

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS INC.,

Petitioner

IPR2024-00503

U.S. Patent No. 7,035,279

**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	9
IV.	RELEVANT LEGAL STANDARDS	10
V.	OVERVIEW OF THE '279 PATENT	12
VI.	CLAIM CONSTRUCTION	15
VII.	IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE ...	16
VIII.	THE COMBINATION OF KOVVALI AND KALMAN RENDERS OBVIOUS CLAIMS 1-5, 9-18.....	17
A.	Summary of Kovvali	17
B.	Summary of Kalman	21
C.	Reasons to Combine Kovvali with Kalman	24
D.	Detailed Analysis of Claims.....	27
1.	Claim 1	27
2.	Claim 2.....	87
3.	Claim 3.....	89
4.	Claim 4.....	95
5.	Claim 5.....	100
6.	Claim 9.....	103
7.	Claim 10.....	104
8.	Claim 11.....	107
9.	Claim 12.....	113
10.	Claim 13.....	115
11.	Claim 14.....	117
12.	Claim 15.....	122
13.	Claim 16.....	124

14.	Claim 17.....	124
15.	Claim 18.....	126
IX.	THE COMBINATION OF KOVVALI, KALMAN, AND MASTERS RENDERS OBVIOUS CLAIM 6	126
A.	Summary of Masters	126
B.	Reasons to Combine Kovvali, Kalman, and Masters.....	130
C.	Detailed Analysis of Claims.....	132
1.	Claim 6.....	132
X.	THE COMBINATION OF KOVVALI, KALMAN, AND RFC2892 RENDERS OBVIOUS CLAIMS 7-8.....	139
A.	Summary of RFC2892	139
B.	Reasons to Combine Kovvali, Kalman, and RFC2892.....	140
C.	Detailed Analysis of Claims.....	140
1.	Claim 7.....	140
2.	Claim 8.....	148
XI.	CONCLUSION	150

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. 7,035,279 (“the ’279 patent”) to Bruckman.

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 1-18 (“the Challenged Claims”) of the ’279 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 ’279 patent, Ex.1001;
- b. the prosecution history of the ’279 patent (“’279 File History”), Ex.1002;
- c. U.S. Patent No. 6,920,113 to Kovvali et al. (“Kovvali”), Ex.1005;

- d. U.S. Patent No. 6,680,912 to Kalman et al. (“Kalman”), Ex.1006;
 - e. U.S. Patent No. 5,920,697 to Masters et al. (“Masters”), Ex.1007;
- and
- f. Tsiang, D. et al., Request for Comments 2892 (“RFC2892”),
“The Cisco SRP MAC Layer Protocol” (Aug. 2000), Ex.1018.

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 No. 6,839,767, Ex.1008;
- b. U.S. Patent No. 6,424,624, Ex.1009;
- c. U.S. Patent No. 6,157,955, Ex.1010;
- d. U.S. Patent No. 6,072,773, Ex.1015;
- e. U.S. Patent No. 6,111,673, Ex.1016;
- f. U.S. Patent No. 6,161,137, Ex.1017;
- g. Bradner, S., Request for Comments 2026 (“RFC2026”), “The Internet Standards Process – Revision 3” (Oct. 1996), Ex.1019; and
- h. U.S. Patent No. 6,738,819 , Ex.1020.

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. A POSITA in the field of the '279 patent, as of its priority date of January 9, 2001, would have been someone knowledgeable and familiar with network flow allocation in a ring topology. Such a POSITA would have a master's degree in computer science, computer engineering, electrical engineering, or equivalent training, and approximately two years of experience working in the field of network communications and resource allocations. Additional work experience can substitute for specific educational background, and vice versa.

21. 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 '279 patent.

IV. RELEVANT LEGAL STANDARDS

22. I am not an attorney. In preparing and expressing my opinions and considering the subject matter of the '279 patent, I am relying on certain basic legal principles that Cisco's counsel has explained to me.

23. I understand that prior art to the '279 patent includes patents and printed publications in the relevant art that predate the priority date of the '279 patent. For purposes of this Declaration, I am applying January 9, 2001, as the priority date of the '279 patent.

24. I have been informed by Cisco's counsel that a claimed invention is unpatentable under 35 U.S.C. § 103 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.


25. I have been further informed by Cisco's counsel that there are several recognized 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.

26. Also, I have been informed and understand that obviousness does not require physical combination/bodily incorporation, but rather consideration of what the combined teachings would have suggested to a POSITA at the time of the alleged invention.

V. OVERVIEW OF THE '279 PATENT

27. The '279 patent relates generally to “flow allocation in a ring topology.” Ex.1001, Title. According to the '279 patent, a network ring is a communication network comprising a plurality of nodes interconnected by links in a ring topology. Ex.1001, 2:19-21, 4:36-38. Each node can communicate with every other node over either a counterclockwise or a clockwise path around the ring. Ex.1001, 4:40-42. Figure 1 of the '279 patent illustrates network 20 that includes nodes 22 interconnected by links 1-6 to form counterclockwise (inner ring 24) and clockwise (outer ring 26) paths:



Ex.1001, Fig. 1 (annotated)

28. Figure 1 also depicts the communication network may include a dispatcher 28. The dispatcher may be an entity that is separate from the network nodes or alternatively the dispatcher may be associated with one of the nodes (e.g. a software process running on the node). Ex.1001, 4:47-51. The dispatcher is responsible for allocating resources (e.g. bandwidth) to nodes 22 as required. Ex.1001, 4:46-55. The dispatcher also determines on which of the rings (i.e. counterclockwise or clockwise) the data flow is to be directed. Ex.1001, 6:29-31.

29. When a node receives a request from one of its clients to send data to another node, the node first checks to determine whether it has sufficient resources already allocated to it in order to service the request. Ex.1001, 6:11-14. If the resources allocated to the node are inadequate, the node submits a request to the dispatcher for additional resources. Ex.1001, 6:14-17.

30. If resources are available (both sufficient unallocated resources left in the specific node budget, and sufficient unused resources left on the appropriate segments of at least one of the rings) the dispatcher determines on which of the rings to allocate the resources and allocates one or more additional “quanta” of resources to the node on the selected ring. Ex.1001, 6:29-31, 6:19-33. The node then carries out the data transmission. Ex.1001, 6:33-34.

31. On the other hand, if the node determines that the resources allocated already to the node are sufficient, the node carries out the requested transmission “without having to communicate with the dispatcher.” Ex.1001, 6:48-51.

32. When both rings have sufficient resources available on all relevant links, the dispatcher chooses one of the rings on which to direct the data flow. Ex.1001, 8:1-6. “[T]he dispatcher proceeds to choose the ring that has, generally speaking, the lower level of current resource allocation.” Ex.1001, 8:3-5. In one embodiment, the ’279 patent describes that relative resource “usage levels” can be

compared for all relevant segments on both rings to select a path for the data flow.
Ex.1001, 8:9-11.

33. As I will explain below, these concepts were well known as of the priority date of the '279 patent.

VI. CLAIM CONSTRUCTION

34. It is my understanding that in order to properly evaluate the '279 patent, the terms of the claims must first be interpreted. It is my understanding that for the purposes of this *inter partes* review, the claims are to be construed under the so-called *Phillips* standard, under which claim terms 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.

35. I have reviewed the entirety of the '279 patent, as well as its prosecution history. It is my opinion that for purposes of applying the prior art presented herein to evaluate the patentability of the Challenged Claims no claim terms require construction. Accordingly, I have analyzed the claims consistent with their ordinary and customary meaning as would be understood by a POSITA.

VII. IDENTIFICATION OF HOW THE CLAIMS ARE UNPATENTABLE

36. The discussion in this Declaration provides a detailed analysis of how the asserted prior art references teach each limitation of the Challenged Claims.

37. 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. In addition to various U.S. Patents, I have reviewed and appropriately cite to a number of “Request for Comments” (RFC) documents published by the Internet Engineering Task Force (IETF). RFCs were well-known as of the priority date of the ’279 patent, and it is my opinion that they were “printed publications” for the purposes for which they are cited throughout. *See, e.g.*, RFC 2026, 6, (“RFCs can be obtained from a number of Internet hosts...”), 8 (IETF documents are “readily available to a wide audience”), 26 (IETF participants “shall publicly announce...every activity” relating to the standardization process). Exhibit 1018 is a document published by the IETF describing the Spatial Reuse Protocol, and a POSITA in the field of network communications, including ring networks and resource allocation methodologies, would have reasonably relied upon such documents. POSITAs frequently cited these same documents before the priority date of the ’279 patent. Ex.1005, 7:1-4 (citing RFC 1483). The ’279 patent itself refers to several RFC documents, including RFC2892. Ex.1001, 1:21-24 (“One solution that has been proposed to meet these needs is the Spatial Reuse Protocol (SRP), which is

described by Tsang et al., in Request for Comments (RFC) 2892 of the Internet Engineering Task Force (IETF). This document which is available at www.ietf.org/rfc.html, is incorporated herein by reference.”; “using methods known in the art, as specified by SRP (RFC 2892)”); *see also* Ex.1001, 1:56-59 (citing RFC 2328), 5:53-58 (citing RFC 2005), 6:1-4 (citing RFC 2814).


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

VIII. THE COMBINATION OF KOVVALI AND KALMAN RENDERS OBVIOUS CLAIMS 1-5, 9-18

A. Summary of Kovvali


39. Like the '279 patent, Kovvali describes a network comprising a plurality of “nodes” interconnected by “links” in a “ring” topology. Ex.1005, 8:19-21. Nodes communicate with each other by passing data over dynamic channels around a ring. Ex.1005, 8:39-41. The ring may be of various types including, for example, a UPSR (unidirectional path switched ring) or a BLSR (bi-directional line switched ring). Ex.1005, 8:30-36.

40. Kovvali’s Figure 4 illustrates a plurality of nodes 420 (labeled A-D) coupled by optical links 422 to form a UPSR, in which all communication between nodes travels in one direction. Ex.1005, 8:30-32.



Ex.1005, Fig. 4 (annotated)


41. Applying Kovvali's teaching of a BLSR, in which "communication travels between nodes on the ring in both directions, with two unidirectional optical links joining each pair of neighboring nodes," Ex.1005, 2:12-16, modified Figure 4 illustrates a plurality of nodes (labeled A-D) coupled by optical links to form "counterclockwise" (blue) and "clockwise" (orange) paths. See Ex.1005, 8:32-36.



Ex.1005, Fig. 4 (modified and annotated)

42. In Figure 4, Kovvali illustrates that node A may be assigned to be an “arbiter” node that acts as a centralized bandwidth manager for bandwidth requests. Ex.1005, 9:1-6, 7:32-33, 4:47-50. Like the dispatcher in the '279 patent, the arbiter node is responsible for assigning resources, such as “dynamic channels” to client nodes and adjusting their allocated rates. Ex.1005, 9:8-11. After the network ring is provisioned, the arbiter node provides “initial data rate allocations for each of the dynamical channels” it assigned. Ex.1005, 12:5-6.

43. The dynamic channel allocations are in “units of columns” of a dynamic frame. Ex.1005, 12:6-7. In one example, Kovvali describes the dynamic frame as including three sections: (A) a three-column overhead section 520 for transmitting dynamic bandwidth management information between nodes 420 and the arbiter node, Ex.1005, 9:49-50, (B) a TDM (time division multiplexed) section 530 for fixed-rate channels, Ex.1005, 9:62-64, and (C) an elastic section 540 for dynamic channels, Ex.1005, 9:65-10:3, as illustrated in Figure 5A:



Ex.1005, Fig. 5A (annotated)

44. Kovvali describes that the dynamic channels in the elastic section can be dynamically allocated to accommodate “bursty data” that may exhibit time-varying rates. Ex.1005, 9:65-10:1, 8:49-53. A client node may request to modify the capacity of assigned channels using control channels in the dynamic frame overhead. Ex.1005, 10:14-18. For example, a client node may send a request for additional

bandwidth to the arbiter node by setting “a request section (REQ)” in the dynamic frame overhead for a bursty channel. Ex.1005, 10:35-38. “The frame with the REQ value passes from the client node to the arbiter node, which extracts the REQ value.” Ex.1005, 12:28-33. The request may be for zero or more columns of bandwidth in the dynamic frame. Ex.1005, 10:39-54.

45. “The arbiter node then determines how to increase or decrease allocations to one or more dynamic channels to accommodate the request.” Ex.1005, 12:35-37. The arbiter node indicates the changes in allocations by setting a “grant section (GNT)” in the dynamic frame overhead for the corresponding bursty channel, indicating zero or more columns of allocated bandwidth. Ex.1005, 10:37-38, 55-60.

B. Summary of Kalman

46. Kalman relates to directing traffic in a ring network topology. Kalman describes a technique for selecting “the preferred output direction for traffic from the source to any destination” on a bi-directional ring network. Ex.1006, 8:5-7. Like the '279 patent, “[t]he disclosed network includes two rings, wherein a first ring transmits data in a ‘clockwise’ direction (or direction 0), and the other ring transmits data in a ‘counterclockwise’ direction (or direction 1 opposite to direction 0).” Ex.1006, 2:61-65. Although there are two paths between any two nodes (i.e. clockwise and counterclockwise), “[a] given packet/flow between two nodes is

transmitted in only a single direction around the network (even when there is a span fault) and is removed from the ring by the destination node.” Ex.1006, 4:40-43.

47. Figure 3 illustrates an example “where node 0 transmits information to node 5 in only the [counterclockwise] direction.... A transmission from node 5 to node 0 would only go through nodes 6 and 7 in the opposite direction. This allows for optimized ring capacity utilization since no capacity is set aside for protection.” Ex.1006, 4:43-48. At node 5, the transmitted information is removed from the ring. The portion of bandwidth used for the transmission from node 0 to node 5 is thus available for other transmissions on the counterclockwise path from node 5 back to node 0. Kalman describes this as “spatial reuse of capacity.” Ex.1006, 4:58-67.




FIG. 3

Ex.1006, Fig. 3 (annotated)

48. Typically, the “least-cost physical route” is selected as the preferred direction for a given data flow. Ex.1006, 4:49-50. Kalman describes that “[f]or each destination node j , a cost is assigned to each output direction (0 and 1)” from the source node on the ring. Ex.1006, 7:50-51. A preferred direction is selected “based on the direction with the lowest cost.” Ex.1006, 7:51-53.

49. The computation of cost from each source node to each destination node “may take into account a variety of factors, including the number of hops from

source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction.” Ex.1006, 7:66-8:5. “Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.” Ex.1006, 8:5-7.

C. Reasons to Combine Kovvali with Kalman

50. A POSITA would have combined Kovvali with Kalman because they both relate to efficient transmission of data in a ring network. Ex.1005, Abstract (“A method and apparatus . . . to efficiently carry both data and voice traffic . . . in a SONET/SDH fiber-optic ring topology.”); Ex.1006, 4:20-23 (“The purpose of the invention described herein is to achieve fast protection in a ring network while providing for efficient network capacity utilization.”).

51. A POSITA would have been motivated to combine Kovvali with Kalman and doing so would have been within a POSITA’s skillset. The combination would have been obvious because it represents merely the use of Kalman’s known technique for selecting a preferred path on a ring network to improve Kovvali’s multi-path ring network for transmitting dynamic frames in the same way. Faced with the choice of transmitting a dynamic frame over one of two paths (e.g. clockwise path or counterclockwise path), a POSITA would look to Kalman for details on how to make the path selection.

52. The combination would have also been obvious because the application of Kalman's technique (selecting a preferred path on a ring network) to improve Kovvali's method (transmitting dynamic frames on a multi-path ring) ready for improvement would yield predictable results. Here, the improvement provided to Kovvali's multi-path ring network by Kalman's path selection technique is efficient transmission of dynamic frames in a ring network using path cost.

53. Kovvali's multi-path ring network would also be improved because the use of a ring selection technique by a centralized bandwidth manager (i.e. Kovvali's "arbiter" node) would allow path selection that accounts for dynamic frame transmissions among all the network nodes across the entire network. By centralizing path selection decisions at the arbiter, more efficient utilization of available network capacity is achieved.

54. For example, referring to Kovvali's modified and annotated Figure 4 discussed in Section VIII.A, in response to a first request to transmit a first dynamic channel from Node B to Node D, the lowest cost path calculated for transmission may be the clockwise (blue) path through Node C. Kalman teaches that the lowest cost path is "typically" used for transmission of "a given packet/flow" between nodes, Ex.1006, 4:49-50, 4:40-45, and thus the arbiter would select this clockwise path for transmitting its data flow. Generally, however, the arbiter would expect to receive more than one request for a dynamic channel, such as a second request for a

second dynamic channel to be transmitted between Node B and Node C. In such a scenario, when selecting a path for transmission of a particular dynamic channel, the arbiter would consider other factors such as the priority of data to be transmitted. For example, if the data associated with the first request has low priority and is delay tolerant, the arbiter may choose to send the respective dynamic channel over a more congested path thereby leaving available paths with less congestion for other requests that may have higher priority and are less delay tolerant. *See, e.g.*, Ex.1006, 15:15-19, 16:49-50 (noting how some “lower priority traffic” may be “delayed” to leave bandwidth available to support transmission requirements of higher-priority “protected” traffic). Thus, by maintaining information about the current amount of bandwidth allocation throughout the network (e.g. congestion), the centralized arbiter would make use of network capacity efficiently by selecting available paths to service multiple requests according to the particular characteristics of the data they transmit.

55. Further, a POSITA would have understood that engineers working on bandwidth allocation schemes would consider a myriad of network configurations to use network resources efficiently. For example, it was well known that, “[t]he final aim of a congestion scheme is to optimize the overall network performances. Schemes that optimize the throughput of one user or one trunk, do not always maximize the overall performance.” Ex.1009, 6:1-5. Thus, a POSITA would have

been motivated to implement Kalman's ring selection technique in Kovvali's centralized bandwidth manager (the "arbiter" node) to optimize the use of network resources for all dynamic channels instead of just one.

56. A POSITA would have had a reasonable expectation of success in combining the teachings of Kovvali and Kalman. Kalman describes that its technique for selecting a preferred path for transmitting data on a network ring may be applied "to any network transmitting frames, cells, or using any other protocol." Ex.1006, 3:38-41. Kalman also describes that, like Kovvali, the nodes of the network ring are interconnected by fiber optic cables. Ex.1006, 11:37-42. Because Kovvali also describes a ring network transmitting dynamic frames between nodes over an optical network, a POSITA would have understood readily that the path selection techniques described in Kalman are applicable to Kovvali. A POSITA also would have had an expectation of success from the combination because of the overall similarities between the network architectures and technologies described in Kovvali and Kalman.

D. Detailed Analysis of Claims

1. Claim 1

- a. [1.0] *In a communication network that includes a plurality of nodes interconnected by links, the nodes including at least first and second nodes, a method for routing a data flow, comprising:***

57. To the extent that the preamble is limiting, Kovvali renders it obvious.

58. Kovvali teaches “a method for communicating between **multiple nodes** of **a communication system** in which a series of fixed-length frames pass between the nodes of the system.” Ex.1005, 4:55-58.

59. Kovvali’s communication network includes nodes, such as IP-based routers, that route a data flow (e.g. data, voice, and/or video traffic) in a SONET/SDH fiber optic ring topology:

A second approach to passing data over SONET infrastructure is called Dynamic Packet Transfer (DPT). This approach uses IP as the basic multiplexing technology instead of ATM. Like the ATM VP Ring approach, the DTP approach assumes replacement of the existing transport infrastructure at each SONET node with **IP-based routers**. DPT defines a layer 2 protocol called Spatial Reuse Protocol (SRP), which is used to transport IP datagrams between nodes in a ring. SRP allows reuse of ring bandwidth across different sections of a ring of nodes. A supplementary algorithm called SRP Fairness algorithm (SRP-fa) was also defined to achieve fair sharing of bandwidth, between various nodes in a ring topology.


Ex.1005, 4:9-21.

In a general aspect, the invention provides a method and apparatus for using SONET technology **to efficiently carry both data and voice traffic** by Add/Drop

Multiplexing of both isochronous traffic, **such as voice and video**, and non-isochronous traffic, **such as data, in a SONET/SDH fiber-optic ring topology**. A portion of a SONET/SDH frame structure (STS-1, STM-1 etc.) is shared by a **number of nodes connected to an optical SONET/SDH ring** (either directly or sitting behind a SONET/ADM).

Ex.1005, 4:35-43.

60. Kovvali depicts in Figure 4 and accompanying text an embodiment that includes a number of SONET nodes (labeled A-D) that are coupled over optical links in a SONET ring:



Ex.1005, Fig. 4 (annotated)

[...]

Referring to FIG. 4, in one embodiment, a **number of SONET nodes 420** (labeled A-D) are **coupled over optical links 422** in a SONET ring 410. Optical links 422 carry signals according to a standard framing format, in this embodiment according to STS-1 framing using OC-1 signaling (see STS-1 block 250 and OC-1 block 280 shown in FIG. 2). In alternative embodiments, optical links 422 may carry signals at a higher standard rate, for example using STS-3 framing with OC-3 signaling, or

STS-12 framing with OC-12 signaling. The following description focuses on STS-1 of STS-3 payloads, however, the approach is applicable to all SONET rates. In FIG. 4, SONET ring 410 is a UPSR (unidirectional path switched ring), in which all communication between nodes travels in one direction. In alternative embodiments, other types of rings, such as a BLSR (bi-directional line switched ring) in which communication between nodes on the ring travels in both clockwise and counterclockwise directions, are used.


Ex.1005, 8:19-36.

61. Therefore, Kovvali teaches a “*communication network that includes a plurality of nodes interconnected by links, the nodes including at least first and second nodes, a method for routing a data flow*” as recited in the preamble.

b. [1.1] *allocating a respective number of quanta of one or more resources in the network to each of the nodes;*

62. Kovvali renders this limitation obvious.

63. First, Kovvali describes that when the network ring is provisioned, one node is assigned as an “arbiter” node, and the remaining nodes are “client” nodes:



Ex.1005, Fig. 4 (annotated)

$$[\dots]$$

In order to dynamically manage the TDM and dynamic channels, **one node 420 is assigned to be the arbiter node** at the time the ring is provisioned. **The remaining nodes** that make use of the TDM and dynamic channels **are client nodes**. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. Requests to assign dynamic channels for communication between particular nodes are sent from a client nodes to the arbiter. The arbiter responds to the requests by assigning dynamic

channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.

Ex.1005, 9:1-11.

64. Second, Kovvali describes that the arbiter node is responsible for assigning “dynamic channels” to particular nodes after the network ring is provisioned. In particular, during configuration of the network ring the arbiter node provides initial data rate allocations to each node’s assigned dynamic channel (*“allocate a respective number of . . . one or more resources in the network to each of the nodes”*):

In a next, optional, step of configuration, **the arbiter node assigns a number of dynamic channels 436 to couple particular nodes (step 930). This involves the arbiter node notifying each node of the dynamic channel which it terminates, and providing the initial data rate allocations for each of the dynamic channels.** The allocations are in units of columns of the dynamic frame. Note that certain dynamic channels initially may be allocated no capacity (zero columns). The client nodes **420** compute the starting column for the dynamic channels by determining the end of TDM section **530**, and then summing the allocated number of columns for all other dynamic channels that have lower identifying numbers to determine the start of that channel's allocation relative to the end of the TDM section.

Ex.1005, 12:1-14.


65. The arbiter node also assigns dynamic channels and adjusts their allocated rates upon request from respective client nodes:

In order to dynamically manage the TDM and dynamic channels, one node 420 is assigned to be the arbiter node at the time the ring is provisioned. The remaining nodes that make use of the TDM and dynamic channels are client nodes. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. Requests to assign dynamic channels for communication between particular nodes are sent from a client nodes to the arbiter. **The arbiter responds to the requests by assigning dynamic channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.**

Ex.1005, 9:1-11.

66. Third, Kovvali teaches that each dynamic channel is allocated zero or more columns (“*quanta*”) in a frame structure. Ex.1005, 10:3-5.

67. In particular, Kovvali describes an STS-1 frame that includes a Synchronous Payload Envelope (SPE) for carrying data between nodes. Ex.1005, 2:34-54. Figure 3 depicts an embodiment of two STS-1 frames, each comprising 90 columns by 9 rows of 8-bit bytes:



Ex.1005, Fig. 3 (annotated)

[...]

Referring to FIG. 3, in the case of STS-1 framing, each of a sequence of STS-1 frames 310 carries an SPE 330. **Each STS-1 frame is typically conceptualized as an array of 90 columns by 9 rows of 8-bit bytes.** As a signal, the rows are concatenated one after another to form a linear signal. Each SPE 330 holds 87 columns by 9 rows, with the remaining 3 columns of the STS-1 frame being used for control and overhead data. In transmission, an SPE is not necessarily aligned with an STS-1 frame. Each STS-1 same has a line and section overhead section 320 that includes a pointer to the first byte of the SPE.


Ex.1005, 2:55-65.

68. Each SPE holds 87 columns by 9-rows of 8-bit bytes:

Each SPE 330 is 87 columns by 9 rows in size. The first column of the SPE is an STS path overhead 340. When the SPE holds seven VT groups 350, the VT groups each use 12 columns; 2 extra columns are left unused to fill the 87 column SPE. If the content of the SPE is not a set of VT groups, then the fill remaining 86 columns can be used for data.

Ex.1005, 2:66-3:5.

69. Kovvali depicts in Figure 5A an embodiment of an STS-1 SPE:



Ex.1005, Fig. 5A (annotated)

70. As shown in Figure 5A, the dynamic frame includes three sections: (A) a three-column overhead section 520 for transmitting dynamic bandwidth management information between nodes 420 and the arbiter node, (B) a TDM (time division multiplexed) section 530 for fixed-rate channels, and (C) an elastic section 540 for dynamic channels:

Referring still to FIG. 5A, **dynamic frame 510 includes three sections. A dynamic frame overhead section 520,**

in this embodiment made up of three columns, is used to transport dynamic bandwidth management information between nodes 420 and the arbiter node. For instance, requests for bandwidth from client nodes and responses from the arbiter node pass in dynamic frame overhead section 520. In alternative embodiments, the byte positions of these sections and fields within the sections may differ.

A second section, **TDM section 530**, is used for statically configured isochronous circuits between nodes. In one alternative, TDM section 530 is reserved for traditional virtual tributary multiplexing, with the minimum granularity of one VT1.5. In another alternative, TDM section 530 includes circuits that are allocated in units of one column. In either alternative, the number of columns for TDM traffic is determined by static provisioning of the fixed-rate channels.

A third section, **elastic section 540**, carries channels that are dynamically allocated and that typically carry bursty data, therefore benefiting from dynamic allocation of their rates to accommodate variations in their rate. Elastic section 540 is provisioned to include a fixed number of dynamic channels. At any time, each dynamic channel is allocated zero or more columns of elastic section 540, and is used for communication between two or more nodes 420. Each of the dynamic channels is associated with a

particular control channel in dynamic frame overhead
section 520.

Ex.1005, 9:48-10:7.

71. Each dynamic channel is allocated zero or more columns of resources:

A third section, **elastic section 540**, carries channels that **are dynamically allocated** and that typically carry bursty data, therefore benefiting from dynamic allocation of their rates to accommodate variations in their rate. Elastic section 540 is provisioned to include a fixed number of dynamic channels. **At any time, each dynamic channel is allocated zero or more columns** of elastic section 540, and is used for communication between two or more nodes 420. Each of the dynamic channels is associated with a particular control channel in dynamic frame overhead section 520.

Ex.1005, 9:65-10:7.

Compared to the ATM VP approach, less bandwidth is lost to overhead. Each ATM cell uses a 5 byte cell header and a 48 byte cell payload (almost 10% overhead). When IP data is encapsulated, for example using LLC/SNAP headers (defined in RFC 1483) still more bandwidth is lost to overhead. In contrast, the invention can use a byte or less of overhead for each dynamic channel per STS-N frame, and **the dynamic channel can be allocated multiple columns in the frame**. This can amount to less

than 1% overhead when STS-3c frames are used according to this invention.

Ex.1005, 6:66-7:8.

In a next, optional, step of configuration, the arbiter node assigns a number of dynamic channels 436 to couple particular nodes (step 930). This involves the arbiter node notifying each node of the dynamic channel which it terminates, and providing the initial data rate allocations for each of the dynamic channels. **The allocations are in units of columns of the dynamic frame.** Note that certain dynamic channels initially may be allocated no capacity (zero columns). The client nodes 420 compute the starting column for the dynamic channels by determining the end of TDM section 530, and then summing the allocated number of columns for all other dynamic channels that have lower identifying numbers to determine the start of that channel's allocation relative to the end of the TDM section.

Ex.1005, 12:1-14.

72. The columns of dynamic channel resources correspond to bandwidth:

In a general aspect, the invention provides a method and apparatus for using SONET technology to efficiently carry both data and voice traffic by Add/Drop Multiplexing of both isochronous traffic, such as voice and video, and non-isochronous traffic, such as data, in a SONET/SDH fiber-

optic ring topology. A portion of a SONET/SDH frame structure (STS-1, STM-1 etc.) is shared by a number of nodes connected to an optical SONET/SDH ring (either directly or sitting behind a SONET/ADM). **The inherent column structure of the SONET/SDH frame is used to define dynamic channels (flexible bandwidth channels) between the nodes. Each dynamic channel can be assigned a priority and optionally a minimum bandwidth.** The channel to column mapping is varied dynamically for data traffic by a centralized bandwidth manager which serves as an arbiter for requests for bandwidth. Both isochronous traffic and best effort data traffic may be carried in an STS-1 frame thereby allowing over-subscribing and statistical multiplexing of data channels over the unassigned portion of the STS-1 frame.

Ex.1005, 4:35-54.

Referring still to FIG. 5A, each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to **request additional bandwidth** for the corresponding bursty channel from the arbiter. GNT527 is set by the arbiter node to **grant bandwidth** to the requested channel. The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel. In this embodiment, in which control

channel 524 is 8 bits, REQ 526 is 3 bits and GNT 527 is 5 bits, the values of REQ and GNT are encoded as follows:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

73. Kovvali’s teaching of allocating columns of resources (i.e. bandwidth) to dynamic channels assigned to nodes is much like the example of allocating resources in the ’279 patent specification where resources include “bandwidth” and “any other allowable factor that may affect the ability of a network link to convey a data flow between nodes”:

It is a further object of some aspects of the present invention to provide a method of data flow allocation that makes optimal use of resources that are available on different network links. **The term “resources” in this context, and throughout the specification and claims of the present patent application, refers not only to**

bandwidth, but also to processing power used to transmit, receive, encrypt or decrypt data, as well as to any other allowable factor that may affect the ability of a network link to convey a data flow between nodes.

Ex.1001, 2:9-18. Accordingly, Kovvali renders obvious “*allocating a respective number ... of one or more resources in the network to each of the nodes.*”

74. A POSITA would have found it obvious that allocating columns of resources (i.e. bandwidth) is allocating “*quanta*” of resources because each column is the smallest discrete unit of bandwidth that can be allocated to a dynamic channel. In particular, with reference to Figure 5A, each column of resources allocated to a dynamic channel corresponds to 9-rows of 8-bit bytes. Ex.1005, 2:55-65, 2:66-3:5. Thus, resources are allocated to dynamic channels in columns that each correspond to 72 bits of bandwidth. Accordingly, in an exemplary scenario where 100 bits of data are signaled between nodes using a dynamic channel, two columns (or 144 bits) of resources would be allocated. In a separate exemplary scenario where 130 bits of data are signaled between nodes using a dynamic channel, two columns (or 144 bits) of resources would be allocated. Because two columns of resources would be allocated when either 100 or 130 bits of data are signaled between nodes, the frequency with which allocation operations must be performed would be reduced. Indeed, two columns of resources would be used to transmit between from 73 up to

144 bits of information without changing the number of columns allocated to the dynamic channel.

75. Kovvali's teaching of allocating columns of resources to dynamic channels is much like the example of allocating resources in "*quanta*" in the '279 patent specification where resources are allocated in quanta to limit the frequency with which allocation and deallocation operations must be performed:

In addition to maintaining segment availability information, the dispatcher is configured with the following parameters for each of the nodes with respect to each of the resources:

node_budget_In—Maximum amount of the resource that the node can receive on the inner ring.

node_budget_Out—Maximum amount of the resource that the node can receive on the outer ring.

node_used_In—Amount of the resource currently allocated to the node on the inner ring.

node_used_Out—Amount of the resource currently allocated to the node on the outer ring.

node_quantum—Amount by which to increase or decrease node_used_In or Out at each resource allocation or deallocation. **Resources are allocated in quanta in order to limit the frequency with**

which allocation and deallocation operations must be performed.

Other parameters used by the dispatcher in its allocation decisions are noted below.

Ex.1001, 5:14-33.

76. Therefore, Kovvali's disclosure of assigning "dynamic channels" to the nodes in the network ring where the resources in the dynamic channels are allocated in "units of columns," renders obvious *"allocating a respective number of quanta of one or more resources in the network to each of the nodes."*

- c. [1.2] *receiving a request at the first node to use a portion of the resources so as to carry the data flow between the first and second nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links;*

77. Kovvali renders this limitation obvious.

78. **First**, as discussed in [1.1] "dynamic channels" are used for communication between two or more nodes (*"carry data flow between the first and second nodes"*):

A third section, **elastic section 540**, carries channels that **are dynamically allocated** and that typically carry bursty data, therefore benefiting from dynamic allocation of their rates to accommodate variations in their rate. Elastic section 540 is provisioned to include a fixed number of dynamic channels. At any time, **each dynamic channel** is

allocated zero or more columns of elastic section 540, and **is used for communication between two or more nodes 420**. Each of the dynamic channels is associated with a particular control channel in dynamic frame overhead section 520.

Ex.1005, 9:65-10:7.

79. Second, each of Kovvali's nodes function as an add/drop multiplexer (ADM) for accommodating dynamic channels. In particular, each node extracts data from dynamic channels that terminate at the node, and multiplexes outbound data from local dynamic channels onto the ring:

In operation, **each node 420 functions as an add/drop multiplexer for the TDM and dynamic channels. That is, the node** extracts data for TDM and dynamic channels that terminate at the node, and **multiplexes outbound data onto the ring for those channels**. Based on the configuration information received by the nodes, each node is aware of which TDM and dynamic channels are added/dropped at that node.


Ex.1005, 12:20-27.

Referring to FIG. 1, communication according to the SONET standard makes use of a ring architecture in which a number of communication nodes are connected by optical links to form a ring. **A SONET ring 110 typically has a number of nodes each of which includes an**

add/drop multiplexer (ADM) 120. Each of the nodes are coupled to two neighboring node by optical paths 122. Communication passes around the ring in a series of synchronous fixed-length data frames. **Each ADM 120 is configured at the time the ring is provisioned to** pass a portion of the communication on the ring without modifying it, to extract (drop) a portion of the communication destined for that node, and **to add outbound communication leaving the node to the optical path.** The dropped and added communication passes between the ADM and local communication equipment, such as a multiplexer 130, which multiplexes a number of separate traffic streams 134. For example, the a separate communication streams 134 may be 1.5 Mb/s (DS1) data stream on which separate 64 kb/s (DS0) telephone channels that are multiplexed. The DS1 data stream is multiplexed onto the optical path and passed between particular nodes on the SONET ring. Typically, SONET ring, 110 is provisioned to provide fixed rate bidirectional communication streams, also known as virtual paths, between pairs of communication steams 134, each of which is coupled to a different ADM 120 on the ring. In operation, the virtual paths coupling different communication streams 134, including their allocated data rates, typically remains fixed for long periods of time.


Ex.1005, 1:48-2:9.

80. Each node that terminates (e.g. that adds) a dynamic channel multiplexes outbound data that was received locally for transmission onto the dynamic channel. Ex.1005, 14:33-37. Kovvali depicts in Figure 4 an example of Node B and Node C each receiving data from local dynamic channels 436 for transmission. Similarly, Nodes B and C perform corresponding actions for receiving data from the ring, e.g., extracting inbound data from the dynamic channels and transmitting it to the local dynamic channels 436. This bidirectional data communication capability is depicted in Fig. 4 through the use of double-headed arrows (green):



Ex.1005, Fig. 4 (annotated)

81. Figure 6 depicts another example where a node 420 receives outgoing packet traffic from local dynamic channels 436 for transmission:



Ex.1005, Fig. 6 (annotated)

82. When a node receives data from local dynamic channels, it processes the data so that the data may be multiplexed onto the outgoing link. In Figure 6, Kovvali illustrates a packet processor 644 that provides an interface to local dynamic channels 436. Ex.1005, 14:27-29. Packet processor 644 includes buffers and other mechanisms to implement appropriate queuing, scheduling and policing processes necessary to handle outgoing packet traffic. Ex.1005, 14:29-33. A Tx (transmit)

packet processor 634 multiplexes the local dynamic channels that are passed from packet processor 644 onto Tx bus 614, and then to SONET framer 610 where they are multiplexed onto the ring optical link. Ex.1005, 14:33-37.

Referring to FIG. 6, a SONET node 420 (FIG. 4), according to the invention, includes several components. **A SONET framer 610 is coupled to optical links 422, which couple the node to its neighboring nodes. SONET framer 610 performs synchronous byte multiplexing/demultiplexing** into 125 microsecond STS-1 frames as dictated by the SONET/SDH standards. SONET framer 610 extracts the receive clock (RxC) and performs SONET/SDH overhead processing (Section/Line Overhead) for the received frames. The SONET framer extracts the Synchronous Payload Envelope (SPE) and places the SPE byte stream on a Rx bus 612. SONET framer 610 also generates a SPE Start signal in order to trigger the back end circuitry to process the payload portion of the byte stream on Rx bus 612. SONET framer 610 also computes the VT pointer and generates a VT Start signal to trigger processing of the incoming Virtual Tributaries.

The back-end circuitry includes a TDM Receive (Rx) processor 622 and a TDM transmit (Tx) processor 624. A slot table 626 holds data identifying the positions of the TDM channels that are added/dropped at this node. TDM

Rx processor 622 passes the TDM channels to TDM Tx processor, as well as a TDM tributary processor 642. TDM tributary processor 642 provides an interface to TDM channels 434. TDM Tx processor 624 multiplexes the TDM channels that pass through the node (i.e., that are not added/dropped at the node) and TDM channels that are passed from TDM tributary processor 642. The multiplexed TDM data is passed to Tx bus 614 and then to SONET framer 610 where it is multiplexed onto the outbound optical link. Dynamic channel data is passed from Rx bus 612 to an Rx packet processor 632, which extracts the elastics section of the frame. Rx packet processor 632 passes the channels that terminate at the node to a **packet processor 644 which provides an interface to dynamic channels 436. Packet processor 644 includes buffers and other mechanisms to implement appropriate queuing, scheduling and policing processes necessary to handle incoming and outgoing packet traffic. A Tx packet processor 634 multiplexes** dynamic channels that are not added/dropped at this node and **dynamic channels that are passed from packet processor 644 onto Tx bus 614, and then to SONET framer 610 where it is multiplexed onto the outbound optical link.**

Ex.1005, 13:62-14:37.

83. A POSITA would have understood that the receipt of data from local dynamic channels by a node (e.g. Node B or Node C) for transmission to another node renders obvious “*receiving a request at the first node to use a portion of the resources*” because it is a request for the node to process (e.g. queue, schedule, police) and multiplex the data onto the portion of network resources allocated by the arbiter (as discussed in [1.1]) to the node for the respective dynamic channels in the dynamic frame.

84. Third, Kovvali describes that while Figure 4 depicts a UPSR (unidirectional path switched ring), in which all communication between nodes travels in one direction, alternative embodiments include other types of rings, such as BLSR (bi-directional line switching ring), in which communications between nodes on the ring travels in both clockwise and counterclockwise directions over two unidirectional links joining each pair of neighboring nodes:

Referring to FIG. 4, in one embodiment, a number of SONET nodes 420 (labeled A-D) are coupled over optical links 422 in a SONET ring 410. Optical links 422 carry signals according to a standard framing format, in this embodiment according to STS-1 framing using OC-1 signaling (see STS-1 block 250 and OC-1 block 280 shown in FIG. 2). In alternative embodiments, optical links 422 may carry signals at a higher standard rate, for example using STS-3 framing with OC-3 signaling, or

STS-12 framing with OC-12 signaling. The following description focuses on STS-1 of STS-3 payloads, however, the approach is applicable to all SONET rates. **In FIG. 4, SONET ring 410 is a UPSR (unidirectional path switched ring), in which all communication between nodes travels in one direction. In alternative embodiments, other types of rings, such as a BLSR (bi-directional line switched ring) in which communication between nodes on the ring travels in both clockwise and counterclockwise directions, are used.**

Ex.1005, 8:19-36.

Referring to FIG. 9, once each node is able to identify the dynamic frame in STS frames arriving at that node, the nodes begin configuration step 720. During this step, the nodes exchange messages over OAM&P channel 528 and control channels 524. First, each node 420 determines its node ID, which is a unique numerical identifier on SONET ring 410 for that node (step 910). At the initiation of this step, the arbiter node has already been selected, for instance during or prior to provisioning of the ring. The node ID is assigned in sequence beginning at 1 for the arbiter node, and incremented in sequence for nodes as the ring is traversed in a predetermined direction. In this embodiment, the arbiter node sends a ID assignment message over OAM&P channel 528 to its neighboring node. The neighboring node receives the message and


assigns its node ID. That neighboring node then increments the ID in the message and forwards the message to its neighbor. This continues until the entire ring is traversed and the arbiter node receives the ID assignment message from the last node. In this embodiment, SONET ring 410 is traversed in the direction of the working channels. **In alternative embodiments which make use of a bi-directional ring**, by convention, the ring is traversed in the clockwise direction. Alternative embodiments may use other approaches to assigning unique identifiers to the nodes. For instance, the IDs may be assigned prior to or during the provisioning phase.

Ex.1005, 11:25-50.

A variety of types of SONET rings are used. In a UPSR (unidirectional path switched ring), all communication between nodes travels in one direction around the ring. In another type of ring, **a BLSR (bi-directional line switched ring), communication travels between nodes on the ring in both directions, with two unidirectional optical links joining each pair of neighboring nodes.**

Ex.1005, 2:10-16.

85. Applying Kovvali's teachings of a BLSR (bi-directional line switched ring), a POSITA would understand that Figure 4 could be modified to show each pair of nodes connected with two unidirectional optical links traveling in opposite directions:



Ex.1005, Fig. 4 (modified and annotated)

86. As shown in modified Figure 4, above, the resulting BLSR network includes two paths for transmitting data between nodes. The first path includes the sequence of links between any two nodes in the clockwise direction (orange). The second path includes the sequence of links between any two nodes in the counterclockwise direction (blue).

87. Therefore, Kovvali's disclosure of a node receiving, processing, and multiplexing data from local dynamic channels onto a ring network that is connected to other nodes with two unidirectional optical links traveling in opposite directions,

renders obvious “*receiving a request at the first node to use a portion of the resources so as to carry the data flow between the first and second nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links.*”

- d. **[1.3] and in response to the request: if the resources already allocated to the first node on one of the paths are sufficient, directing the data flow from the first node to the second node over the one of the paths without requesting additional resources;**

88. Kovvali renders this limitation obvious.

89. First, as discussed in [1.1] Kovvali describes that an arbiter node assigns “dynamic channels” to particular nodes in the network with initial data rate allocations during a configuration step:

In a next, optional, step of configuration, **the arbiter node assigns a number of dynamic channels 436 to couple particular nodes (step 930). This involves the arbiter node notifying each node of the dynamic channel which it terminates, and providing the initial data rate allocations for each of the dynamic channels.** The allocations are in units of columns of the dynamic frame. Note that certain dynamic channels initially may be allocated no capacity (zero columns). The client nodes **420** compute the starting column for the dynamic channels by determining the end of TDM section **530**, and then summing the allocated number of columns for all other

dynamic channels that have lower identifying numbers to determine the start of that channel's allocation relative to the end of the TDM section.

Ex.1005, 12:1-14.

90. As discussed in [1.1], the arbiter node may also assign dynamic channels and adjust their allocated rates upon request from respective client nodes:

In order to dynamically manage the TDM and dynamic channels, one node 420 is assigned to be the arbiter node at the time the ring is provisioned. The remaining nodes that make use of the TDM and dynamic channels are client nodes. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. Requests to assign dynamic channels for communication between particular nodes are sent from a client nodes to the arbiter. **The arbiter responds to the requests by assigning dynamic channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.**

Ex.1005, 9:1-11.

91. Second, as discussed in [1.2], Kovvali teaches that each node that terminates (e.g. that adds or drops) a dynamic channel multiplexes outbound data that was received locally for transmission onto the dynamic channel. Ex.1005, 12:20-27, 13:63-14:37, and 1:48-2:9.

92. In particular, Kovvali teaches that after receiving an allocation for a dynamic channel, a node determines the portion of the frame that is allocated to the dynamic channel. The node adds (or multiplexes) data for transmission into the portion of the frame that is already allocated to the dynamic channel, and then transmits the frame onto the network:

The method includes determining at a first node an allocation of a portion of the second part of the frames for each of the communication channels and then broadcasting the allocation from the first node to the other nodes including passing information in the first part of one or more frames. The method also includes, at each of the second and the third nodes, receiving control information in the first part of a sequence of frames and processing the received control information to determine the portion of a frame that is allocated to the first communication channel. **The second node receives a first frame from the network, adds data for transmission to the third node into the portion of the first frame that is allocated to the first communication channel and transmits the first frame onto the network.** The third node receives the first frame from the network, and extracts the data from the portion of the first frame that is allocated to the first communication channel.

Ex.1005, 5:66-6:14.

93. A POSITA would have understood that adding (or multiplexing) data for transmission into the portion of a frame that is already allocated to a dynamic channel, renders obvious “*if the resources already allocated to the first node on one of the paths is sufficient, directing the data flow to the second node . . . without requesting additional resources*” because the data is included in only the portion of the frame that is already allocated to the dynamic channel without requesting additional resources, such as bandwidth. As discussed further at [1.4] below, Kovvali explains that a node may take action (e.g. request additional bandwidth) if the data to be transmitted onto a dynamic channel exceeds the allocated bandwidth; it would have been obvious to a POSITA, therefore, that no action would be needed or taken (e.g. the node would not request additional bandwidth) when the existing allocation is sufficient (e.g., the data to be transmitted fits within the number of columns in the transmission frame already allocated, as discussed at [1.1]). As discussed at [1.3], Kovvali’s arbiter makes “initial data rate allocations” before any node issues a bandwidth request. Ex.1005, 12:1-6. Thus, it would have been obvious to a POSITA for Kovvali’s nodes to have “*resources already allocated*” and available for use without requiring a request for additional resources.

94. Third, as discussed in [1.2] the BSLR network taught by Kovvali includes two paths for transmitting between nodes: a clockwise path and a counterclockwise path. Further, as discussed in [1.1] dynamic channels are used for

communication between two or more nodes. Thus, a POSITA would understand that Kovvali's disclosure of transmitting a frame "onto the network" after multiplexing data for a dynamic channel into a frame renders obvious "*directing the data flow from the first node to the second node over the one of the paths without requesting additional resources,*" because the frame transmitted on the network would use either the clockwise or the counterclockwise path.

95. Therefore, Kovvali's disclosure of a node receiving outbound data that was received locally for transmission onto the dynamic channel, multiplexing the outbound data for transmission into the portion of the frame that is already allocated to the dynamic channel without requesting additional resources, and transmitting the frame onto a BSLR network, renders obvious "*and in response to the request: if the resources already allocated to the first node on one of the paths are sufficient, directing the data flow from the first node to the second node over the one of the paths without requesting additional resources.*"

- e. **[1.4] *and if the resources already allocated to the first node are insufficient, performing the following steps:***

96. Kovvali explains that data transmitted on the dynamic channels may be "bursty" in nature, and thus some steps are taken when a large burst of data is to be transmitted. Ex.1005, 4:30-34. It would have been obvious to a POSITA that for such bursty traffic, the data to be transmitted would occasionally exceed the bandwidth allocated for it (and thus, "*the resources already allocated to the first*

node are insufficient”). As explained further below, the prior art renders obvious performing the recited steps. *See* [1.5], [1.6], [1.7] and [1.8] *infra*.

f. **[1.5] *requesting an increase in an allocation to the first node of the requested resources;***

97. Kovvali renders this limitation obvious.

98. **First**, Kovvali teaches that data traffic may be “bursty” in nature and exhibit time-varying rates:

Traditional SONET equipment has several desirable features, including interoperability and speedy fault recovery. The SONET framing structure, and the defined transport and path overhead bytes provide the basis for interoperability, and provide mechanisms for fault management and protection switching. However, the SONET framing structure, virtual tributary approach and payload mapping are not well suited to carrying data at rates that do not match the SONET signal hierarchy, or that exhibit **bursty or time-varying rates**.

Ex.1005, 4:25-34.

In addition to passing data over standard rate virtual tributaries 432 between nodes 420, in accordance with the invention, SONET ring 410 passes data over TDM (isochronous time division multiplexed) channels 434 and dynamic channels 436 between nodes. TDM channels 434 and dynamic channels 436 are not limited to rates in the SONET signal hierarchy shown in FIG. 2. For example,

TDM channels 434 may have rates lower than the 1.5 Mb/s of a VT1.5 virtual tributary. In particular, a TDM channel may use as little as a single STS-1 column yielding a data rate of 576 Kb/s. Dynamic channels 436 can be repeatedly assigned to communication between different nodes after SONET ring 410 is provisioned, and furthermore, the data rate assigned to any particular dynamic channel 436 can be dynamically changed, for example to **accommodate time-varying data rate requirements of bursty communication between nodes.**

Ex.1005, 8:37-53.

The LAN backbone and access technologies that are used to pass data to and from the SONET infrastructure generally support different data rates than those of the PDH digital signal hierarchy described above. The technologies in use today and those emerging for the future include Ethernet (10 Mb/s), Fast Ethernet (100 Mb/s), Gigabit Ethernet (1 Gb/s), Frame Relay, ATM, DSL and Cable. The variety of rates and formats of data streams that are mapped to legacy Virtual Tributary rates can result in gross inefficiencies of “stranded” bandwidth that is left unused in a SONET Frame. In addition, **the existing synchronous digital signal hierarchy was designed for static allocations of data rates over data streams between nodes and hence is not particularly flexible to deal with the bursty nature of data traffic.**

For example, a SONET ring may be provisioned according to the maximum data rates guaranteed between particular nodes, making it difficult to make use of unused capacity between one pair of nodes for data communication among another pair.

Ex.1005, 3:21-39.

99. Second, to accommodate the time-varying data rate requirements of bursty communication between nodes, Kovvali describes the data rate assigned to any particular dynamic channel can be dynamically changed:

In addition to passing data over standard rate virtual tributaries 432 between nodes 420, in accordance with the invention, SONET ring 410 passes data over TDM (isochronous time division multiplexed) channels 434 and dynamic channels 436 between nodes. TDM channels 434 and dynamic channels 436 are not limited to rates in the SONET signal hierarchy shown in FIG. 2. For example, TDM channels 434 may have rates lower than the 1.5 Mb/s of a VT1.5 virtual tributary. In particular, a TDM channel may use as little as a single STS-1 column yielding a data rate of 576 Kb/s. Dynamic channels 436 can be repeatedly assigned to communication between different nodes after SONET ring 410 is provisioned, and furthermore, **the data rate assigned to any particular dynamic channel 436 can be dynamically changed, for example to**

**accommodate time-varying data rate requirements of
bursty communication between nodes.**

Ex.1005, 8:37-53.


100. Third, Kovvali teaches that to change a data rate allocated to a dynamic channel assigned to it, a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth. Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38.

Referring still to FIG. 5A, **each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel from the arbiter. GNT527 is set by the arbiter node to grant bandwidth to the requested channel.** The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel. In this embodiment, in which control channel 524 is 8 bits, REQ 526 is 3 bits and GNT 527 is 5 bits, the values of REQ and GNT are encoded as follows:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

101. Figure 5A depicts an embodiment of the dynamic frame that includes a request section (REQ) and a grant section (GNT) as part of the dynamic frame overhead section of the frame:



Ex.1005, Fig. 5A (annotated)

[...]

Referring still to FIG. 5A, dynamic frame 510 includes three sections. A dynamic frame overhead section 520, in

this embodiment made up of three columns, is used to transport dynamic bandwidth management information between nodes 420 and the arbiter node. For instance, **requests for bandwidth from client nodes and responses from the arbiter node pass in dynamic frame overhead section 520.** In alternative embodiments, the byte positions of these sections and fields within the sections may differ.

Ex.1005, 9:48-56.


102. Thus, if the bandwidth of a dynamic channel already assigned to a node becomes insufficient because the data rate requirement on the dynamic channel has increased (e.g. it has experienced a burst of data to be transmitted), it would have been obvious that a client node may send a request to the arbiter node for additional bandwidth so that the increased data rate requirement may be accommodated.

103. Therefore, Kovvali's description of a client node sending a request to an arbiter node for additional bandwidth to accommodate a dynamic channel with an increased data rate requirement, renders obvious "*if the resources already allocated to the first node are insufficient . . . requesting an increase in allocation to the first node of the requested resources.*"

- g. ***[1.6] determining, for the links comprised by each of the paths, respective levels of use of the requested resources due to communications in progress over the network;***

104. The combination of Kovvali and Kalman renders this limitation obvious.

105. First, as discussed in [1.2], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. As shown in Figure 4, below, the resulting BLSR network includes two paths for transmitting data between nodes. The first path includes the sequence of links between any two nodes in the clockwise direction (blue). The second path includes the sequence of links between any two nodes in the counterclockwise direction (orange). It would have been obvious to a POSITA that when the arbiter receives a request to increase the bandwidth allocated for communications from a first node to a second node, the arbiter would have two path options to choose from for potentially accommodating that bandwidth increase request.



Ex.1005, Fig. 4 (modified and annotated)

106. Second, Kalman discloses a technique for selecting a “preferred direction” on a ring for transmitting data between nodes. Ex.1006, 7:51-53. A preferred direction for traffic between two nodes is selected based on the direction with the lowest cost:

For each destination node j , a cost is assigned to each output direction (0 and 1) from node 0 on the ring. **A preferred direction for traffic from nodes 0 to j is selected based on the direction with the lowest cost.** For simplicity, the mechanism for reassigning costs to the path

to each destination node for each output direction from node 0 operates with a constant number of operations, irrespective of the current condition of the ring. (The mechanism may be further optimized to always use the minimum possible number of operations, but this will add complexity to the algorithm without significantly increasing overall protection switching speed.) The mechanism for reassigning an output direction to traffic packets destined for a given node based on the path cost minimizes the time required to complete this reassignment.

Ex.1006, 7:50-63.

107. In particular, Kalman describes a network that includes two rings, where a first ring transmits data in clockwise direction, and the second ring transmits data in a counterclockwise direction. Ex.1006, 2:56-67. During operation, data between nodes flows on the ring that provides the lowest-cost path to the destination node:

A network protection and restoration technique is described that efficiently utilizes the total bandwidth in the network to overcome the drawbacks of the previously described networks, that is not linked to a specific transport protocol such as SONET, and that is designed to meet the Telcordia 50 millisecond switching requirement. **The disclosed network includes two rings, wherein a**

first ring transmits data in a “clockwise” direction (or direction 0), and the other ring transmits data in a “counterclockwise” direction (or direction 1 opposite to direction 0). Additional rings may also be used. The traffic is removed from the ring by the destination node.

During normal operations (i.e., all spans operational and undegraded), data between nodes flows on the ring that provides the lowest-cost path to the destination node. If traffic usage is uniformly distributed throughout the network, the lowest-cost path is typically the minimum number of hops to the destination node. Thus, both rings are fully utilized during normal operations. Each node determines the lowest-cost path from it to every other node on the ring. To do this, each node must know the network topology.

Ex.1006, 2:56-67.

The purpose of the invention described herein is to achieve fast protection in a ring network while providing for efficient network capacity utilization. **Certain aspects of the preferred embodiment are:**

- a. Transmission of a given packet between two nodes in only one direction around the ring (rather than in both directions as is done in SONET UPSR).**

- b. Differentiation between “protected” and “unprotected” traffic classes.
- c. A fast topology communication mechanism to rapidly communicate information about a span break to all nodes in the ring.
- d. A fast re-routing/routing table update mechanism to re-route paths impacted by a span break the other direction around the ring.
- e. An optional interim wrapping mechanism that may be used to further increase protection switching speed.

Ex.1006, 4:20-37.

108. For example, Kalman depicts in Figure 3 an embodiment where a given packet/flow between two nodes (specifically, from Node 0 to Node 5) is transmitted in only a single direction around the network:




FIG. 3

Ex.1006, Fig. 3 (annotated)

109. In Figure 3, node 0 transmits information to node 5 in only one direction via the counterclockwise path (blue) (i.e. the path traversing node 7 and node 6), and not the clockwise path (orange) (i.e. the path traversing node 1, node 2, node 3, and node 4). Kalman explains that a transmission from node 5 to node 0 would only go via the clockwise path (orange) (i.e. the path traversing node 6 and node 7), and not the counterclockwise path (blue) (i.e. the path traversing node 4, node 3, node 2, and node 1).

A given packet/flow between two nodes is transmitted in only a single direction around the network (even when there is a span fault) and is removed from the ring by the destination node, **as is shown in FIG. 3 where node 0 transmits information to node 5 in only the direction indicated by the thick arrows. A transmission from node 5 to node 0 would only go through nodes 6 and 7 in the opposite direction.** This allows for optimized ring capacity utilization since no capacity is set aside for protection.

Ex.1006, 4:40-48.

110. Kalman discloses that the path selected for transmitting data is the path with the lowest cost, which may be determined in various ways, such as evaluating the length of the transmission path (measured by the number of “hops” between the two nodes). Kalman also explains other ways of evaluating or measuring the “cost” of a potential transmission path:

The least-cost physical route is typically used for protected traffic. This is often the shortest-hop physical route. For example, a transmission from node 0 to node 2 would typically be transmitted via node 1. The shortest-hop physical route corresponds to the least-cost route when traffic conditions throughout the network are relatively uniform. If traffic conditions are not uniform,

the least-cost physical route from node 0 to node 2 can instead be the long path around the ring.

Ex.1006, 4:49-57.

For each destination node *j*, a cost is assigned to each output direction (0 and 1) from node 0 on the ring. **A preferred direction for traffic from nodes 0 to *j* is selected based on the direction with the lowest cost.** For simplicity, the mechanism for reassigning costs to the path to each destination node for each output direction from node 0 operates with a constant number of operations, irrespective of the current condition of the ring. (The mechanism may be further optimized to always use the minimum possible number of operations, but this will add complexity to the algorithm without significantly increasing overall protection switching speed.) The mechanism for reassigning an output direction to traffic packets destined for a given node based on the path cost minimizes the time required to complete this reassignment.

Ex.1006, 7:50-63.

A table is maintained at each node with the columns Destination Node, direction 0 cost, and direction 1 cost. An example is shown as Table 1. **The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node *j* may take into account a variety of factors,** including the number of hops from source to

destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction. **Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.** The example given below assumes that the costs correspond only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

Ex.1006, 7:64-8:16.

The preferred direction is that with the lower cost to reach destination node j. In the event that the costs to reach node j on direction 0 and on direction 1 are equal, than either direction can be selected. (Direction 0 is selected in this example.) The normal operational cost for each physical route (source to destination) is computed from the link status table shown in Table 2.

Ex.1006, 8:34-40.

111. Kalman teaches that one of the variety of factors that may be taken into account when computing cost is “level of traffic congestion”:

A table is maintained at each node with the columns Destination Node, direction 0 cost, and direction 1 cost. An example is shown as Table 1. **The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node j may take into account a variety of factors, including** the number of hops from source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and **the level of traffic congestion in that direction.** Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly. The example given below assumes that the costs correspond only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

Ex.1006, 7:64-8:16.

112. A POSITA would have been familiar that “traffic congestion” was a common way to refer to the level of use of link bandwidth and/or other resources in

a network by the traffic being sent over the network. *See, e.g.*, Ex.1008, 8:34-36 (“A congestion condition occurs when the sources 12, 13, attempt to exceed the finite bandwidth of the resources along the paths 10, 12”); Ex.1009, 1:36-39 (“Congestion is generally defined as a condition during which network performance is degraded due to saturation of network resources such as communication links, processor cycles, memory buffers, etc.”). Thus, a POSITA would have found it obvious for the “level of traffic congestion” disclosed by Kalman to include the level of use of link bandwidth and/or other resources by the dynamic frame resources described in Kovvali.

113. Third, Kalman discloses that the lowest cost route (e.g. the preferred direction) may be identified by maintaining a table indicating the cost of transmitting traffic from a source node to a destination node in both output directions:

**A table is maintained at each node with the columns
Destination Node, direction 0 cost, and direction 1 cost.**

An example is shown as Table 1. The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node j may take into account a variety of factors, including the number of hops from source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction. Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly. The example

given below assumes that the costs correspond only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

Ex.1006, 7:64-8:16.

114. Kalman's Table 1 depicts one example of such a table, showing the cost assigned in each output direction (0 and 1) (i.e. clockwise and counterclockwise, respectively) from source node 0 to each destination node j :

A table is maintained at each node with the columns Destination Node, direction 0 cost, and direction 1 cost. **An example is shown as Table 1. The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node j** may take into account a variety of factors, including the number of hops from source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction. Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly. The example given below assumes that the costs correspond

only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

For a representative ring with a total of 8 nodes (in clockwise order 0, 1, 2, 3, table's normal operational setting at node 0 is:

TABLE 1

<u>Preferred direction table at node 0</u>			
Destination Node	Direction 0 cost	Direction 1 cost	Preferred Direction
1	1	7	0
2	2	6	0
3	3	5	0
4	4	4	0
5	5	3	1
6	6	2	1
7	7	1	1

Ex.1006, 7:64-8:33.

115. Kalman describes that based on the cost associated with sending traffic in a particular direction, the preferred output direction may be selected, such as the direction with the lower cost to reach the destination node:

A table is maintained at each node with the columns Destination Node, direction 0 cost, and direction 1 cost. An example is shown as Table 1. The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node j may take into account a variety of factors, including the number of hops from source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction. **Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.** The example given below assumes that the costs correspond only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

Ex.1006, 7:64-8:16.

The preferred direction is that with the lower cost to reach destination node j. In the event that the costs to reach node j on direction 0 and on direction 1 are equal, than either direction can be selected. (Direction 0 is selected in this example.) The normal operational cost for

each physical route (source to destination) is computed
from the link status table shown in Table 2.

Ex.1006, 8:34-40.

116. As described in more detail in Section IX.C, a POSITA would have been motivated to combine the teachings of Kovvali and Kalman, and doing so would have been within the POSITA's skillset. A POSITA would have appreciated that by having Kovvali's arbiter direct data transmission toward the least-occupied or less congested portion of the ring network (e.g., to the clockwise or counterclockwise portion that has more unallocated bandwidth), more efficient utilization of the available network capacity would be achieved. Alternatively, a POSITA would have appreciated that having Kovvali's arbiter direct data transmission toward the most-occupied or more congested portion of the ring network (e.g., to the clockwise or counterclockwise portion that has more allocated bandwidth) would leave a larger contiguous block of unallocated space on the other portion of the network (e.g., in the opposite direction), thus making it possible to accommodate a larger bandwidth request that may arrive in the future. While either approach would have its respective benefits, consistently applying either strategy would depend upon Kovvali's arbiter node maintaining information about the current amount of bandwidth allocations throughout the network (*"determining, for*

the links comprised by each of the paths, respective levels of use of the requested resources”).

117. Therefore, the combination of Kovvali and Kalman render obvious “*determining, for the links comprised by each of the paths, respective levels of use of the requested resources due to communications in progress over the network.*”

- h. [1.7] *selecting which of the paths is to carry the data flow responsive to the determined levels of use of the requested resources on the links comprised in each of the paths;***

118. See analysis at [1.6]. As discussed in [1.6], the combination of Kovvali and Kalman teach selecting the lowest cost path that accounts for the “level of traffic congestion” on the links of the path.

119. Therefore, the combination of Kovvali and Kalman render obvious “*selecting which of the paths is to carry the data flow responsive to the determined levels of use of the requested resources on the links comprised in each of the paths.*”

- i. [1.8] *and increasing the allocation of the resources to the first node on the selected path by a predetermined quantum.***

120. Kovvali renders this limitation obvious.

121. As discussed in [1.5], Kovvali teaches that a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth.

Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38.

122. The values in REQ and GNT may be encoded to reduce the number of bits needed for the control channel according to a predetermined amount of bandwidth (e.g., a specified number of columns in the dynamic frame structure) requested:

Referring still to FIG. 5A, each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel from the arbiter. GNT 527 is set by the arbiter node to grant bandwidth to the requested channel. **The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel.** In this embodiment, in which control channel 524 is 8 bits, REQ 526 is 3 bits and GNT 527 is 5 bits, **the values of REQ and GNT are encoded as follows:**

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

123. A POSITA would have understood that when the requested bandwidth is granted by the arbiter, the bandwidth (e.g., the number of columns) allocated to the client node is increased (or decreased) by the requested amount. For example, the existing allocation may have been 5 columns. If the value of REQ sent by the client node is “100” (corresponding to a predetermined request for 1 additional column of bandwidth), and the value of GNT returned by the arbiter is “00110” (corresponding to a total allocation of 6 columns), the additional resources allocated to the client node for sending the dynamic channel are increased by 1 column of additional bandwidth. Because the increased allocation was equal to the amount of increase earlier requested by the node, the increased allocation would have been “*by a predetermined quantum*” as claimed. It would have been obvious to a POSITA for

Kovvali's arbiter to respond to an allocation request by adjusting the allocation by the amount requested when the network had sufficient capacity to accommodate the request. Again, in such instances, the allocation adjustment would be "*by a predetermined quantum*" (e.g., the requested amount) as claimed. More generally, Kovvali's table of potential REQ and GNT values shows that all bandwidth requests and grants are made in terms of a number of dynamic frame columns, and thus, the allocations are for predetermined, quantized amounts of bandwidth. That is, Kovvali's signaling mechanism for bandwidth requests and grants does not allow for allocations to be made on a per-bit or per-byte basis; instead, as discussed at [1.1], requests and allocations are made in 72-bit increments (corresponding to the number of bits in one dynamic frame column).

124. Therefore, Kovvali renders obvious "*increasing the allocation of the resources to the first node on the selected path by a predetermined quantum.*"

2. Claim 2

- a. [2.0] *A method according to claim 1, wherein the network comprises an Internet Protocol (IP) network.***

125. Kovvali renders this limitation obvious.

126. Kovvali describes that one of the predominant network protocols for carrying data on ring networks includes the IP (Internet Protocol):

Some approaches to the data transport problems are based
on the **predominant networking protocols that are used**

to carry data on networks today, namely IP (Internet Protocol) and ATM (Asynchronous Transfer Mode). In particular, two types of approaches to passing data over SONET infrastructure have been proposed. A common theme in the two solutions, which are outlined below, **is the retention of the ring topology of interconnection of the SONET ring,** and support for redundancy and protection switching in the infrastructure.

Ex.1005, 3:40-49.

127. In particular, Kovvali discloses that one approach to passing data over SONET infrastructure is called Dynamic Packet Transfer (DPT) that uses the IP-based routers to transport IP datagrams between nodes in a ring:

A second approach to passing data over SONET infrastructure is called Dynamic Packet Transfer (DPT). This approach uses IP as the basic multiplexing technology instead of ATM. Like the ATM VP Ring approach, **the DTP approach assumes replacement of the existing transport infrastructure at each SONET node with IP-based routers. DPT defines a layer 2 protocol called Spatial Reuse Protocol (SRP), which is used to transport IP datagrams between nodes in a ring.** SRP allows reuse of ring bandwidth across different sections of a ring of nodes. A supplementary algorithm called SRP Fairness algorithm (SRP-fa) was also defined

to achieve fair sharing of bandwidth, between various nodes in a ring topology.

Ex.1005, 4:9-21.

128. It would have been obvious to a POSITA for Kovvali's ring network to employ Internet Protocol, for example by using the Dynamic Packet Transfer capability discussed in Kovvali. A POSITA would have understood that a ring network that uses the IP networking protocol renders obvious an "*Internet Protocol (IP) network.*"

129. Employing Internet Protocol (IP) for data transmissions on Kovvali's ring network also would have been obvious because Kovvali specifically focuses on handling "bursty" traffic, and Internet Protocol (IP) was well-known as a "bursty" data communications protocol. Ex.1005, 8:47-53; Ex.1015, 12:50-51 ("inherently bursty traffic (such as IP—Internet Protocol—traffic)"); Ex.1016, 1:35-36 ("packet-switched (that is, bursty) IP traffic"); Ex.1017, 2:16-19 (describing "the typically 'bursty' nature of transmissions between nodes of" a "TCP/IP network").

130. Therefore, Kovvali renders obvious "*wherein the network comprises an Internet Protocol (IP) network.*"


3. Claim 3

- a. **[3.0]** *A method according to claim 1, wherein the nodes are interconnected in a ring, and wherein the plurality of paths comprises a first path traversing the ring in one direction and a second path traversing the ring in the opposite direction, and wherein selecting which of the*

paths is to carry the data flow comprises selecting the first or the second path.

131. The combination of Kovvali and Kalman render this limitation obvious.

132. First, as discussed in [1.2], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. As shown in Figure 4, below, the resulting BLSR network includes two paths for transmitting data between nodes. The “*first path*” includes the sequence of links between any two nodes in the clockwise direction (orange). The “*second path*” includes the sequence of links between any two nodes in the counterclockwise direction (blue). The two paths are in “*opposite direction[s]*.”



Kovvali at Fig. 4 (modified and annotated)

133. Thus, Kovvali renders obvious “wherein the nodes are interconnected in a ring, and wherein the plurality of paths comprises a first path traversing the ring in one direction and a second path traversing the ring in the opposite direction.”

134. **Second**, as discussed in [1.6], Kalman discloses the first path or the second path is selected as the preferred direction for transmitting data based on which of the two paths has the lowest cost:

The least-cost physical route is typically used for protected traffic. This is often the shortest-hop physical route. For example, a transmission from node 0 to node 2

would typically be transmitted via node 1. The shortest-hop physical route corresponds to the least-cost route when traffic conditions throughout the network are relatively uniform. If traffic conditions are not uniform, the least-cost physical route from node 0 to node 2 can instead be the long path around the ring.

Ex.1006, 4:49-57.

For each destination node *j*, a cost is assigned to each output direction (0 and 1) from node 0 on the ring. **A preferred direction for traffic from nodes 0 to *j* is selected based on the direction with the lowest cost.** For simplicity, the mechanism for reassigning costs to the path to each destination node for each output direction from node 0 operates with a constant number of operations, irrespective of the current condition of the ring. (The mechanism may be further optimized to always use the minimum possible number of operations, but this will add complexity to the algorithm without significantly increasing overall protection switching speed.) The mechanism for reassigning an output direction to traffic packets destined for a given node based on the path cost minimizes the time required to complete this reassignment.

Ex.1006, 7:50-63.

A table is maintained at each node with the columns Destination Node, direction 0 cost, and direction 1 cost.

An example is shown as Table 1. The computation of the cost on a direction from node 0 (assuming node 0 as the source) to node j may take into account a variety of factors, including the number of hops from source to destination in that direction, the cumulative normalized bit error rate from source to destination in that direction, and the level of traffic congestion in that direction. **Based on these costs, the preferred output direction for traffic from the source to any destination can be selected directly.** The example given below assumes that the costs correspond only to the normalized bit error rate from source to destination in each direction. The cost on a given link is set to 1 if the measured bit error rate is lower than the operational bit error rate threshold. Conveniently, if all links are fully operational, the cumulative cost from node 0 to node j will be equal to the number of hops from node 0 to node j if there is no traffic congestion. Traffic congestion is not taken into account in this example.

Ex.1006, 7:64-8:16.

The preferred direction is that with the lower cost to reach destination node j. In the event that the costs to reach node j on direction 0 and on direction 1 are equal, than either direction can be selected. (Direction 0 is selected in this example.) The normal operational cost for each physical route (source to destination) is computed from the link status table shown in Table 2.

Ex.1006, 8:34-40.

For a representative ring with a total of 8 nodes (in clockwise order 0, 1, 2, 3, table's normal operational setting at node 0 is:

TABLE 1

Preferred direction table at node 0			
Destination Node	Direction 0 cost	Direction 1 cost	Preferred Direction
1	1	7	0
2	2	6	0
3	3	5	0
4	4	4	0
5	5	3	1
6	6	2	1
7	7	1	1

Ex.1006, 8:17-33.

135. As discussed in [1.6] and in Section VIII.C, a POSITA would have been motivated to combine the teachings of Kovvali and Kalman, and doing so would have been within the POSITA's skillset.


136. Therefore, Kovvali and Kalman render obvious "*wherein the nodes are interconnected in a ring, and wherein the plurality of paths comprises a first path traversing the ring in one direction and a second path traversing the ring in the opposite direction, and wherein selecting which of the paths is to carry the data flow comprises selecting the first or the second path.*"

4. Claim 4

- a. **[4.0]** *A method according to claim 3, wherein selecting the first or the second path comprises conveying the data flow over an inner or outer data link ring within the network provided by a Spatial Reuse Protocol (SRP).*

137. Kovvali and Kalman each render this limitation obvious.

138. First, as discussed in [1.2], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. As shown in Figure 4, below, the resulting BLSR network includes two paths for transmitting data between nodes. The “*first path*” includes the sequence of links between any two nodes in the clockwise direction (orange) (“*outer data link ring*”). The “*second path*” includes the sequence of links between any two nodes in the counterclockwise direction (blue) (“*inner . . . data link ring*”).



Ex.1005, Fig. 4 (modified and annotated)

139. As discussed in [2.0], Kovvali discloses that one approach to passing data over SONET infrastructure is called Dynamic Packet Transfer (DPT) that uses the IP-based routers to transport data in IP datagrams between nodes. The “DPT defines a layer 2 protocol called Spatial Reuse Protocol (SRP),” that allows reuse of ring bandwidth across different sections of a ring of nodes:

A second approach to passing data over SONET infrastructure is called Dynamic Packet Transfer (DPT). This approach uses IP as the basic multiplexing technology instead of ATM. Like the ATM VP Ring

approach, the DTP approach assumes replacement of the existing transport infrastructure at each SONET node with IP-based routers. DPT defines a layer 2 protocol called Spatial Reuse Protocol (SRP), which is used to transport IP datagrams between nodes in a ring. SRP allows reuse of ring bandwidth across different sections of a ring of nodes. A supplementary algorithm called SRP Fairness algorithm (SRP-fa) was also defined to achieve fair sharing of bandwidth, between various nodes in a ring topology.

Ex.1005, 4:9-21.

140. It would have been obvious to a POSITA for Kovvali's ring network to transport IP datagrams using the Spatial Reuse Protocol that Kovvali describes. Thus, Kovvali renders obvious "*conveying data flow . . . within the network provided by a Spatial Reuse Protocol (SRP).*"

141. Kalman also teaches the use of "spatial reuse" when delivering data flow between two nodes in a network. In particular, Kalman discloses that when data is transmitted along a preferred path in a network, the data is removed from the ring by the destination node, which ensures that traffic does not use more capacity than is necessary to deliver it to the destination node:

A given packet/flow between two nodes is transmitted in only a single direction around the network (even when there is a span fault) and is removed from the

ring by the destination node, as is shown in FIG. 3 where node 0 transmits information to node 5 in only the direction indicated by the thick arrows. A transmission from node 5 to node 0 would only go through nodes 6 and 7 in the opposite direction. This allows for optimized ring capacity utilization since no capacity is set aside for protection.

The least-cost physical route is typically used for protected traffic. This is often the shortest-hop physical route. For example, a transmission from node 0 to node 2 would typically be transmitted via node 1. The shortest-hop physical route corresponds to the least-cost route when traffic conditions throughout the network are relatively uniform. If traffic conditions are not uniform, the least-cost physical route from node 0 to node 2 can instead be the long path around the ring.

The removal of packets from the ring by the destination node ensures that traffic does not use more capacity than is necessary to deliver it to the destination node, thus enabling increased ring capacity through spatial reuse of capacity. An example of spatial reuse is the following. If 20% of span capacity is used up for traffic flowing from node 0 to node 2 via node 1, then the removal of this traffic from the ring at node 2 means that the 20% of span capacity is now available for any traffic flowing on any of the other

spans in the ring (between nodes 2 and 3, nodes 3 and 4, etc.).

Ex.1006, 4:40-67.

142. A POSITA would thus understand that a ring network that enables increased ring capacity through “spatial reuse” when passing data between nodes renders obvious “*conveying data flow . . . within the network provided by a Spatial Reuse Protocol (SRP).*”

143. A POSITA would have had an expectation of success in implementing a *Spatial Reuse Protocol (SRP)* in the context of the Kovvali-Kalman combination. Notably, such protocols had already been developed and defined by standards bodies such as the Internet Engineering Task Force, which had issued RFC2892 entitled “The Cisco SRP MAC Layer Protocol.” Ex.1018, 1. RFC2892 defines the “‘Spatial Reuse Protocol’ (SRP) for use with ring based” networks like those in Kovvali and Kalman.

144. It also would have been obvious to a POSITA to consider Kovvali’s clockwise and counterclockwise ring paths to correspond to “*inner*” and “*outer*” rings, or vice versa. The terms “inner” and “outer” ring were simply used by network engineers to give names to the two separate paths around a bidirectional ring network; the terms were used somewhat like “clockwise” and “counterclockwise” to distinguish between the two paths. As the ’279 patent confirms, labeling the paths

around a ring network as “inner” and “outer” is arbitrary. Ex.1001, 1:27-28; *see also id.* 2:23-25 (“a clockwise or a counterclockwise direction, or in other words,... over an inner or an outer ring”).

145. Therefore, Kovvali and Kalman each render obvious “*wherein selecting the first or the second path comprises conveying the data flow over an inner or outer data link ring within the network provided by a Spatial Reuse Protocol (SRP).*”

5. Claim 5

- a. [5.0] *A method according to claim 1, wherein the one or more resources comprise a link bandwidth.***

146. Kovvali renders this limitation obvious.

147. First, as discussed in [1.1], Kovvali teaches that the arbiter node is responsible for assigning dynamic channels to “couple” particular nodes in the network ring. Ex.1005, 12:1-14, 9:1-11. Each dynamic channel may be assigned bandwidth dynamically by a “centralized bandwidth manager which serves as an arbiter for requests for bandwidth”:

In a general aspect, the invention provides a method and apparatus for using SONET technology to efficiently carry both data and voice traffic by Add/Drop Multiplexing of both isochronous traffic, such as voice and video, and non-isochronous traffic, such as data, in a SONET/SDH fiber-optic ring topology. A portion of a SONET/SDH frame

structure (STS-1, STM-1 etc.) is shared by a number of nodes connected to an optical SONET/SDH ring (either directly or sitting behind a SONET/ADM). The inherent column structure of the SONET/SDH frame is used to define dynamic channels (flexible bandwidth channels) between the nodes. **Each dynamic channel can be assigned a priority and optionally a minimum bandwidth. The channel to column mapping is varied dynamically for data traffic by a centralized bandwidth manager which serves as an arbiter for requests for bandwidth.** Both isochronous traffic and best effort data traffic may be carried in an STS-1 frame thereby allowing over-subscribing and statistical multiplexing of data channels over the unassigned portion of the STS-1 frame.

Ex.1005, 4:35-54.

Referring still to FIG. 5A, dynamic frame 510 includes three sections. A dynamic frame overhead section 520, in this embodiment made up of three columns, is used to transport dynamic bandwidth management information between nodes 420 and the arbiter node. For instance, **requests for bandwidth from client nodes and responses from the arbiter node** pass in dynamic frame overhead section 520. In alternative embodiments, the byte positions of these sections and fields within the sections may differ.

Ex.1005, 9:48-56.

148. Second, as discussed in [1.0], the nodes are arranged in a ring structure and are “coupled over optical links”:

Referring to FIG. 4, in one embodiment, **a number of SONET nodes 420** (labeled A-D) are **coupled over optical links 422** in a SONET ring 410. Optical links 422 carry signals according to a standard framing format, in this embodiment according to STS-1 framing using OC-1 signaling (see STS-1 block 250 and OC-1 block 280 shown in FIG. 2). In alternative embodiments, optical links 422 may carry signals at a higher standard rate, for example using STS-3 framing with OC-3 signaling, or STS-12 framing with OC-12 signaling. The following description focuses on STS-1 or STS-3 payloads, however, the approach is applicable to all SONET rates. In FIG. 4, SONET ring 410 is a UPSR (unidirectional path switched ring), in which all communication between nodes travels in one direction. In alternative embodiments, other types of rings, such as a BLSR (bi-directional line switched ring) in which communication between nodes on the ring travels in both clockwise and counterclockwise directions, are used.

Ex.1005, 8:19-36.

149. It would have been obvious to a POSITA that allocating bandwidth to a dynamic channel between two nodes effectively allocates that same bandwidth to

the dynamic channel on the corresponding links between the two nodes. This is because—as Kalman confirms—the dynamic channel that couples the nodes also occupies the allocated portion of bandwidth on the link (or links) along the path from the transmitting node to the receiving node. *See* Ex.1006, 4:40-66, Fig.3 (discussing how, with spatial reuse, data transmissions only occupy bandwidth on the links between the transmitter and receiver, leaving the corresponding bandwidth on other links of the ring available for other transmissions).

150. Therefore, Kovvali renders obvious “*wherein the one or more resources comprise a link bandwidth.*”

6. Claim 9

- a. **[9.0]** *A method according to claim 1, wherein selecting which of the paths is to carry the data comprises verifying that a sufficient amount of the requested resources is available to carry the data flow on every one of the links comprised in the selected path.*

151. Kovvali and Kalman render this limitation obvious.

152. As discussed in [7.0] and [8.0], a POSITA would have found it obvious that when the resources requested for transmitting a dynamic channel exceeds the available capacity on any link of a path, the arbiter would either deny the request (e.g. GNT 00000) or return a lower value up to “[a]ll available columns” (e.g. GNT 11111) for the channel in the frame so that the maximum number of columns that can be allocated to the dynamic channel are not exceeded on any of the links in the

path (“*verifying that a sufficient amount of the requested resources is available to carry the data flow on every one of the links comprised in the selected path.*”). In other words, Kovvali and Kalman render obvious for the arbiter to deny a request for bandwidth that cannot be accommodated. In order to deny such requests, it would have been obvious to a POSITA for the arbiter to check whether there is sufficient bandwidth from the source to the destination of the data flow, including checking for sufficient bandwidth on each link that the data flow will traverse. As Kalman teaches, spatial reuse protocol allows a single portion of a dynamic frame to be used for different purposes (different transmissions) on different links around the network ring. *See* Ex.1006, 4:48-67. Thus, it would have been obvious for Kovvali’s arbiter node to maintain the bandwidth allocations it has made for each link of the network, and for the arbiter node to confirm when making a new bandwidth allocation that the relevant links have sufficient capacity.

153. Thus, the combination of Kovvali and Kalman render obvious, “*wherein selecting which of the paths is to carry the data comprises verifying that a sufficient amount of the requested resources is available to carry the data flow on every one of the links comprised in the selected path.*”

7. Claim 10

- a. **[10.0]** *A method according to claim 1, wherein selecting which of the paths is to carry the data flow comprises selecting the one of the paths having the lowest level of a*

predetermined measure of use of the requested resources.

154. The combination of Kovvali and Kalman render this limitation obvious.

155. As discussed in [1.6] and [1.7], the combination of Kovvali and Kalman render obvious a dispatcher that selects the lowest cost path that accounts for a variety of factors including the “level of traffic congestion” (“*predetermined measure of use of the requested resources*”) on the links of the path. It would have been obvious to a POSITA that the level of traffic congestion is a “*predetermined*” measure of use because the level of traffic congestion would need to be measured before it would be used in the cost calculation.

156. A POSITA would have also found it obvious that where the value of all the factors in the lowest cost path calculation are the same except for the “level of traffic congestion,” the path with the lowest “level of traffic congestion” would be selected because it has the lowest cost.

157. For example, Kalman describes in one example that the computation of cost for a particular path may take into account a variety of factors, including the number of hops from source to destination, the cumulative normalized bit error rate, and the level of traffic congestion. Ex.1006, 7:66-8:5. In the exemplary scenario where the number of hops between the source and destination node are the same in both directions (e.g. they are equidistant) and that the cumulative normalized bit error rate is the same, then the level of traffic congestion would be a determining

factor leading to the selection of one of the paths. For example, the path with the lowest level of traffic congestion may be selected because it has the lowest cost.

158. Further, Kalman describes that the computation of cost for a particular path “may” take into account a variety of factors and provides examples where the cost is based on just one factor (e.g. packet loss rate, average packet queuing delay, or variation of packet queuing delay):

3. The method of claim 1 wherein said costs are based upon a **packet loss rate**.

4. The method of claim 1 wherein said costs are based upon an **average packet queuing delay** at the routing switch receiving packets from the link.

5. The method of claim 1 wherein said costs are based upon the **variation of packet queuing delay** at the routing switch receiving packets from the link.

Ex.1006, 7:66-8:1, 16:36-44. A POSITA would have understood that “average packet queuing delay” and “variation of packet queuing delay” are each examples of a level of traffic congestion. For example, when a network is congested, packets may be “dropped” or “delay[ed].” Ex.1009, 5:65-55; Ex.1008, 3:53-57. As a network became busier, more packets would be dropped or delayed. Thus, Kalman’s measures of “packet loss rate” and “queueing delay” would have been measurements of network usage. Accordingly, a POSITA would have found it obvious that the

calculation of cost could also be based on the level of traffic congestion (or another measure of resource use) alone.


159. Therefore, the combination of Kovvali and Kalman render obvious, “*wherein selecting which of the paths is to carry the data flow comprises selecting the one of the paths having the lowest level of a predetermined measure of use of the requested resources.*”

8. Claim 11

- a.** [11.0] *A method according to claim 1, wherein receiving the request comprises choosing a dispatcher within the network to manage allocation of the resources, wherein the dispatcher receives and processes the request.*

160. Kovvali renders this limitation obvious.

161. First, as discussed in [1.1], Kovvali describes that when the network ring is provisioned, one node is assigned as an “arbiter” node, and the remaining nodes are “client” nodes:



Ex.1005, Fig. 4 (annotated)

[...]

In order to dynamically manage the TDM and dynamic channels, **one node 420 is assigned to be the arbiter node** at the time the ring is provisioned. **The remaining nodes** that make use of the TDM and dynamic channels **are client nodes**. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. Requests to assign dynamic channels for communication between particular nodes are sent from a client nodes to the arbiter. The arbiter responds to the requests by assigning dynamic

channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.

Ex.1005, 9:1-11.

162. Accordingly, a POSITA would have found it obvious that the step of “*receiving the request*” includes assigning one of the nodes in the ring as an “arbiter” node.

163. Second, Kovvali describes that the “arbiter” (“*dispatcher*”) may be a centralized bandwidth manager that processes requests for dynamic channel resources:

In a general aspect, the invention provides a method and apparatus for using SONET technology to efficiently carry both data and voice traffic by Add/Drop Multiplexing of both isochronous traffic, such as voice and video, and non-isochronous traffic, such as data, in a SONET/SDH fiber-optic ring topology. A portion of a SONET/SDH frame structure (STS-1, STM-1 etc.) is shared by a number of nodes connected to an optical SONET/SDH ring (either directly or sitting behind a SONET/ADM). The inherent column structure of the SONET/SDH frame is used to define dynamic channels (flexible bandwidth channels) between the nodes. Each dynamic channel can be assigned a priority and optionally a minimum bandwidth. The channel to column mapping is varied dynamically for data traffic by **a centralized bandwidth manager which**

serves as an arbiter for requests for bandwidth. Both isochronous traffic and best effort data traffic may be carried in an STS-1 frame thereby allowing over-subscribing and statistical multiplexing of data channels over the unassigned portion of the STS-1 frame.

Ex.1005, 4:35-54.

164. The arbiter in Kovvali maintains a database of all active dynamic channels, including their terminal nodes and current bandwidth allocation:

The arbiter maintains a database of all active TDM and dynamic channels. In particular, **for each dynamic channel, the arbiter maintains a list of all terminal nodes for that channel, as well as the current allocation for that channel.**

Ex.1005, 13:9-12.

165. Third, as discussed in [1.1] and [1.5], Kovvali describes that the arbiter node receives requests from client nodes to assign dynamic channels and adjust their allocated rates:

In order to dynamically manage the TDM and dynamic channels, one node 420 is assigned to be the arbiter node at the time the ring is provisioned. The remaining nodes that make use of the TDM and dynamic channels are client nodes. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. **Requests to assign dynamic channels for communication between**

particular nodes are sent from a client nodes to the arbiter. The arbiter responds to the requests by assigning dynamic channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.

Ex.1005, 9:1-11.

166. In particular, as described in [1.5] Kovvali teaches that to change a bandwidth allocation to a dynamic channel assigned to it, a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth. Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38.

Referring still to FIG. 5A, **each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel from the arbiter. GNT 527 is set by the arbiter node to grant bandwidth to the requested channel.** The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel. In this embodiment, in which control channel 524 is 8 bits, REQ 526 is 3 bits and GNT

527 is 5 bits, the values of REQ and GNT are encoded as follows:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

Referring still to FIG. 5A, dynamic frame 510 includes three sections. A dynamic frame overhead section 520, in this embodiment made up of three columns, is used to transport dynamic bandwidth management information between nodes 420 and the arbiter node. For instance, **requests for bandwidth from client nodes and responses from the arbiter node pass in dynamic frame overhead section 520**. In alternative embodiments, the byte positions of these sections and fields within the sections may differ.

Ex.1005, 9:48-56.


167. Therefore, Kovvali's disclosure of choosing an arbiter that manages network bandwidth, receives requests for dynamic bandwidth, and responds with bandwidth allocation grants, renders obvious, *"wherein receiving the request comprises choosing a dispatcher within the network to manage allocation of the resources, wherein the dispatcher receives and processes the request."*

9. Claim 12

- a. **[12.0] *A method according to claim 11, wherein choosing the dispatcher comprises choosing one of the nodes to act as the dispatcher.***

168. Kovvali renders this limitation obvious.

169. As discussed in as discussed in [1.1] and [11.0], Kovvali describes that when the network ring is provisioned, one node is assigned as an "arbiter" node, and the remaining nodes are "client" nodes:



Ex.1005, Fig. 4 (annotated)

[...]

In order to dynamically manage the TDM and dynamic channels, **one node 420 is assigned to be the arbiter node** at the time the ring is provisioned. **The remaining nodes** that make use of the TDM and dynamic channels **are client nodes**. In FIG. 4, node A is the arbiter node and the remaining nodes are client nodes. Requests to assign dynamic channels for communication between particular nodes are sent from a client nodes to the arbiter. The arbiter responds to the requests by assigning dynamic

channels to particular nodes, and by adjusting the allocated rates for the dynamic channels.

Ex.1005, 9:1-11.

170. Therefore, Kovvali's disclosure of choosing one of the nodes in the ring to the arbiter node ("*dispatcher*") renders obvious, "*wherein choosing the dispatcher comprises choosing one of the nodes to act as the dispatcher.*"

10. Claim 13

- a. **[13.0]** *A method according to claim 11, wherein requesting the increase in the allocation comprises requesting and receiving the allocation from the dispatcher.*

171. Kovvali renders this limitation obvious.

172. As discussed in [1.5] Kovvali teaches that to change a bandwidth allocation to a dynamic channel assigned to it, a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth. Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38.

Referring still to FIG. 5A, **each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to request additional bandwidth for**

the corresponding bursty channel from the arbiter. GNT527 is set by the arbiter node to grant bandwidth to the requested channel. The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel. In this embodiment, in which control channel 524 is 8 bits, REQ 526 is 3 bits and GNT 527 is 5 bits, the values of REQ and GNT are encoded as follows:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

Referring still to FIG. 5A, dynamic frame 510 includes three sections. A dynamic frame overhead section 520, in this embodiment made up of three columns, is used to transport dynamic bandwidth management information between nodes 420 and the arbiter node. For instance, **requests for bandwidth from client nodes and responses from the arbiter node pass in dynamic frame**

overhead section 520. In alternative embodiments, the byte positions of these sections and fields within the sections may differ.

Ex.1005, 9:48-56.

173. Therefore, Kovvali renders obvious “*wherein requesting the increase in the allocation comprises requesting and receiving the allocation from the dispatcher.*”

11. Claim 14

- a. **[14.0]** *A method according to claim 1, and comprising, upon release of the requested resources: comparing an amount of the resources that have been released to a predetermined threshold; and if the amount is greater than the predetermined threshold, deallocating the predetermined quantum from the first node.*

174. Kovvali renders this limitation obvious.

175. First, Kovvali teaches that a particular dynamic channel may pass data for a set of “sub-channels”:

Each channel is associated with a set of two or more node IDs, corresponding to the nodes which terminate that channel. **The nodes that terminate a particular channel may in fact use that channel to pass data for a set of sub-channels.** Nodes which do not add or drop that channel are in general not aware of the number of sub-channels, or how those sub-channels are multiplexed onto the channel. In this way, the number of channels needed is

related to the number of unique sets of terminal nodes, not to an underlying number of communication sessions passing between nodes.

Ex.1005, 12:66-13:8.

176. Kovvali describes that sub-channels pass separate data streams between terminating nodes of a dynamic channel:

Referring to FIG. 11, the requesting node communicates a request to establish a new communication session to the arbiter over the OAM&P channel (step 1110). The request identifies the terminal nodes of the session and the priority of the session. The arbiter node receives the request and first determines whether the request has the same terminal nodes and priority as an already assigned dynamic channel (step 1120). If it does not, the arbiter node picks the next available channel number that is not currently assigned to a communication session and assigns sub-channel number of 0 to the new stream and communicates the assignment to the requesting node over the OAM&P channel (step 1130). **If the requested terminal nodes and priority are the same as an already assigned dynamic channel, the arbiter node assigns a new sub-channel** number to the new stream and communicates the assignment to the requesting node over the OAM&P channel (step 1140). The sub-channel number is not “visible” on the ring, and is only communicated to the terminating nodes in order for

them to correctly multiplex and demultiplex traffic streams on separate sub-channels on a common dynamic channel.

Ex.1005, 13:30-49.

177. As discussed in [1.5], data traffic may be “bursty” in nature with time-varying data rate requirements. Ex.1005, 4:25-34, 8:37-53. Thus, a POSITA would have found it obvious that due to the bursty nature of data traffic, an existing sub-channel may no longer be needed (e.g. a burst of data packets to transmit on the subchannel have subsided) and may be released:

Referring to FIG. 7, operation of the system proceeds through several steps. First, SONET ring 410 is provisioned using standard SONET provisioning techniques (step 710). At the end of this step, each node 420 is able to identify dynamic frame 520 in the STS frames that it receives. Next, a configuration step is carried out in which the node ids of the nodes which are participating in the communication using the dynamic frames are determined, and other initial configuration of the TDM and dynamic channels takes place (step 720). After the configuration step is completed, nodes can communicate over TDM and dynamic channels that were set up in the configuration step. **In operation, from time to time, nodes may initiate a bandwidth allocation step in which** the bandwidth allocated to an existing channel is

requested to be changed (step 730), or a channel assignment step in which a channel is assigned to communication among particular nodes, or **an existing channel is released if it is no longer needed** (step 740).

Ex.1005, 10:66-11:16.

178. Second, as discussed in [1.5], to accommodate the time-varying data rate requirements of bursty communication between nodes, Kovvali describes the data rate assigned to any particular dynamic channel can be dynamically changed. Ex.1005, 4:47-50. A client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37, 46-55.

179. It would have been obvious to a POSITA that upon release of a sub-channel, the corresponding bandwidth requirement of the dynamic channel would be reduced by the resources allocated to the subchannel. In this scenario a client node would determine whether to request a decrease in the dynamic channel allocation. Doing so would reduce the resources that would be unnecessarily allocated for a sub-channel that no longer needs them.

180. Third, as discussed in [1.1] and [1.8], allocations of bandwidth for dynamic channels are in units of columns. Ex.1005, 12:1-14. A column corresponds to a certain number of rows and bytes, for example 9-rows of 8-bit bytes. Ex.1005, 2:55-65, 2:66-3:5. Thus, as discussed in [1.1], resources are allocated to dynamic channels in columns that each correspond to 72 bits of information.

181. Accordingly, before making a request for a decrease in the dynamic channel allocation, a node would first compare the bits of data that have been released to the number of bits in a column (“*a predetermined threshold*”) to determine whether a decrease of an entire column (or entire columns) is warranted. For example, in an exemplary scenario where the released subchannel carried 100 bits of information, the node would compare this value with the 72 bits of information in a column. Because the 100 bits of information released is greater than the 72 bits of information in a column, a node would request a decrease of an entire column of bandwidth. Doing so ensures that the decrease in the dynamic channel bandwidth allocation is attributed only to the released sub-channel, and thus will not affect other sub-channels that may be using shared columns of the dynamic frame.

182. Fourth, as discussed in [1.8], in response to a client node request (REQ) to deallocate bandwidth, the arbiter responds by setting a grant section (GNT) in the dynamic frame specifying the total amount of bandwidth granted to the dynamic channel. Ex.1005, 10:35-38. A POSITA would have understood that when the request for decreased bandwidth is granted by the arbiter, the bandwidth (e.g., the number of columns) allocated to the client node is decreased by the requested amount.

183. A POSITA would have further understood that the requested amount may also be “*the predetermined quantum*” by which the allocation was previously

increased as discussed in [1.8]. When the bursty nature of data traffic has subsided, the increased allocation of resources would have no longer been needed. In this scenario, the additional resources allocated for sending the bursty data on the dynamic channel may be deallocated.

184. Continuing the example discussed in [1.8], the existing allocation may have been 6 columns. If the value of REQ sent by the client node is 001 (corresponding to a request for a deallocation of 1 column of bandwidth), and the value of GNT returned by the arbiter is 00101 (corresponding to a total allocation of 5 columns), the resource deallocated to the client node for sending the dynamic channel is 1 column.

185. Therefore, Kovvali renders obvious “*comparing an amount of the resources that have been released to a predetermined threshold; and if the amount is greater than the predetermined threshold, deallocating the predetermined quantum from the first node.*”

12. Claim 15

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

186. See analysis at [1.0].

b. [15.1] *a plurality of nodes, comprising at least first and second nodes; a plurality of links, interconnecting the nodes to provide communications therebetween;*

187. See analysis at [1.0].

- c. ***[15.2] and a dispatcher, which is adapted to allocate a respective number of quanta of one or more resources in the network to each of the nodes,***

188. See analysis at [1.1] and [11.0].

- d. ***[15.3] wherein the first node is adapted, upon receiving a request to use a portion of the resources so as to carry the data flow between the first and second nodes over one of a plurality of paths therebetween, each such path comprising a respective sequence of the links,***

189. See analysis at [1.2].

- e. ***[15.4] to direct the data flow from the first node to the second node over one of the paths without requesting additional resources if the resources already allocated to the first node on the one of the paths are sufficient,***

190. See analysis at [1.3].

- f. ***[15.5] and to request an increase in an allocation of the requested resources to the first node if the resources already allocated to the first node are insufficient,***

191. See analysis at [1.4] and [1.5].

- g. ***[15.6] and wherein in response to the increase requested by the first node, the dispatcher is adapted to determine, for the links comprised in each of the paths, respective levels of use of the requested resources due to communications in progress over the network***

192. See analysis at [1.6].

- h. ***[15.7] and to select which of the paths is to carry the data flow responsive to the determined levels of use of the requested resources on the links comprised in each of the paths,***

193. See analysis at [1.7].

- i. ***[15.8] and to increase the allocation of the resources to the first node on the selected path by a predetermined quantum.***

194. See analysis at [1.8].

13. Claim 16

- a. ***[16.0] A network according to claim 15, wherein the dispatcher is adapted to select for carrying the data flow the one of the paths that has the lowest level of a predetermined measure of use of the requested resources.***

195. See analysis at [10.0].

14. Claim 17

- a. ***[17.0] A network according to claim 16, wherein the measure of use comprises, for each of at least two of the paths, a maximum level of use of at least one of the requested resources, taken over all of the links comprised in the paths.***

196. Kalman renders this limitation obvious.

197. As discussed in [10.0], Kalman provides examples where the cost is based on just one factor (e.g. packet loss rate, average packet queuing delay, or variation of packet queuing delay), Ex.1006, 16:37-41, and that a POSITA would have found it obvious that the calculation of cost could be based on the level of traffic congestion alone.

198. Referring to Kalman's Figure 3 (discussed in [1.6]), for a given packet/flow between two nodes (e.g. Node 0 to Node 5), in an exemplary scenario there may be congestion on the counterclockwise path (blue) between Node 0 and

Node 7, corresponding to a link cost of 10, and no congestion on the remaining links on the counterclockwise path (i.e. Node 7 to Node 6, and Node 6 to Node 5). On the clockwise path (orange), there may be congestion between Node 1 and Node 2, corresponding to a link cost of 25, and no congestion on the remaining links of the clockwise path (i.e. Node 0 to Node 1, Node 1 to Node 2, Node 2 to Node 3, Node 3 to Node 4, and Node 4 to Node 5).




FIG. 3

Ex.1006, Fig. 3 (annotated)

199. Applying Kalman's teachings, the counterclockwise path would be selected because the maximum link cost on that path (e.g. 10) is lower than the maximum link cost on the clockwise path (e.g. 25), resulting in the total path cost for the counterclockwise path (blue) being the lower than the path cost for the clockwise path.

200. Therefore, Kalman renders obvious "*wherein the measure of use comprises, for each of at least two of the paths, a maximum level of use of at least one of the requested resources, taken over all of the links comprised in the paths.*"

15. Claim 18


- a.** **[18.0]** *A network according to claim 15, wherein the first node is adapted, upon release of the requested resources, to compare an amount of the resources that have been released to a predetermined threshold, and if the amount is greater than the predetermined threshold, to deallocate the predetermined quantum from the first node.*

201. See analysis at [14.0].

IX. THE COMBINATION OF KOVVALI, KALMAN, AND MASTERS RENDERS OBVIOUS CLAIM 6

A. Summary of Masters

202. Like Kalman, Masters relates to choosing one of a plurality of routes between nodes based on total cost. Ex.1007, 3:34-37. Figure 7 illustrates an example of a network where messages may be routed across various communication paths between nodes, labeled C-G. Ex.1007, 13:57-67.



Ex.1007, Fig. 7 (annotated)

203. In Figure 7, Masters depicts two possible routes for sending a message to site E from site C, “which reflect taking different routes having different intermediate hops.” Ex.1007, 14:32-37. One possible route (blue) “travels from node C to F to E,” and another route (orange) travels “from nodes C to F to G to E.” Ex.1007, 14:37-41.

204. “Each potential route is composed of one or more connectors specified in a sequential order. The message is passed sequentially from one connector to the next connector in the route until the message reaches its destination site. Each connector provides direct connectivity between two sites and has an associated

individual cost.” Ex.1007, 3:40-45. In Figure 7, the first exemplary route (blue) from node C to node E uses connectors Vcf and Vfe, and the second exemplary route (orange) uses connectors Vcf, Vfg, and Vge. Ex.1007, 14:33-34, Fig. 8.

205. Masters describes that a “repository” including a routing table of message routing information is used to identify the total cost associated with each of the respective plurality of routes between connectors. Ex.1007, 3:27-39, 4:63-5:9. The route chosen is based on its associated cost, for example, “the lowest cost route,” Ex.1007, 3:27-39, 7:29-38. “The total cost of a route is equal to the individual costs of its constituent connectors, plus an optional increment amount for each intermediate site encountered along the route.” Ex.1007, 3:45-47. Figure 8 provides an example of a routing table with the total cost of various routes from node C to nodes D, E, F, and G. For example, the total cost associated with sending a message along the first exemplary route (blue) using connectors Vcf and Vfe is 13. The total cost associated with sending a message along the second exemplary route (orange) from node C to E using connectors Vcf, Vfg, and Vge is 16.

Site C's Routing Table

<u>Destination</u>	<u>Route</u>	<u>Total Cost</u>
D	Vcd	1
D	Vcf, Vfe, Ved	21
D	Vce, Ved	108
E	Vcd, Vde	12
E	Vcf, Vfe	13
E	Vcf, Vfg, Vge	16
F	Vcf	9
F	Vcd, Vde, Vef	18
F	Vce, Vef	106
G	Vcf, Vfg	12
G	Vcf, Vfe, Veg	18
G	Vcd, Vde, Vef, Vfg	21
G	Vce, Veg	105
Internet	Vcf, Vfg, Vgi	14
.	.	.
.	.	.
.	.	.

Ex.1007, Fig. 8 (annotated)

206. The value assigned to cost may be “any other measure of resources” including the amount of “processing time” that will be expended if the site connector is used, Ex.1007, 6:2-8, or the “processing power used to send a message by that connector.” Ex.1007, 16:37-38. “One of the plurality of routes is chosen based on its associated total cost, for example, the lowest cost route, and the message is sent from the originating site to the destination site using the chosen route.” Ex.1007, 3:34-37.

B. Reasons to Combine Kovvali, Kalman, and Masters

207. A POSITA would have been motivated to combine Masters with Kovvali and Kalman, because Masters also relates to efficient transmission of data in a communication network. Ex.1007, 15:7-10 (“to determine the best route for a message to take to reach a final destination” based on “cost information”); *see also* Section VIII.C (describing reasons to combine Kovvali and Kalman).

208. A POSITA would have been motivated to combine Masters with Kovvali and Kalman, and doing so would be within a POSITA’s skillset. The combination would have been obvious because it represents merely the use of Master’s known technique for selecting a preferred route based on the cost of “processing power” to improve Kalman’s total cost calculation in the same way. When processing power of a node affects the total cost of selecting a path for transmitting data, a POSITA would look to Masters for details on how to account for the cost of processing power.

209. The combination would have also been obvious because the application of Masters technique (selecting a preferred route on a network based on total cost that accounts for “processing power”) to improve Kalman’s method (selecting a preferred path on a ring network based on total cost) ready for improvement would yield predictable results. Here, the improvement provided to Kalman’s total cost

calculation is accounting the cost of “processing power” that may impact the total cost of a particular path.

210. It was known in the art that “the processing power available at a switch or router is limited.” Ex.1010, 3:16-17. While that processing power is typically “dedicated to the primary business of the switch/router—switching or routing,” other obligations can sometimes cause “performance drops.” Ex.1010, 3:18-20. Furthermore, the same level of processing power may not be available for all ports on a switch or router. Ex.1010, 3:22-24 (“Putting significant processing power on all the ports of a switch or router is not cost-effective.”). As a result, it would have been obvious to a POSITA for Kovvali’s arbiter node to take into account the available processing power of the nodes in the network when evaluating potential paths for network traffic flowing between two nodes.

211. A POSITA would have had a reasonable expectation of success in combining the teachings of Masters and Kalman. Kalman describes that the computation of cost associated with sending data between nodes “may take into account a variety of factors.” Ex.1006, 7:66-8:1. Masters simply contemplates that the cost of a route may include the “processing power” used by the connectors to send the data. Ex.1007, 16:33-38. Masters also explicitly suggests the combination by stating that its techniques “find applicability in any network that transmits information between different nodes, or locations,” Ex.1007, 3:65-67, and thus a

POSITA would have understood that Masters technique for including “processing power” in a total cost calculation for selecting a preferred path for sending data between nodes would have been readily applicable to the combined network of Kovvali and Kalman.

C. Detailed Analysis of Claims

1. Claim 6

- a. **[6.0]** *A method according to claim 1, wherein the one or more resources comprise a processing power associated with each of the links.*

212. The combination of Kovvali, Kalman, and Masters renders this limitation obvious.

213. First, as discussed in [1.6], Kovvali teaches a BLSR network where each pair of nodes are connected with two unidirectional optical links traveling in opposite directions. Ex.1005, 2:10-16. Further, as discussed in [1.6], Kalman discloses a technique for selecting a preferred direction for transmitting data between nodes based on the direction with the lowest cost. Ex.1006, 7:50-63.

214. Second, Masters also describes a technique for selecting a route among a plurality of routes based on lowest cost:

According to another aspect of the invention, messages may be routed in the electronic messaging system by maintaining and using a repository of message routing information which forms the basis for message routing decisions. The repository includes an identity of at least

one destination site, **a plurality of routes for reaching the destination site, and a total cost associated with each of the respective plurality of routes. One of the plurality of routes is chosen based on its associated total cost, for example, the lowest cost route,** and the message is sent from the originating site to the destination site using the chosen route. This route may traverse one or more intermediate sites before it reaches the destination site.

Ex.1007, 3:27-39.

215. In particular, Master describes that a route between sites is composed of one or more “connectors.” Ex.1007, 3:40-41. A message is passed along the route from one connector to another connector until the message reaches its destination site. Ex.1007, 3:41-43. “Each connector provides direct connectivity between two sites and has an associated individual cost.” Ex.1007, 3:43-45.

Each potential route is composed of one or more connectors specified in a sequential order. The message is passed sequentially from one connector to the next connector in the route until the message reaches its destination site. Each connector provides direct connectivity between two sites and has an associated individual cost. The total cost of a route is equal to the individual costs of its constituent connectors, plus an optional increment amount for each intermediate site

encountered along the route. The cost associated with a particular connector and/or route represents a predetermined measure of resources (e.g., monetary expense, bandwidth usage, processing time) for using that connector.

Ex.1007, 3:40-52.

One component of the directory service is the site connector--a software object (defined as a discrete entity, comprising data and processes that act on that data, for performing a specific task) that facilitates intersite communications by configuring and managing RPC connections across different sites within a local area network ("LAN"). **Each instance of the site connector object identifies one and only one remote site, and describes the connectivity between the local site and the designated remote site.** In other words, a separate site connector object is provided for each remote site to which a potential connection exists. Generally speaking, a site connector is used to connect together sites of the same type, for example, those operating under the same vendor's messaging software and/or using the same addressing scheme.

Another class of object potentially present in the directory service is a gateway connector object, which is similar to a site connector object except that a site connector describes and manages connections to other sites of the

same type within the same LAN, whereas a gateway connector object is typically used to reach a remote foreign messaging system--i.e., one that is provided by another vendor or that utilizes a different addressing scheme. In addition, whereas a site connector ordinarily facilitates only RPC connectivity between sites, a gateway connector may provide a broader assortment of connectivity mechanisms such as telephone lines, networks and the like.

Each instance of the gateway connector object ordinarily identifies only one foreign system, and thus describes the connectivity between the local site and that foreign system, but a single gateway connector may be used to reach more than one remote system, if desired. In addition, a gateway connector may be used to interconnect messaging sites of the same type, in a manner similar to a site connector object.

Ex.1007, 5:24-58.

216. Masters describes that a cost may be associated with the amount of resources that will be expended in using a particular site connector. Ex.1007, 5:66-6:1. The cost may correspond to “the amount of processing time” that will be expended if the site connector is used, Ex.1007, 6:2-8, or the “processing power used to send a message by that connector,” Ex.1007, 16:33-38.

FIG. 2 is a simplified representation of the directory service showing two different classes of connector objects--the site connector object 903 and the gateway connector object 923. The site connector object 903 includes its common name 905 within the directory service so that the site connector object may be located by other objects and a display name 904, e.g., ASCII encoded text which is readable by a human. An optional cost parameter 906 may be included which specifies a predetermined measure of resources that will be expended in using the site connector. **The value assigned to the cost parameter 906** may correspond to the monetary expense of using this connector (for example, the charge that would be imposed by the communications service provider) or it **may correspond to any other measure of resources--for example, the amount of processing time or bandwidth that will be expended if the site connector is used.**

Ex.1007, 5:59-6:8

11. The method of claim 9 in which the **individual cost associated with a particular connector** represents a predetermined measure of resources for using that connector.

12. The method of claim 11 in which the predetermined measure of resources comprises at least one of (i) monetary expense incurred by using the connector; (ii) amount of bandwidth used to send a message by that

connector; or (iii) **measure of processing power used to send a message by that connector.**

Ex.1007, 16:30-38.

217. Accordingly, it would have been obvious to a POSITA that when the cost of a link between two connectors includes “the processing power used to send a message by that connector,” the processing power is “*associated with*” the link directly connecting two sites.

218. Masters also describes that a “repository” including a routing table of message routing information for identifying the total cost associated with each of the plurality of routes between connectors, Ex.1007, 3:27-39, 4:63-5:9, is used for choosing a route based on its associated cost (e.g. lowest cost route), Ex.1007, 3:27-39, 7:29-38:

According to another aspect of the invention, messages may be routed in the electronic messaging system by maintaining and using a repository of message routing information which forms the basis for message routing decisions. **The repository includes an identity of at least one destination site, a plurality of routes for reaching the destination site, and a total cost associated with each of the respective plurality of routes. One of the plurality of routes is chosen based on its associated total cost, for example, the lowest cost route, and the message is sent from the originating site to the destination**

site using the chosen route. This route may traverse one or more intermediate sites before it reaches the destination site.

Ex.1007, 3:27-39.

In one embodiment of the messaging system architecture, the ability to communicate with an arbitrarily large number of messaging sites in an integrated and cohesive manner may be achieved through the use of a "directory service"--a **repository of information** that is distributed across the sites in an organization and which defines the set of remote sites that are reachable from any particular site and how those remote sites may be reached. The directory service includes **a routing table which lists potential routes for reaching remote sites**, software components that facilitate connections between sites, and a database that stores information concerning the attributes of the individual servers and end-users of the e-mail system, including permissions and the like.

Ex.1007, 4:63-5:9.

Each entry in the routing table 40 has three fields: a final destination 46 which represents the address space for which the message is bound, a route 44 by which the message may travel to reach that destination, and an associated total cost ("TC") 42 for using that route. A route is a series of one or more consecutive communication paths which traverse different address

spaces and over which a message will travel to reach its destination. The route field 44 in a routing table will list, in proper routing path order, the name of connector leading out of the local site followed by the names of subsequent connectors in remote sites along which the message will travel en route to its final destination.

Ex.1007, 7:16-27.

219. A POSITA would have found it obvious that calculating the total cost for each of a plurality of routes based on “the processing power used to send a message” by the respective connectors in those routes, and selecting the “lowest cost route,” renders obvious determining respective levels of resource usage (as discussed above in [1.6]) and selecting a path responsive to the determined levels of use (as discussed in [1.7]) based on “*a processing power associated with each of the links.*”

X. THE COMBINATION OF KOVVALI, KALMAN, AND RFC2892 RENDERS OBVIOUS CLAIMS 7-8

A. Summary of RFC2892

220. RFC2892 specifies a “Spatial Reuse Protocol” (SRP) for use with ring based media. Ex.1018, 1. RFC2892 describes a “SRP Fairness Algorithm (SRP-fa)” for allocating bandwidth fairly around the ring. Ex.1018, 9. When there is no congestion, “[e]ach node can send traffic up to max_usage. Max_usage is a per node

parameter that limits the maximum amount of low priority traffic a node can send.”
Ex.1018, 29.

B. Reasons to Combine Kovvali, Kalman, and RFC2892

221. A POSITA would have been motivated to combine RFC2892 with Kovvali and Kalman, because RFC2892 also relates to “efficient use of bandwidth” in a ring network. Ex.1018, 1. As discussed in [4.0], Kovvali and Kalman each describe “spatial reuse” for transmitting traffic around the ring. In particular, Kovvali describes explicitly a “Spatial Reuse Protocol (SRP)” allowing reuse of ring bandwidth across different sections of a ring of nodes. Ex.1005, 4:17-18. Kovvali describes further that “[a] supplementary algorithm called SRP Fairness algorithm (SRP-fa) was also defined to achieve fair sharing of bandwidth between various nodes in a ring topology.” Ex.1005, 4:18-21. A POSITA carrying out the teachings of SRP and SRP-fa in Kovvali would look to known documents such as RFC2892 that specify the SRP and SRP-fa respectively.

C. Detailed Analysis of Claims

1. Claim 7

- a. *[7.0] A method according to claim 1, wherein selecting which of the paths is to carry the data flow comprises comparing an amount of the one or more resources requested to a resource budget assigned to the first node, and permitting the data flow only if allocating the*

requested resources will not cause a total of the resources allocated to the first node to exceed the budget.

222. The combination of Kovvali, Kalman, and RFC2892 renders this limitation obvious.

223. First, as discussed in [1.1], an arbiter node assigns “dynamic channels” to particular nodes and adjusts the allocated rates for the dynamic channels upon request by client nodes. Ex.1005, 9:1-11. Each dynamic channel is allocated zero or more columns in a dynamic frame. Ex.1005, 10:3-5.

224. Second, RFC2892 describes a “SRP Fairness Algorithm (SRP-fa)” in which nodes are assigned parameters that limit the amount of low priority traffic that each node can send. Ex.1018, 9, 29. In particular, RFC2892 describes that according to the SRP-fa, a variable “allowed_usage” is set that defines “the current maximum transmit usage for that node.” Ex.1018, 29. When a network is not congested, the value of “allowed_usage” is allowed to build up periodically “up to” the constant “max_usage” which is a “per node parameter” that limits the maximum amount of low priority traffic a node can send. Ex.1018, 29.

3.3. SRP Fairness Algorithm (SRP-fa) Overview

If a node experiences congestion, then it will advertise to upstream nodes via the opposite ring the value of its transmit usage counter. The usage counter is run through a low pass filter function to stabilize the feedback. Upstream nodes will adjust their transmit rates so as not to

exceed the advertised values. Nodes also propagate the advertised value received to their immediate upstream neighbor. Nodes receiving advertised values who are also congested propagate the minimum of their transmit usage and the advertised usage.

Congestion is detected when the depth of the low priority transit buffer reaches a congestion threshold.

Usage messages are generated periodically and also act as keepalives informing the upstream station that a valid data link exists.

Ex.1018, 9.

6. SRP-fa Rules Of Operation

The SRP-fa governs access to the ring. The SRP-fa only applies to low priority traffic. High priority traffic does not follow SRP-fa rules and may be transmitted at any time as long as there is sufficient transit buffer space is sufficient transit buffer space.

The SRP-fa requires three counters which control the traffic forwarded and sourced on the SRP ring. The counters are `my_usage` (tracks the amount of traffic sourced on the ring), `forward_rate` (amount of traffic forwarded on to the ring from sources other than the host) and **`allowed_usage` (the current maximum transmit usage for that node).**

With no congestion all nodes build up allowed usage periodically. Each node can send up to max_usage. Max_usage is a per node parameter than limits the maximum amount of low priority traffic a node can send.

When a node sees congestion it starts to advertise its my_usage which has been low pass filtered (lp_my_usage).

Congestion is measured by the transit buffer depth crossing a congestion threshold.

A node that receives a non-null usage message (rcvd_usage) will set its allowed usage to the value advertised. However, if the source of the rcvd_usage is the same node that received it then the rcvd_usage shall be treated as a null value. When comparing the rcvd_usage source address the ring ID of the usage packet must match the receiver's ring ID in order to qualify as a valid compare. The exception is if the receive node is in the wrap state in which case the usage packet's ring ID is ignored.

Ex.1018, 28-29.

225. In another example, RFC2892 describes pseudo-code defining a variable "MAX_ALLOWANCE" that sets the maximum allowed usage for a node:

6.1. SRP-fa pseudo-code

A more precise definition of the fairness algorithm is shown below:

Variables:

lo_tb_depth	low priority transit buffer depth
my_usage	count of octets transmitted by host
lp_my_usage	my_usage run through a low pass filter
my_usage_ok	flag indicating that host is allowed to transmit
allow_usage	the fair amount each node is allowed to transmit
fwd_rate	count of octets forwarded from upstream
lp_fwd_rate	fwd_rate run through a low pass filter
congested	node cannot transmit host traffic without the TB buffer filling beyond its congestion threshold point.
rev_usage	the usage value passed along to the upstream neighbor

Constants:

MAX_ALLOWANCE = configurable value for max allowed usage for this node.

[. . .]

THESE ARE UPDATED EVERY CLOCK CYCLE:

```
=====

my_usage      is incremented by 1 for every octet that
               is transmitted by the host (does not
               include data transmitted from the
               Transit Buffer).

fwd_rate      is incremented by 1 for every octet that
               enters the Transit Buffer

if ((my_usage < allow_usage) &&
    !((lo_tb_depth > 0) && (fwd_rate <
my_usage)) &&
    (my_usage < MAX_ALLOWANCE))
    // true means OK to send host packets
    my_usage_ok = true;
```

Ex.1018, 30-31.

226. In the SRP-fa pseudo-code example described in RFC2892 above, at each clock cycle, the variable “my_usage” (amount of data transmitted by host) is compared to the constant “MAX_ALLOWANCE” (maximum allowed usage for this node) to check for the condition that “my_usage” is less than MAX_ALLOWANCE. Only if this condition is true, is it “OK to send host packets.”

227. A POSITA would have thus found it obvious that “max_usage” corresponds to a “*resource budget assigned to the first node*” because a node can only send “up to” max_usage which is the value for the “maximum amount of low

priority traffic” a node can send. Similarly, a POSITA would have found it obvious that “MAX_ALLOWANCE” corresponds to a “*resource budget assigned to the first node*” because it is the value configured for the “max[imum] allowed usage for this node.”

228. Third, as discussed in [1.5], Kovvali teaches that to change a bandwidth allocation to a dynamic channel assigned to it, a client node may request increased (or decreased) bandwidth for a corresponding bursty channel from the arbiter. Ex.1005, 10:35-37. In particular, the client node sets a request section (REQ) within the dynamic frame overhead specifying the amount of requested bandwidth. Ex.1005, 10:35-37. The arbiter responds by setting a grant section (GNT) within the dynamic frame specifying the amount of bandwidth granted. Ex.1005, 10:37-38.

Referring still to FIG. 5A, **each control channel 524 includes two sections, a request section (REQ) 526 and a grant section (GNT) 527. REQ 526 is set by a client node 420 (FIG. 4) to request additional bandwidth for the corresponding bursty channel from the arbiter. GNT527 is set by the arbiter node to grant bandwidth to the requested channel.** The values in REQ and GNT are encoded thereby reducing the number of bits needed for the control channel. In this embodiment, in which control channel 524 is 8 bits, REQ 526 is 3 bits and GNT 527 is 5 bits, the values of REQ and GNT are encoded as follows:

REQ	Requested Bandwidth
000	No bandwidth required
001	Deallocate 1 column
010	Deallocate 3 columns
011	Deallocate 6 columns
100	Allocate 1 additional column
101	Allocate 3 additional columns
110	Allocate 6 additional columns
111	Allocate 10 additional columns
GNT	Granted Bandwidth
00000	No Bandwidth allocated
00001	Numerical value of the number
-11110	of columns allocated
11111	All available columns allocated to the channel

Ex.1005, 10:33-62.

229. It would have been obvious to a POSITA that an arbiter would compare the total bandwidth requested for a dynamic channel with the maximum allowed usage for the node (e.g. “max_usage” or “MAX_ALLOWANCE”) so that the arbiter does not grant bandwidth that exceeds a maximum amount (e.g., the maximum number of dynamic channel columns) allowable for the node. It would have been obvious to a POSITA for Kovvali’s arbiter not to allocate more bandwidth than the maximum value allowed so that network bandwidth can be allocated fairly, as discussed in RFC2892 relating to the SRP Fairness Algorithm.

230. Therefore Kovvali renders obvious “*wherein selecting which of the paths is to carry the data flow comprises comparing an amount of the one or more resources requested to a resource budget assigned to the first node, and permitting*

the data flow only if allocating the requested resources will not cause a total of the resources allocated to the first node to exceed the budget.”

2. Claim 8

- a. [8.0]** *A method according to claim 7, wherein comparing the amount of the one or more resources comprises comparing the amount of each of the resources requested to the resource budget assigned for each of the resources, and wherein permitting the data flow comprises permitting the flow only if all of the resources requested for at least one of the paths are within the budget.*

231. The combination of Kovvali, Kalman, and RFC2892 renders this limitation obvious.

232. As discussed in [7.0], Kovvali and RFC2892 teach comparing the total bandwidth requested for a dynamic channel (“*each of the resources requested*”) with the maximum allowed bandwidth usage (e.g. “max_usage” or “MAX_ALLOWANCE”) for the node (“*resource budget assigned for each of the resources*”). A POSITA would have found it obvious that when the resources requested for transmitting a dynamic channel exceeds the maximum allowed bandwidth usage, the arbiter would either deny the request (e.g. GNT 00000) or return a lower value up to “[a]ll available columns” (e.g. GNT 11111) for the channel in the frame so that the maximum number of columns that the node’s bandwidth usage does not exceed the configured maximum value (e.g. “max_usage” or “MAX_ALLOWANCE”).

233. A POSITA would have also found it obvious that if a dynamic channel had a minimum data requirement that exceeds the maximum allowed bandwidth usage for the node (*“the budget”*), the arbiter would deny the request for resources. Denying the request would ensure the network only delivers an expected quality of service according to the minimum data requirement of the dynamic channel.

234. It would have been obvious to a POSITA for Kovvali’s arbiter node to take into account quality of service requirements of different data flows because it was well known in the art that network traffic may have “different QoS [quality of service]” parameters with varying bandwidth requirements. Ex.1020, 1:59-61; 1:61-67 (“For satisfying services with QoS parameters...traffic management mechanisms are required to ensure that sufficient resources exist to meet the service objectives.... Prominent among those mechanisms is the AC [admission control] function which limits the volume of the traffic accepted to the network to allow the various QoS objectives to be satisfied.”); 2:34-36 (“A QoS manager keeps track of the resource availability of the network, and accepts a new request only if the requested resource is available.”); 2:42-49 (“initializing a QoS manager at an admit limit (AL) representing a bandwidth utilization capacity allocated to a service offered on the weakest link of the network; receiving a service request including a service type and a required bandwidth; comparing the AL with the required bandwidth and accepting the request whenever the required bandwidth does not exceed the AL, and rejecting

the request whenever the required bandwidth exceeds the AL"). Accordingly, Kovvali and RFC2892 render obvious *"permitting the flow only if all of the resources requested for at least one of the paths are within the budget."*

235. Therefore, Kovvali, Kalman, and RFC2892 render obvious *"wherein comparing the amount of the one or more resources comprises comparing the amount of each of the resources requested to the resource budget assigned for each of the resources, and wherein permitting the data flow comprises permitting the flow only if all of the resources requested for at least one of the paths are within the budget."*

XI. CONCLUSION

236. 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: 1/26/2024

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