

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

Patent Owner Demonstratives IPR2023-00402 U.S. Patent No. 8,830,821 June 12, 2024

James Carmichael
Steve McBride
Mitch Yang
Carmichael IP, PLLC

Michael Ng
George Stamatopoulos
Kobre & Kim LLP

DEMONSTRATIVE EXHIBIT -
NOT EVIDENCE

Overview

1.The '821 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

1.The '821 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

The '821 Patent

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

(10) **Patent No.:** **US 8,830,821 B2**
(45) **Date of Patent:** **Sep. 9, 2014**

(54) **METHOD FOR SUPPORTING MPLS
TRANSPORT PATH RECOVERY WITH
MULTIPLE PROTECTION ENTITIES**

(56) **References Cited**
U.S. PATENT DOCUMENTS

(75) Inventors: **Daniel Cohn**, Raanana (IL); **Rafi Ram**,
Herzliya (IL)

6,314,093 B1 11/2001 Mann et al.
7,561,534 B2 7/2009 Alicherry et al.
7,961,602 B2 * 6/2011 Tschio 370/227
2005/0185654 A1 8/2005 Zadikian et al.
2007/0064702 A1 * 3/2007 Bates et al. 370/392
2010/0157794 A1 * 6/2010 Nakash 370/228
2010/0309778 A1 * 12/2010 Young 370/219

(73) Assignee: **Orckit-Corrigent 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 324 days.

(21) Appl. No.: **13/311,128**

(22) Filed: **Dec. 5, 2011**

(65) **Prior Publication Data**
US 2012/0327762 A1 Dec. 27, 2012

Related U.S. Application Data

(60) Provisional application No. 61/499,943, filed on Jun.
22, 2011.

(51) **Int. Cl.**
H04J 1/16 (2006.01)
H04L 12/28 (2006.01)

(52) **U.S. Cl.**
USPC **370/217**; **370/219**; **370/227**; **370/228**;
370/392

(58) **Field of Classification Search**
USPC **370/219**, **227-228**, **392**
See application file for complete search history.

S. Bryant, et al; MPLS-TP Linear Protection; Nokia Siemens Net-
works, Html markup produced by rfcmarkup 1.104, available from
http://tools.ietf.org/tools/rfcmarkup; Jun. 2, 2011.
Haiyan Zhang et al; Linear Protection Switching in MPLS-TP;
Telecom Italia; Html markup produced by rfcmarkup 1.104, avail-
able from http://tools.ietf.org/tools/rfcmarkup; Jul. 12, 2010.


* cited by examiner

Primary Examiner — Jung Park
Assistant Examiner — Chuong M Nguyen

(74) *Attorney, Agent, or Firm* — Peter A. Nieves; Sheehan
Phinney Bass + Green PA

(57) **ABSTRACT**
A system and method for selecting MPLS network transport
entities between a first endpoint and a second endpoint are
presented. Working and protection entities are selected from
a set of available entities based upon minimizing the prob-
ability of simultaneous failure and/or minimizing a cost func-
tion base upon one or more metrics.

20 Claims, 6 Drawing Sheets



(10) **Patent No.:** **US 8,830,821 B2**
(45) **Date of Patent:** **Sep. 9, 2014**

(57) **ABSTRACT**

A system and method for selecting MPLS network transport entities between a first endpoint and a second endpoint are presented. Working and protection entities are selected from a set of available entities based upon minimizing the probability of simultaneous failure and/or minimizing a cost function base upon one or more metrics.


Ex. 1001, p. 1

FIELD OF THE INVENTION

The present invention relates to telecommunications and more particularly, is related to high availability of telecommunication entities.


Ex. 1001, 1:17-19

The '821 Patent



Ex. 1001, Fig. 1; POR, 8-10

The '821 Patent




Ex. 1001, Fig. 1; POR, 8-10

A first exemplary embodiment of an entity protection system **100** is shown in FIG. 1. The first embodiment for executing the functionality may be a computer, as described below. An entity protection switch **130** switches between a working entity **122** and a protection entity **124** from a set of entities **120** or entity descriptors represented within an entity selection logic block **110**. Each entity **120** represents a network route between a first endpoint and a second endpoint, and may include entity cost data associated with each entity, for example, interior gateway protocol (IGP) data, traffic engineering (TE) data, and probability of entity failure. An overall cost logic block **140** uses results from a concurrent failure probability matrix **150** and an entity cost determination block **160** to select the working entity **122** and the protection entity **124** from the set of entities **120**, as described below. The concurrent failure probability matrix **150** examines the set of entities **120** and calculates the probability of concurrent failure for each entity pair. For example, a pair having multiple spans and/or nodes in common may be more likely to simultaneously fail than a pair having no spans or nodes in common. Similarly, an entity traversing a network element having a high mean-time-between-failures value may be less likely to fail than an entity traversing a network element having a low mean-time-between-failures value.

Ex. 1001, 3:30-53; POR, 8, 48-49

The '821 Patent




Ex. 1001, Fig. 1; POR, 8-10

The working entity **122** and the protection entity **124** are inputs to an entity protection switch **130**. The entity protection switch **130** may employ, for example, an MPLS linear protection scheme such as those defined in IETF draft-ietf-mpls-tp-linear-protection, IETF draft-zulr-mpls-tp-linear-protection-switching, and ITU-T G.8131 T-MPLS Linear Protection. However, a person having ordinary skill in the art will appreciate that other protection switching mechanisms may be employed within the scope of this disclosure. The entity protection switch **130** may select either the working entity **122** or the protection entity **124** to be the active entity **126**.


Ex. 1001, 3:54-65, POR, 8

The '821 Patent



Ex. 1001, Fig. 2; POR, 8-9

The '821 Patent





Ex. 1001, Fig. 3; POR, 8

FIG. 3 is a flowchart expanding on the entity selection process as shown by block 230 (FIG. 2). An entity pair is selected, as shown by block 310. The entity pair may be selected based upon the overall cost function, as determined above, as shown by block 220 (FIG. 2). Cost functions are described further below. One entity of the entity pair may be selected as a working entity, as shown by block 320. Similarly, a protection entity may be selected from the entity pair, as shown by block 330. An active entity may be selected from the working entity and the protection entity, as shown by block 340. The protection entity may be selected before the working entity, or the working entity may be selected before the protection entity. The active entity may be selected after either or both the working entity and/or the protection entity have been selected. It should be noted that entity pair reselection, as triggered by a reselection event, may be substantially similar to initial entity selection.



Ex. 1001, 4:49-65; POR, 8

The '821 Patent



Ex. 1001, Figs. 3 and 4A; POR, 10



The '821 Patent



An entity selection method may generally minimize an overall cost function. The cost function for an entity pair is calculated, as shown in block 220 (FIG. 2), and an entity pair may be selected based in part upon the calculated cost function in block 310 (FIG. 3). Flowcharts expanding upon selecting a pair based upon the overall cost function of block 310 (FIG. 3) are shown by FIGS. 4A, 4B, and 4C. A first embodiment of the overall cost minimization minimizes the probability of an over-50 ms traffic hit, as shown by FIG. 4A. The entity pair may be selected based upon each entity of the pair having the lowest possible probability of individual failure. Alternatively, the entity pair may be selected to minimize the probability of a span or node failure simultaneously affecting the working and protection entities, as shown by block 440, and further described below.

Ex. 1001, Figs. 4A, 4C, 5:33-47, POR, 10

The '821 Patent





A second embodiment of overall cost minimization shown in FIG. 4B, minimizes entity cost according to one or more predefined metrics (block 450), for example, interior gateway protocol (IGP) or traffic engineering (TE) data, for the selected working entity. When network efficiency optimization is desired, the MPLS linear protection block may be configured to be revertive, and therefore, in the absence of entity failures, the working entity will be active.

Since the first overall cost minimizing embodiment and the second overall cost minimizing embodiment may produce conflicting results, a third overall cost minimizing entity selection embodiment, shown in FIG. 4C, may minimize an operator-defined function of the simultaneous failure probability (SFP) and the entity cost. An example of this function would be to minimize SFP (block 410) by selecting working and protection entity pairs with the lowest SFP, and, if there are multiple pairs with the same SFP within an error margin, as shown by block 420, choosing the pair that minimizes entity IGP cost (block 430). It should be noted that the term “simultaneous failure” encompasses concurrent failures of

Ex. 1001, Figs. 4A, 4C, 5:48-5:67, POR, 10

The '821 Patent



The probability of simultaneous failure in entities i and j can be expressed as shown in Eq. 1,



$$SFP_{ij} = \sum_k P(C_{ij}^k) \quad \text{Eq. 1}$$

where C_{ij}^k is the k -th element common to entities i and j , and $P(x)$ is the element failure probability function. A metric related to the probability of an entity failure is known as the availability of the entity. Entity availability measures the probability that an end-to-end connection will survive the failure of one or more spans or nodes between the endpoints.

Ex. 1001, Figs. 4A, 4C, 6:5-6:19, POR, 10

Ex. 1001, p. 3



The '821 Patent



Entity elements may include spans and/or nodes. Span failure probabilities may be assigned, for example, manually by a network operator based on network knowledge. Alternatively, span failure probabilities may be automatically calculated by the network, for example, by a management entity or by the endpoints themselves, based on span parameters such as span length, which in some scenarios can be automatically detected by the endpoints, for example, based on measured received optical power. Node failure probabilities may be configured by the operator, typically based on vendor-provided mean time between failure (MTBF) parameters, and on infrastructure failure probability, for example, power failure probability.

Ex. 1001, Figs. 4A, 4C, 6:20-33; POR, 10

The '821 Patent



As discussed above, entity selection logic may select a pair of entities that minimize a combined cost function, or cost logic, of the entity cost and the SFP. As shown in FIG. 1, this may be calculated by populating entries, for example, entity descriptors **120**, in the entity selection logic **110**, for example, containing SFP values for each entity pair as calculated by a concurrent failure probability matrix **150**, and the entity cost vector **160**, for example, containing entity cost for each entity, to a combined cost logic function **140** that yields a combined cost metric for each entity pair.

Ex. 1001, Figs. 4A, 4C, 6:34-43; POR, 10

The '821 Patent

1. An entity selection method performed by a network device, comprising the steps of:
providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
determining an overall cost for each entity pair of said plurality of entities;
selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and
if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,
wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1

The '821 Patent

14. A system for selecting entities within an MPLS network, comprising:
a data structure comprising a plurality of transport entity descriptors;
an entity protection switch configured to switch between a working entity and a protection entity; and
digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity;
logic configured to determine an entity cost of said plurality of transport entity descriptors: and
logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,
wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 14

1.The '821 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

Petitioner's Constructions Are Untimely

37 C.F.R. § 42.104(b) *Identification of challenge.* Provide a statement of the precise relief requested for each claim challenged. The statement must identify the following:

...

How the challenged claim is to be construed. ...

Sur-reply, 2

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

Petition, 12-13; Sur-reply, 2

Petitioner's Constructions Are Untimely

[0076], [0142]-[0145], Claim 1, Claim 11. Consistent with the analysis at [1.1]-[1.2], it would have been obvious for the “plurality of candidate path pairs” (pairs of primary and protection LSPs) to correspond to the provided plurality of LSPs of the MPLS architecture. Ex.1003, ¶¶102-103.

Petition, 35; Sur-reply, 2

First, as discussed at [1.2], [2.1], [5.1], Doshi's path calculation considers disjointedness of the working LSP (“*working entity*”) and the protection LSP (“*protection entity*”). Further, as discussed at [14.3], a POSITA would have found

Petition, 80; Sur-reply, 2

The “selection limitation”

1. An entity selection method performed by a network device, comprising the steps of:
providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
determining an overall cost for each entity pair of said plurality of entities;
selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and
if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,
wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1; POR, 27

The “selection limitation”

1. An entity selection method performed by a network device, comprising the steps of:
providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
determining an overall cost for each entity pair of said plurality of entities;
selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and
if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,
wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1; POR, 27-28

The “selection limitation”

1. An entity selection method performed by a network device, comprising the steps of:

- providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
- determining an overall cost for each entity pair of said plurality of entities;
- selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and

if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1; POR, 27-28

The “selection limitation”

9. The method of claim 4, wherein said overall cost function comprises:
selecting a subset of entity pairs wherein each entity pair of said subset has substantially minimum probability of a concurrent failure of said protection entity and said working entity; and
if said subset comprises at least two entity pairs, selecting an entity pair from said subset that minimizes an entity cost function.

Ex. 1001, Cl. 9; Reply, 7-8; Sur-reply, 5

The “selection limitation”






FIG. 2

Ex. 1001, Figs. 2, 3, 4C; Sur-Reply, 3-4

The “selection limitation”



The “selection limitation”





FIG. 4A

Ex. 1001, Figs. 3 and 4A; POR, 10; Sur-reply, 5-6

The “selection limitation”





FIG. 4B

Ex. 1001, Figs. 3 and 4A; POR, 10; Sur-reply, 5-6

The “selection limitation”






FIG. 4C

Ex. 1001, Figs. 3 and 4A; POR, 10; Sur-reply, 5-6

The “selection limitation”

69. This is not what the '740 patent teaches, nor what claims 1 and 17 require. In contrast to Doshi, a POSITA reading and practicing claim 1 of the '821 patent would understand that in a network with $K=10$ entities there are 90 corresponding candidate path pairs to select from, assuming the order is not interchangeable (i.e., treat pairs “AB” and “BA” as not interchangeable, or $\{A,B\} \neq \{B,A\}$). Thus, in the '821 Patent, the number of N possible pairs of primary/restoration paths in a network with K entities available to be paired is $N=K*(K-1)$. N is the number of pairs, K is the number of entities that can make up the pairs, and $K-1$ the factor by which the number of available entities must be multiplied to account for the fact that each of the K entities has $K-1$ potential pairings (assuming it is the primary entity, each of the K entities can be paired with each other entity in the plurality as a restoration entity). For $K=10$, $N=90$; thus, a POSITA would understand that practicing claim 1 of the '821 patent provides nine times the number of potential pairs to choose from as compared to Doshi's method.

The “selection limitation”



“working entity” versus “protection entity”

The multiprotocol label switching-transport profile (MPLS-TP) framework requires support of the transport entity or trail recovery. For example, in the event of a span or a node failure of a label switch path (LSP), pseudowire (PW) or sub network connection (SNC), the entity should recover within 50 ms of the failure. This is known as a sub-50 ms recovery, or a sub-50 ms traffic hit. References to a transport entity in this document are applicable to LSP, PW and SNC.

Ex. 1001, 1:38-45; Sur-reply, 6-7

“working entity” versus “protection entity”

Several MPLS-TP mechanisms exist to accomplish a sub-50 ms node or span failure recovery by defining a second transport entity, the protection entity, between the same endpoints as the original entity, the working entity, where the working and protection entities do not share the span or nodes against whose failure protection is required. When a failure is detected at the working entity, the endpoints switch traffic to the protection entity. Traffic recovery is expected to occur within 50 ms of fault detection. These mechanisms rely on proactive connectivity verification between the endpoints for fast failure detection, and support both 1:1 and 1+1 protection architectures.

Ex. 1001, 1:46-57; Sur-reply, 6-7

“working entity” versus “protection entity”

14. A system for selecting entities within an MPLS network, comprising:
a data structure comprising a plurality of transport entity descriptors;
an entity protection switch configured to switch between a working entity and a protection entity; and
digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity;
logic configured to determine an entity cost of said plurality of transport entity descriptors; and
logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,
wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 14; POR, 48; Sur-reply, 6, 8-9

“working entity” versus “protection entity”

Several MPLS-TP mechanisms exist to accomplish a sub-50 ms node or span failure recovery by defining a second transport entity, the protection entity, between the same endpoints as the original entity, the working entity, where the working and protection entities do not share the span or nodes against whose failure protection is required. When a failure is detected at the working entity, the endpoints switch traffic to the protection entity. Traffic recovery is expected to occur within 50 ms of fault detection. These mechanisms rely on proactive connectivity verification between the endpoints for fast failure detection, and support both 1:1 and 1+1 protection architectures.

Examples of such mechanisms can be found in IETF draft-ietf-mpls-tp-linear-protection, IETF draft-zulr-mpls-tp-linear-protection-switching, and ITU-T G.8131 T-MPLS Linear Protection. However, in network environments where span or

Ex. 1001, 1:46-61; Sur-reply, 6-8

“working entity” versus “protection entity”




FIG. 3 is a flowchart expanding on the entity selection process as shown by block 230 (FIG. 2). An entity pair is selected, as shown by block 310. The entity pair may be selected based upon the overall cost function, as determined above, as shown by block 220 (FIG. 2). Cost functions are described further below. One entity of the entity pair may be selected as a working entity, as shown by block 320. Similarly, a protection entity may be selected from the entity pair, as shown by block 330. An active entity may be selected from the working entity and the protection entity, as shown by block 340. The protection entity may be selected before the working entity, or the working entity may be selected before the protection entity. The active entity may be selected after either or both the working entity and/or the protection entity have been selected. It should be noted that entity pair reselection, as triggered by a reselection event, may be substantially similar to initial entity selection.

Ex. 1001, Fig. 3, 4:49-65; Sur-reply, 6-7

“working entity” versus “protection entity”

3/20/24, 3:04 PM

draft-zulr-mpls-tp-linear-protection-switching-01

MPLS Working Group
Internet Draft
Intended status: Standards Track

Haiyan Zhang
Huawei
Igor Umansky
Alcatel-Lucent
Han Li
China Mobile
Jeong-dong Ryoo
ETRI
Alessandro D'Alessandro
Telecom Italia

Expires: January 12, 2011

July 12, 2010

Linear Protection Switching in MPLS-TP
draft-zulr-mpls-tp-linear-protection-switching-01.txt

Abstract

This document specifies a linear protection switching mechanism for MPLS-TP. This mechanism supports 1+1 unidirectional/bidirectional protection switching and 1:1 bidirectional protection switching. It is purely supported by MPLS-TP data plane, and can work without any control plane.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunications Union Telecommunications Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by the ITU-T.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [RFC 78](#) and [RFC 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress".

Zhang, et al.

Expires January 12, 2011

[Page 1]

<https://datatracker.ietf.org/doc/html/draft-zulr-mpls-tp-linear-protection-switching-01>

Ex.1043 / IPR2023-00402 / Page 1 of 27
CISCO SYSTEMS, INC. v. ORCKIT CORPORATION


1/27

- 1:n architecture

Details for the 1:n protection switching architecture will be provided in a future version of this draft.

Ex. 1043, p.1; Sur-reply, 9

“working entity” versus “protection entity”



Ex. 1001, Fig. 1; POR, 8-9

1.The '821 Patent

2.Claim Construction

3.The Challenged Claims Are Patentable

The Challenged Claims Are Patentable

Grounds	Claims	Basis
#1	1-7, 9-11, 13, and 17-20	35 U.S.C. § 103 (Pre-AIA) over Doshi and Guichard
#2	8 and 12	35 U.S.C. § 103 (Pre-AIA) over Doshi, Guichard, and Huang
#3	14-16	35 U.S.C. § 103 (Pre-AIA) over Doshi, Guichard, and Xu

Petition, 14

The Challenged Claims Are Patentable

1. The “selection limitation” (“selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost”)
 - Independent Claim 1 & all claims depending therefrom (2-13); independent Claim 17 & all claims depending therefrom (18-20)
2. The “reselection limitation” (“reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair ...”)
 - All claims
3. The probability of “concurrent failure” limitation (“digital logic configured to determine a probability of concurrent failure”)
 - Independent Claim 14 & all claims depending therefrom (15-16)
4. Claims 2, 13 & 18

POR, 26-27

Doshi/Guichard Fail to Disclose the Selection Limitation

1. An entity selection method performed by a network device, comprising the steps of:

- providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
- determining an overall cost for each entity pair of said plurality of entities;
- selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and

if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair, wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1; POR, 27-28

Doshi/Guichard Fail to Disclose the Selection Limitation

determining an overall cost for each entity pair of said plurality of entities;

selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and

Ex. 1001, Cl. 1; POR, 27-28

Doshi/Guichard Fail to Disclose the Selection Limitation


A plurality of transport entities between a first endpoint and a second endpoint are provided, as shown by block 210. Each possible entity pairing is examined for cost in turn, starting with a first entity pair, as shown by block 215. An overall cost for each pair of the plurality of transport entities is determined, as shown by block 220. The overall cost may be determined based on the probability of simultaneous failure of the working and protection entities (SFP), as described below, the entity cost, or a combination of SFP and entity cost. The entity cost may be defined as the sum of the link metrics, for example IGP or TE metrics, along the entity. Three embodiments of overall cost determination are described in more detail below (FIGS. 4A, 4B and 4C).

Ex. 1001, 4:11-23; POR, 8-9; Sur-reply, 3-4

Doshi/Guichard Fail to Disclose the Selection Limitation

69. This is not what the '740 patent teaches, nor what claims 1 and 17 require. In contrast to Doshi, a POSITA reading and practicing claim 1 of the '821 patent would understand that in a network with $K=10$ entities there are 90 corresponding candidate path pairs to select from, assuming the order is not interchangeable (i.e., treat pairs "AB" and "BA" as not interchangeable, or $\{A,B\} \neq \{B,A\}$). Thus, in the '821 Patent, the number of N possible pairs of primary/restoration paths in a network with K entities available to be paired is $N=K*(K-1)$. N is the number of pairs, K is the number of entities that can make up the pairs, and $K-1$ the factor by which the number of available entities must be multiplied to account for the fact that each of the K entities has $K-1$ potential pairings (assuming it is the primary entity, each of the K entities can be paired with each other entity in the plurality as a restoration entity). For $K=10$, $N=90$; thus, a POSITA would understand that practicing claim 1 of the '821 patent provides nine times the number of potential pairs to choose from as compared to Doshi's method.

Doshi/Guichard Fail to Disclose the Selection Limitation



In the '821 Patent:

Number of pairs ("N") is computed is given by the following formula based on the number of paths ("K"):

$$N = K * (K - 1)$$

For $K=10 \rightarrow N = 90$

For $K=3 \rightarrow N = 6$


But in Doshi:

Number of pairs is limited to K, thus:

$K=10 \rightarrow N = 10$ pairs

$K=3 \rightarrow N = 3$ pairs

Doshi/Guichard Fail to Disclose the Selection Limitation



In the '821 Patent:

Number of pairs ("N") is computed is given by the following formula based on the number of paths ("K"):

$$N = K * (K - 1)$$

For K=10 → N = 90

For K=3 → N = 6

But in Doshi:

Number of pairs is limited to K or L, thus:

$K \text{ (\& } L \leq K) \rightarrow N = K \text{ or } N = L$

Thus, N= 10 or N = 3 pairs

Doshi/Guichard Fail to Disclose the Reselection Limitation

1. An entity selection method performed by a network device, comprising the steps of:
providing a plurality of multi protocol label switching (MPLS) transport entities between a first endpoint and a second endpoint;
determining an overall cost for each entity pair of said plurality of entities;
selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost; and
if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair,
wherein said entity pair reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.

Ex. 1001, Cl. 1; POR, 34

Doshi/Guichard Fail to Disclose the Reselection Limitation

Indisputably, Guichard does *not* disclose the use of entity pairs or of a protection entity; and Doshi does not disclose reoptimization of either an entity pair or of a protection entity. Petitioner's challenge based on Guichard and Doshi thus rests on two steps: *first*, combining Guichard's primary path reoptimization with Doshi's entity pair to produce a system in which the primary path is reoptimized, and, *second*, adapting this combination to practice reoptimization of both a primary path (as disclosed in Guichard) and a protection path (which is not disclosed in either Doshi or Guichard).

POR, 35; Ex. 2025, ¶ 83-84.


No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

Reoptimizing a Traffic Engineering LSP

Network state keeps changing. Links and nodes fail and recover; new TE LSPs are set up while others may be torn down. Consequently, for any TE LSP, a more optimal path may appear in the network. It is then highly desirable to detect the existence of such a path and reoptimize a TE LSP along a better path when it becomes available. Consider the example shown in Figure 2-15.

Ex. 1006, p. 79; Petition, 20-21

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection



No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

After you have determined the network recovery objectives, you should evaluate the various available network recovery mechanisms while keeping in mind various criteria. The convergence time is obviously the first one that comes to mind. How much time is required to reroute the traffic upon a network element failure? This is clearly a critical aspect to consider, but certainly not the only one. The following are a few other important aspects:

- **The scope of recovery**—Determines whether the recovery mechanism can handle link failure, Shared Risk Link Group (SRLG) failure, or node failure. As a reminder, the concept of SRLG relates to the notion of multiple links sharing a common resource whose failure would provoke all the links to fail simultaneously. For example, consider the case of multiple links routed across the same fiber. A fiber cut would imply the failure of all the links. We say that they share the same SRLG.
- **QoS during failure**—Does the alternate path (usually called a backup path) provide an equivalent QoS?
- **Network overhead**—This relates to the number of extra state required in the network by the recovery mechanism.
- **Cost**—There is a very wide range of network recovery mechanisms whose costs vary by several orders of magnitude. For instance, an optical 1+1 protection mechanism, although very efficient in terms of rerouting time and QoS along the backup path, requires additional equipment for the traffic replication and some extra bandwidth dedicated to protection (actually as much bandwidth as the protected bandwidth). On the other hand, other recovery mechanisms such as IP are usually cheaper both in terms of extra equipment and network resources (of course, they may not provide the same level of performance).
- **Network stability**—Does the network recovery under some circumstances (such as when very fast recovery is required) potentially lead to network instability when faced with frequent path changes?

This list is not meant to be exhaustive. It simply provides some of the most important aspects to evaluate when considering a network recovery mechanism. These different aspects are illustrated throughout this book in the various case studies.

In some cases, it might be desirable to elect to use a combination of recovery mechanisms, although such an option requires some extra care to avoid race conditions between recovery mechanisms and double protection, which may lead to a lack of optimality in terms of backup resources. For instance, a common approach consists of combining SONET protection with fine-tuning Interior Gateway Protocol (IGP) rerouting.

Ex. 2028, p. 96; POR, 21; Sur-reply, 13-14

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

Protection Versus Restoration

There are actually two families of network recovery mechanisms—protection and restoration. A protection mechanism relies on the precomputation (and signaling) of a backup (alternate) path before any failure. In contrast, a restoration mechanism requires the computation of the

backup path only after the failure has occurred; in other words, the backup path is computed on the fly. In terms of convergence time, protection mechanisms are usually faster but also require more network overhead because extra state is required.

Ex. 2028, p. 96-97; Sur-reply, 14

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

Publication Classification

- (51) **Int. Cl.**
G06F 11/00 (2006.01)
- (52) **U.S. Cl.** **370/216**

(57) **ABSTRACT**

A resilient load balancing method uses fixed paths and a fixed path-splitting strategy to enable ingress routers to efficiently reroute traffic after a failure. An off-line management system computes a set of fixed paths and a set of splitting ratios for routing demand from ingress routers to egress routers, with sufficient capacity to meet demands under each failure scenario. That data is then used by the ingress router to reroute demand after observing a failure.

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

[0030] In the most general sense, the present disclosure considers path protection and traffic engineering together to develop an optimal failure recovery plan. This is in contrast to most existing methods which consider them separately. The fast failure-recovery mechanism disclosed herein offers nearly optimal load balancing while using a static set of paths. Multipath routing between each pair of ingress and egress routers is utilized. Such end-to-end routing has several benefits. First, these routes do not need to change when a failure is detected, which saves time, reduces overhead, and improves path stability. Second, end-to-end load balancing spreads the traffic in the network more effectively than local rerouting. Third, it enables faster recovery and lower protocol overhead than conventional link-state routing protocols, like Open Shortest Path First (OSPF), which not only rely on flooding link-state advertisements and re-computing shortest paths but also suffer from transient forwarding loops. A fourth benefit is the solution's simplicity; most of the functionality is incorporated in the Network-Management Software (NMS) rather than the network elements. The NMS is responsible for selecting the end-to-end paths and calculating the path-splitting parameters for each router. The solution has a minimalist control plane, used only for failure detection, which leads naturally to a simpler network where smart, expensive routers are replaced with cheaper routers and/or switches with a limited feature set.

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

[0049] The set of failure scenarios S is a well defined “set” of “subsets” of links or nodes that are modeled for network restoration and resulting network availability. That is, these subsets of links and nodes are the deterministic failures that we want to assess how the network responds against. Of course, any subset has a non-zero probability of simultaneous failure, but because of the monstrous size of “all” subsets, for practicality the set is usually restricted to the most probable. The most probable are typically potential single upper layer failures (the individual links and nodes (singleton sets)) and in addition potential failure of a single link or node in a lower layer network. Because multiple links at the upper layer network of interest can route over a single lower layer link, a lower layer failures causes simultaneous failure of multiple upper layer links. Such groups of links are sometimes called “shared risk groups.”

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

[0056] The capability of the underlying routers, in part, determines the optimization problem the NMS solves. On one extreme, the network can support an optimal configuration of paths and splitting ratios for every network failure scenario $s \in S$. While not scalable in practice, the solution to this optimization problem serves as a performance baseline and as a way to compute a suitable set of paths P_d for each demand d . A more practical alternative, however, can comprise each ingress router u_d storing splitting ratios for every detectable failure $o \in O_d$. After observing the path failure(s), router u_d can switch to the new splitting configuration for the remaining paths. Alternatively, a single splitting configuration can be used across all failures. In this approach, router u_d simply renormalizes the splitting percentages for the active paths.

[0057] Algorithms for three scenarios to be solved by the NMS are discussed. Several of the optimization problems are NP-hard (as proven in the Appendix), and resolved with heuristics (presented later in the disclosure) that achieve nearly optimal performance in practice.

Ex. 1025 [Xu], [0056]-[0057]; POR, 43

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

[0010] In contrast to the level of protection provided by SONET, a typical mesh network provides a “path redial” protection scheme. That is, the call setup process for a connection is automatically re-initiated in the event of a failure. More particularly, when a fault is discovered in a single link between two nodes along a path between a source and a destination, the source node of the connection using the path determines an alternate path through the network and switches the traffic that was using the connection to the alternate path. The alternate path is chosen to avoid use of the faulty link.

[0011] When considering the granularity of choice of protection schemes that a network service provider may offer to customers, the SONET 1:1 protection scheme may be considered to be the best, though at a cost. The path redial protection scheme may be seen to be less desirable, as there are circumstances where a backup path may not be available and the time required to determine a backup path may be significant when compared to the SONET protection scheme. In summary then, the SONET 1:1 protection scheme may be considered to be a “platinum” level of protection service, the path redial protection scheme may be considered to be a “silver” level of protection service and no protection at all may be considered to be a “bronze” level of protection. It would then be in the best interest of network service providers if there were a “gold” level of protection service that could be offered to customers.

No Reason to Create or Adapt the Doshi/Guichard Combination to Practice Reselection

[0012] When considering the qualities of such a “gold” level of service protection, it would be preferable to have a faster switching time than path redial and better use of the reserved protection bandwidth than SONET. It would also be desirable that the “gold” level of service protection is interoperable with LDP and/or CR-LDP protocols.

[0013] The “gold” level of service protection should have support for maintenance switching, i.e., switching to avoid a node taken out of service for planned maintenance. For support by this “gold” level of service protection, the maintenance switching should be deterministic and independent of maintenance elsewhere in the network. That is, the route that any working traffic will take around a node taken out of service for planned maintenance should be known ahead of time so as to avoid maintenance taking place at other nodes.

[0014] The “gold” level of service protection should also have an ability to protect any single link failure. Additionally, the “gold” level of service protection should be arranged such that the network size is not limited by the signaling channel bandwidth. For example, the use of bytes K1 and K2 in the line overhead for automatic protection switching in SONET is known to limit the size of SONET rings to 16 nodes.

Ex. 1007 Huang], [0010], [0011]; POR, 43

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation

14. A system for selecting entities within an MPLS network, comprising:

- a data structure comprising a plurality of transport entity descriptors;

- an entity protection switch configured to switch between a working entity and a protection entity; and

- digital logic configured to select said working entity and said protection entity from said plurality of transport entity descriptors, comprising: logic configured to determine a probability of concurrent failure of said working entity and said protection entity;

- logic configured to determine an entity cost of said plurality of transport entity descriptors; and

- logic configured to reselect said working entity and said protection entity from said plurality of transport entity descriptors upon a reselection event,

wherein said reselection event is selected from a group consisting of adding an entity to said plurality of transport entities, removing an entity from said plurality of transport entities, an operational status change for one of said plurality of transport entities, and a change in overall cost for one of said plurality of transport entities.


Ex. 1001, Cl. 1; POR, 47-48

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation

restoration capacity. Normally the extra capacity is sized for recovering from only a single failure at a given time in the network. This is due to the fact that the probability of occurrence of more than one failure during the time interval required to repair a failure is very insignificant, and therefore most of the time it is sufficient to cover just a single network failure. This also provides opportunity to share restoration capacity between disjoint network failures. To understand sharing, realize that if two primary paths carrying traffic are routed disjointly, then a failure affecting one of them will not affect the other. Now, if the corresponding protection paths are routed such that they share common links then the restoration capacity on those links can be shared for restoring traffic between any failures along their primary routes. In general, sharing can be achieved by:

Ex. 1005, [0040]; Petition, 81; POR, 49

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation



Ex. 1005, Fig. 11; Reply, 23-24; Sur-reply, 19

[0005] In accordance with a first aspect of the disclosure, a method is provided for balancing load in a network after failure of a link. The network includes a plurality of interconnected vertices and edges $e \in E$ and a set of traffic demands $d \in D$. Each demand d defines a flow requirement h_d from an ingress router u_d to an egress router v_d . The network additionally includes a set of failure states s .

Ex. 1025 [Xu], [0005]; Sur-reply, 19-20

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation

[0002] Distributed computing networks, such as the Internet, comprise a vast number of routers and links connecting Network Access Devices (NADs). Routers are the electronic devices used to connect NADs to each other, or more likely to private or public networks. Links are the medium used to make the connection between routers, and ultimately users of a network. Examples of links are Ethernet cable, coaxial cable, Digital Subscriber Line, fiber-optic cable and a plurality of wireless technologies, just to name a few. The arrangement of routers and links, combined with support technologies and hardware, allows almost any machine (e.g. client, server, etc.) to access any machine in the same network. The exact path that a network communication session takes, i.e. the specific routers and links traversed, is flexible and may be determined by the Network Management System (NMS), and chosen depending on a variety of conditions, including proximity, availability and speed.

Ex. 1025 [Xu], [0002]; Sur-reply, 20

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation

[0006] At an ingress router u_d , traffic is routed by automatically balancing traffic load over a predefined set of paths P_d from the ingress router u_d to the egress router v_d according to a set of splitting ratios α_p^o . Each splitting ratio defines a fraction of the demands d to be transmitted over a path p of the set of paths P_d in a case of a detectable failure state o corresponding to a subset of a set of modeled failure states s . The set of modeled failure states are states that cause at least one path p of the predefined set of paths P_d to fail. The set of modeled failure states includes a no-failure state in which the network has no failures. Each modeled failure state has at least a predetermined probability of occurrence.

[0008] The method may also include pre-computing, in an off-line management system, the paths P_d and splitting ratios α_p^o . The step of pre-computing the fixed paths P_d may comprise obtaining an optimal flow on each edge for all modeled failure states s ; and using a decomposition algorithm to compute the corresponding fixed paths P_d through the network.

Ex. 1025 [Xu], [0006], [0008]; POR, 52

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation

[0030] In the most general sense, the present disclosure considers path protection and traffic engineering together to develop an optimal failure recovery plan. This is in contrast to most existing methods which consider them separately. The fast failure-recovery mechanism disclosed herein offers nearly optimal load balancing while using a static set of paths. Multipath routing between each pair of ingress and egress routers is utilized. Such end-to-end routing has several benefits. First, these routes do not need to change when a failure is detected, which saves time, reduces overhead, and improves path stability. Second, end-to-end load balancing spreads the traffic in the network more effectively than local rerouting. Third, it enables faster recovery and lower protocol overhead than conventional link-state routing protocols, like Open Shortest Path First (OSPF), which not only rely on flooding link-state advertisements and re-computing shortest paths but also suffer from transient forwarding loops. A fourth benefit is the solution's simplicity; most of the functionality is incorporated in the Network-Management Software (NMS) rather than the network elements. The NMS is responsible for selecting the end-to-end paths and calculating the path-splitting parameters for each router. The solution has a minimalist control plane, used only for failure detection, which leads naturally to a simpler network where smart, expensive routers are replaced with cheaper routers and/or switches with a limited feature set.

Ex. 1025 [Xu], [0030]; POR, 52


[0035] The network architecture of the present disclosure uses multiple paths between each ingress-egress router pair in the network. Using pre-established end-to-end paths allows fast failure recovery as the ingress router can shift the load away from the failed paths, avoiding dynamic path recalculation. Using multiple paths also allows the ingress router to balance the load in the network which helps to reduce network traffic congestion. By utilizing this architecture, the ingress router has a simple static configuration that determines the traffic-splitting ratios among the available paths, while intermediate routers merely forward packets over pre-established paths. As a result, the router is a much simpler device that does not need to collect congestion feedback, participate in a routing protocol, interact with the management system upon failure detection, or solve any computationally difficult problems.

Ex. 1025 [Xu], [0035]; POR, 52

[0049] The set of failure scenarios S is a well defined “set” of “subsets” of links or nodes that are modeled for network restoration and resulting network availability. That is, these subsets of links and nodes are the deterministic failures that we want to assess how the network responds against. Of course, any subset has a non-zero probability of simultaneous failure, but because of the monstrous size of “all” subsets, for practicality the set is usually restricted to the most probable. The most probable are typically potential single upper layer failures (the individual links and nodes (singleton sets)) and in addition potential failure of a single link or node in a lower layer network. Because multiple links at the upper layer network of interest can route over a single lower layer link, a lower layer failures causes simultaneous failure of multiple upper layer links. Such groups of links are sometimes called “shared risk groups.”


Ex. 1025 [Xu], [0049]; POR, 53-54; Sur-reply, 21

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation



Ex. 1025 [Xu], Fig. 1; Reply, 26

Doshi/Guichard/Xu Fail to Disclose the Concurrent Probability Limitation



Ex. 1025 [Xu], Fig. 1; POR, 24-25

Claims 2 and 18


2. The method of claim 1, wherein said step of selecting an entity pair further comprises:
selecting a working entity from said entity pair; and
selecting a protection entity from said entity pair.

18. The non-transitory computer readable media of claim 17, wherein said step of selecting an entity pair further comprises:
selecting a working entity from said plurality of transport entities;
selecting a protection entity from said plurality of transport entities; and
selecting an active entity from the set consisting of said working entity and said protection entity.


Ex. 1001 Cl. 2, 18

Claims 2 and 18

'821 Patent Claim 2



Doshi



Claim 13

13. The method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.

Reliance on background knowledge of a POSITA to supply a missing limitation “should generally only be done when ‘the [missing] limitation in question [is] unusually simple and the technology particularly straightforward’ and that ‘references to common sense ‘cannot be used as a wholesale substitute for reasoned analysis and evidentiary support.’”

Koninklijke Philips N.V. v. Google LLC, 948 F.3d 1330, 1338 (Fed. Cir. 2020) (quoting *Arendi*, 832 F.3d at 1361-1362)

POR, 57-58

“Absent some articulated rationale, a finding that a combination of prior art would have been ‘common sense’ or ‘intuitive’ is no different than merely stating the combination ‘would have been obvious.’”

In re Van Os, 844 F.3d 1359, 1361 (Fed. Cir. 2017)

POR, 59