

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS INC.,

Petitioner

IPR2024-01450

U.S. Patent No. 7,596,088

**DECLARATION OF NADER F. MIR, PH.D.,
UNDER 37 C.F.R. § 1.68 IN SUPPORT OF PETITION
FOR *INTER PARTES* REVIEW**

TABLE OF CONTENTS

I.	INTRODUCTION	4
II.	QUALIFICATIONS AND PROFESSIONAL EXPERIENCE	6
III.	LEVEL OF ORDINARY SKILL IN THE ART	10
IV.	RELEVANT LEGAL STANDARDS	12
V.	BACKGROUND	14
VI.	CLAIM CONSTRUCTION	15
VII.	EXPLANATION OF HOW THE CLAIMS ARE UNPATENTABLE.....	15
VIII.	DZIONG AND SATHYANATH RENDER OBVIOUS CLAIMS 1-4, 6, 8-11, 13, AND 15-20	16
A.	Summary of Dziong	16
B.	Summary of Sathyanath	19
C.	Combining Dziong and Sathyanath	20
D.	Detailed Analysis of Claims.....	21
1.	Claim 1	21
2.	Claim 2	73
3.	Claim 3	79
4.	Claim 4	81
5.	Claim 4	85
6.	Claim 6	89
7.	Claim 8	92
8.	Claim 9	103
9.	Claim 10	107
10.	Claim 11	107
11.	Claim 13	108
12.	Claim 15	109
13.	Claim 16	112

14.	Claim 17.....	113
15.	Claim 18.....	116
16.	Claim 19.....	117
17.	Claim 20.....	118
IX.	DZIONG, SATHYANATH, AND JENSEN RENDER OBVIOUS CLAIMS 7 AND 14.....	119
A.	Summary of Jensen	119
A.	Combining Jensen with Dziong and Sathyanath.....	119
1.	Claim 7.....	120
2.	Claim 14.....	126
X.	CONCLUSION.....	127

I, Nader F. Mir, Ph.D., do hereby declare as follows:

I. INTRODUCTION

1. I am making this declaration at the request of Cisco Systems, Inc. in the matter of the *Inter Partes* Review of U.S. Patent No. 7,596,088 (“the ’088 patent”) to Agmon et al.

2. I am being compensated for my work in this matter at my standard hourly rate. I am also being reimbursed for reasonable and customary expenses associated with my work and testimony in this proceeding. My compensation is not contingent on the outcome of this matter or the specifics of my testimony.

3. I have been asked to provide my opinions regarding whether the subject matter of claims 1-4, 6-11, and 13-20 (“the Challenged Claims”) of the ’088 patent would have been obvious to a person having ordinary skill in the art (“POSITA”) at the time of the alleged invention, in light of the prior art. It is my opinion that the Challenged Claims would have been obvious to a POSITA.

4. In the preparation of this declaration, I have studied:

- a. the ’088 patent, Ex. 1001;
- b. the prosecution history of the ’088 patent (“’088 File History”), Ex. 1002;
- c. U.S. Patent Application Publication No. 2005/0220026 to Dziong et al. (“Dziong”), Ex. 1005;

d. U.S. Patent Application Publication No. 2004/01900445 to Dziong and Sathyanath (“Sathyanath”), Ex. 1006; and

e. U.S. Patent Application Publication No. 2003/0198213 to Jensen et al. (“Jensen”), Ex. 1007;

5. In forming the opinions expressed below, I have considered: the documents listed above; the relevant legal standards, including the standard for obviousness; and my own knowledge and experience based upon my work in the field of network communications and security as described below, as well as portions of the following additional materials:

a. U.S. Patent Application Publication No. 2003/0172180 to Reynders et al., Ex. 1010;

b. U.S. Patent No. 6,256,309 to Daley et al., Ex. 1011;

c. U.S. Patent No. 6,850,976 to Xia (“Xia”), Ex. 1033;

d. U.S. Patent Application Publication No. 2005/0220123 to Wybenga et al. (“Wybenga”), Ex. 1034; and

e. Internet Engineering Task Force Request for Comments 3031 (“RFC3031”), Ex.1039.

6. Unless otherwise noted, all **emphasis** in any quoted material has been added. Claim terms are *italicized*.

II. QUALIFICATIONS AND PROFESSIONAL EXPERIENCE

7. My complete qualifications and professional experience are described in my *Curriculum Vitae*, a copy of which can be found in Exhibit 1004. The following is a brief summary of my relevant qualifications and professional experience.

8. I am a tenured professor in the Electrical Engineering Department of San Jose State University, California, where I also held the position of Associate Chair from 2006 through 2008. I have been a faculty member at San Jose State University since 2001.

9. In 1995, I received a Ph.D. in electrical engineering from Washington University, St. Louis. I received an M.S. degree in electrical engineering in 1990, also from Washington University. In 1985, I received a B.S. in electrical engineering from Polytechnic University.

10. I have over 30 years of experience with computer networks and communication systems, including both industry and academic experience. Before my current position at San Jose State University, I was an Assistant Professor in the Department of Electrical and Computer Engineering at the University of Kentucky. I joined San Jose State University in 2001 as an Associate Professor. Since 2005, I have been a Full Professor in the Electrical Engineering Department at San Jose State University, where I perform research and teach in the areas of

computer networks, networking devices, switches, routers, virtualization, virtual switches, cloud computing, software-defined networking in clouds, and cloud data centers.

11. I hold several technical editorial positions for various journals, including IEEE Communication Magazine. As a Technical Editor of IEEE Communications Standards Magazine for Cloud and Edge Computing, I am responsible for accepting or rejecting scientific articles submitted to the journal in the areas of computer networking focusing on cloud systems. From 2005 through 2020, I held the technical editorial position for IEEE Communications Magazine. I am a senior member of the IEEE and have served as a member of the technical program committees and the steering committees for a number of major IEEE communications and networking conferences.

12. For more than 30 years, I have conducted research in the areas of computer and communication networks; networking devices; protocols including (but not limited to) packet-switched networks, cloud and edge computing, software-defined networking in cloud systems, integrated voice, video, data center networks, TCP/IP, client/server operations, content delivery networks, multimedia and streaming, client/server, networking traffic flow analysis, wireless/mobile networks, networking devices, network security, and network virtualization, among others.

13. I have received a number of university awards as well as national and international awards. In particular, I have received a number of research grants from private, state, and governmental funding agencies for conducting research in the fields of computers and communication networks. I am also the recipient of several university teaching recognition awards from San Jose State University; several research excellence awards; and the school's published book award for the year.

14. Before my positions in academia, I worked in various positions in the telecommunication and computer network industries. From 1985 to 1988, I was a Telecomm R&D Engineer at the Telecommunication Research & Development Center (TRDC) in Surrey, England. In particular, I headed engineering projects in the area of telecommunication engineering. I also participated in the design of a high-speed switch for digital Private Branch Exchange (PBX) and SS7 networks. In 1991, I worked as an Electrical Engineering Consultant for Bussmann Company, where I instructed engineers on electrical engineering topics and prepared them for the National Principles and Practice of Engineering (PE) exam. From 1989 to 1995, I worked as a Computer/Communications Chip-Design Engineer at the Computer & Communications Research Center in St. Louis, Missouri. In that position, I designed a VLSI gigabit switch control plane chip for a high-speed computer network using CMOS technology. From 1995 to 1996, I was

a Research Scientist Engineer at Advanced Telecommunications Institute, Stevens Institute of Technology in Hoboken, New Jersey. There I headed analyzing communications and high-speed computer networking projects.

15. Since obtaining my Ph.D., I have served in various consulting positions in the area of computer networking technology and communication engineering, including for CITS Inc. in San Jose, California (2012-2013), Morrison & Foerster LLP in San Francisco, California (2015-2016), and Butler and Snow in Jackson, Mississippi (2017-2018).

16. I have authored or co-authored over 100 articles in peer-reviewed journals, conference proceedings, texts, and monographs in the aforementioned areas of computer networking, communications, and related fields. I have published two textbooks, one of which is entitled “Computer Communication Networks.” This textbook has been adopted by numerous universities worldwide.

17. I am a named inventor on one U.S. patent, which relates to a packet-switching network and method of routing packets.

18. I have served as an expert witness and technical consultant in litigation and *inter partes* review matters concerning computer networks, wireless networks at both the protocol and hardware levels, and telecommunications.

III. LEVEL OF ORDINARY SKILL IN THE ART

19. I understand there are multiple factors relevant to determining the level of ordinary skill in the pertinent art, including (1) the levels of education and experience of persons working in the field at the time of the invention; (2) the sophistication of the technology; (3) the types of problems encountered in the field; and (4) the prior art solutions to those problems.

20. The '088 patent “relates generally to communication networks, and particularly to methods and systems for selecting routes for communication connections that share bandwidth allocations in ring configurations.” '088 patent, 1:6-9. The '088 patent refers to a variety of known networking technologies and standards that the patent presumes the reader is already familiar with, including:

- Resilient Packet Ring (RPR), standardized in IEEE 802.17;
- Multiprotocol Label Switching (MPLS) protocol, described and standardized in various Internet Engineering Task Force (IETF) publications, including Request for Comments (RFC) 3031, RFC3036, and RFC3212;
- Resource Reservation Protocol (RSVP), including MPLS traffic engineering extensions to RSVP, described and standardized in IETF publications including RFC2205 and RFC3209;

- Quality of Service frameworks such as the Integrated Services (IntServ) protocol architecture, described and standardized in IETF publications including RFC1633;
- Synchronous Optical Network (SONET) technology used in ring networks;
- Internet Protocol (IP) routing protocols, such as Open Shortest Path First (OSFP) and Intermediate System to Intermediate System (IS-IS) protocols;
- Traffic Engineering extensions to OSFP and IS-IS; and
- Asynchronous Transfer Mode (ATM) protocol.

See '088 patent, 1:10-61, 2:4-6, 5:2-3, 5:20-26, 6:23-26.

21. A POSITA in the field of the '088 patent, as of its filed date of January 24, 2006, would have been someone knowledgeable and familiar with the above list of technologies and protocols. Such a POSITA would have a bachelor's degree in computer science, computer engineering, electrical engineering, or equivalent training, and approximately four years of experience working in the field of network communications, including experience with MPLS tunnels and bandwidth allocation. Additional work experience can substitute for specific educational background, and vice versa. In 2006, my own level of education and experience exceeded that which I have identified for a POSITA. Thus, as of the priority date of the '088 patent, I was qualified as a person of ordinary skill.

22. For purposes of this Declaration, in general, and unless otherwise noted, my statements and opinions, such as those regarding my own experience and what a POSITA would have understood or known generally (and specifically related to the references I consulted herein), reflect the knowledge that existed in the relevant field as of the priority date of the '088 patent.

IV. RELEVANT LEGAL STANDARDS

23. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '088 patent, I am relying on certain basic legal principles that Cisco's counsel has explained to me.

24. I understand that prior art to the '088 patent includes patents and printed publications in the relevant art that predate the priority date of the '088 patent. For purposes of this Declaration, I am applying January 24, 2006, as the priority date of the '088 patent.

25. I have been informed by Cisco's counsel that a claimed invention is unpatentable for obviousness if the differences between the claimed invention and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a POSITA. I have also been informed by Cisco's counsel that the obviousness analysis considers factual inquiries, including the level of ordinary skill in the art, the scope and content of the prior art, and the

differences between the prior art and the claimed subject matter. In forming the opinions expressed in this Declaration, I have considered these factual inquiries.

26. I have been further informed by Cisco's counsel that the disclosures of multiple prior art references may be considered in combination, but that the rationale for doing so should be expressly laid out. I understand that patent law recognizes multiple rationales for combining references or modifying a reference to show obviousness. These rationales include: (a) combining prior art elements according to known methods to yield predictable results; (b) simple substitution of one known element for another to obtain predictable results; (c) use of a known technique to improve a similar device (method, or product) in the same way; (d) applying a known technique to a known device (method, or product) ready for improvement to yield predictable results; (e) choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; and (f) some teaching, suggestion, or motivation in the prior art that would have led a POSITA to modify the prior art or to combine prior art teachings to arrive at the claimed invention. I understand that these enumerated rationales are just examples and that the combination of multiple references can be supported by any reason that has a rationale and factual underpinning.

V. BACKGROUND

27. The '088 patent presumes the reader is familiar with various networking concepts, which would have been part of the background knowledge of a POSITA. For example, a POSITA would have known and understood at least the following networking technologies:

- Multi-Protocol Label Switching (MPLS), described in RFC3031 and RFC3036 (see '088 patent, 1:19-30);
- Resource ReSerVation Protocol (RSVP), described in RFC2205 (see '088 patent, 1:46-52);
- Traffic Engineering extensions to RSVP, known as RSVP-TE, described in RFC3209 (see '088 patent, 1:42-46);
- Ring-topology networks, such as those used with Synchronous Optical Network (SONET) networks (see '088 patent, 6:23-26); and
- MPLS applied in the context of a ring network (see '088 patent, 4:64-67).

28. A POSITA also would have been familiar with the basic components of equipment capable of communicating on a network and performing routing protocols. For example, a POSITA would have known that device that communicates over a network must have a network interface to facilitate those

communications. Similarly, a device that performs a path-selection process must have a processor that is configured to perform the steps of the path-selection process. Most commonly, this would have been a general-purpose computer processor, and the path-selection process would be written as software executed by the processor.

VI. CLAIM CONSTRUCTION

29. It is my understanding that to properly evaluate the '088 patent, the terms of the claims must first be interpreted. I understand that in the proceeding to which this Declaration is being submitted, the words of the claims are given their ordinary and customary meaning as would have been understood by a POSITA in light of the specification and prosecution history, unless the inventor has set forth a special meaning for a term. I have also been informed that claim terms only need to be construed to the extent necessary to resolve the obviousness inquiry. I have reviewed the entirety of the '088 patent, as well as its prosecution history. It is my opinion that the obviousness of the claims can be evaluated without construing any claim terms.

VII. EXPLANATION OF HOW THE CLAIMS ARE UNPATENTABLE

30. The discussion in this Declaration provides a detailed analysis of how the asserted prior art references disclose or render obvious each limitation of the Challenged Claims.

31. As part of my analysis, I have considered, and discuss in detail, the scope and content of the prior art and any differences between the alleged invention and the prior art.

32. It is my opinion that the alleged invention recited in the Challenged Claims would have been obvious in view of the prior art discussed below and the knowledge of a POSITA.

VIII. DZIONG AND SATHYANATH RENDER OBVIOUS CLAIMS 1-4, 6, 8-11, 13, AND 15-20

A. Summary of Dziong

33. U.S. Patent Application Publication No. 2005/0220026 to Dziong et al. (Ex. 1005, “Dziong”) was filed April 2, 2004, and published October 6, 2005. Dziong is titled “Calculation of Link-Detour Paths in Mesh Networks.”

34. Like the '088 patent, Dziong recognized the potential value of exploiting “bandwidth-sharing opportunities.” Dziong, [0029]; *see* '088 patent at 5:30-51. In one aspect, Dziong generally relates to protecting communication paths against potential failures by establishing alternate pathways that allow traffic to be rerouted around a failed node or link. Dziong, [0006]-[0008]. Dziong explains that there existed a variety of technologies for protecting communication paths against potential failures, but these existing technologies made inefficient use of bandwidth:


[0036] Another problem, associated with some current link-based recovery mechanisms is a failure to take advantage of recovery bandwidth-sharing opportunities. For example, consider topological rings 102 and 104 depicted in FIG. 1. Ring 102 (A-B-C-D-A) (with two units of capacity) and ring 104 (B-C-F-E-B) (with four units of capacity) have the link between nodes B and C (i.e., link B-C) in common. Note that, in this example, the capacity of a ring is limited to the capacity of the link in the ring with the lowest capacity. In accordance with the current SONET and SDH ring standards, ring 102 can use one unit of bandwidth for working traffic, while reserving one unit of bandwidth for protection. Similarly, ring 104 can use two units of bandwidth for working traffic, while reserving two units of bandwidth for protection. Since link B-C is common to both rings and thus carries three units of working traffic, it should reserve three units of protection bandwidth to protect against failures of other links in the two rings. This equates to providing sufficient recovery bandwidth on link B-C to accommodate a failure of at least one link (other than B-C) in each of ring 1 and ring 2 simultaneously.

[0037] However, modem-day networks have very high reliability and typically a very fast repair interval (i.e., the time it takes to recover from a single failure, restore service, fix the failure, and switch back to the original configuration--if that is part of the protocol, or at least reserve new recovery paths based on the modified configuration). In the present invention, this reliability is taken into account by assuming that, since the probability of experiencing a second failure during the recovery

interval following an initial failure is insignificant, the probability of two or more co-existing failures can essentially be ignored.

[0038] Considering this, reserving separate capacity for each ring, ring 102 and ring 104 in our example, is wasteful of resources. In the above scenario, this equates to the assumption that the reserved bandwidth in the network need only accommodate a failure of a link of ring 102 or a link of ring 104, but not both simultaneously. With this assumption, the bandwidth reserved on link B-C to cover a single failure on this two-ring network need only be two units (as opposed to three). In the event of a failure of any one of the other links of ring 102, one unit of the reserved bandwidth along link B-C can be used for recovery purposes. Similarly, in the event of a failure of any one of the other links of ring 104, both units of the reserved bandwidth along link B-C can be used for recovery purposes. Thus, the recovery bandwidth reserved on link B-C is shared between the two rings, yielding a more efficient use of network resources.

Dziong, [0036]-[0038].



Dziong, Fig. 1.

35. Dziong uses a variety of terms to refer to a path that provides protection in the event of a failure on another path, including “recovery paths” ([0009]), “protection paths” ([0009]), and “link-detour paths” ([0009]). To the extent that Dziong differentiates among these terms, their differences are not relevant to comparing Dziong’s technology to the ’088 patent. *See* Dziong, [0083] (“the terms ‘restoration’ and ‘protection’ are often used interchangeably in the art”), [0034] (“The present invention can use either or both types of recovery mechanism though protection is preferred.”).

B. Summary of Sathyanath

36. U.S. Patent Application Publication No. 2004/0190445 to Dziong and Sathyanath (Ex.1006, “Sathyanath”) was filed September 26, 2003, and published

September 30, 2004. Sathyanath is titled “Restoration path calculation in mesh networks.”

37. Like the '088 patent and Dziong, Sathyanath describes improving network performance and capacity through the “efficient sharing of protection capacity.” Sathyanath, [0012]. Sathyanath expands on Dziong’s disclosure by providing further details regarding a “link-cost function” used to evaluate the cost of a potential link or path. *See* Dziong, [0104], (“...the recovery path is calculated using the... link-cost function...”), [0113] (“the link-cost function can take into account several metrics...”). Specifically, Sathyanath explains that a link-cost function may compute “the incremental bandwidth that needs to be reserved” for a recovery path. Sathyanath, [0205]. Sathyanath’s link-cost function leads to “choosing a restoration path that would require the minimum amount of additional bandwidth to be reserved along its links.” Sathyanath, [0200].

C. Combining Dziong and Sathyanath

38. The first paragraph of Dziong’s specification expressly refers to, and incorporates by reference, Sathyanath by referring to Sathyanath’s patent application number (10/673,057). Dziong, [0001]; Sathyanath, title page item (21). Thus, a person reading Dziong would have been expressly instructed to refer also to the disclosure of Sathyanath, and as a result, would have found it obvious to consider them together. In particular, a POSITA evaluating Dziong would have

found Sathyanath's additional examples to be informative of the way in which Dziong contemplated choosing recovery paths to maximize the amount of bandwidth sharing in a network. *See* Dziong, [0060] ("maximizing sharing of link-detour bandwidth").

39. A POSITA also would have found it obvious to combine features from the various examples provided in Dziong and Sathyanath. Indeed, they generally encourage this, as they both explain that their described embodiments are not "mutually exclusive" and that they contemplate "modifications of the described embodiments." Dziong, [0027], [0196]; Sathyanath, [0037], [0295].

D. Detailed Analysis of Claims

1. Claim 1

- a. [1.0] *A method for establishing a connection through a communication network that includes multiple segments, comprising:*

40. To the extent that the preamble is limiting, Dziong and Sathyanath render it obvious.

41. Dziong describes "**methods** for calculating and distributing link-protection parameters." Dziong, [0010]; *see also* Dziong, claim 1 ("A link-based recovery **method for a mesh network**, the method comprising...").

42. Dziong describes establishing a link-detour path for a connection: [0058] Embodiments of the present invention may include a distributed method for calculating primary and link-detour paths in a mesh

network. This method improves the number of connections admitted to a network and reduces the probability of crank-backs due to unavailable or overutilized link-detour paths.... A source node may assume that bandwidth for a new connection is available, and then start to **signal to establish the primary path and link-detour paths for the connection.**

Dziong, [0058].


43. Dziong's link-detour path is itself a "*connection*." Dziong also refers to the link-detour path as a protection path or recovery path. *See* Dziong [0058], [0008] ("A link-detour path is that portion of a link-based recovery path that corresponds to the failure."), [0009] ("recovery paths (also known as protection paths)"). Note that Sathyanath also introduces the further synonymous term "restoration path." Sathyanath, [0046] ("the process of 'restoration' is synonymous in this specification with the process of 'protection.'"), [0045] ("Sharing can be achieved... along the restoration path...").

44. Sathyanath further describes establishing a protection path ("*establishing a connection*") in a network using label-switched paths provided by MPLS protocol:

[0250] In case of the centralized server approach, a centralized server computes the additional protection bandwidth needed on each link along the protection path to admit a new connection (LSP). This information is then conveyed as part of network signaling to **establish the protection LSP.**

Sathyanath, [0250]. Since Dziong does not specify a protocol used to establish link-detour paths, it would have been obvious to use MPLS as described by Sathyanath.


45. Dziong and Sathyanath provide multiple examples of a “communication network that includes multiple segments.” For example, Dziong illustrates in Figure 5 (below) that there are multiple links providing communication paths from a source node to destination node. Dziong, [0070]. These links, labeled with the letters A through Q, are “segments.” See ’088 patent, 4:34-35 (“network links, referred to as segments”).



Dziong, Fig. 5 (annotated).

46. Dziong illustrates another “communication network that includes multiple segments” in Figure 6a (below). In the figure, four nodes (labeled A through D)¹ are interconnected by links (“multiple segments”).

¹ Note that in contrast to Dziong’s Fig. 5 example, where the labeling letters correspond to links), the labeling letters in Dziong’s Fig. 6a example correspond to nodes.



Dziong, Fig. 6a (annotated).


[0089] In network 602, each solid line represents all of the demands that have the same link-detour path, which is represented by a corresponding broken line. For example, solid line 606 in SRLG 4 corresponds to one or more demands totaling 15 units of bandwidth between nodes A and D. Broken lines 608 between nodes A and B, nodes B and C, and nodes C and D correspond to the common link-detour path for those demands.

Dziong, [0089].

47. Dziong's example of the recovery path A-B-C-D (broken lines 608 that "correspond to a link-detour path") is a "connection through a communication network that includes multiple segments." Dziong explains that the example network in Fig. 6a is meant to illustrate the path protection concepts that Dziong discloses. Dziong, [0092] ("The example shows one-link-long connections and their protection paths that, in effect, illustrate link-based protection concepts.").

Thus, it would have been obvious to a POSITA for the recovery path A-B-C-D to have been established using the recovery path selection techniques described in Dziong and Sathyanath.

48. Sathyanath shows in Fig. 17 yet another example of “a communication network that includes multiple segments.” In the figure, each of the links L1 through L6 is a “segment[.]” As I annotate below, the links are arranged in pairs (L1-L2, L3-L4, and L5-L6) to form multiple pathways from the source (R1) to the destination (R5). Each pathway has “multiple segments.”



Sathyanath, Fig. 17 (annotated).

[0215] As an example, consider the partial network 1700 of FIG. 17. Consider the primary path L1-L2 with four units of bandwidth. Also, assume that there are two possible disjoint routes (with respect to the

primary path) to the destination, namely route A (L3-L4) and route B(L5-L6).

Sathyanath, [0215].

49. In summary, Dziong and Sathyanath describes methods for establishing a recovery path in a network that includes multiple nodes and links among those nodes, which renders obvious “*method for establishing a connection through a communication network that includes multiple segments*” as claimed.

b. [1.1] *identifying two or more candidate routes for a new connection having a requested bandwidth to be established through the network;*

50. Dziong and Sathyanath render this limitation obvious.

51. Dziong describes evaluating and choosing from among multiple “*candidate routes*” in the form of potential recovery paths (a.k.a. link detour “LD” paths):

3. The invention of claim 2, wherein computing an LD path comprises: computing a sum of the link costs associated with each link j in **two or more candidate LD paths** for link i of the first primary path to determine an LD path link cost sum; and **choosing the LD path** for the link i of the first primary path to be the candidate LD path with the smallest LD path link cost sum.

Dziong, claim 3.


52. Dziong further describes “choosing the path with the lowest cost.”

Dziong, [0065], [0104], [0168]. The act of “choosing” a path shows that there would have been multiple potential paths previously identified.

53. Dziong further renders obvious “*identifying two or more candidate routes*” through its discussion of Figure 5. In the figure, which I annotate below, Dziong describes multiple recovery paths that may be considered to protect against a potential failure of link A:

[0070] FIG. 5 illustrates a simple network with both path-based and link-based recovery paths. For path-based recovery, two link- and node-disjoint paths between the source and destination nodes are shown. One is primary path A-B-C-D and the other is recovery path E-F-G-H-I. For link-based recovery, each link on the primary path has its own link-detour (LD) path that is defined by its source node (the link's upstream node), its destination node (the link's downstream node), and a set of transit nodes. For example, in the case of a failure of link A, demands that were carried on that link can be rerouted to LD path J-K. Alternatively, though much less efficiently, demands on link A could have been rerouted along LD path E-F-G-H-I-D-C-M-L. Other alternative routes could also be used, contingent on any hop limit imposed on the network.

Dziong, [0070].




Dziong, Fig. 5 (annotated).

54. In the event of a failure of link A, Dziong contemplates a link-based recovery path that substitutes link detour path J-K for the original link A. Thus, first recovery path from the source to the destination would be J-K-B-C-D. Dziong also describes a potential path-based recovery path, E-F-G-H-I. These are each “candidate routes.” Dziong further describes yet another link-detour path, E-F-G-H-I-D-C-M-L, for routing around a link A failure. Thus, a third source-to-destination “candidate route” would be E-F-G-H-I-D-C-M-L-B-C-D (which Dziong acknowledges would be “much less efficient[.]”). Dziong, [0070].

55. Dziong also illustrates multiple “candidate routes” in Figure 6a, which I have annotated below. As discussed at [1.0], the link from node A to node D is referred to as shared-risk link group (SRLG) 4. To protect against a potential failure of this link, it would have been immediately obvious to a POSITA that there are two potential recovery paths from node A to node D: first, the recovery path 608 (through nodes A-B-C-D, marked with a dotted blue line below) (see Dziong,

bandwidth”), [0105] (“in some cases, a request for a protection path for a link can be rejected due to lack of available bandwidth on a candidate protection link”).

57. Dziong shows in Figs. 6b that, indeed, 15 units of protection bandwidth are ultimately reserved on link A-B to protect against a potential failure of SRLG 4 (i.e., the link A-D). Thus, it would have been obvious to a POSITA that when SRLG 4 was created and Dziong’s method was used to establish a recovery path to protect against a failure of SRLG 4, the recovery path was “*requested*” to have a bandwidth of 15 units:



Dziong, Fig. 6b (annotated).

58. As I discussed at [1.0], Dziong renders obvious “*establishing a connection*” in the form of a recovery path. Dziong explains, for example, that a

recovery path (a.k.a. protection path) is set up to protect against a failure of a primary path providing a new service:

[0052] FIG. 3 depicts an exemplary process for achieving link-based protection at the demand level. As shown, in step 302, **a working path for a new service is computed along with protection paths for each link in the working path.** In order to admit a new service into the network, there should be sufficient capacity in the network to admit the new service along the working path and also guarantee the service's recovery from any single failure along its route.

Dziong, [0052].

59. Dziong's protection path for a new service renders obvious a "*new connection... through the network.*" The '088 patent confirms that a path set up to protect against a failure elsewhere in the network is a "*new connection.*" See '088 patent, 5:5-8 ("...the new connection may be set up to protect the existing connection against failure of a node or segment along its route.").

60. Sathyanath provides still further examples showing the obviousness of the "*identifying*" limitation. For example, Sathyanath describes using extensions to the well-known OSPF protocol to share information used to identify an optimal protection path.² Sathyanath's reference to an "optimal" path shows that there

² As I noted previously, Sathyanath uses the terms "restoration path" and "protection path" interchangeably. Sathyanath, [0078] ("protection path (also known as 'restoration path')").

would have been other, suboptimal paths, and hence, “*two or more candidate routes*”:


[0187] The following describes how the concepts of shared restoration maybe applied to EP networks by appropriate extension of the OSPF/OSPF-TE protocols. These extensions allow OSPF/OSPF-TE to disseminate sharing information to nodes in an IP network and enable nodes to **identify**, for each new service, an optimal combination of a primary path and **a restoration path**, from source to destination, such that the amount of bandwidth reserved on the restoration path is minimized.

Sathyanath, [0187].

61. Sathyanath’s Fig. 17 specifically shows an example network with a primary path L1-L2 and two candidate recovery paths:

[0215] As an example, consider the partial network 1700 of FIG. 17. Consider the primary path L1-L2 with four units of bandwidth. Also, assume that there are **two possible disjoint routes** (with respect to the primary path) to the destination, namely **route A (L3-L4) and route B (L5-L6)**.

Sathyanath, [0215].



Sathyanath, Fig. 17 (annotated).

62. Since the primary path L1-L2 has “four units of bandwidth” (Sathyanath, [0215]), it would have been obvious to a POSITA for the potential recovery routes to have an equal amount of “*requested bandwidth*.” Sathyanath confirms that the recovery routes are requested to have four units of bandwidth by stating that route B requires “no additional units of bandwidth... since 4 units have been reserved already.” Sathyanath, [0222].

63. In summary, Dziong teaches evaluating multiple potential protection paths to protect against a failure on a primary path providing a new service at a specified bandwidth, which renders obvious “*identifying two or more candidate*

routes for a new connection having a requested bandwidth to be established through the network.”

- c. **[1.2] providing a sharing definition specifying one or more other connections, which have respective bandwidths and which are permitted to share a bandwidth allocation with the new connection,**

64. Dziong and Sathyanath render this limitation obvious.

65. Dziong repeatedly describes the benefits of sharing bandwidth among multiple recovery paths:

[0036] Another problem, associated with some current link-based recovery mechanisms is a failure to **take advantage of recovery bandwidth-sharing opportunities.**

Dziong, [0036].

[0060] In embodiments of the present invention, link-detour path calculation involves **maximizing sharing of link-detour bandwidth.**

Dziong, [0060].

66. To achieve this bandwidth sharing, Dziong describes how nodes will advertise to one another information about their connections, including information about what bandwidth is available for sharing:

[0054] Next, in step 306, sharing between disjoint link failures is achieved by determining, via bookkeeping information, the amount of protection bandwidth that would need to be reserved on each link for recovering demands affected by any single other link failure in the network. Recovery of each other link may require a different amount of

capacity. On each link that is part of a recovery path, the maximum of required recovery capacities required on that link is calculated. This maximum is then reserved on the link if sufficient capacity exists on the link. For a distributed implementation, **each node keeps track of this sharing information** for each of its incident links. **Signaling is used to update this sharing information** with admittance of every new demand into the network. In the case of SONET/SDH networks, the reservation information is kept in terms of time slots associated with the demands, though other schemes are possible.

Dziong, [0054].

The recovery-path calculation algorithm makes use of information including how much bandwidth can be-shared at each link in the network. This information is obtained by **advertising, across the network, the amount of bandwidth reserved for recovery** on each link and by bookkeeping, in each node, information about all recovery paths that would be activated when a protected link fails.

Dziong, [0060].

67. Dziong's discussion of advertising shared bandwidth information renders obvious "*providing a sharing definition.*" Dziong explains that each node, upon receipt of sharing information from other nodes, performs "bookkeeping" so that its local store of bandwidth sharing information includes "information about all recovery paths" in the network. Dziong, [0060]. Thus, it would have been obvious to a POSITA that Dziong's bandwidth sharing information includes information about "*one or more other connections.*" Dziong also states that the


advertised information includes the “amount of bandwidth reserved for recovery on each link,” and thus, Dziong teaches that these “*other connections...have respective bandwidths.*”

68. Dziong illustrates in Fig. 6b a “bandwidth reservation table 604 for link A-B of FIG.6(a).” Dziong, [0088]. Thus, Fig. 6b illustrates the information that Dziong contemplates being advertised throughout the network (“*a sharing definition*”). See Dziong, [0064] (“each node in the network advertises to other nodes in the network the amount of protection bandwidth it currently has reserved on each of its incident links”). The bandwidth reservation table shows that the three protection paths on link A-B (which protect against failures of SRLG 3, 4, and 5) “*are permitted to share a bandwidth allocation*”:

[0090] Bandwidth reservation table 604 describes the current state of the bandwidth of link A-B in terms of cross-connection (XC) units. In particular, "P-XC" represents the number of XC units in link A-B that currently support demands (i.e., 12 units in this example). "SRLG 3 failure" represents the number of XC units in link A-B that are reserved to protect demands on SRLG 3 (i.e., 48 units). Similarly, "SRLG 4 failure" and "SRLG 5 failure" represent the numbers of XC units in link A-B that are reserved to protect demands on SRLG 4 (i.e., 15 units) and SRLG 5 (i.e., 36 units), respectively. Note that link A-B is not part of the link-detour path for the other 24 XC units of demand on SRLG 4 (i.e., represented by solid line CE).

[0091] Table 604 is based on the assumption that **protection bandwidth is shared** for SRLG-disjoint failures. As such, **RSRV-XC** represents the **actual amount of bandwidth that needs to be reserved on link A-B to protect against any one SRLG failure** in network 602. In the current example, this corresponds to the maximum of the protection bandwidths required by each of SRLGs 3, 4, and 5 protected by link A-B (i.e., 48 units). Thus, link A-B provides protection bandwidth that is shared between disjoint SRLGs 3, 4, and 5.

Dziong, [0090]-[0091].




Dziong, Fig. 6b (annotated).

69. In addressing limitation [1.2], I discussed how Dziong's method of reserving shared bandwidth would have been used to create a recovery path to

protect against a failure of link A-D (i.e., SRLG 4). In this example, the recovery path A-B-C-D is the “*new connection.*” Before SRLG 4 was established,³ it would have been obvious to a POSITA that the bandwidth reservation table for link A-B would include protection bandwidth reservations for only SRLG 3 and SRLG 5. I have prepared below an example of how Dziong’s bandwidth reservation table for link A-B would look like in this situation. The protection bandwidth reservations for SRLG3 and SRLG 5 correspond to the claimed “*other connections... which are permitted to share a bandwidth allocation.*”

³ Dziong does not expressly state the order in which the three protection paths across link A-B were established. There are only a handful of ways that they could have been sequenced, and in any event, my analysis here does not meaningfully depend on that order. Because of Dziong’s teaching of bandwidth sharing among protection paths, the order of their establishment is not important—regardless of the sequence, the total reserved bandwidth (after all three protection paths are established) would be 48 units, as shown in Fig. 6b. I have simply chosen to treat the SRLG 4 protection path as a last-to-be-established path to further explain how Dziong’s bandwidth sharing concept would work.



Dziong, Fig. 6b (modified to remove protection bandwidth allocated for SRLG 4).

70. Thus, Dziong illustrates in Fig. 6b information about the protection bandwidth that is reserved on link A-B. Dziong explains that this information is advertised to other nodes in the network. Dziong, [0064] (“each node in the network advertises to other nodes in the network the amount of protection bandwidth it currently has reserved on each of its incident links”). Thus, it would have been obvious to a POSITA for this modified bandwidth reservation table to have been advertised across Dziong’s network in Fig. 6a.

71. My analysis above treats the “*sharing definition*” as reflecting the information that would be advertised across the network before the claimed “*new connection*” is established. I note, however, that the language of claim 1 continues in the next limitation ([1.3]) to recite that the *sharing definition* is “*such that a total bandwidth allocated jointly to the new connection and to the other connections*” has a certain value. To the extent that the claimed “*sharing definition*” is interpreted as reflecting the state of bandwidth allocations after the “*new connection*” is established, Dziong still renders the claim limitation obvious. Dziong teaches that bandwidth sharing information is “flooded to the network” “upon the occurrence of a large change in services on the network or other periodic timer.” Dziong, [0099]. This could include, for example, upon the establishment of “a new primary path in the network.” Dziong, [0099]. Thus, it would have been obvious to a POSITA for the bandwidth reservation information in Fig. 6b (without modification) to be advertised across the network following the establishment of SRLG 4 and its corresponding recovery path.

72. Dziong provides relatively few details about the data format that would be used to advertise the bandwidth sharing information. Sathyanath, however, provides further details about this. Like Dziong, Sathyanath contemplates that “sharing information... is advertised, using a data network, to other nodes in the network.” Sathyanath, [0084]. Sathyanath describes a variety of ways that


bandwidth sharing information maintained and shared, including an “aggregate node-link vector representation” (exemplified in Fig. 6). Sathyanath, [0102]-[0103]. To facilitate bandwidth sharing, “each node disseminates aggregate link-node vectors for its connected links.” Sathyanath, [0105]. The aggregate link-node vector provides “information about which links can share the restoration bandwidth.” Sathyanath, [0122]. Thus, Sathyanath’s aggregate link-node vector also corresponds to the claimed “*sharing definition*.”

73. Sathyanath describes an example of an aggregate link-node vector in conjunction with Figures 5 and 6:

[0021] FIG. 5 depicts an exemplary shared mesh data network supporting multiple services via disjoint primary and restoration paths.

[0022] FIG. 6 depicts node-link aggregate vector and node-aggregate vector representations of sharing information.

Sathyanath, [0021]-[0022].




Sathyanath, Fig. 5.

74. More specifically, Fig. 6 shows an exemplary “aggregate link-node vector V_{nla} for link L5 associated with the exemplary network 500 of FIG. 5.” Sathyanath, [0103]. The aggregate link-node vector includes information on other connections and their respective bandwidths:

[0104] For this example, it is assumed that the bandwidth demand of each service is of unit value, however, extension to the more general case of non-unit bandwidth demands would be understood to one skilled in the art. As illustrated, the value of the vector in the position corresponding to node NA is 1. This indicates that, for a failure of node NA, an aggregate bandwidth demand of 1 (associated with protection path R1) will be placed on link L5 corresponding to one (unit bandwidth) service. Similarly, the 2 in the position of V_{nla} corresponding to node NB indicates that for a failure of node NB, link

L5 will be required to support two aggregate (unit bandwidth) services paths (i.e., associated with protection paths R1 and R2). Similarly, **potential failures of links L1, L2, L3, and L4 correspond to aggregate bandwidth reservation values in V_{nla} of 1, 2, 3, and 1, corresponding to restoration paths R1, R1+R2, R2+R3+R4, and R1, respectively.** Note that the bandwidth that is reserved on link L5 is the maximum of the bandwidth requirements associated with any single network component failure. In the example network of FIG. 5, the reservation bandwidth is thus $\max \{1, 2, 1, 2, 3, 1\}=3$.

Sathyanath, [0104].



Sathyanath, Fig. 6 (partial, annotated).

75. Thus, the restoration paths identified in the aggregate link-node vector are “other connections, which have respective bandwidths.” Since the actually reserved bandwidth is the “maximum of the bandwidth requirements associated with any single network component failure” it would have been obvious that all of

the restoration paths share a single bandwidth reservation. Sathyanath, [0104].

Sathyanath confirms this by explaining instances where “the restoration path for the new service can share the bandwidth reserved for the restoration paths of the existing services.” Sathyanath, [0120]. Thus, the restoration paths “*are permitted to share a bandwidth allocation.*”

76. Sathyanath describes other data formats, including a “node-aggregate vector” (exemplified in Fig. 6), a “binary node-link vector” (exemplified in Fig. 7), a “node-binary vector” (exemplified in Fig. 7), and “alternative representations of the link-state sharing information... as would be understood to one skilled in the art.” Sathyanath, [0102]-[0103], [0106]-[0109], [0111]-[0112], [0114]; *see also* [0226]-[0243] (describing “Formats for Sharing” data via a link-state advertisement protocol). These other data formats further render obvious the claimed “*sharing definition.*”

77. Finally, Sathyanath describes a “TE/Share database” that is used to store bandwidth sharing information at each node. The information in the TE/Share database is advertised throughout the network:

[0194] In this scheme, in order to support sharing/restoration, every node maintains a local sharing information database (e.g., the **TE/Share database**). This database contains a data structure for each link L that the node owns. The data structure for each link contains information (e.g., bandwidth, link-id) about all the other links in the

network for which link L provides restoration capacity. Since this could be a large amount of information for a large network, only **changes in this database are advertised.**

Sathyanath, [0194].

78. Sathyanath provides an example set of entries for a “TE/Share database” for the network shown in Fig. 17. *See* Sathyanath, [0215]-[0220]. As I discuss further in my analysis of [1.3] below, the entries identify the amount of shareable protection bandwidth and the links (“*other connections*”) for which that bandwidth is already reserved (“*which have respective bandwidths and are permitted to share...*”). Thus, Sathyanath’s TE/Share database is a further example of a “*sharing definition*” as claimed.

79. In summary, Dziong and Sathyanath describe numerous examples of advertising across the network the amount of shared protection bandwidth allocated to recovery paths, which renders obvious “*providing a sharing definition specifying one or more other connections, which have respective bandwidths and which are permitted to share a bandwidth allocation with the new connection.*”

- d. **[1.3] such that a total bandwidth allocated jointly to the new connection and to the other connections in any given segment is equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment;**

80. Dziong and Sathyanath render this limitation obvious.


81. Dziong states that the “amount of bandwidth actually reserved on each incident link is the maximum of the bandwidths required to recover from any of the failures.” Dziong, [0063]. Dziong’s “amount of bandwidth actually reserved” corresponds to the claimed “*total bandwidth allocated*.” Dziong describes how reserving this amount of protection bandwidth “guarantee[s] the service’s recovery from any single failure along its route.” Dziong, [0052]. Thus, a POSITA would have understood that the allocated bandwidth is the maximum required by any of the protection paths, i.e., “*equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment*.”

82. With respect to the example of Figs. 6a-6b, Dziong explains that recovery bandwidth can be shared among different shared-risk link groups (SRLGs):

[0088] To achieve bandwidth efficiency better than that of rings, recovery bandwidth sharing is considered between different single-event failures. In other words, the **bandwidth reserved for recovery of a particular SRLG can be shared with bandwidth reserved for recovery of other disjoint SRLGs**, since it is assumed that only one failure at a time will occur. FIG. 6(a) illustrates an exemplary optical/SONET network 602 and FIG. 6(b) shows bandwidth reservation table 604 for link A-B of FIG. 6(a).

Dziong, [0088].

83. Fig. 6b shows that the total amount of recovery bandwidth that is actually allocated on link A-B (a “*given segment*”) is the largest bandwidth among the protection bandwidths for SRLGs 3, 4, and 5. Again, this is because Dziong teaches that “protection bandwidth is shared for SRLG-disjoint failures.” Dziong, [0091]. Continuing with the example of how Dziong’s recovery path for SRLG 4 would have been established, I have annotated Fig. 6b to show the recited claim limitations. The 15 units of protection bandwidth for SRLG 4 corresponds to the “*bandwidth allocated to the new connection.*” The 48 units of protection bandwidth for SRLG 3 and the 36 units for SRLG 5 correspond to “*bandwidth allocated... to the other connections.*” The “*total bandwidth allocated*” is the maximum of these protection bandwidths, i.e., 48 units.



Dziong, Fig. 6b (annotated).

[0091] Table 604 is based on the assumption that protection bandwidth is shared for SRLG-disjoint failures. As such, **RSRV-XC** represents the actual amount of bandwidth that needs to be reserved on link A-B to protect against any one SRLG failure in network 602. In the current example, this corresponds to the maximum of the protection bandwidths required by each of SRLGs 3, 4, and 5 protected by link A-B (i.e., 48 units). Thus, link A-B provides protection bandwidth that is shared between disjoint SRLGs 3, 4, and 5.

Dziong, [0091].

84. Fig. 6b also shows another bandwidth reservation, P-XC, which is allocated 12 units of bandwidth: “‘P-XC’ represents the number of XC [cross-connect] units in link A-B that currently support demands (i.e., 12 units in this example).” Dziong, [0090]. Because the P-XC reservation is being used to support current traffic demands, a POSITA would have understood that P-XC is, essentially, a primary path reservation. Accordingly, the P-XC reservation represents bandwidth that cannot be shared. Thus, P-XC is not an “*other connection[]*” as claimed, since it is not “*permitted to share a bandwidth allocation with the new connection.*” See [1.2].

85. Sathyanath provides further examples showing a “*total bandwidth allocated*” as claimed. Sathyanath’s Fig. 5 example shows that the bandwidth reserved on link L5 is the “maximum of the bandwidth requirements associated with any single network component failure.” Sathyanath, [0104]. A new service and corresponding restoration path are then added to the network:

[0117] To understand the operation and use of the sharing information vectors, **consider the addition of a new service** to exemplary network 500 of FIG. 5 as illustrated in FIG. 8. A primary path node-link vector V_{pnl} for the new service 802 of FIG. 8 is illustrated in FIG. 9. Sathyanath, [0117].

86. As part of adding the new service, Sathyanath describes checking to see if the existing bandwidth reservation on link L5 is sufficient for both the pre-

existing restoration paths (“*the other connections*”) and the new restoration path (“*the new connection*”):

[0120] One way to make this determination is to add the primary path node-link vector V_{pnl} associated with the new service path to the aggregate node-link vector V_{nla} that was associated with the previous aggregate reservation on target link L5. This addition is done on an element-by-element basis to form a sum vector V_{snl} . The maximum value of this resulting sum vector is then compared to the current reservation bandwidth of link L5. **If the maximum value is equal to the current restoration bandwidth, then the restoration path for the new service can share the bandwidth reserved for the restoration paths of the existing services.** Otherwise, supporting the new restoration path will require reservation of an additional bandwidth unit. In the case of the exemplary new service 802 of FIG. 8, the sum vector V_{snl} resulting from the vector addition of the primary path node-link vector V_{pnl} of FIG. 9 and the aggregate node-link vector V_{nla} of FIG. 6 is {1,3,1,3,3,1}, where the maximum value is three, which is equal to the current reservation bandwidth on link L5. Thus, the restoration path for this new primary path may be added to link L5 without the reservation of additional bandwidth on link L5.

Sathyanath, [0120].

87. This example shows that the bandwidth allocated for link L5 (the “*total bandwidth allocated*” in a “*given segment*”) is the maximum required to restore any one service (“*a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment*”).

88. Sathyanath's example of Fig. 17 also shows this limitation.

Sathyanath examples that on link L4 (a "given segment"), there is an existing bandwidth allocation of four units to protection against failures on links L8 and L9 ("other connections"):

[0215] As an example, consider the partial network 1700 of FIG. 17. Consider the primary path L1-L2 with four units of bandwidth. Also, assume that there are two possible disjoint routes (with respect to the primary path) to the destination, namely route A (L3-L4) and route B(L5-L6). Further, assume that the entries of the current TE/Share database (DB) are as shown below:

[0216] Entry for L3 in TE/Share DB→4:: L1, L8; 5::L10

[0217] Entry for L4 in TE/Share DB→4:: L8, L9

[0218] Entry for L5 in TE/Share DB→4:: L9, L10

[0219] Entry for L6 in TE/Share DB→4:: L8, L10

[0220] The P-set for this example is {L1, L2}, the first R-set, R_A -set={L3, L4} and the second R-set, R_B -set={L5, L6}. The B-Li-sets for restoration path A are B-L3-set={L1, L8, L10} and B-L 4-set ={L8, L9}. Similarly, the B-Li-sets for restoration path B are B-L5-set={L9, L10}, B-L6-set ={L8, L10}. Evaluating the first restoration path, it is noted that none of the B-Li-sets are null, thus each of the links within the first restoration path offers some potential for sharing. Here, for example, the entry for link L3 indicates that L3 needs to reserve four units of bandwidth to protect L1 and L8 and five units of bandwidth to protect L10. Note that links L8, L9, and L10 are outside of the partial network 1700 shown in FIG. 17, but are still considered in the above

example to levy protection requirements on the links as indicated below. Note that the actual reserved bandwidth on link L3 would be the maximum of those two bandwidths, namely five units. Similarly, **the entry for L4 indicates that L4 needs to reserve four units of bandwidth to protect against failures on L8 and L9.** The entries for links L5 and L6 are interpreted similarly.

Sathyanath, [0215]-[0220].

89. If a new recovery path with 4 units of bandwidth were established over link L4, no new bandwidth would need to be allocated. Thus, this example shows that the “*total bandwidth allocated*” would be “*equal to a maximum of the requested bandwidth*” (four units protecting L1-L2) “*and of the bandwidths of the other connections*” (four units protecting L8 and L9):

[0221] In operation, source node R1 would compare the primary path L1-L2 against the entries for every link in a restoration path (e.g., $R_A\text{-set}=\{L3, L4\}$) and repeat the comparison for all possible restoration paths (e.g., $R_A\text{-set}$ and $R_B\text{-set}$). Upon comparison with path L3-L4, it is found that link L1 is already shared and that, if node R1 were to choose this restoration path, then the reservation bandwidth on L3 would need to be increased by 3 i.e., $4+4-5$) additional units of bandwidth. However, no additional bandwidth would need to be reserved on L4.

Sathyanath, [0221].

90. In Sathyanath’s example of Fig. 17, the path L5-L6 is ultimately chosen as the “*preferred route*,” as discussed further below for limitations [1.4] and [1.5]. However, links L5 and L6 in the TE/Share database have similar

bandwidth reservations. *See* Sathyanath, [0218]-[0219]. Sathyanath explains that neither L5 nor L6 require any additional bandwidth allocation to provide a recovery path for the primary path L1-L2. Sathyanath, [0222]. Thus, the same conclusion is true for links L5 and L6: their *total bandwidth allocated*” is “*equal to a maximum of the requested bandwidth*” (four units protecting L1-L2) “*and of the bandwidths of the other connections*” (the existing recovery paths on links L5 or L6).

91. In summary, Dziong and Sathyanath describe numerous examples of allocating an amount of recovery bandwidth on a link that is the maximum of the bandwidths of a new recovery path and existing recovery paths, which renders obvious “*that a total bandwidth allocated jointly to the new connection and to the other connections in any given segment is equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment.*”

- e. **[1.4] for each of the two or more candidate routes, calculating a cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route, responsively to the sharing definition;**

92. Dziong and Sathyanath render this limitation obvious.

93. Dziong describes calculating a link cost value, where cost is inversely proportional to the amount of bandwidth sharing:

[0065] In step 406, an LD path for each link in a set of candidate primary paths for a new demand is calculated in such a way that sharing

in the network is maximized (e.g., by **summing the cost of each link** in each candidate LD path, where **the cost is an inverse function of the degree to which the protection bandwidth for that link can be shared** with other links in the network, and by choosing the path with the lowest cost).

Dziong, [0065].

94. Since Dziong describes “cost” as being inversely related to the amount of bandwidth sharing, it would have been obvious to a POSITA that Dziong’s “cost” is directly related to the amount of non-shared bandwidth. Since Dziong prefers bandwidth sharing whenever possible (*see, e.g.,* [0060]: “maximizing sharing of link-detour bandwidth”), it would have been obvious to a POSITA that the only reason that recovery bandwidth would not be shared would be when the amount of recovery bandwidth required is more than the existing recovery bandwidth reservation. Thus, it would have been obvious to a POSITA for Dziong’s “cost” calculation to reflect the “*cumulative additional bandwidth reservation to be allocated to the new connection.*” Since Dziong describes “summing the cost of each link” in a candidate path ([0065]), it would have been obvious to calculate the “*cumulative additional bandwidth... in the segments along the candidate route.*” Since the shared bandwidth information is what Dziong describes advertising (see discussion above regarding “sharing definition” in [1.2]),

it would have been obvious that Dziong's cost calculation is performed
“*responsively to the sharing definition.*”

95. Dziong further describes a potential “link-cost function”:


[0113] Depending on the path calculation objective, **the link-cost function can take into account several metrics including:** administrative weight (which can be considered as a link bandwidth unit cost), available bandwidth, and delay. In the following discussion, maximization of bandwidth utilization is the focus, where the metrics of importance are administrative weight, **available bandwidth**, and **available shared bandwidth**.

Dziong, [0113].

96. Dziong notes that the factors considered by the link-cost function may depend on the desired “objective.” As noted previously, one potential objective is the “maximizing [of] sharing of link-detour bandwidth.” Dziong, [0060].

97. Dziong provides a specific example, in connection with Figure 9, which illustrates the obviousness of this claim limitation. Multiple concepts are all packed into this one figure, so I have annotated the figure below in a series of images to help explain its various ideas. First, “boxes 902, 904, and 906 represent the use of bandwidth on three links (A, B, and C) along a link detour path.” Dziong, [0132]. The link detour path protects against a failure along path “908 associated with a primary path link i.” Dziong, [0132]. As shown below, Figure 9 shows both the primary path (dashed orange line passing through link i), the

recovery path (dashed purple line passing through links A-B-C) and the bandwidths of links i, A, B, and C.



Dziong, Fig. 9 (annotated).

98. Dziong’s illustration of the bandwidth for link i shows that bandwidth on each link is divided into bandwidth allocated for primary paths (at the top), bandwidth allocated for recovery (“restoration”) paths (at the bottom), and unallocated (available) bandwidth in the middle. Dziong explains that links A, B, and C each have some bandwidth reserved for existing recovery paths:

Each of boxes 902, 904, and 906 also represent the use of each associated links bandwidth for primary, available, and recovery

bandwidth categories. For example, as represented by box 902, **link A** has 8 STS-1s reserved for protection (916). **Link B** has 15 STS-1s reserved for protection (918) and **link C** has 10 STS-1s reserved (920).

Dziong, [0132].




FIG. 9

Dziong, Fig. 9 (annotated).

99. A POSITA would have been familiar with Dziong's reference to "STS-1." STS-1 is an acronym for "Synchronous Transport Signal 1," which is a basic unit of transmission in the context of SONET. An STS-1 provides 51.84 Mbit/s of bandwidth. Other common SONET transmission signals are generally

multiples of the STS-1 bandwidth. For example, an STS-3 (which Dziong also references at [0134]) corresponds to a bandwidth that is exactly three times the size of an STS-1, that is, a bandwidth of 155.52 Mbit/s. As can be seen from Figure 9 and Dziong's discussion, Dziong generally treats the STS-1 as a unit of bandwidth. *See* Dziong, [0132] ("5 STS-1 units of bandwidth.... Link B has 15 STS-1s reserved"), [0134] ("link B where 15 units are already reserved").


100. Dziong shows in Fig. 9 that the recovery path does not require any new recovery bandwidth allocations because all of the required recovery bandwidth is shared with the existing reservations. Specifically, the new recovery path requires 5 units of bandwidth, but links A, B, and C all have existing recovery bandwidth reservations that are greater than 5 units. Thus, Dziong's cost for the recovery path A-B-C would be zero. This also shows that Dziong's cost of protection path calculation is made "*responsively to the sharing definition*," since it refers to the existing recovery path reservations.

101. Dziong also discusses, in the context of Figure 9, the potential cost of a recovery path for alternative primary path link bandwidths. Dziong expressly teaches that the cost of a recovery path for either 1 or 3 units of bandwidth would be zero:

As illustrated, a [new] demand of either STS-1 or STS-3 can be accommodated without requiring any additional bandwidth for

recovery on the detour links A, B, and C. **Therefore the CPP [cost of protection path] value for those demands is shown as zero.**

Dziong, [0134].



Dziong, Fig. 9 (partial, annotated).

102. Fig. 9 also illustrates an example of how Dziong contemplates calculating a cost of protection path. For the example of a new protection path having 12 units of bandwidth (i.e., for an STS-12), Dziong states that Link A would require 9 additional bandwidth units, Link B would require 2 units, and Link C would require 7 units:

For an STS-12 demand, 9 additional STS-1s are needed on link A for recovery since three STS-1s are already reserved and can be shared (assuming the reservation is for a disjoint link recovery). Similarly, 2 additional STS-1s are needed on link B where 15 units are already reserved (10 above the 5 STS-1 demand of link i), and 7 additional STS-

1s are required on link C. These numbers are used with the respective weights to calculate CPPi value for a STS-12 demand.

Dziong, [0134].

103. Dziong shows calculating the cost of protection path (CPP) value in Fig. 9:

$$(STS\ 12:\ CPP = [9*AW1 + 2*AW2 + 7*AW3]/12)$$

Dziong, Fig. 9 (partial).

104. For clarity, the Fig. 9 equation calculates a protection path cost as:

$$CPP = [9*AW1 + 2*AW2 + 7*AW3]/12$$

where the three “AW” values represent “administrative weights.” Dziong, [0131].

The use of administrative weights, however, is optional: “In some cases, the link costs of the LD [link detour] paths can be a function of an administrative weight for the LD path links.” Dziong, [0012]. Removing the administrative weights, as Dziong suggests, would leave the equation for calculating a cost of protection path as the sum of the additional recovery bandwidth reservation required on each link (9+2+7), divided by the total recovery bandwidth (12):

$$CPP = [9 + 2 + 7]/12$$

105. Dziong does not explain why its cost of protection path equation divides by the total recovery bandwidth. Regardless, the numerator in Dziong’s

calculation is the sum of the additional recovery bandwidth reservation required on each link (9+2+7). Thus, Dziong renders obvious the claimed concept of *“calculating a cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route.”*

106. I also note that Dziong explains that a cost of protection path that is negative indicates that the requested recovery bandwidth is unavailable. In other words, the requested recovery bandwidth cannot be accommodated. For example, the recovery path A-B-C in Fig. 9 has insufficient capacity to protect a primary path of 48 units or 192 units bandwidth:

Since the shared bandwidth plus the **available capacity of link B is insufficient** (see box 904) to accommodate an STS-48 demand, **the CPPi value associated with STS-48 is negative**. Similarly, the CPPi value is negative for an STS-192 demand.

Dziong, [0134].

107. Thus, it would have been obvious to a POSITA that a cost of protection path of zero is the minimum cost possible for a viable recovery path. Dziong describes and illustrates in Fig. 9 that the cost of protection path along links A-B-C for a recovery path having between 1 unit and 5 units of bandwidth is zero.

108. Sathyanath provides further details and examples of link-cost functions. “One metric that can be used to evaluate a restoration path is the overall

amount of additional bandwidth that must be reserved in the network to support protection of the new service using that restoration path.” Sathyanath, [0200].

Sathyanath describes evaluating “a number of candidate restoration paths R_1 , R_2 , R_3 , ..., R_N ” (“*each of the two or more candidate routes*”). Sathyanath, [0201].

When bandwidth sharing is possible, Sathyanath describes calculating a link cost as the amount of additional bandwidth that needs to be reserved, which may be zero:

[0203] (b) Sharing possible--.... Thus, by definition, some protection bandwidth sharing is possible on link L_i . There are two cases.

[0204] In the first case, the bandwidth $Br(i)$ that is presently reserved on link L_i is greater than or equal to B_p . In this case, no additional bandwidth need be reserved to protect the new service, and **the cost of using this link of the candidate restoration path is zero.**

[0205] In the second case, B_p is greater than the presently reserved bandwidth $Br(i)$ on link L_i , and **the incremental bandwidth that needs to be reserved** is $B_p - Br(i)$. If, however, B_p is greater than the total capacity of the link, then the link cannot be used to protect the service, and the cost of link L_i is preferably set to a sufficiently high value to exclude this link and, in fact, this entire candidate restoration path, from further consideration as a restoration path for this service on this primary path. Alternatively, this "total capacity" constraint may be applied to exclude such paths from being included when the set of candidate restoration paths is originally generated.

Sathyanath, [0203]-[0205].

109. For each candidate path, Sathyanath considers “the set of links L_i ” that make up the candidate path (“*the segments along the candidate route*”). Sathyanath, [0201]. If the recovery bandwidth already reserved on a link L_i is at least as much as is required to protect the new service, then “no additional bandwidth need be reserved” and “the cost of using this link [L_i] of the candidate restoration path is zero.” Sathyanath, [0204]. If additional recovery bandwidth is required and available, however, then the cost of link L_i is “the incremental bandwidth that needs to be reserved” (“*additional bandwidth reservation to be allocated to the new connection*”). Sathyanath, [0205]. The total cost for a recovery path “is calculated as a sum of the costs of its constituent links” (“*calculating a cumulative additional bandwidth reservation to be allocated*”). Sathyanath, [0209]. Since the cost of each link is the incremental bandwidth that needs to be reserved on that link, the overall recovery path cost corresponds to a “*cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route.*”

110. Sathyanath describes calculating link and path recovery costs with respect to its example network shown in Fig. 17. As discussed previously, Fig. 17 shows a primary route L_1 - L_2 and two candidate recovery routes, route A (L_3 - L_4) and route B (L_5 - L_6):

[0215] As an example, consider the partial network 1700 of FIG. 17. Consider the primary path L1-L2 with four units of bandwidth. Also, assume that there are two possible disjoint routes (with respect to the primary path) to the destination, namely route A (L3-L4) and route B (L5-L6).

Sathyanath, [0215].

111. Sathyanath explains that link L3 would require 3 additional units of bandwidth. This is because link L3 “needs to reserve four units of bandwidth to protect L1 and L8 and five units of bandwidth to protection L10.” Sathyanath, [0220]. So while link L3 has five units of recovery bandwidth reserved, four of those units are for an existing connection that traverses link L1. If link L1 failed, link L3 would need a total of 8 units of bandwidth to act as a recovery path for both the existing 4-unit connection and the new 4-unit primary path. Meanwhile, link L4 would require no additional bandwidth. Thus, it would have been obvious to a POSITA that the total incremental bandwidth required for recovery route A is 3 units.

Upon comparison with path L3-L4, it is found that link L1 is already shared and that, if node R1 were to choose this restoration path, then the reservation bandwidth on L3 would need to be increased by 3 i.e., 4+4-5) additional units of bandwidth. However, no additional bandwidth would need to be reserved on L4.

Sathyanath, [0221].

112. In contrast, no additional bandwidth is required on links L5 or L6.


This is because there are already 4 units of bandwidth reserved on these links, and none of those existing reservations is to protect connections on links L1 or L2.

Sathyanath, [0218]-[0219]. Thus, the existing and new recovery paths can share the 4 units of bandwidth already allocated, and total incremental bandwidth required for recovery route B is 0 units:

[0222] Upon comparison with restoration path L5-L6, it is found that the intersection of P-set {L1, L2} and R_B-set {L8, L9, L10} is null, and no additional units of bandwidth would need to be reserved (since 4 units have been reserved already). Therefore, under such circumstances, the restoration path L5-L6 would be considered to be the better restoration path.

Sathyanath, [0222].

113. I have illustrated below Sathyanath's Fig. 17 to show the recovery path costs for the candidate routes. Sathyanath's discussion of the Fig. 17 example shows, for each of potential recovery routes A and B (*"for each of the two or more candidate routes"*), calculating the incremental amount of bandwidth to be reserved on the route links for existing recovery paths and a new recovery path (*"calculating a cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route"*) using information from the TE/Share database (*"responsively to the sharing definition"*).



Sathyanath, Fig. 17 (annotated).

114. In summary, Dziong and Sathyanath teach calculating, for multiple candidate recovery paths, a cost of protection path value by summing the amount of additional recovery bandwidth that must be reserved on each link based on shared bandwidth information, which renders obvious “*for each of the two or more candidate routes, calculating a cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route, responsively to the sharing definition.*”

- f. [1.5] *selecting a preferred route having a minimum value of the cumulative additional bandwidth reservation out of the two or more candidate routes; and*

115. Dziong and Sathyanath render obvious this limitation.

116. As I discussed regarding limitation [1.4], Dziong and Sathyanath describe calculating a cost of protection path as a number, which is a “*value*.” The cost of protection path is calculated as the sum of additional recovery bandwidth reservations required across all links of a candidate recovery path (“*cumulative additional bandwidth reservation*”).

117. Dziong and Sathyanath both describe selecting a recovery path that has the least additional bandwidth required. For example, Dziong teaches “choosing” the path “with the smallest LD [link-detour] path link cost sum.” Dziong, claim 3. Sathyanath similarly describes to “select[...] the restoration path with the lowest path cost.” Sathyanath, Abstract. Sathyanath further explains:

[0200] When a new service is added to a network, the source node or a network management entity has the task of determining appropriate primary and restoration paths for that service. Assuming a set of candidate primary paths can be calculated for the service as well as a set of potential restoration paths for one or more of those primary paths, it is then of interest to determine, for each candidate primary path, a restoration path that offers the best sharing opportunity. One metric that can be used to evaluate a restoration path is the overall amount of additional bandwidth that must be reserved in the network to support protection of the new service using that restoration path. By **choosing a restoration path that would require the minimum amount of additional bandwidth to be reserved along its links**, an operator can reduce the overall impact to the available capacity in the network.

Sathyanath, [0200].

118. Thus, it would have been obvious to a POSITA that Dziong and Sathyanath teach “*selecting a preferred route having a minimum value... out of the two or more candidate routes.*”

119. Sathyanath illustrates this principle in action through its example in Fig. 17. As discussed at [1.4], candidate route A (L3-L4) requires 3 additional units of bandwidth, while candidate route B (L5-L6) requires 0 additional units of bandwidth. The recovery path cost of route B (0 units) has a “*minimum value*” of all the potential recovery paths. Sathyanath explains that route B is therefore the better route:

[0222] Upon comparison with restoration path L5-L6, it is found that the intersection of P-set {L1, L2} and R_B-set {L8, L9, L10} is null, and no additional units of bandwidth would need to be reserved (since 4 units have been reserved already). Therefore, under such circumstances, **the restoration path L5-L6 would be considered to be the better restoration path.**

Sathyanath, [0222].

120. Thus, it would have been obvious to a POSITA to select route B (L5-L6) as a “*preferred route*” because it is “the better restoration path.”

121. In summary, Dziong and Sathyanath teach choosing as a recovery path the candidate path that requires the least amount of additional bandwidth to be allocated, which renders obvious “*selecting a preferred route having a minimum*

value of the cumulative additional bandwidth reservation out of the two or more candidate routes.”

g. [1.6] *establishing the new connection using the preferred route.*

122. Dziong and Sathyanath render obvious this limitation.

123. As I discussed above for limitation [1.0], Dziong and Sathyanath describe establishing a new recovery path, which corresponds to the claimed “*establishing the new connection.*” As I discussed at [1.5], Dziong and Sathyanath also describe choosing from multiple potential recovery paths a “*preferred route.*” Thus, it would have been obvious to a POSITA for the chosen recovery route to be established so that it could be used to provide recovery services in the event of a failure along a path being protected.

124. For example, Dziong further describes “establishing” a new recovery path, for example, by reserving the bandwidth needed on the links that make up the recovery path:

"Protection" typically refers to a recovery mechanism where the paths are computed and reserved in advance of a failure.

Dziong, [0034].

[0058] Embodiments of the present invention may include a distributed method for calculating primary and link-detour paths in a mesh network. This method improves the number of connections admitted to a network and reduces the probability of crank-backs due to unavailable

or overutilized link-detour paths. A crankback is the failure to establish either a primary or protection path based on the unavailability of resources that were anticipated to be available. A crankback can occur, for example, during the reservation of bandwidth along a calculated protection path for a link. **A source node may** assume that bandwidth for a new connection is available, and then **start to signal to establish** the primary path and **link-detour paths** for the connection.

Dziong, [0058].

125. Dziong further describes how a recovery path, such as a link-detour path, can be established by sending a reservation setup message. In other words, the reservation setup message is the “signal to establish” a recovery path, as referenced in the Dziong quote above. The reservation setup message is sent along the recovery path, and each node on the recovery path reserves the required amount of bandwidth:

[0175] In the enhanced loop avoidance embodiment, **the reservation setup message is sent along the full-LD path**. The message carries information that classifies nodes as branching, merging, signaling, or transit nodes. The message also carries associated failure IDs, if applicable. The failure IDs are used to avoid having more than one shortened-LD path associated with the set of links in the protected segment associated with the shortened-LD path.

* * *

[0178] In the basic embodiment, branching, merging, and transit nodes reserve cross-connects in anticipation of a failure. In this scheme, **each**

link in the protected segment makes an independent reservation.

This limits, in general, using the same shortened-LD path for all links in the protected segment since each LD path can be different. As a consequence, there is a possibility of partly overlapping protection segments and that may cause problems in a scenario where a second failure occurs.

Dziong, [0175], [0178].

126. Dziong further confirms that the reservation process “*establish[es]*” the recovery path along the preferred route by describing what happens when a failure occurs. No new path set up messages are required. Instead, the branching node (i.e., the node that diverts traffic from the primary path onto the recovery path) simply activates the previously-chosen recovery path:

[0183] In the enhanced embodiment, **the failure message is sent** from the reservation control node to the branching node, then along the shortened-LD path to the merging node, and then to the downstream node. At the branching and merging nodes, the node first checks to see if any of the associated failure IDs are already activated. If so, no action is taken other than confirmation that the recovery is in place. Otherwise, **the reserved cross-connect is activated.**

Dziong, [0183].

127. Sathyanath similarly describes the ““**pre-establishment...** of alternative paths.” Sathyanath, [0046]. Sathyanath specifically describes establishing a protection path as a label-switched path in an MPLS network:

[0250] In case of the centralized server approach, a centralized server computes the additional protection bandwidth needed on each link along the protection path to admit a new connection (LSP). This information is then conveyed as part of network **signaling to establish the protection LSP**.

* * *

[0260] It is assumed that **every protection LSP is established** with an assigned bandwidth of zero or some minimal value sufficient to carry OAM traffic. After a failure, traffic coming into the ingress node for a primary LSP is switched to its corresponding protection LSP, and traffic exiting the egress node for the primary LSP is selected from the protection LSP, instead of from the primary LSP. The aggregate bandwidth for the protected traffic on primary LSPs is reserved in advance as part of the protection pool in the SMDN. Note that the bandwidth of the protection pool is not reassigned to any protection LSP. After the failure, each protection LSP that is providing protection against a specific failure needs to be assigned the same bandwidth and QoS characteristics as were assigned to the primary LSP it protects. This requires changing the bandwidth assignment of each of these protection LSPs. This can be accomplished through RSVP-TE using its bandwidth change procedure (see RFC 3209 for details).

Sathyanath, [0250], [0260].


128. In summary, Dziong and Sathyanath teach establishing a chosen recovery path, which renders obvious “*establishing the new connection using the preferred route.*”

2. Claim 2

- a. [2.0] *The method according to claim 1, wherein at least some of the segments are arranged in a ring configuration, and*

129. Dziong and Sathyanath render obvious this limitation.

130. Dziong illustrates a network with a “ring configuration” in Fig. 6a. As discussed at [1.0], each link between two nodes corresponds to a “segment[.]”



Dziong, Fig. 6a (annotated).

131. Sathyanath also describes applying its techniques to a network that “has a ring topology.” Sathyanath, [0248]. A POSITA would have found it obvious


for a network with a “ring topology” to have “*at least some... segments... arranged in a ring configuration.*”

132. In summary, Dziong and Sathyanath describe applying their bandwidth-sharing techniques to a ring network, which renders obvious “*wherein at least some of the segments are arranged in a ring configuration.*”

- b. [2.1] *wherein identifying the two or more candidate routes comprises identifying first and second routes traversing respective first and second ringlets of the ring configuration.*

133. Dziong and Sathyanath render obvious this limitation.

134. As discussed at [2.0], Dziong’s example network in Fig. 6a is a ring network. Dziong shows that the largest ring (around the perimeter of the nodes A-B-C-D) also includes a bisecting connection between nodes A and C. Thus, the network also includes two smaller rings, A-B-C and A-C-D. These smaller rings are “*first and second ringlets.*”




Dziong, Fig. 6a (annotated).

135. As I discussed regarding the Fig. 6a example for limitation [1.1], there are two potential link-detour paths to be used in the event of a failure of primary path A-D. First, there is link-detour path A-B-C-D, which traverses the path A-B-C corresponding to the “*first ringlet*.” Second, there is link-detour path A-C-D, which corresponds exactly to the “*second ringlet*.” Thus, these two candidate routes are “*first and second routes traversing respective first and second ringlets of the ring configuration.*”

136. I note that the '088 patent specification describes the concept of “ringlets” as corresponding to the clockwise and counterclockwise directions on a ring:

In the exemplary configuration of FIG. 1, four network nodes denoted A, B, C and D are arranged in a ring configuration 26, such as a resilient packet ring (RPR) configuration as defined in the IEEE 802.17 standard cited above. Ring 26 comprises two ringlets, namely a **clockwise (CW) ringlet 28** and a **counterclockwise (CCW) ringlet 32**. Segments 25 that belong to ring 26 are referred to as ring segments.


'088 patent, 4:37-43.



'088 Patent, Fig. 1 (annotated).

137. If the claimed “ringlets” are interpreted in a limiting way as corresponding to clockwise and counterclockwise directions on a ring network, Dziong still renders this limitation obvious. Dziong shows in its example of Fig. 2,

for example, a ring network with traffic (represented by arrows 204 and 208) moving in both a clockwise and a counterclockwise direction:



Dziong, Fig. 2 (annotated).

138. Dziong describes applying its bandwidth-sharing concepts to the example of Fig. 2:

[0045] ...For example, **consider a network as illustrated in FIG. 2** where nodes A, B, C, D, and E are connected in a ring topology, forming ring A-B-C-D-E-A. Assume that nodes C, D, and E have no other connected links or nodes in the network other than those illustrated. This means that there is just one shortest detour (C-B-A-E-D) between nodes C and D that can avoid traversing link C-D.

[0046] Now consider a demand 202 from node A to node B whose primary path, for some reason, is A-E-D-C-B. Assume that the demand is recovered by link-based recovery mechanisms, where the detour for

link D-C is the path D-E-A-B-C. In this case, if link D-C fails, traffic for demand 202 will flow along the primary path segment A-E-D (204), followed by flow along the detour path D-E-A-B-C (206), followed by flow along the primary path segment C-B (208). In this case, traffic will flow along links A-E, E-D, and B-C twice. This backhauling can be avoided if the protection scheme detects the backhaul while computing the protection paths and avoids the backhaul by moving the protection switching functions. **In the present example**, backhauling can be avoided by moving the protection switching function to nodes A and B. To accomplish this, **the present invention can incorporate additional bookkeeping and signaling that allow the computation and selection of the appropriate switching nodes for protection.**

Dziong, [0045]-[0046].

139. Notably, both Dziong and Sathyanath generally contemplate applying their bandwidth-sharing techniques to a wide variety of networks:

[0076] Note that, although optical or SONET network examples are used in various discussion herein, the present invention can be applied to networks based on different technologies and topologies, including wired/wireless, optical/electrical, and mesh topology.

Dziong, [0076].

[0049] The concepts of the present invention are discussed in the context of exemplary shared mesh data networks. However, without loss of generality, they may be applied generally to any mesh network (e.g., Sonet, ATM, frame relay, MPLS) that provides traffic-switching control at the nodes incident to the links in the network.

Sathyanath, [0049].

140. Thus, it would have been obvious to a POSITA to apply the disclosures of Dziong and Sathyanath to a bidirectional ring network having clockwise and counterclockwise “*ringlets*.” It would have been obvious to a POSITA that on a ring network, two potential recovery paths would go in opposite directions around the ring network. Thus, it would have been obvious for the step of identifying potential recovery paths (discussed above for limitation [1.1]) to comprise “*identifying first and second routes traversing respective first and second ringlets of the ring configuration.*”

141. In summary, Dziong and Sathyanath describe identifying paths on ring networks, including paths that traverse a ring in clockwise and counterclockwise directions, which renders obvious “*wherein identifying the two or more candidate routes comprises identifying first and second routes traversing respective first and second ringlets of the ring configuration.*”

3. Claim 3

- a.** **[3.0]** *The method according to claim 1, wherein establishing the new connection comprises establishing a tunnel through the network in accordance with a tunnel-based communication protocol.*

142. Dziong and Sathyanath render obvious this limitation.

143. As I discussed above from limitation [1.0], Dziong and Sathyanath both contemplate establishing a new recovery path as an MPLS label-switched

path, which corresponds to “*establishing the new connection.*” Sathyanath explains in the context of its Fig. 1 that an MPLS label-switched path is a “*tunnel*”:

[0050] As illustrated, SMDN 100 has been provisioned with five primary label-switched paths (LSPs) LSP-1 to LSP-5. **A label-switched path is effectively a tunnel** between two nodes that carries service traffic according to a predetermined route. For example, **LSP-1 is a tunnel** from node N3 to node N5 that follows the path N3-N2-N5. (LSP-1 could alternatively be described by the links it traverses, i.e., L5-L4, or by the nodes and links, i.e., N3-L5-N2-L4-N5; however, for this discussion, in cases where no ambiguity will arise from the node-only notation, that notation will be used.)

Sathyanath, [0050].

144. Thus, a POSITA would have recognized that MPLS is a “*tunnel-based communication protocol.*” The MPLS architecture allows packets to be encapsulated with a label and then routed via the label instead of whatever headers were on the packet before it was encapsulated. This effectively makes it seem as though two routers are consecutive routers (i.e., adjacent to one another), even though they not; rather, they are connected by a “tunnel”:

Sometimes a router Ru takes explicit action to cause a particular packet to be delivered to another router Rd, even though Ru and Rd are not consecutive routers on the Hop-by-hop path for that packet, and Rd is not the packet's ultimate destination. For example, this may be done by encapsulating the packet inside a network layer packet whose destination address is the address of Rd itself. This creates a “tunnel”

from Ru to Rd. We refer to any packet so handled as a "Tunneled Packet".

RFC3031, §3.27, p. 32. This understanding is consistent with the '088 patent, which acknowledges that MPLS is a "tunneled protocol[]." '088 patent, 1:19-21.

145. In summary, Dziong and Sathyanath describe establishing a new recovery path as an MPLS label-switched path, which renders obvious "*wherein establishing the new connection comprises establishing a tunnel through the network in accordance with a tunnel-based communication protocol.*"

4. Claim 4

- a. [4.0] *The method according to claim 1, wherein the new connection and at least one of the other connections specified in the sharing definition traverse at least one common segment along the preferred route,***


146. Dziong and Sathyanath render obvious this limitation.

147. As I discussed previously with respect to limitations [1.0] through [1.2] and [1.5], Dziong provides an example embodiment in Figs. 6a-6b that demonstrates the recited claim limitations:


- "*new connection*" corresponds to a protection path from node A to node D;
- "*other connections*" corresponds to protection paths reserved on link A-B, including SRLG 3 and SRLG 5;

- “*sharing definition*” corresponds to advertised information about connections on link A-B, including information about protection paths for SRLG 3 and SRLG 5; and
- “*preferred route*” corresponds to protection path A-B-C-D.

148. Fig. 6a shows that link A-B is one “*segment*” along the “*preferred route*” A-B-C-D. Fig. 6b shows that the protection paths for SRLG 3 and SRLG 5 also traverse link A-B.



Dziong, Fig. 6a (annotated).



Dziong, Fig. 6b (annotated).

149. Link A-B is *"at least one common segment along the preferred route."* As discussed at [1.2], the information in Dziong's Fig. 6b corresponds to the claimed *"sharing definition."* Thus, Dziong's example in Fig. 6b shows that the new protection path for SRLG 4 (*"the new connection"*) and the existing protection paths for SRLG 3 and SRLG 5 (*"and at least one of the other connections specified in the sharing definition"*) are routed along link A-B (*"traverse at least one common segment along the preferred route"*).

150. Sathyanath's Fig. 17 example also demonstrates this limitation. As discussed above for limitation [1.5], Sathyanath explains that route B on links L5-

L6 is the “better restoration path,” and thus, corresponds to the claimed “*preferred route*.” Links L5 and L6 have existing recovery bandwidths reserved for other links, such as links L9 and L10:

Further, assume that the entries of the current TE/Share database (DB) are as shown below:

[0216] Entry for L3 in TE/Share DB → 4:: L1, L8; 5::L10

[0217] Entry for L4 in TE/Share DB → 4:: L8, L9

[0218] Entry for L5 in TE/Share DB → 4:: L9, L10

[0219] Entry for L6 in TE/Share DB → 4:: L8, L10

[0220] ... Here, for example, the entry for link L3 indicates that L3 needs to reserve four units of bandwidth to protect L1 and L8 and five units of bandwidth to protect L10. Note that links L8, L9, and L10 are outside of the partial network 1700 shown in FIG. 17, but are still considered in the above example to levy protection requirements on the links as indicated below. Note that the actual reserved bandwidth on link L3 would be the maximum of those two bandwidths, namely five units. Similarly, the entry for L4 indicates that L4 needs to reserve four units of bandwidth to protect against failures on L8 and L9. The entries for links L5 and L6 are interpreted similarly.

Sathyanath, [0215]-[0220].

151. Thus, link L5 is an example of a “*common segment along the preferred route*” that is traversed by both the “*new connection*” (new recovery path L5-L6 to protect primary path L1-L2) and “*at least one of the other connections specified in the sharing definition*” (recovery paths for links L9 and L10).


152. In summary, Dziong and Sathyanath describe how a new recovery path can share bandwidth reserved on links for other, existing recovery paths, which renders obvious “*wherein the new connection and at least one of the other connections specified in the sharing definition traverse at least one common segment along the preferred route.*”

5. Claim 4

- a. [4.1] and wherein establishing the new connection comprises allocating the maximum of the requested bandwidth and of the bandwidths of the at least one of the other connections in the at least one common segment to the new connection and to the at least one of the other connections.**

153. Dziong and Sathyanath render obvious this limitation.

154. Continuing with the analysis from limitation [4.0] of Dziong’s Fig. 6b, Dziong illustrates allocating 15 units of bandwidth for the new protection path for SRLG 4 (“*the new connection*”). The existing protection paths for SRLG 3 and SRLG 5 (“*the at least one of the other connections*”) are allocated 48 units and 36 units, respectively. The maximum of these values is 48 units, which is the amount that Dziong shows being allocated for link A-B (“*the at least one common segment*”):



Dziong, Fig. 6b (annotated).

155. Dziong explains that the bandwidth reserved for these protection paths is equal to the “maximum” needed to protect any one path. The reserved bandwidth is shared among the protection paths, and thus “*allocate[ed]... to*” all of them:

[0091] Table 604 is based on the assumption that protection bandwidth is shared for SRLG-disjoint failures. As such, RSRV-XC represents the actual amount of bandwidth that needs to be reserved on link A-B to protect against any one SRLG failure in network 602. In the current example, this corresponds to the **maximum of the protection bandwidths required by each of SRLGs 3, 4, and 5 protected by link A-B (i.e., 48 units)**. Thus, link A-B provides **protection bandwidth that is shared** between disjoint SRLGs 3, 4, and 5.

Dziong, [0091].

156. Sathyanath's Fig. 17 example also demonstrates this limitation. As discussed above for limitation [4.0], Sathyanath explains that route B on links L5-L6 is the "better restoration path," and thus is where a "*new connection*" is established to protect primary path L1-L2. As discussed at [1.0], this "*new connection*" would have a "*requested bandwidth*" of 4 units to match the bandwidth of the L1-L2 path being protected. Sathyanath, [0215]. Sathyanath's exemplary bandwidth sharing information (the TE/Share database) shows that Links L5 and L6 already have 4 units of bandwidth reserved to protect other links:

Further, assume that the entries of the current TE/Share database (DB) are as shown below:

[0216] Entry for L3 in TE/Share DB → 4:: L1, L8; 5:: L10

[0217] **Entry for L4 in TE/Share DB → 4:: L8, L9**

[0218] **Entry for L5 in TE/Share DB → 4:: L9, L10**

[0219] Entry for L6 in TE/Share DB → 4:: L83, L10

157. I note that Sathyanath explains how to interpret the TE/Share database entries with reference to the entries for links L3 and L4, which like links L4 and L5, each have 4 units of bandwidth reserved. A POSITA would have understood from Sathyanath's explanation that the number appearing in a TE/Share database entry between the arrow (→) and the double-colon (::) represents the amount of bandwidth reserved for protection:

[0220] ... Here, for example, the entry for **link L3 indicates that L3 needs to reserve four units of bandwidth** to protect L1 and L8 and five units of bandwidth to protect L10. Note that links L8, L9, and L10 are outside of the partial network 1700 shown in FIG. 17, but are still considered in the above example to levy protection requirements on the links as indicated below. Note that the actual reserved bandwidth on link L3 would be the maximum of those two bandwidths, namely five units. Similarly, the entry for L4 indicates that **L4 needs to reserve four units of bandwidth to protect against failures** on L8 and L9. **The entries for links L5 and L6 are interpreted similarly.**

Sathyanath, [0215]-[0220].

158. In short, the TE/Share database shows that the “*bandwidths of the at least one of the other connections*” on links L5 and L6 are 4 units.

159. Sathyanath further explains that the amount of reserved bandwidth for protection on links L5 and L6 remains the same—at 4 units—when the new protection path for L1-L2 is created:

[0222] Upon comparison with restoration path L5-L6, it is found that the intersection of P-set {L1, L2} and R_B-set {L8, L9, L10} is null, and **no additional units of bandwidth would need to be reserved** (since 4 units have been reserved already). Therefore, under such circumstances, the restoration path L5-L6 would be considered to be the better restoration path.

Sathyanath, [0222].

160. The amount of bandwidth reserved on links L5 and L6 (4 units) is equal to the “*maximum*” of the requested bandwidth (4 units) and the amount of bandwidth previously allocated for other protection paths (4 units). Since Sathyanath repeatedly discusses sharing bandwidth among recovery paths, it would have been obvious that this bandwidth is then reserved for both the new protection path and the existing protection paths (“*to the new connection and to the at least one of the other connections*”).

161. Thus, Dziong and Sathyanath describe allocating shared bandwidth among new and existing recovery paths, where the total recovery bandwidth allocated is the maximum of the amount of bandwidth needed for the new recovery path and the amount of bandwidth needed for the existing recovery paths, which renders obvious “*allocating the maximum of the requested bandwidth and of the bandwidths of the at least one of the other connections in the at least one common segment to the new connection and to the at least one of the other connections.*”

6. Claim 6

- a.** [6.0] *The method according to claim 1, wherein providing the sharing definition comprises distributing the sharing definition to nodes of the network using a reservation protocol.*

162. Dziong and Sathyanath render obvious this limitation.

163. As I discussed above at [1.2], Dziong describes advertising information the amount of bandwidth that may be shared among recovery paths:

[0060] In embodiments of the present invention, link-detour path calculation involves maximizing sharing of link-detour bandwidth. The recovery-path calculation algorithm makes use of information including how much bandwidth can be-shared at each link in the network. This information is obtained by **advertising, across the network, the amount of bandwidth reserved for recovery on each link** and by bookkeeping, in each node, information about all recovery paths that would be activated when a protected link fails.

* * *

[0062] The primary path is calculated by taking into account the link cost and constraints that take into account the costs and constraints of link-detour paths for each link in the primary path. The link-detour path cost and constraints for each link are **distributed to each node by an advertising protocol**.

Dziong, [0060], [0062].

164. Dziong further explains that the bandwidth-sharing advertisements can be distributed across the network using a link-state advertisement protocol:

[0056] Finally, in step 310, link status and sharing information is passed to nodes in the network using an appropriate link state advertisement (LSA)-based routing protocol.

* * *

[0098] An alternative to centralized solution is a distributed implementation where paths are calculated in their respective source nodes. Such a solution has the advantage of being much more scalable and resilient to network element failures. While a distributed

implementation avoids signaling to a centralized server, it requires an advertisement protocol that distributes information about link states across the network. Such **link-state advertisement (LSA) protocols are usually already present** in communications networks to support primary path calculations. However, **some extensions may be required** in order to advertise information about link bandwidth sharing capabilities, and, in the case of link recovery, link-recovery costs.

Dziong, [0056], [0098].

165. Sathyanath provides further details about distributing bandwidth-sharing information to other network nodes. For example, Sathyanath explains that RSVP-TE protocol may be used to distribute this information:

[0292] While this invention has been described with respect to dissemination of sharing information by specific mechanisms including OSPF, **alternative means for dissemination of information within mesh networks could also be used** (e.g., **RSVP-TE**, TCP/IP connections, constrained routing label distribution protocol (CR-LDP)), and SNMP) as would be understood to one skilled in the art.

Sathyanath, [0292].

166. A POSITA would have been familiar with RSVP-TE, which is defines Traffic Engineering extensions to Resource ReSerVation Protocol (RSVP). *See* Sathyanath, [0056]. Thus, RSVP-TE is an example of a “*reservation protocol*.” Notably, RSVP-TE is the same protocol contemplated by the ’088 patent for distributing a “*sharing definition*.” ’088 patent, 7:10-12.

167. In summary, Dziong and Sathyanath teach distributing bandwidth sharing information to other network nodes using RSVP-TE, which renders obvious “*wherein providing the sharing definition comprises distributing the sharing definition to nodes of the network using a reservation protocol.*”

7. Claim 8

168. Claim 8 is an independent claim, and its limitations largely mirror those recited for claim 1 analyzed above. Thus, the bulk of claim 8’s features would have been obvious from the disclosures of Dziong and Sathyanath for the reasons I have already discussed. I address below the specifically recited features of claim 8 that are different relative to claim 1.

a. [8.0] *A network node, comprising:*

169. Dziong and Sathyanath render obvious this limitation.

170. Dziong describes how its recovery-path-selection process may be performed by either a network node or on a specialized server:

[0096] Recovery-path calculation can be centralized or distributed. In a centralized solution, the path calculation is performed in a specialized server that keeps track of all connection states. When a new connection demand arrives at a source node, a request is sent to the server to calculate the primary path and associated protection path(s). Once the server calculates the paths, this information is sent back to the source node....

[0098] An alternative to centralized solution is a distributed implementation where paths are calculated in their respective source nodes. Such a solution has the advantage of being much more scalable and resilient to network element failures.

Dziong, [0096], [0098].

171. Sathyanath similarly explains that recovery paths may be calculated and chosen on either a network node or on a centralized server:

[0078] ... It is assumed that for a particular protection path, **the upstream node to the link computes and updates the sharing information** (for an optical network, this may also include selection of the time slots, wavelength, and ports) and then passes this information to the downstream node so that both nodes connected to link have the same view of resources and sharing for the link....

[0080] ... In the case of a centralized architecture,... the complete sharing-information table resides on a centralized server and maybe made available to each node through a centralized command signaling structure (e.g., SNMP). In the case of a distributed architecture, similar capability can be achieved at the cost of periodically exchanging the information of the complete sharing-information table between nodes. Thus, a centralized approach provides savings in terms of this inter-nodal exchange of information, while a distributed approach provides additional robustness (e.g., the elimination of a single point of failure) and scalability.

* * *

[0250] In case of the centralized server approach, **a centralized server computes the additional protection bandwidth needed** on each link along the protection path to admit a new connection (LSP). This information is then conveyed as part of network signaling to establish the protection LSP.

Sathyanath, [0078], [0080], [0250].

172. A POSITA would have understood that Sathyanath's discussion of a "centralized server" generally corresponds to, and provides further details for, Dziong's "specialized server."

173. A node in one of the Dziong's or Sathyanath's exemplary networks is a "*network node*." Alternatively, Dziong's specialized server is also a "*network node*." Thus, Dziong and Sathyanath disclose multiple examples of a "*network node*."

b. [8.1] *a network interface, which is arranged to communicate with other network nodes over a communication network that includes multiple segments; and*

174. Dziong and Sathyanath render obvious this limitation.

175. As I discussed above for limitation [1.0], Dziong and Sathyanath describe a variety of example networks comprised of multiple links among nodes. Such networks correspond to the claimed "*communication network that includes multiple segments*."

176. Sathyanath explains that a network node has an interface coupling the node to a network link:

By convention, the **node** that owns a link is the one that is **connected to the outgoing interface side of the link**, and that node is the one that makes advertisements about the link.

Sathyanath, [0194]. Thus, it would have been obvious to a POSITA for a network node to include a “*network interface*” allowing it to connect to a network link and communicate with other nodes in a network. Indeed, a network node must have a network interface to do so, and thus it was common and conventional to do so.

177. Dziong and Sathyanath confirm that nodes in their networks are “*arranged to communicate with other network nodes.*” For example, in Dziong’s Fig. 6a example, link A-B supports a current communication demand of 12 unit:

[0090] Bandwidth reservation table 604 describes the current state of the bandwidth of link A-B in terms of cross-connection (XC) units. In particular, “P-XC” represents the number of XC units in **link A-B that currently support demands (i.e., 12 units** in this example).

Dziong, [0090].

178. Thus, node A is currently in communication with node B.

179. More generally, Dziong and Sathyanath both generally relate to communicating over a network:

[0004] The present invention relates to **communications networks**, and, more specifically, to mechanisms for calculating recovery paths in mesh communication networks.

Dziong, [0004].

[0004] The present invention relates to **network communications** and, more specifically, to restoration path sharing.

Sathyanath, [0004].

180. Thus, it would have been obvious to a POSITA that the nodes in their networks are “*arranged to communicate with other network nodes.*”

181. In summary, Dziong and Sathyanath describe networks of nodes that communicate with one another, which renders obvious for a network node to include “*a network interface, which is arranged to communicate with other network nodes over a communication network that includes multiple segments.*”

c. [8.2] a processor, which is arranged

182. Dziong and Sathyanath render obvious this limitation.

183. As I discussed above for limitation [8.0], the claimed “*network node*” corresponds to a node in Dziong’s or Sathyanath’s networks, or to Dziong’s example of a specialized server. Any of these nodes would have required a processor to perform the algorithms and calculations described in Dziong and Sathyanath. *See, e.g.,* Xia, 4:50-54 (“**components found in a conventional router, such as a processor**, a real-time operating system, and programs and hardware that implement the routing and switching functions”); Wybenga, [0052]-

[0053] (“Inbound **network processor 221 operates under the control of control software** (not shown) stored in memory 330. Similarly, outbound network processor 223 operates under the control of control software (not shown) stored in memory 340.... Network processors 221 and 223 in router 100 share routing information...”).

184. Alternatively, it would have been obvious for a network node or a specialized server to have a “*processor*” for performing the bandwidth-sharing processes described in Dziong and Sathyanath. Route-selection processes (like those of Dziong and Sathyanath) were commonly implemented as software executing on a processor. *See, e.g.*, Wybenga, [0010] (“Traditionally, **a single processor is used to forward all packets** in a router or switch.”); Xia, 3:5-7 (“Routers 102-108 may be, for example, dedicated hardware routers, or **general purpose computers running routing software...**”). Thus, it would have been obvious to a POSITA for the nodes and servers described by Dziong and Sathyanath to include a “*processor*.”

- d. **[8.3] *to identify two or more candidate routes for a new connection having a requested bandwidth to be established through the network,***

185. This limitation is substantially identical to the text of limitation [1.1]. The only minor difference is one of grammar: whereas [1.1] recites “*identifying*,” this limitation recites “*to identify*.” This difference does not affect my analysis, so

this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- e. **[8.4] *to accept a sharing definition specifying one or more other connections, which have respective bandwidths and which are permitted to share a bandwidth allocation with the new connection***

186. This limitation is very similar to the text of limitation [1.2]. In contrast to the previous limitation, however, the difference is more than mere grammar: whereas [1.2] recites “***providing a sharing definition,***” this limitation recites “***to accept a sharing definition.***”

187. Both Dziong and Sathyanath contemplate that their nodes and servers would both send and receive bandwidth sharing information. Indeed, it would make no sense for their nodes to “advertise” bandwidth sharing information throughout the network if the other network nodes did not then “*accept*” this information. For example, Dziong explains how each node will keep track of bandwidth sharing information about all recovery paths, which the node obviously would have obtained from other nodes’ advertisements:

[0054] Next, in step 306, sharing between disjoint link failures is achieved by determining, via bookkeeping information, the amount of protection bandwidth that would need to be reserved on each link for recovering demands affected by any single other link failure in the network. Recovery of each other link may require a different amount of capacity. On each link that is part of a recovery path, the maximum of

required recovery capacities required on that link is calculated. This maximum is then reserved on the link if sufficient capacity exists on the link. For a distributed implementation, **each node keeps track of this sharing information for each of its incident links. Signaling is used to update this sharing information** with admittance of every new demand into the network.

Dziong, [0054].

[0060] In embodiments of the present invention, link-detour path calculation involves maximizing sharing of link-detour bandwidth. The recovery-path calculation algorithm makes use of information including how much bandwidth can be-shared at each link in the network. This information is obtained by advertising, across the network, the amount of bandwidth reserved for recovery on each link and by **bookkeeping, in each node, information about all recovery paths** that would be activated when a protected link fails.

Dziong, [0060].

188. Sathyanath specifically describes a node updating its local database on bandwidth sharing information, then sharing the information to a downstream. It is obvious that the downstream node “accept[s]” the bandwidth sharing information since the two nodes then have “the same view of resources and sharing for the link”:

[0078] In steady-state operation of a shared mesh data network, it is assumed that once a new demand (i.e., service request) is received at a node, a route computation routine is invoked that calculates both a

primary path (also known as "working path") as well as a protection path (also known as "restoration path") for the demand. Signaling (e.g., SNMP, PNNI, or RSVP-TE signaling) is used to establish the primary path for the service and to reserve bandwidth for the protection path. The signaling along the protection path carries the information of the primary path in terms of its demand bandwidth as well as nodes and links it contains. For each of the links incident to a node in the protection path, **the node updates its local reserved bandwidth database** (corresponding to its portion of TABLE 1). This update involves incrementing, by the requested demand bandwidth, the value in each entry of each relevant row that corresponds to a link or node in the primary path (i.e., those links or nodes that, upon failure, would cause a disruption in a service along the primary path and thereby levy a bandwidth demand on the links in the corresponding protection path). The update also involves calculating the maximum value of each updated row. Note that a link connecting two nodes can be a downstream link for each node in different directions. It is assumed that for a particular protection path, **the upstream node to the link computes and updates the sharing information** (for an optical network, this may also include selection of the time slots, wavelength, and ports) and **then passes this information to the downstream node so that both nodes connected to link have the same view of resources and sharing for the link**. This scheme is specifically useful in cases where primary and protection paths are bi-directional.

Sathyanath, [0078].

189. Sathyanath further describes updating a table of bandwidth sharing information either “centrally” (e.g., at Dziong’s specialized server) or at other nodes in the network:

[0079] As an example, consider the arrival of a request for shared mesh protection service between nodes N4 and N2 of SMDN 100. Assume that, at the arrival of the request, the network was in a state captured by FIG. 1 and TABLE 1 of FIG. 2. Further, assume that LSP-11 (i.e., N4-L3-N5-L4-N2) and LSP-12 (i.e., N4-L1-N1-L2-N2) are the primary an[d] protection paths, respectively, determined to service this request. As before, signaling along the primary path is used to establish or provision the primary LSP. Also, as before, signaling along the protection path is used to carry the route information as well as the bandwidth demand (assumed in this example to be one unit) of this new service to all the nodes (i.e., N4, N1, and N2) in protection path LSP-12 for the new primary path LSP-11. As a result, **rows L1 and L2 of TABLE 1 will be updated (centrally or, in the case of a distributed architecture, at nodes N4 Oust row L1), N1 (both rows L1 and L2), and N2 Oust row L2))**. The update result is captured in TABLE 2 of FIG. 3. Note that an additional one unit of bandwidth is now needed on both links L1 and L2 to guarantee the recovery of the new connection request (LSP-11) from a failure along its working route (i.e., N4-L3-N5-L4-N2).

Sathyanath, [0079].

190. In summary, Dziong and Sathyanath suggest or describe updating a local database of bandwidth sharing information using information advertised by other nodes in the network, which renders obvious “*to accept a sharing definition.*”

191. The remainder of this limitation is the same as recited in limitation [1.2] and would have been obvious for the reasons I discussed previously.

- f.** **[8.5]** *such that a total bandwidth allocated jointly to the new connection and to the other connections in any given segment is equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment,*

192. This limitation is identical to the text of limitation [1.3]. Thus, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- g.** **[8.6]** *to calculate for each of the two or more candidate routes a cumulative additional bandwidth reservation to be allocated to the new connection in the segments along the candidate route, responsively to the sharing definition,*

193. This limitation is substantially identical to the text of limitation [1.4]. The only difference is a minor change in wording and grammar: whereas [1.4] recites “*for each of the two or more candidate routes, calculating,*” this limitation recites “*to calculate for each of the two or more candidate routes.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- h. [8.7] *to select a preferred route having a minimum value of the cumulative additional bandwidth reservation out of the two or more candidate routes, and***

194. This limitation is substantially identical to the text of limitation [1.5].

The only difference is a minor change in grammar: whereas [1.5] recites “*selecting*,” this limitation recites “*to select*.” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- i. [8.8] *to establish the new connection using the preferred route.***

195. This limitation is substantially identical to the text of limitation [1.6].

The only difference is a minor change in grammar: whereas [1.6] recites “*establishing*,” this limitation recites “*to establish*.” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

8. Claim 9

- a. [9.0] *The node according to claim 8, wherein at least some of the segments are arranged in a ring configuration associated with the node, and***

196. This limitation is substantially identical to the text of limitation [2.0].

The only difference is the minor addition of the language that the *ring configuration* is “*associated with the node*.” As shown in Dziong’s Fig. 6a example discussed above for limitation [2.0], all of the nodes A through D are on

the ring. Thus, the “*ring configuration*” of Dziong’s Fig. 6a is “*associated*” with each of the nodes.

197. The remainder of this limitation is the same as recited in limitation [2.0] and would have been obvious for the reasons I discussed previously.

b. [9.1] *wherein the processor comprises a ring-level connection admission control (CAC) module*

198. Dziong and Sathyanath render obvious this limitation.

199. As I discussed regarding limitation [8.0], Dziong describes making bandwidth sharing decisions at a “specialized server.” *See, e.g.*, Dziong, [0096]. As I discussed for limitation [8.2], it was both common and necessary for such a server to have a “processor” for performing the steps of Dziong’s bandwidth-sharing process.

200. Dziong describes admitting a new service to the network, including both the new primary (working) path and recovery path:

[0052] FIG. 3 depicts an exemplary process for achieving link-based protection at the demand level. As shown, in step 302, a working path for a new service is computed along with protection paths for each link in the working path. **In order to admit a new service into the network**, there should be sufficient capacity in the network to admit the new service along the working path and also guarantee the service's recovery from any single failure along its route.

Dziong, [0052].

201. Sathyanath further describes performing “admission control” for new network connections:

[0248] In case of static provisioning, a pool of bandwidth on each link is set aside in advance of service provisioning specifically for the purpose of recovering traffic affected by a failure. If the amount of protection bandwidth is known in advance, then static allocation becomes an option.... By using static allocation, in applying shared mesh protection methodology, bookkeeping is minimized. The **call admission control (i.e., the process that accepts a new call or data communication request into the network)** for the working traffic (LSP) typically will only admit those calls that will be sufficiently protected by the available bandwidth along the protection path (given at most a single failure in the network).

Sathyanath, [0248].

202. Since Dziong and Sathyanath both contemplate checking whether there is sufficient bandwidth available to provide a requested service, it would have been obvious to a POSITA for their bandwidth-sharing processes to include performing “*connection admission control.*”

203. As I discussed previously regarding limitation [2.0], Dziong and Sathyanath describe ring network embodiments. Where Dziong’s specialized server makes recovery path selections for a ring network, it would have been obvious for the specialized server to perform “*ring-level connection admission control.*” The portion of the processor in the specialized server that would perform

this function corresponds to the claimed “*ring-level connection admission control (CAC) module.*”

204. In summary, Dziong and Sathyanath describe a specialized server that makes connection admission decisions for a ring network, which renders obvious “*wherein the processor comprises a ring-level connection admission control (CAC) module.*”

- c. **[9.2] *which is arranged to identify first and second candidate routes traversing respective first and second ringlets of the ring configuration***

205. This limitation is very similar to the text of limitation [2.1]. The only difference is a minor change in grammar: whereas [2.1] recites “*identifying,*” this limitation recites “*which is arranged to identify.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- d. **[9.3] *and to select the preferred route traversing one of the first and second singlets.***⁴

206. This limitation substantially repeats the features recited in limitations [9.2] and [8.7]. First, limitation [8.7] recites “*to select a preferred route... out of the two or more candidate routes.*” Second, limitation [9.2] recites that the candidate routes traverse “*first and second ringlets.*” In the context of a ring

⁴ The word “singlets” in this claim is an obvious typographical error. I have analyzed the claims as if it recited “*ringlets.*”

network where the candidate routes are ringlets, it would have been obvious to a POSITA that the recovery path selection process of Dziong and Sathyanath would choose a “*preferred route*” that traverses one of the ringlets.

9. Claim 10

- a. **[10.0]** *The node according to claim 8, wherein the new connection comprises a tunnel established through the network in accordance with a tunnel-based communication protocol.*

207. This limitation is substantially identical to the text of limitation [3.0]. The only minor difference is one of grammar: whereas [3.0] recites “*establishing the new connection comprises establishing a tunnel,*” this limitation recites “*the new connection comprises a tunnel established.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

10. Claim 11

- a. **[11.0]** *The node according to claim 8, wherein the new connection and at least one of the other connections specified in the sharing definition traverse at least one common segment along the preferred route,*

208. This limitation is identical to the text of limitation [4.0]. Thus, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- b. **[11.1]** *and wherein the processor is arranged to allocate the maximum of the requested bandwidth and of the bandwidths of the at least one of the other connections*

in the at least one common segment to the new connection and to the at least one of the other connections.

209. This limitation is substantially identical to the text of limitation [4.1].

The only minor difference is one of grammar: whereas [4.1] recites “*wherein establishing the new connection comprises allocating,*” this limitation recites “*wherein the processor is arranged to allocate.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

11. Claim 13

- a. [13.0] *The node according to claim 8, wherein the processor is arranged to accept the sharing definition using a reservation protocol.***

210. This limitation is very similar to the text of limitation [6.0]. There are two minor differences, both of which relate to differences between independent claims 1 and 8 that I have already discussed.

211. First, as I already discussed regarding limitation [8.2], claim 8 is directed to a network node that includes a “*processor*” that is “*arranged*” to perform process steps that largely mirror the limitations of claim 1. Thus, claim 13 refers to the “*processor*” of its parent claim, claim 8.

212. Second, as I already discussed regarding limitation [8.4], claim 8 recites “*to accept*” a sharing definition, whereas claim 1 recites “*providing*” a sharing definition.

213. These differences do not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above for limitations [6.0], [8.2], and [8.4].

12. Claim 15

214. Claim 15 is an independent claim, and its limitations largely mirror those recited for claim 1 analyzed above. Thus, the bulk of claim 15's features would have been obvious from the disclosures of Dziong and Sathyanath for the reasons I have already discussed. I address below the specifically recited features of claim 15 that are different relative to claim 1.

a. [15.0] *A communication network, comprising:*

215. This limitation recites a portion of the text of limitation [1.0]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

b. [15.1] *a plurality of network segments*

216. This limitation is substantially similar to a portion of the text of limitation [1.0]. Where claim limitation [1.0] recited "*a communication network that includes multiple segments,*" claim limitations [15.0] and [15.1] refer to a "*communication network comprising: a plurality of network segments.*" The minor differences in wording do not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

c. [15.2] *a network node*

217. This limitation is identical to the text of limitation [8.0]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

d. [15.3] *a network interface, which is arranged to communicate with other nodes in the network; and*

218. This limitation is substantially similar to the text of limitation [8.1]. The differences between limitation [8.1] and [15.3] relate to the fact that claim 8 introduces the features of a “*communication network*” having “*multiple segments*” in limitation [8.3], whereas those concepts are recited in claim 15 through the text of limitations [15.0] and [15.1]. These minor differences do not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

e. [15.4] *a processor, which is arranged*

219. This limitation is identical to the text of limitation [8.2]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- f. ***[15.5] to identify two or more candidate routes for a new connection having a requested bandwidth to be established through the network,***

220. This limitation is identical to the text of limitation [8.3]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- g. ***[15.6] to accept a sharing definition specifying one or more other connections, which have respective bandwidths and which are permitted to share a bandwidth allocation with the new connection***

221. This limitation is identical to the text of limitation [8.4]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- h. ***[15.7] such that a total bandwidth allocated jointly to the new connection and to the other connections in any given segment is equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment,***

222. This limitation is identical to the text of limitations [1.3] and [8.5]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- i. ***[15.8] to calculate for each of the two or more candidate routes a cumulative additional bandwidth reservation to be allocated to the new connection in the segments***

along the candidate route, responsively to the sharing definition,

223. This limitation is identical to the text of limitation [8.6]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

j. **[15.9]** *to select a preferred route having a minimum value of the cumulative additional bandwidth reservation out of the two or more candidate routes, and*

224. This limitation is identical to the text of limitation [8.7]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

k. **[15.10]** *to establish the new connection using the preferred route.*

225. This limitation is identical to the text of limitation [8.8]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

13. Claim 16

a. **[16.0]** *The network according to claim 15, wherein at least some of the segments are arranged in a ring configuration associated with the network node, and*

226. This limitation is substantially similar to the text of limitation [2.0] and identical to the text of limitation [9.0]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- b. **[16.1] *wherein the processor comprises a ring-level connection admission control (CAC) module,***

227. This limitation is identical to the text of limitation [9.1]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- c. **[16.2] *which is arranged to identify first and second candidate routes traversing respective first and second ringlets of the ring configuration,***

228. This limitation is identical to the text of limitation [9.2]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- d. **[16.3] *and to select the preferred route traversing one of the first and second ringlets.***

229. This limitation is identical to the text of limitation [9.3] (apart from the typographical error “*singlets*” in [9.3] that I previously explained). Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

14. Claim 17

230. Claim 17 is an independent claim, and its limitations largely mirror those recited for claim 1 analyzed above. Thus, the bulk of claim 17’s features would have been obvious from the disclosures of Dziong and Sathyanath for the reasons I have already discussed. I address below the specifically recited features of claim 17 that are different relative to claims 1, 9, and 15.

- a. **[17.0] *A computer software product used in a network node in a communication network that includes multiple segments,***

231. This limitation reiterates much of the text of limitation [1.0], which is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

232. Regarding the language “*computer software product*,” I previously discussed for limitations [8.0] and [8.2] that it would have been obvious to a POSITA to implement the bandwidth sharing processes in Dziong and Sathyanath using a general-purpose computer processor. A POSITA would have understood that the implementation of those bandwidth sharing processes would take the form of processor instructions, i.e., computer software. Thus, it would have been obvious to implement route selection and bandwidth allocation processes as a “*computer software product*.”

- b. **[17.1] *the product comprising a tangible storage medium, in which program instructions are stored, which instructions, when read by a computer, cause the computer***

233. This limitation simply recites basic and well-known features relating to how computers store and use software. A POSITA would have been familiar with these basic computer concepts, such as having a processor execute stored instructions. Accordingly, this limitation would have been obvious to a POSITA.

- c. ***[17.2] to identify two or more candidate routes for a new connection having a requested bandwidth to be established through the network,***

234. This limitation is identical to the text of limitation [8.3]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- d. ***[17.3] to accept a sharing definition specifying one or more other connections, which have respective bandwidths and which are permitted to share a bandwidth allocation with the new connection***

235. This limitation is identical to the text of limitation [8.4]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- e. ***[17.4] such that a total bandwidth allocated jointly to the new connection and to the other connections in any given segment is equal to a maximum of the requested bandwidth and of the bandwidths of the other connections in the given segment,***

236. This limitation is identical to the text of limitations [1.3] and [8.5]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- f. ***[17.5] to calculate for each of the two or more candidate routes a cumulative additional bandwidth reservation to be allocated to the new connection in the segments***

along the candidate route, responsively to the sharing definition,

237. This limitation is identical to the text of limitation [8.6]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

g. ***[17.6] to select a preferred route having a minimum value of the cumulative additional bandwidth reservation out of the two or more candidate routes, and***

238. This limitation is identical to the text of limitation [8.7]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

h. ***[17.7] to establish the new connection using the preferred route.***

239. This limitation is identical to the text of limitation [8.8]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

15. Claim 18

a. ***[18.0] The product according to claim 17, wherein at least some of the segments are arranged in a ring configuration associated with the node, and***

240. This limitation is substantially similar to the text of limitation [2.0] and identical to the text of limitation [9.0]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- b. **[18.1] *wherein the instructions cause the computer to identify first and second candidate routes traversing respective first and second ringlets of the ring configuration, and***

241. This limitation is substantially similar to the text of limitations [2.1] and [9.2]. The only minor difference is one of grammar: whereas [2.1] recites “*wherein identifying the two or more candidate routes comprises identifying,*” this limitation recites “*wherein the instructions cause the computer to identify.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- c. **[18.2] *to select the preferred route traversing one of the first and second ringlets.***

242. This limitation is identical to the text of limitation [9.3] (apart from the typographical error “*singlets*” in [9.3] that I previously explained). Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

16. Claim 19

- a. **[19.0] *The product according to claim 17, wherein the new connection comprises a tunnel established through the network in accordance with a tunnel-based communication protocol.***

243. This limitation is substantially similar to the text of limitation [3.0]. The only minor difference is one of grammar: whereas [3.0] recites “*wherein establishing the new connection comprises establishing a tunnel,*” this limitation

recites “*wherein the new connection comprises a tunnel established.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

17. Claim 20

- a. **[20.0]** *The product according to claim 17, wherein the new connection and at least one of the other connections specified in the sharing definition traverse at least one common segment along the preferred route,*

244. This limitation is identical to the text of limitation [4.0]. Accordingly, this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

- b. **[20.1]** *and wherein the instructions cause the computer to allocate the maximum of the requested bandwidth and of the bandwidths of the at least one of the other connections in the at least one common segment to the new connection and to the existing connection.*

245. This limitation is substantially similar to the text of limitation [4.1]. The only minor difference is one of grammar: whereas [4.1] recites “*wherein establishing the new connection comprises allocating,*” this limitation recites “*wherein the instructions cause the computer to allocate.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong and Sathyanath for the reasons I discussed above.

**IX. DZIONG, SATHYANATH, AND JENSEN RENDER OBVIOUS
CLAIMS 7 AND 14**

A. Summary of Jensen

246. U.S. Patent Application Publication No. 2003/0198213 to Jensen et al. (Ex.1007, “Jensen”) was filed May 3, 2001, and published October 23, 2003. Jensen is titled “Communications Networks.”

247. Jensen describes designing a network to improve its performance, with a particular emphasis on improving the distribution of routes through a network. *See* Jensen, [0040]. Jensen explains that load-balancing can further improve a network’s performance by distributing traffic across available routes:

Load-balancing algorithms can improve the traffic distribution in the network. Using a simple load-balancing algorithm (to break ties on equal length routes, find the resource with the highest utilisation, choose the route with the lowest of the two) for both node and links showed that traffic could be balanced across the most highly utilised nodes or link.

Jensen, [0041].

A. Combining Jensen with Dziong and Sathyanath

248. As I discuss in more detail below, Sathyanath describes a route-selection process that, when there are multiple candidate routes that all have the same lowest cost, chooses among those routes at random. Jensen describes a more strategic approach to selecting a route in this situation by considering resource utilization and choosing a route with the lowest resource utilization. Thus, Jensen’s

approach improves network performance by better distributing traffic loads in a network. Accordingly, it is my opinion that a POSITA would have found it obvious to combine the concepts disclosed by Jensen with those disclosed by Dziong and Sathyanath.

1. Claim 7

- a. [7.0] *The method according to claim 1, wherein selecting the preferred route comprises reverting to a lower priority route selection method when the calculated cumulative additional bandwidth reservation is equal for all of the two or more candidate routes.***

249. As I discussed previously for limitation [1.5], Dziong and Sathyanath provide various examples and illustrations of “*selecting a preferred route.*” The examples I discussed previously did not involve multiple candidate routes that would have an equal cost, i.e., where “*the calculated cumulative additional bandwidth reservation is equal for all of the two or more candidate routes.*” As I discuss below, Jensen describes a suitable “*lower priority route selection method*” to balance traffic loads among multiple, equally preferable candidate routes.


250. Sathyanath describes with respect to Fig. 15 a process of iterating through a list of candidate routes to identify which of them has the lowest cost (and thus would require the minimum amount of additional bandwidth reservation). As each candidate route is evaluated, Sathyanath’s process tests to see whether it has a

lower cost than the previously-identified lowest-cost route. If so, then the process updates its identified lowest-cost route with the new, even-lower cost route:

[0209] The above information is consolidated into an exemplary procedure for determining a minimum-cost restoration path given a particular primary path as illustrated by FIG. 15. As shown, in step 1500, a primary path is selected for a new service of bandwidth B_p , and the set of links P -set on that primary path is initialized. Next, in step 1502, N candidate restoration paths are chosen from amongst the possible alternative paths in the network that are disjoint from the primary path. In step 1504, restoration path index j is initialized to zero, and the variables `Lowest_Cost` and `BestRPathID` are initialized to a sufficiently high value (e.g., infinity) and 0, respectively. Next, in step 1506, the cost `Cost_Rj` of the restoration path is calculated as a sum of the costs of its constituent links (more detail is shown in FIG. 16 described below), and the calculated `Cost_Rj` is compared with the current value of `Lowest_Cost` in the test of step 1508.

[0210] **If `Cost13 Rj` is less than `Lowest_Cost`, then, in step 1510, `Lowest_Cost` is set equal to `Cost_Rj`, and `BestRPathID` is set equal to restoration path index j .** Next, or if the test of step 1508 fails, in step 1512, index j is incremented, and the test of step 1514 is then performed to determine whether all the candidate restoration paths have been considered. If they have, then the procedure finishes in step 1516. If they have not, then the procedure continues with the calculation of the cost of the next candidate restoration path in step 1506.

Sathyanath, [0209]-[0210].



Comparison step checks only for a lower cost and evaluates as “false” for a candidate route having the same cost as the current lowest cost route

Sathyanath, Fig. 15 (annotated).

251. Notably, Sathyanath’s process in Fig. 15 checks in step 1508 whether the current-iteration route has a lower cost than the previously-identified lowest-cost route. If the two routes had the same cost, the comparison in step 1508 would evaluate as false, and no updating of the identified lowest-cost route would occur. Thus, when a set of candidate routes has more than one candidate route having the same lowest cost, Sathyanath’s process implicitly chooses one of those routes—specifically, whichever route is in the set of routes is evaluated first.

252. Sathyanath does not discuss or describe the order in which candidate routes would be evaluated in the Fig. 15 process. Thus, a POSITA would have understood that when there are multiple least-cost candidate routes, Sathyanath's process essentially chooses one of those routes at random, since the order in which the routes appear in the set of candidate routes is unspecified.

253. Selecting one route out of multiple candidate routes was a known technique for load-balancing traffic. For example, prior art U.S. Patent No. 6,256,309 explains that "load balancing is provided by choosing a random number between all equal-cost paths for a particular connection." US6256309, 9:63-65. As another example, U.S. Patent Publication 2003/0172180 states that "A pseudo-random number generator may be used to break ties in the case of equal cost links." US20030172180, [0063].

254. While randomly choosing among candidate routes will, indeed, distribute the traffic load across the candidate routes, a POSITA would have recognized that the traffic load would be better balanced across the candidate routes by choosing a route in a more strategic way. The choosing process does not need to be complicated. Jensen describes a "simple load-balancing algorithm"—an algorithm so simple, in fact, that Jensen explains it as a parenthetical aside:

Load-balancing algorithms can improve the traffic distribution in the network. Using a **simple load-balancing algorithm (to break ties on equal length routes, find the resource with the highest utilisation,**

choose the route with the lowest of the two) for both node and links showed that traffic could be balanced across the most highly utilised nodes or links. In effect, the two biggest hubs in the network were balanced. This is not surprising as many equal-length alternative routes traverse both hubs: these are the nodes or links that get balanced.

Jensen, [0041].

255. It would have been obvious to a POSITA to improve Sathyanath's random-selection process with a strategic route selection method like Jensen's simple load-balancing algorithm. In the combination, Jensen's load-balancing algorithm would be employed in the event that two or more candidate routes were calculated to require an equivalent amount of additional bandwidth reservation (i.e., the candidate routes were of equal cost). A POSITA would have recognized that the combination would have a predictable and beneficial result: since Jensen's method identifies resources with high utilization and chooses a route with the lowest utilization, Jensen's algorithm would steer recovery paths away from routes experiencing high utilization. In other words, Jensen's algorithm would prefer routes with low utilization. By making more use of under-utilized routes, Jensen's load-balancing algorithm would improve the balance of traffic on links and routes within a network. Jensen further explains that improving the traffic balance would have secondary beneficial effects, such as to "mitigate the effects of poor network design." Jensen, [0039]. Thus, Jensen's load-balancing algorithm would improve

the balance of traffic in Dziong's and Sathyanath's networks in the same that it does in Jensen's own networks.

256. As Jensen describes its load-balancing algorithm as "simple," it would have been within the ordinary level of skill in the art. Since the need for load-balancing was a known problem in choosing routes through a network, it would have been obvious to incorporate Jensen's load-balancing approach into the route-selection processes in Dziong and Sathyanath when there are multiple candidate routes that would require the same amount of additional bandwidth to be reserved. Invoking Jensen's load-balancing algorithm in such situations corresponds to *"reverting to a lower priority route selection method when the calculated cumulative additional bandwidth reservation is equal for all of the two or more candidate routes."*

257. In summary, Sathyanath describes choosing among multiple equivalent candidate routes at random, while Jensen describes more thoughtfully choosing among candidates by considering a secondary criterion such as resource utilization, which renders obvious *"wherein selecting the preferred route comprises reverting to a lower priority route selection method when the calculated cumulative additional bandwidth reservation is equal for all of the two or more candidate routes."*

2. Claim 14

- a. [14.0]** *The node according to claim 8, wherein the processor is arranged to revert to a lower priority route selection method when the calculated cumulative additional bandwidth reservation is equal for all of the two or more candidate routes.*

258. This limitation is substantially similar to the text of limitation [7.0].

The only difference is a minor change in wording and grammar: whereas [7.0] recites “*wherein selecting the preferred route comprises reverting,*” this limitation recites “*wherein the processor is arranged to revert.*” This difference does not affect my analysis, so this limitation is obvious from the disclosures of Dziong, Sathyanath, and Jensen for the reasons I discussed above.

X. CONCLUSION

259. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true, and that these statements were made with knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code.

Dated: 9/20/2024

Nader F. Mir
Nader F. Mir, Ph.D.