 60 in network element 32 determines a downstream mapping of each downstream frame received through connection 48 to the appropriate backplane trace, using a suitable mapping function. The mapping performed by module 60 should naturally consider the user port to which the frame is addressed, in order to send the frame to the line card serving this user port. Module 60 controls multiplexer 44 in order to send each frame over the appropriate backplane trace. In addition to determining the appropriate line card for each frame (i.e., determining over which LAG group 58 to send the frame), the downstream mapping also balances the load of frames within each LAG group, responsively to attributes of the downstream frames.

For example, the downstream mapping can be implemented in two sequential stages. First, for each Ethernet frame, module 60 determines the appropriate line card to which the frame should be sent, depending on the destination user port. Then, within the backplane traces belonging to the LAG group of the appropriate line card, module 60 selects a particular trace and controls MUX 44 to send the frame over the selected trace. Alternatively, any other suitable mapping method can be used by module 60.


Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 5:23-67

The “selecting” limitation

Control module 60 may be implemented in hardware, as software code running on a suitable processor, or as a combination of hardware and software elements. Module 60 may comprise an independent module, or be integrated with other components of network element 32. Different embodiments of network element 32 may comprise any number of line cards 40. Each line card may serve any number of user ports 24 and any number of backplane traces 56. In general, there need not be a direct relation between the number of user ports and the number of backplane traces. Each user port and each backplane trace may have any suitable bandwidth. User ports can have equal or different bandwidths. Similarly, the backplane traces within each LAG group 58 can have equal or different bandwidths.


Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 6:1-14

The “selecting” limitation




Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, Fig. 1

The “selecting” limitation




Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, Fig. 2.

The “selecting” limitation



Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, Fig. 3

The “selecting” limitation



Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, Fig. 5.

The “selecting” limitation

In the system configuration of FIG. 2, consider the downstream frames arriving from network 28, via network processor 38, to multiplexer 44 and addressed to aggregated user port 64. In general, the processing of these frames in network element 32 comprises two consecutive mapping operations. First, each frame is mapped to one of user ports 24 in external LAG group 68. Determining the user port implicitly determines through which line card 40 the frame will pass. Then, the same frame is mapped to one of backplane traces 56 in the appropriate LAG group 58 that serves the selected line card. In some embodiments, in order to reduce the hardware complexity and/or computational complexity of network element 32, control module 60 performs a single combined mapping operation that combines the two mapping operations described above.

Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 6:50-65

The “selecting” limitation

FIG. 4 is a flow chart that schematically illustrates a method for single-stage downstream hashing, in accordance with an embodiment of the present invention. Although the following description demonstrates the method using the simplified system configuration of FIG. 3 above, the method can also be applied in any other system configuration comprising two stages of link aggregation, such as the configurations discussed in the descriptions of FIGS. 1 and 2 above.

The method begins with control module 60 determining a hashing size parameter denoted N_{bpow} at a hash size definition step 80. N_{bpow} is defined as $N_{bpow} = N_{extp} \cdot N_{bpt}$ wherein N_{extp} denotes the number of user ports in external LAG group 68, and N_{bpt} denotes the number of backplane traces 56 in each LAG group 58. For simplicity, it is assumed that all line cards 40 having a user port in external LAG group 68 have the same number of backplane traces 56 connecting them to MUX 44. In the present example, $N_{extp} = 4$, $N_{bpt} = 4$, therefore $N_{bpow} = 16$.

Paper 21 at 21-32; Paper 32 at 2-8;
Ex1001, 7:21-30

“arranged in parallel”

1. A method for communication, comprising:
coupling a network node to one or more interface modules
using a first group of first physical links arranged in parallel, 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;
coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Paper 32 at 8-13;
Ex1001, Cl. 1

“arranged in parallel”

- Petitioner’s construction: (Paper 27 [Reply] at 15)


Petitioner maintains that no construction of “*parallel*” is required. If the Board determines otherwise, however, then “*parallel*” should be construed to include what is described in the ’740 patent: lines that transfer data between a user node and a communication network are in “parallel.” The ’740 patent explains that

- PO’s construction: no construction needed – plain and ordinary meaning is apparent to a POSITA. EX2024 [Gomez Decl], ¶¶80-96

Paper 32 at 8-13

“arranged in parallel”

A POSITA would understand that the relevant endpoints for *internal backplane traces 56* are the *internal* components at which those terminate: *line cards 40* and *MUX 44*:



See, e.g., EX2024, ¶¶86-88.

“arranged in parallel”

22	Q.	I want to invite your attention to
23		Paragraph 14 of your declaration.
24	A.	Okay.

Sur-Reply at 8-13; Ex 2026 [Houh Dep. Tr.], 66:22-24

14. As explained by the prior art, “[l]ink 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.” Bruckman, [0002]. “Link

Paper 32 at 8-13; Ex. 1003 [Houh Decl.], par. 14

“arranged in parallel”

14. As explained by the prior art, “[l]ink 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.” Bruckman, [0002]. “Link

Sur-Reply at 8-13; Ex. 1003, ¶ 14

Link aggregation (LAG) 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 (referred to as the “LAG group”). Traffic transmitted between the end-

Paper 32 at 8-13; Ex 1001, 1:11-14

“arranged in parallel”

8 Q. Do you agree with the statement from
9 Dr. Bruckman?
10 A. I mean, I don't think I have any cause to --
11 to disagree specifically with this statement.
12 Q. Is this had a comprehensive definition of link
13 aggregation?
14 A. I think it's just a statement that it says. I
15 wouldn't necessarily characterize it the way you are.
16 Q. You mean you wouldn't characterize it as a
17 definition?
18 A. Yeah, I wouldn't characterize it the way you
19 characterized it in your question. It's simply a
20 statement. And, as I said, I have no reason to disagree
21 with what it's saying here.
22 Q. To be clear, I'm just asking if you consider
23 this to be a definition of link aggregation.
24 Do you consider this statement from Bruckman
25 to be a definition of link aggregation?

“arranged in parallel”

1	A.	I think it's just a statement. I'm not saying
2	it's a definition here.	It's a description of something
3	that Bruckman is saying that link aggregation can be,	
4	and it's simply an example given in Bruckman. And	
5	I'm -- I'm not -- I don't have any words around that	
6	that says this is what the definition is in my	
7	declaration.	

14. As explained by the prior art, “[l]ink 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.” Bruckman, [0002]. “Link

Paper 32 at 8-13; Ex. 2026, 68:1-7; Ex. 1003, par. 14


“arranged in parallel”

of a single user port 24. For this purpose, in the exemplary embodiment of FIG. 2, a number of user ports 24 are configured to form an aggregated user port 64. User ports 24 forming port 64 are configured as an Ethernet LAG group, referred to as an external LAG group 68. Thus, the bandwidth of port 64 is generally equal to the sum of bandwidths of the individual user ports 24 in external LAG group 68, in both upstream and downstream directions.

External LAG group 68 may comprise any number of user ports 24, belonging to any number of line cards 40. In general, a particular line card 40 may have some of its user ports 24 assigned to an external LAG group and other user ports 24 used individually or assigned to another external LAG group.

An external multiplexer (MUX) 72 performs the multiplexing and de-multiplexing (mapping) functions of the external LAG group. MUX 72 is external to network element 32 and is often located in user equipment 76, separate and distinct from network element 32. Aggregated port 64 is typically connected on the downstream side to a user node, such as a layer 2 or layer 3 switch (not shown).


LAG 68 with “parallel” links that terminate at two different line cards



Ex.1001, Fig. 2 (annotated); Ex.1025, ¶14.

Paper 32 at 8-13; Paper 27 at p. 15; Ex. 1001, 6:21-38


“arranged in parallel”



- 1 Q. Let's go to page 17 of your declaration.
- 2 A. Okay.
- 3 Q. There's a figure, Figure 2.
- 4 A. Yes.
- 5 Q. That figure's from the '740 patent, right?
- 6 A. That's what it's labeled, yes.
- 7 Q. And on that figure there is something referred
- 8 to as a network element labeled 32. Correct?
- 9 A. I see that.
- 10 Q. And within that network element 32 there are a
- 11 number of line cards labeled 40. Correct?
- 12 A. Yes, I see that.
- 13 Q. And let's take the first of those line cards,
- 14 line card 1.
- 15 Do you see that?
- 16 A. Yes.
- 17 Q. There are several line cards that start with
- 18 the number 1 and end in N. Correct?
- 19 A. Yes.
- 20 Q. The first line card is connected to a number
- 21 of lines shown in Figure 2 on the right-hand side of the
- 22 line card.
- 23 Do you see that?
- 24 A. Yes.
- 25 Q. Those lines represent backplane traces.

Paper 32 at 8-13; Ex 1003, p. 17; Ex 2026, 59:1-25


“arranged in parallel”



- 1 Correct?
- 2 A. If you look at the -- the specification of the
- 3 '740, the -- the items labeled 56 are often described
- 4 as -- as backplane traces.
- 5 Q. Okay. You can also look at your declaration
- 6 at Paragraph 25, which is right under the figure I was
- 7 pointing to, where you state: The interface modules are
- 8 connected to the communication network 28 using
- 9 backplane traces 56.
- 10 Do you see that?
- 11 A. That's what I say, yes.
- 12 Q. Okay, great.
- 13 Now, these backplane traces 56 on the
- 14 left-hand side connect to different line cards, right?
- 15 A. Yes, yeah, the -- the 56, yes.
- 16 Q. And different sets of backplane traces connect
- 17 to different line cards. Correct?
- 18 A. In this example figure, that's -- that's
- 19 what's shown.
- 20 Q. Okay. I want to direct your attention to the
- 21 very top set of backplane traces that are connected to
- 22 line card 1.
- 23 Do you see that?
- 24 A. Yes.
- 25 Q. And what's shown here on Figure 2 is four

Paper 32 at 8-13; Ex 1003, p. 17; Ex 2026, 60:1-25

“arranged in parallel”



'740 patent, Fig. 2 (annotated).

1 backplane traces that are connected to line card 1.

2 Do you see that?

3 A. Yes.

4 Q. And those four backplane traces are connected

5 also to a different component in network element 32 that

6 is labeled 44 and referred to as MUX.

7 Do you see that?

8 A. Yes.

9 Q. Is that component a multiplexer, in your

10 understanding?

11 A. The specification does indicate that 44 is a

12 multiplexer. And MUX is -- is the standard abbreviation

13 I've seen used before.

14 Q. Now, let me ask you this, Dr. Houh.

15 This top group of backplane traces that

16 connect on the one hand to line card 1 and on the other

17 hand to MUX 44, is this set of backplane traces arranged

18 in parallel?

19 A. I mean, certainly the figure suggests that

20 these are -- are parallel traces.

21 Q. What in the figure suggests to you that these

22 are parallel traces?


23 A. Well, there are four separate lines, and

24 they're -- they're -- they're parallel lines as drawn

25 on -- in the figure.

Paper 32 at 8-13; Ex 1003, p. 17; Ex 2026, 61:1-25

“arranged in parallel”



'740 patent, Fig. 2 (annotated).

- 1 Q. And the fact that they're parallel lines as
2 drawn would lead you to believe that these are parallel
3 traces. Correct?
- 4 A. I mean, coupled with the context of the
5 specification, I mean, I -- I think, you know, the
6 figures are read and described to the -- the POSITA
7 reviewing the patent.
- 8 Q. Does the fact that these four backplane traces
9 all start at the same line card and end at the same
10 multiplexer also suggest that these are parallel
11 backplane traces?
- 12 A. I think it's certainly consistent with things
13 that are described in the -- in the specification as
14 well.

Paper 32 at 8-13; Ex 1003, p.17, Ex 2026, 62:1-14

“arranged in parallel”

15 Q. If these backplane traces had different
16 starting and ending points, would that suggest to you
17 that they are parallel traces?

18 ATTORNEY CLEMENTS: Objection. Form.

19 A. I don't think that alone would -- would -- oh,
20 wait. So you're saying if the -- if the trace had
21 different starting points and end points? I mean, they
22 still could be parallel as laid out on a board, but if
23 the communications are going to different places,
24 then -- then -- then if communications over those traces
25 are going to different places, I don't think that's what
1 is thought of as -- as -- as parallel links or traces,
2 or links.


3 BY ATTORNEY STAMATOPOULOS:

4 Q. And when you say "that is thought of," you
5 mean by a person of ordinary skill in the art?

6 A. Yes.

Paper 32 at 8-13; Ex
1003 at p. 17; Ex
2026, 62:15-63:6

“arranged in parallel”



The '740 Patent

1.The '740 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

The Challenged Claims Are Patentable

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

Paper 2, p. 21

The Challenged Claims Are Patentable


1. No disclosure of “coupling one of the one or more interface modules to a communication network using a second group of links arranged in parallel” (the “arranged in parallel” limitation)
 - All independent claims and all claims depending therefrom
2. No disclosure of “selecting, in a single computation... a first physical link out of the first group and a second physical link out of the second group” (the “two-link-with-one-computation selection” limitation)
 - Independent claims 1, 4-5, 13, 17, 19-20, and 28 and all claims depending therefrom
3. No disclosure of “selecting from two or more backplane traces” (the “backplane trace selection” limitation)
 - Independent Claim 14, 15, 29, and 30 & all claims depending therefrom

Bruckman/Basso do not disclose a second group of links arranged in parallel


1. A method for communication, comprising:
coupling a network node to one or more interface modules
using a first group of first physical links arranged in parallel, 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;
coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Ex. 1001, Cl. 1


Bruckman/Basso do not disclose a second group of links arranged in parallel



Bruckman/Basso do not disclose a second group of links arranged in parallel



Bruckman/Basso do not disclose a second group of links arranged in parallel



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.


89. A POSITA would understand that in the context of the claims, providing the backplane traces 56 as parallel physical links enables a network element “in which there is no direct relationship or connection between user ports and backplane traces.” And a POSITA would understand that traces that each lead to a different sets of ports, such that the ports (or first physical links) in a group can only directly connect to a single trace, are not arranged in parallel. Ex2028 [All About Circuits], 2 (“In a parallel circuit, all components share the same electrical nodes.”); Ex2032 [Basic Electricity], 1 (“A parallel circuit is constructed by connecting the terminals of all the individual load devices”).

Bruckman/Basso do not disclose two selections/ a backplane trace selection

1. A method for communication, comprising:
coupling a network node to **one or more interface modules**
using a first group of first physical links arranged in parallel, 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;
coupling **each of the one or more interface modules** to a communication network using a **second group of second physical links arranged in parallel**, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Ex. 1001, Cl. 1

Bruckman/Basso do not disclose two selections/ a backplane trace selection



Bruckman/Basso do not disclose two selections/ a backplane trace selection

15 Q. And so you need a mechanism to determine which
16 of the physical links will transmit a particular frame
17 when presented with an option between several different
18 physical links that connect between two end points,
19 right?

20 A. I mean, I do talk about such a mechanism in
21 Paragraph 15.

Bruckman/Basso do not disclose two selections/ a backplane trace selection

15. When using link aggregation, a mechanism is used to determine which of the physical links will transmit a particular frame. This mechanism includes “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.” Bruckman, [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.” Bruckman, [0005]. These distribution algorithms may include, for example, a hash function. *See* Bruckman, [0006]-[0012].

Bruckman/Basso do not disclose two selections/ a backplane trace selection

22 Q. What -- do you need a similar mechanism when
23 there is only one physical link available for a
24 connection between two end points?

25 ATTORNEY CLEMENTS: Objection. Form.


1 A. I don't think it would be necessary in such a
2 case, but you could certainly have a distributor
3 function that -- that distributes frames among one link.
4 But it wouldn't be necessary to do so --

Bruckman/Basso do not disclose two selections/ a backplane trace selection

7 Q. Why would it not be necessary?

8 A. Because there's -- you -- the -- the purpose
9 of -- one of the purposes of doing such aggregation is
10 to achieve higher bandwidth by using two parallel paths.
11 So you'd -- you'd want to leverage the fact that you can
12 use two parallel paths to go faster than one path. But
13 if there's one path, you don't really have a choice if
14 there's one physical link between -- you know, to the
15 other element. You don't really have a choice. Doesn't
16 mean you couldn't distribute among one link, but it's
17 just not -- wouldn't do anything to speed it up or speed
18 up the data transfer.

Bruckman/Basso do not disclose two selections/ a backplane trace selection



Bruckman/Basso do not disclose making two selections with a single computation

1. A method for communication, comprising:
coupling a network node to one or more interface modules using a first group of first physical links arranged in parallel, 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;
coupling each of the one or more interface modules to a communication network using a second group of second physical links arranged in parallel, 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;
receiving a data frame having frame attributes sent between the communication network and the network node;
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
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.

Ex. 1001, Cl. 1

Bruckman/Basso do not disclose making two selections with a single computation

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.

Bruckman/Basso do not disclose making two selections with a single computation

TABLE I

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)
}
```

Bruckman/Basso do not disclose making two selections with a single computation


```
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 provides limited description of how the hash computation is used 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.

Bruckman/Basso do not disclose making two selections with a single computation



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

Bruckman/Basso do not disclose making two selections with a single computation



Bruckman/Basso do not disclose making two selections with a single computation

[0009] A network device, e.g., router, may comprise 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. These ports may be connected to another one or more network devices. The switch fabric may be configured to direct incoming packets of data to particular blades where one or more of the network processors in the recipient blade may be configured to process the received packets.

Bruckman/Basso do not disclose making two selections with a single computation

[0010] Upon a network processor receiving a packet of data, the network processor may index into a table, commonly referred to as a forwarding table, to determine the table associated with a particular logical interface as well as the next destination address. The forwarding table may comprise a plurality of entries where each entry may comprise information indicating a particular table associated with a particular logical interface as well as the next destination address. Each logical interface may be associated with a table storing a plurality of entries containing blade/port combinations as discussed further below. In one embodiment, an entry may be indexed in the forwarding table using a destination address in the received packet header.

Bruckman/Basso do not disclose making two selections with a single computation

[0011] A hash function may then be performed on the received packet to generate a hash value. In one embodiment, a hash function may be performed on the source and destination address in the packet header to generate a hash value.

Bruckman/Basso do not disclose making two selections with a single computation

[0041] In step 504, a hash function may be performed on the received packet to generate a hash value. In one embodiment, a hash function may be performed on the source and destination address in the packet header to generate a hash value. It is noted that in other embodiments a hash function may be performed on different fields, e.g., port, type of service, in the received packet to generate a hash value.

Bruckman/Basso do not disclose making two selections with a single computation

[0042] In step 505, the hash value generated in step 504 may be used to index into the table associated with a particular logical interface 405 determined in step 503. Upon indexing into the table associated with the logical interface 405 determined in step 503, an appropriate blade 402/port 404 combination may be identified in step 506 to transmit the received packet of data as explained below.


Bruckman/Basso do not disclose two selections/ a backplane trace selection

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

- coupling the user ports to one or more user interface modules;
- coupling each user interface module to the communication network via a backplane using two or more backplane traces arranged in parallel;
- receiving data frames sent between the user ports and the communication network, the data frames having respective frame attributes;
- 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
- 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.

Paper 21 at 63-69; Paper 27 at 30-31

Bruckman/Basso do not disclose two selections/ a backplane trace selection



Paper 21 at 63-69; Paper 32 at 30-31