

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.

Petitioner

v.

PORTSMOUTH NETWORK CORP.

Patent Owner

Case No. IPR2024-00954

U.S. Patent No. 8,199,637

DECLARATION OF HENRY H. HOUH, PH.D.

ON BEHALF OF PETITIONER

TABLE OF CONTENTS

	<u>Page</u>
I. Background.....	1
II. Relevant Law	9
A. Claim Construction	9
B. Obviousness.....	10
C. Ordinary Skill in the Art.....	12
III. Summary of Opinions.....	13
IV. Overview of the '637 Patent.....	13
A. Purported Invention Of The '637 Patent.....	13
B. Prosecution History	20
V. Claim Construction.....	24
VI. Level of Ordinary Skill in the Art	24
VII. Overview of the Prior Art.....	25
A. Mitchell	25
VIII. Specific Grounds for Challenge	33
A. Ground 1: Mitchell Renders Obvious Claims 1, 7, 9-11, 13, 15- 17, 19, 21, 23, and 25	33
1. Independent Claim 1	33
2. Independent Claim 16.....	75
3. Dependent Claim 17	85

4.	Dependent Claims 7, 13, 19, and 23	86
5.	Dependent Claims 9, 15, and 25	88
6.	Dependent Claims 10 and 21	90
7.	Dependent Claim 11	92
B.	Ground 2: Mitchell and Bhate Render Obvious Claims 2, 6, 12, 18, and 22	96
1.	Dependent Claim 2	96
2.	Dependent Claims 6, 12, 18, And 22	105
C.	Ground 3: Mitchell And Raahemi Render Obvious Claims 3 And 4	107
1.	Dependent Claim 3	107
2.	Dependent Claim 4	117
D.	Ground 4: Mitchell And Huai Render Obvious Claim 5	120
1.	Dependent Claim 5	120
E.	Ground 5: Mitchell And Yamauchi Render Obvious Claims 8, 14, 20, And 24	125
1.	Dependent Claims 8, 14, 20, And 24	125
F.	Ground 6: Mitchell And McGee Render Obvious Claims 9, 15, And 25	129
1.	Dependent Claims 9, 15, And 25	129
IX.	Conclusion	132
X.	Availability for Cross-Examination	132
XI.	Right to Supplement	133

XII. Jurat.....	133
-----------------	-----

I, Henry H. Houh, declare as follows:

1. My name is Henry H. Houh.
2. I have been retained as an expert witness on behalf of Petitioner Cisco Systems, Inc. to provide expert opinions on the validity of U.S. Patent No. 8,199,637 (the “637 patent”).

I. Background

3. I received a Ph.D. in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology (“MIT”) in 1998. I also received a Master of Science degree in Electrical Engineering and Computer Science in 1991, a Bachelor of Science degree in Electrical Engineering and Computer Science in 1989, and a Bachelor of Science degree in Physics in 1990, all from MIT.

4. I am currently self-employed as an independent technical consultant. I am also president of a company, Einstein’s Workshop, that provides supplemental science, technology, engineering, and mathematics (“STEM”) education to children of all ages.

5. I first entered telecommunications in 1987 when I worked as a summer intern at AT&T Bell Laboratories as part of a five-year dual degree program at MIT. I continued to work at AT&T Bell Laboratories as part of this MIT program. While at MIT, I was a teaching assistant (“TA”) in the Electrical Engineering and Computer Science Department’s core Computer Architectures

course. I first was a TA as a senior for a role typically reserved for graduate students. I later became head TA. The course covered various topics in computer architectures, as well as programming. As a TA, I helped write homework assignments, lab assignments (including those that involved programming), and exams. I also taught in the recitation sections.

6. Later, as part of my doctoral research at MIT from 1991-1998, I was a research assistant in the Telemedia Network Systems (“TNS”) group at the Laboratory for Computer Science. The TNS group built a high-speed gigabit network and created applications that ran over the network. Example applications included ones for remote video capture, processing, and display of video on computer terminals. In addition to working on the design of core network components, designing and building the high-speed links, and designing and writing the device drivers for the interface cards, I also set up the group’s web server. Also, I helped to maintain, install, and upgrade the networking devices used within the group, along with other graduate students. Being a networking group, we regularly discussed topic of concern in networking, including network security and the kinds of attacks that were carried out over a network. I used several network scanning applications, including Nmap, to assess the level of security in our own network.

7. I also helped to build the web pages that initiated the above-mentioned video sessions via a web interface. Vice President Al Gore visited our group in 1996 and received a demonstration of—and remotely drove—a radio-controlled toy car with a wireless video camera mounted on it that was built by our group. This toy car device received commands transmitted over a network from a remote computer, and video data from the toy car was transmitted wirelessly then over a computer network back to the user controller. On occasion, we allowed users visiting our web site to drive the toy car from their remote computer while they watched the video on their computer. The video stream was encoded by TNS-designed hardware, streamed over the TNS-designed network, and displayed using TNS-designed software. We also installed a multicast backbone (“MBONE”) router within our group around the mid 1990’s, and I personally accessed multicast video at that time. As part of my consulting work, I have also studied commercial deployments of IPTV and their use of multicast, including the AT&T U-Verse system first deployed around 2006.

8. I defended and submitted my Ph.D. thesis, titled “Designing Networks for Tomorrow’s Traffic,” in January 1998. As part of my thesis research, I analyzed local area and wide area flows to show a more efficient method for routing packets in a network, based on traffic patterns at the time. The traffic flow

data included Ethernet, IP, TCP or UDP, and RTP header information, which I analyzed to come to the conclusions in my thesis.

9. From 1997 to 1999, I was a Senior Scientist and Engineer at NBX Corporation, a start-up that made business telephone systems for streaming packetized audio over data networks instead of using traditional telephone lines. NBX was later acquired by 3Com Corporation.

10. As part of my work at NBX, I designed the core audio reconstruction software algorithms for the telephones, as well as the packet transmission algorithms. I also designed and validated the core packet transport protocol used by the phone system. The protocol was used for all signaling in the phone system, including for the setup of conference calls. The NBX system also featured a computer interface for initiating phone calls, which could also initiate conference calls. The NBX system also supported the Telephony Application Programming Interface (“TAPI”) that allowed other computer programs to integrate with our system telephony features. I also programmed the first prototype of our phone which communicated using IP, and I demonstrated our IP phones working over the Internet when we attended a trade show in California. The phone connected over the Internet to our headquarters in Andover, MA. I also programmed a data capture and analysis system that provided data on the operation of the system.

11. From 1999-2004, I was employed by Empirix or its predecessor company, Teradyne. Empirix was a leader in test tools for telecommunications protocols and systems, providing functional testing tools as well as load testing tools. From 2000-2001, I conceived and built a test platform for testing Voice-over-IP (VoIP). The first application on this new test platform was a cloud emulator for simulating the effects of transmitting VoIP over a busy network. The test platform was based on a network processor chip, which could be programmed to cut-through packets while processing packet data such as various protocol layer headers including addresses included therein and even packet data contents. I also designed a protocol analyzer built on the same platform. The application captured and performed protocol decoding at various layers in the protocol stacks of captured packets, including detailed Ethernet header decoding, IP header decoding, TCP header decoding, UDP header decoding, RTP header decoding, and many other specified protocols. The application was also designed to reconstruct entire conversations that spanned multiple packets.

12. At Empirix, I also rearchitected the design of the Web Test division's core product, e-Test Suite. E-Test Suite was a software program designed to perform functional and load testing of web applications. I was technical lead on the project, and also worked on programming a Javascript interpreter for the product. Javascript is the flexible programming layer that allows users to interact with web

pages through actions such as clicking and dragging that allows web page programmers to extend the functionality of a web page beyond clicking on links. I also designed and prototyped a web server security testing product, which scanned web servers for specific types of vulnerabilities, including looking for administrative logins which could be attacked with password attacks and buffer overflow attacks, and attacks such as SQL query injection attacks. We used the prototype on some customer web sites, with their permission, and successfully detected vulnerabilities.

13. In 2006, as part of my role at BBN Technologies, I helped found PodZinger Inc., now known as RAMP Inc. PodZinger utilized BBN's speech recognition algorithms to search through the spoken words in audio and video segments. While I was Vice President of Operations and Technology, PodZinger followed its initial prototype with a full streaming audio and video search solution. I also created a social networking web site, which BBN sold to a venture-funded startup company. In the process of creating the web site, I designed and specified the authentication and authorization protocols.

14. Until October 2022, I was President of Einstein's Workshop, a STEM education company for kids. At Einstein's Workshop, we taught programming to children of all ages, starting with Scratch programming for young elementary students and advancing through programming in Python and for game

development for high school students. I conceived many of the programming classes and have taught programming to novice and more advanced students. We also created an educational computer-aided design tool, called BlocksCAD. We received a grant, on which I was the principal investigator, to further develop BlocksCAD and successfully met our grant objectives. BlocksCAD was spun-off from Einstein's Workshop and has been part of various incubator/accelerator programs in education including LearnLaunch and Mass Challenge as well as Play Labs at the MIT Media Lab. Today, BlocksCAD is used worldwide in makerspaces as well as in schools throughout the United States.

15. I have reviewed all of the exhibits cited in this declaration, including the following:

Exhibit No.	Description
1001	U.S. Patent No. 8,199,637 B2 to Bruckman <i>et al.</i> ("637 Patent")
1002	Declaration of Dr. Henry Houh
1003	<i>Curriculum Vitae</i> of Dr. Houh
1004	Prosecution History of U.S. Patent No. 8,199,637 B2
1005	U.S. Patent No. 8,208,370 B1 to Mitchell <i>et al.</i> ("Mitchell")
1006	U.S. Patent No. 7,113,699 B1 to Bhate <i>et al.</i> ("Bhate")
1007	U.S. Patent Publication No. 2007/0047556 A1 to Raahemi <i>et al.</i> ("Raahemi")
1008	U.S. Patent No. 6,614,785 B1 to Huai <i>et al.</i> ("Huai")

Exhibit No.	Description
1009	U.S. Patent Publication No. 2005/0073965 A1 to Yamauchi (“Yamauchi”)
1010	U.S. Patent Publication No. 2006/0029097 A1 to McGee <i>et al.</i> (“McGee”)
1011	U.S. Patent Publication No. 2006/0109802 A1 to Zelig <i>et al.</i> (“Zelig”)
1012	U.S. Patent Publication No. 2005/0271035 A1 to Cohen <i>et al.</i> (“Cohen”)
1013	U.S. Patent No. 7,154,900 B1 to Alanqar <i>et al.</i> (“Alanqar”)
1014	U.S. Patent No. 4,646,289 A to Tsiakas <i>et al.</i> (“Tsiakas”)
1015	U.S. Patent No. 7,463,580 B2 to Agmon <i>et al.</i> (“Agmon”)
1016	Docket Control Order, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Dec. 18, 2023)
1017	Notice of Disclosures, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Jan. 8, 2024)
1018	Order Granting Unopposed Motion for Leave to File Amended Complaint for Patent Infringement, <i>Portsmouth Network Corp. v. Cisco Sys., Inc.</i> , No. 2:23-cv-00441 (E.D. Tex. Mar. 12, 2024)
1019	Director’s June 21, 2022, Memorandum on Discretionary Denials

16. I understand the ’637 patent has a filing date of July 22, 2008 and claims priority to a U.S. Provisional Patent Application No. 60/737,733 filed on November 16, 2005. Thus, I have considered certain issues from the perspective of a person of ordinary skill in the art (“POSITA”) as described below as of November 16, 2005, the alleged priority filing date of the ’637 patent.

17. I have been retained by the Petitioner as an expert in the field of information technology and networking. I am working as an independent consultant in this matter and am being compensated at my normal consulting rate of \$680 per hour for my time. My compensation is not dependent on and in no way affects the substance of my statements in this Declaration.

18. I have no financial interest in the Petitioner. I similarly have no financial interest in the '637 patent and have had no contact with any of the named inventors of the '637 patent.

II. Relevant Law

19. I am not an attorney. For the purposes of this declaration, I have been informed about certain aspects of the law that are relevant to my opinions. My understanding of the law is as follows:

A. Claim Construction

20. I have been informed that claim construction is a matter of law and that the final claim construction will ultimately be determined by the Board. For the purposes of my analysis in this proceeding and with respect to the prior art, I have been informed that IPRs are currently reviewed under “the *Phillips* standard.”

21. I have been informed that under the *Phillips* standard, claim terms are given their plain and ordinary meaning as understood by a person of ordinary skill

in the art at the time of the invention in light of the claim language and the patent specification.

22. I have been informed that embodiments described in the specification are generally encompassed by the claims.

23. I have been informed that the patentee can serve as their own lexicographer. As such, if a claim term is provided with a specific definition in the specific, I should interpret that claim term in light of the particular definition provided by the patentee.

B. Obviousness

24. I have been informed and understand that a patent claim can be considered to have been obvious to a person of ordinary skill in the art at the time the application was filed. This means that, even if all of the requirements of a claim are not found in a single prior art reference, the claim is not patentable if the differences between the subject matter in the prior art and the subject matter in the claim would have been obvious to a person of ordinary skill in the art at the time the application was filed.

25. I have been informed and understand that a determination of whether a claim would have been obvious should be based upon several factors, including, among others:

- the level of ordinary skill in the art at the time the application was filed;
- the scope and content of the prior art; and
- what differences, if any, existed between the claimed invention and the prior art.

26. I have been informed and understand that a claim can be unpatentable in light of a single prior art reference if it would have been obvious to modify that reference to come up with the claimed subject matter.

27. I have been informed and understand that the teachings of two or more references may be combined in the same way as disclosed in the claims, if such a combination would have been obvious to one having ordinary skill in the art.

28. In determining whether a combination based on either a single reference or multiple references would have been obvious, it is appropriate to consider, among other factors:

- whether the teachings of the prior art references disclose known concepts combined in familiar ways, which, when combined, would yield predictable results;
- whether a person of ordinary skill in the art would have implemented a predictable variation, and would have seen the benefit of doing so;

- whether the claimed elements represent one of a limited number of known design choices, and would have a reasonable expectation of success by those skilled in the art;
- whether a person of ordinary skill would have recognized a reason to combine known elements in the manner described in the claim;
- whether there is some teaching or suggestion in the prior art to make the modification or combination of elements claimed in the patent; and
- whether the claim applies a known technique that had been used to improve a similar device or method in a similar way.

29. I understand that one of ordinary skill in the art has ordinary creativity and is not an automaton.

30. I understand that in considering obviousness, it is important not to determine obviousness using the benefit of hindsight derived from the patent being considered.

C. Ordinary Skill in the Art

31. I understand that, in the context of an invalidity analysis, a POSITA is a hypothetical person who looks to prior art at the time of the invention.

32. I further understand that the factors that may be considered in determining the level of ordinary skill in the art include: (1) the problems

encountered in the art; (2) the prior art solutions to the problems encountered in the art; (3) the rapidity of innovation; (4) the sophistication of the technology; and (5) the education level of active workers in the field. I understand that these factors need not all be taken into account for the analysis and that one or more of these factors may control.

III. Summary of Opinions

33. It is my conclusion that claims 1-25 (the “Challenged Claims”) of the ’637 patent are rendered obvious by the prior art identified below:

Ground	References	Basis (35 U.S.C.)	Claims
1	Mitchell	§ 103	1, 7, 9-11, 13, 15-17, 19, 21, 23, 25
2	Mitchell and Bhate	§ 103	2, 6, 12, 18, 22
3	Mitchell and Raahemi	§ 103	3, 4
4	Mitchell and Huai	§ 103	5
5	Mitchell and Yamauchi	§ 103	8, 14, 20, 24
6	Mitchell and McGee	§ 103	9, 15, 25

IV. Overview of the ’637 Patent

A. Purported Invention Of The ’637 Patent

34. The ’637 Patent generally describes “[a] method for handling communication failures in a network, which includes provisioning first and second multipoint-to-multipoint communication services over respective first and second


alternative sets of links that connect the endpoints in a network.” EX1001, Abstract.

35. The '637 states that in existing networks that provide primary and backup sets of links providing communication services, “there is no direct physical connection between the primary and backup external networks.” EX1001, 4:1-6. When a failure in the primary set of links occurs, “processes in upper layers are unable to identify this loss of connectivity in a timely manner in order to rapidly transfer the traffic to the backup set,” resulting in a failover that is “often slow, on the order of several minutes.” EX1001, 4:6-11. The '637 Patent explains that because of the lack of direct physical connection between the primary and backup networks, “upper layer protocols can only rely on timeout mechanisms to identify the loss of connectivity,” which results in long timeout intervals due to the complex implementation. EX1001, 4:12-16.

36. Therefore, the '637 Patent's alleged invention provides “improved methods for rapidly initiating the diversion of traffic from the failed primary topology to the backup topology” by propagating failure information to the nodes in the network, deactivating the physical layer of the primary set of links, and failing over to the backup set of links upon detecting a loss of connectivity in the first set. EX1001, 4:17-28. Because the backup set of links is predetermined and the physical layer of the first set of links is deactivated, the described failover

method is allegedly faster than existing failover methods. EX1001, 4:29-37 (“For example, in layer-3 networks, the loss of connectivity identified by the endpoints initiates a routing protocol database synchronization process among the network nodes. As a result of the synchronization process, routing tables used by the nodes for packet forwarding begin to use the newly available links. The resulting traffic diversion process is significantly faster, in comparison to known methods, and is typically completed within a few seconds.”).

37. The '637 Patent depicts an exemplary network in Figure 1, annotated below. EX1001, FIG. 1, 4:41-43 (“FIG. 1 is a block diagram that schematically illustrates a ring-based communication network 20, in accordance with an embodiment of the present invention.”).



EX1001, FIG. 1 (annotated).

38. Primary and backup topologies are shown as primary and backup VPLS instances *VPLS1* and *VPLS2*. EX1001, FIG. 1, 5:11-20 (“In some embodiments, the service provided to the endpoints comprises a virtual private local area network service (VPLS), and the primary and backup topologies are associated with two different VPLS instances, denoted VPLS1 and VPLS2. The user nodes and main sites are associated with both VPLS instances. Each user node and main site is connected to the corresponding ring node using two different interfaces or links, one link carrying the packets associated with each VPLS

instance.”), 5:36-59(“Thus, the endpoints of network 20 are able to communicate via two alternative topologies. The first topology, which uses VPLS1, comprises the ring network and links 44A . . . 44D. The second topology, which uses VPLS2, comprises the ring network and links 48A . . . 48D. The first and second topologies defined above are respectively referred to herein as the “VPLS1 topology” and the “VPLS2 topology” for brevity. Note that some physical elements, such as the elements of the ring network, are common to both topologies. Other elements, such as links 44A . . . 44D and 48A . . . 48D and some of routers 56, are separate. Typically, each endpoint is connected to the primary and backup links using a local routing network. For example, user node 24A is connected to links 44A and 48A using a local routing network 52A. Each local routing network comprises one or more routers 56. In some embodiments, a single redundant router supports both VPLS1 and VPLS2. Routing network 52D, for example, comprises a single router that connects main site 28 to both VPLS instances. Alternatively, the local routing network may comprise a pair of simpler and lower cost routers, with each router connected to a single VPLS. For example, user node 24B is connected via a pair of routers, one router connected to each of links 44B and 48B and to VPLS1 and VPLS2, respectively.”)

39. In Figure 1, **ring nodes 32A-C**, and **32D** are traversed by **VPLS1** and **VPLS2**. EX1001, 5:1-5 (“Network 20 comprises two alternative topologies that

connect the different endpoints. One topology serves as the primary topology, and the other topology serves as backup. Each of the primary and backup topologies is identified with a different service instance.”), 5:11-16 (“In some embodiments, the service provided to the endpoints comprises a virtual private local area network service (VPLS), and the primary and backup topologies are associated with two different VPLS instances, denoted VPLS1 and VPLS2. The user nodes and main sites are associated with both VPLS instances.”).

40. **VPLS1** provides the communication service to **main site 28** and **user nodes 24A-C** (“**endpoints**”) over the **ringlets 36 and 40** (“**ring network 34**”) and **links 44A-D**. **Ring network 34** and **links 44A-D** make up the first set of network links. EX1001, 5:26-38 (“For example, in FIG. 1, user node 24A is connected through a local routing network to ring node 32A using two separate links 44A and 48A. Links 44A and 48A are associated with VPLS1 and VPLS2, respectively. Similarly, endpoints 24B, 24C and 28 are connected to their ring nodes using VPLS1 over links 44B . . . 44D, and using VPLS2 over links 48B . . . 48D, respectively. Links 44A . . . 44D and 48A . . . 48D may comprise, for example, optical fiber links, copper connections, or any other suitable communication link. Thus, the endpoints of network 20 are able to communicate via two alternative topologies. The first topology, which uses VPLS1, comprises the ring network and links 44A . . . 44D.”).

41. **VPLS2** provides the communication service to the **endpoints** over the **ring network 34** and **links 48A-D**. **Ring network 34** and **links 44A-D** make up the second set of network links. EX1001, 5:26-40 (“For example, in FIG. 1, user node 24A is connected through a local routing network to ring node 32A using two separate links 44A and 48A. Links 44A and 48A are associated with VPLS1 and VPLS2, respectively. Similarly, endpoints 24B, 24C and 28 are connected to their ring nodes using VPLS1 over links 44B . . . 44D, and using VPLS2 over links 48B . . . 48D, respectively. Links 44A . . . 44D and 48A . . . 48D may comprise, for example, optical fiber links, copper connections, or any other suitable communication link. Thus, the endpoints of network 20 are able to communicate via two alternative topologies. The first topology, which uses VPLS1, comprises the ring network and links 44A . . . 44D. The second topology, which uses VPLS2, comprises the ring network and links 48A . . . 48D.”).

42. The '637 Patent describes an example where there is a failure in one of the links associated with **VPLS1**. First, a network node (e.g., **ring node 32D**) detects the failure. Then, the detecting network node (e.g., **ring node 32D**) propagates the failure information to the **other nodes**. EX1001, 7:10-15 (“In principle, when a network node detects a local failure or other loss of connectivity in one of the links of the primary topology, the node propagates this information to the other nodes of the primary topology. The detecting node propagates the

information by distributing a message, which is referred to herein as a remote fault indication (RFI).”), 9:13-15 (“When link 44D fails, node 32D detects the failure and sends out an RFI message to the other nodes.”). Each network node receives the failure information and deactivates the physical layer associated with the first set of links. EX1001, 7:16-18 (“Upon receiving an RFI message, **each node** disables the physical layers of the primary topology links that are locally-connected to the node.”), 9:15-16 (“Nodes 32A . . . 32C receive the RFI message and locally shut down links 44A . . . 44C, respectively.”). Finally, the **endpoints** detect a loss of connectivity in response to the deactivation of the physical layer of the first set of links and trigger a failover to **VPLS2**. EX1001, 7:22-25 (“As a result, the endpoints rapidly detect a local loss of connectivity and trigger a topology change, i.e., a diversion of the traffic from the failed primary topology to the backup topology.”).

B. Prosecution History


43. The '637 Patent was filed as U.S. Patent Application No. 12/093,824 on July 22, 2008. EX1004, 1-107. U.S. Patent Application No. 12/093,824 was originally-filed with claims 1-28. EX1004, 36-41. The '637 Patent claims priority to PCT Application No. PCT/IL2006/001283, which was filed November 7, 2006. EX1001, Cover. The '637 Patent claims priority to U.S. Provisional Patent Application No. 60/737,733, which was filed on November 16, 2005. EX1001,

Cover; *see also* EX1004, 1-33. Solely for the purposes of this IPR, I understand the earliest claimed priority date is November 16, 2005.

44. On December 15, 2009, the Examiner issued a Non-Final Office Action rejecting claims 5, 6, 11, 12, 21, and 24 under 35 U.S.C. §112, second paragraph, as being indefinite. The Examiner also rejected claims 1-28 under 35 U.S.C. §§102 and 103 based on the following references: Grenier (US 6,741,553), Le Faucheur (US2006/0215548), Filsfils (US2006/0153200), Quinlan (US2002/0021671), Mark (2003/0128706), and Tsurumi (US2005/0207348)—none of which are relied upon in this instant Petition. EX1004, 126-139; EX1001, Cover.

45. In response, the Applicant canceled claim 5, and amended claims 1, 6-13, 21, and 23. The Applicant amended independent claims 1, 13, and 23 to include the features “propagating failure information by a detecting node to each node of the other nodes of the plurality of network nodes traversed by a first set of links” and “receiving the failure information at a node and deactivating a physical layer of the first set of links connected thereto, thereby causing a loss of connectivity in the first set of links” and argued that the prior art references fail to disclose these features. To support the amendments, the Applicant cited to Figure 3, block 76 (shown below) and paragraph [0030] of the specification (which corresponds to the text reproduced below). EX1004, 155-166.

FIG. 3



EX1004, 33.

Embodiments of the present invention provide improved methods for rapidly initiating the diversion of traffic from the failed primary topology to the backup topology. In these
25 the failed primary topology to the backup topology. In these embodiments, when a network node detects a local failure or other loss of connectivity in the primary topology, the node propagates this information to the other nodes. Upon receiving the propagated failure information, each node deactivates the
30 physical layers of the primary set links that are connected to the node. As a result, the endpoints rapidly detect a local loss of connectivity and cause upper-layer protocols to perform a topology change, i.e., a diversion of traffic to the backup set of links.

EX1004, 13.

46. The Applicant noted that claim 24 was unchanged in their remarks, however, the claim was amended as shown below. EX1004, 161, 163-164.

24. (Original) The network according to claim 23, wherein ~~one of the~~ a first network nodes node is arranged to detect the failure and to send a message indicating the failure to ~~either~~ a second network nodes node, and wherein ~~another of the~~ a third network nodes node, which provides the physical layer resources to ~~the one or more of the~~ links in the first set, is arranged to deactivate the physical layer resources responsively to accepting the message.

EX1004, 161.

47. On June 16, 2010, the Examiner issued a Final Office Action rejecting claims 1-4 and 6-28 under 35 U.S.C. §§102 and 103 based on the same prior art references cited in the Non-Final Office Action. EX1004, 169-182.

48. In response, the Applicant filed a Request for Continued Examination, canceled claims 13-22, added new claims 29-36, and amended claim 23. EX1004, 196-207. The Applicant argued that “claim 1 requires that physical layers of all links connected [to] all other nodes other than the detecting node are to be deactivated.” EX1004, 206.

49. On October 12, 2011, the Examiner issued a Non-Final Office Action rejecting claims 7-10, 12, 24-28, and 33-36 under 35 U.S.C. §112, second paragraph, as being indefinite. EX1004, 231-235.

50. In response, the Applicant amended claims 7-10, 12, 24-28, and 33-36 to provide antecedent basis. EX1004, 244-252. The Examiner issued a Notice of

Allowance on January 18, 2012 and the '637 Patent issued on June 12, 2012 with claims 1-25. EX1001, Cover; EX1004, 256, 278.

51. However, as I explain below, the purported allowable subject matter would have been obvious to a POSITA, at the time of the alleged invention, in view of certain prior art references that the Examiner never considered.

V. Claim Construction

52. I have reviewed the claims, specification, and file history of the '637 patent.

53. I have applied the ordinary and customary meaning of the claim terms as understood by a POSITA in my analysis provided herein, as I have concluded that the claim terms in the '637 Patent do not need to be construed.

VI. Level of Ordinary Skill in the Art

54. In my opinion, a person of ordinary skill in the art ("POSITA") in the field of the '637 patent as of its effective filing date of November 16, 2005, would have had a bachelor's degree in computer science, computer engineering, electrical engineering, or an equivalent field, or approximately two years of experience working in the field of information technology and networking. In my opinion, additional education could make up for less practical experience, and vice versa.

55. I met these requirements for POSITA at the time of the effective filing date of the '637 patent.


VII. Overview of the Prior Art

A. Mitchell

56. Mitchell is directed to a method for a fast failover to a second set of links following a failure in the first set of links in a communications network.

EX1005, Abstract.

57. Figure 4 below illustrates an example of this failover method in a data processing system. EX1005, FIG. 4, 7:34-36 (“FIG. 4 illustrates a data processing system including a primary Ethernet switch network element according to an embodiment of the present invention.”).



EX1005, FIG. 4 (annotated).

58. First, a **failure** occurs in the *link* between network switch 410a and core router 404a. This link is in the *group of primary links*. The **failure** is then

detected by network switch 410a, which propagates the **failure** to all network nodes downstream of the failed *link* (e.g., to network switch 410a, primary Ethernet switch 412a, and endstations 414a-414n). EX1005, 7:36-55 (“In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404 a and 404 b), primary and secondary network switches (410 a and 410 b), and primary and secondary Ethernet switches (412 a and 412 b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414 a-414 n similarly coupled to primary Ethernet switch 412 a via one of a plurality of primary downstream links 416 a-416 n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412 b via one of a plurality of secondary downstream links 418 a-418 n (e.g., via a secondary network interface card or host bus adapter) as shown. Using one or more embodiments of the present invention, a failure such as that illustrated on a link between primary network switch 410 a and core router 404 a may be propagated downstream to one or more of endstations 414. Endstations 414 may include any of a number of network elements (e.g., management modules, processor blades, processor modules, or the like).”).

59. Then, the physical layer of the *group of primary links* is **deactivated** (e.g., *disables*, which is indicated by **X**'s). Next, Mitchell's system quickly fails

over to the *backup set of links* to continue providing the communication services.

EX1005, 3:52-63 (“Using embodiments of the present invention, a link failure may be propagated within a bounded or predetermined period of time. For example, where embodiments of the present invention are implemented using predominantly special-purpose hardware (e.g., application specific integrated circuits, special-purpose processors, programmable logic devices, or the like) or specific (e.g., vectored) interrupts, a link between a network element and a downstream portion of a communications network may be disabled within a period of time substantially less than or equal to 50 milliseconds of detecting a failure of another link between the network element and an upstream portion of the communications network.”), Abstract (“A method and system for fast link failover is disclosed. According to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link or group of links between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 2:24-33 (“Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to

a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained.”).


60. Though the failed *link* in Figure 4 is the *link* between network switch 410a and core router 404a, a **failure** can occur in any of the links in the *group of primary links*. As shown in Figure 2 below, a **failure** occurs in *link 108a* between *primary switch 104a* and upstream network 102. This **failure** results in the propagation of the failure and **deactivation** of the downstream *links 110a-110n* (indicated by **X**'s) and a failover to the *backup set of links*. EX1005, 4:31-6:5 (“FIG. 2 illustrates a data processing system including a primary switch network element according to an embodiment of the present invention. As previously described with respect to FIG. 1, data processing system 200 of the illustrated embodiment of FIG. 2 includes an upstream communications network portion 102, a primary switch 104 a, a secondary switch 104 b, and a number of redundantly-linked network elements (e.g., multi-homed endstations 106 a, 106 b . . . 106 n). Upstream communications network portion 102, (e.g., a network core, wide, metropolitan, or local area network, or the like) including any of a number of network elements is coupled to primary switch 104 a using a primary upstream link 108 a and to secondary switch 104 b using a secondary upstream link 108 b. Multi-

homed endstations 106 a, 106 b . . . 106 n are each similarly coupled to primary switch 104 a via one of a plurality of primary downstream links 110 a-110 n (e.g., via a primary network interface card or host bus adapter) and to secondary switch 104 b via one of a plurality of secondary downstream links 112 a-112 n (e.g., via a secondary network interface card or host bus adapter) as shown. While links 108, 110, and 112 have been illustrated as direct connections between network elements (e.g., upstream communications network portion 102, primary and secondary switches 104, and endstations 106) in alternative embodiments, any of links 108, 110, and 112 may be wireless (e.g., using electro-magnetic, optical, infrared and/or acoustic signals or transmission media). For example, one or more of primary downstream links 110 and secondary downstream links are implemented in one embodiment of the present invention using a wireless local area network (e.g., IEEE 802.11x standard) communication protocol. It will be appreciated that the use of the terms “upstream” and “downstream” within the present description is relative based upon the particular network element considered. For example, a link between primary switch 104 a and endstation 106 a is considered a “downstream” link from the perspective of primary switch 104 a and an “upstream” link from the perspective of endstation 106 a. Similarly, primary switch 104 a may be considered an “upstream” network element from the perspective of endstation 106 a and a “downstream” network element from the perspective of upstream communications

network portion 102. Moreover, a direct physical or virtual link or path need not exist between two network elements for an upstream/downstream relationship to exist and no network element need be coupled to a link for it to be considered “upstream” or “downstream”. Consequently, even after primary downstream links 110 are disabled, an endstation 106 is still considered a downstream network element with respect to primary switch 104 a. In the illustrated embodiment of FIG. 2, primary switch 104 a includes a configuration interface 202 and a link failure propagation module 204. While configuration interface 202 and link failure propagation module 204 have been illustrated in FIG. 2 as included within primary switch 104 a, in an alternative embodiment of the present invention one or more of configuration interface 202 and link failure propagation module 204 may be stored at or executed from a switch or other network element within data processing system 200 which is directly or indirectly coupled to one or more of endstations 106 a-106 n. In one embodiment, configuration interface 202 is used to provide a user interface for configuring link failure propagation module 204 and the operation thereof and link failure propagation module 204 is used to propagate link failure and/or link state as further described herein. According to one embodiment, configuration interface 202 is configured to receive data from a user specifying various configuration parameters such as whether or not link failure propagation is to be enabled or not, and if enabled, how link failure propagation is to be

performed (e.g., automatically or on demand, for one or more individually specified ports, for all ports associated with one or more specified virtual networks (e.g., virtual local area or storage area networks), or for all ports of a switch). In another embodiment of the present invention, configuration parameters additionally include data specifying what action is to be taken when a previously failed link (or a new link) becomes operational. For example, one or more associated endstations may alternatively “fail back” to a previous primary link which becomes operational again following a failure or the endstation(s) may continue to use a “failed over” secondary link or set of links. In operation, data is transmitted between one or more of multi-homed endstations 106 a-106 n and upstream communications network portion 102 using primary downstream links 110 a-110 n, primary switch 104 a, and primary upstream link 108 a. Following a failure of primary upstream link 108 a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104 a) link failure propagation module 204 of primary switch 104 a is used to detect the failure and responsively disable one or more of primary downstream links 110 a-110 n. According to the illustrated embodiment of FIG. 2, all downstream links or ports of primary switch 104 a are disabled in response to the detection of a failure of primary upstream link 108 a. In another embodiment, only those downstream links or ports which are individually predetermined or identified using configuration

interface 202 are disabled, and in yet another embodiment, only those downstream links or ports associated with a virtual network such as a VLAN or VSAN which was predetermined or identified using configuration interface 202 are disabled. To disable one or more links according to embodiments of the present invention, any of a number of techniques may be implemented. For example, according to one embodiment of the present invention, a physical layer protocol circuit or “PHY” associated with one or more ports and/or links is disabled using an administrative command, causing all ports and/or links associated with the PHY to become disabled.”).



EX1005, FIG. 2 (annotated).

VIII. Specific Grounds for Challenge

A. Ground 1: Mitchell Renders Obvious Claims 1, 7, 9-11, 13, 15-17, 19, 21, 23, and 25

1. Independent Claim 1


i. [1pre] A method for handling a communication failure in a network, comprising:

61. To the extent the preamble may be interpreted to be a claim limitation, Mitchell teaches the preamble.

62. For instance, Figure 4 depicts a network that includes various network elements (*e.g.*, network core 402, core routers 404a and 404b, network switches 410a and 410b, primary ethernet switch 412a, secondary ethernet switch 412b) that are connected to endpoints (*e.g.*, multi-homed endstations 414a-414n). EX1005, FIG. 4, 7:34-55 (“FIG. 4 illustrates a data processing system including a primary Ethernet switch network element according to an embodiment of the present invention. In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404 a and 404 b), primary and secondary network switches (410 a and 410 b), and primary and secondary Ethernet switches (412 a and 412 b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414 a-414 n similarly coupled to primary Ethernet switch 412 a via one of a plurality of primary downstream links 416 a-416 n (*e.g.*, via a primary network interface card

or host bus adapter) and to secondary Ethernet switch 412 b via one of a plurality of secondary downstream links 418 a-418 n (e.g., via a secondary network interface card or host bus adapter) as shown. Using one or more embodiments of the present invention, a failure such as that illustrated on a link between primary network switch 410 a and core router 404 a may be propagated downstream to one or more of endstations 414. Endstations 414 may include any of a number of network elements (e.g., management modules, processor blades, processor modules, or the like).”).

63. As I explained above, in Section VIII.A, Figure 4 of Mitchell shows an example of how Mitchell’s fast link failover system handles a link failure, for example, by disabling downstream links in the communications network.



EX1005, FIG. 4 (annotated).

64. Specifically, the data processing system described in Mitchell handles a communication failure (*e.g.*, link failure) by propagating failure information downstream from the failed link and deactivating the other links associated with the failed link. EX1005, 7:36-53. The resulting failover to the second set of links is a fast failover. EX1005, Abstract (“A method and system for fast link failover is disclosed.”), 1:7-9 (“The present invention relates generally to communications networks and more particularly to a method and system for fast link failover.”), 2:23-42 (“Disclosed is a method and system for fast link failover. Using one or more embodiments of the present invention, network connectivity (*e.g.*, data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained. According to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the


communications network using one or more alternate links.”), 3:52-63 (“Using embodiments of the present invention, a link failure may be propagated within a bounded or predetermined period of time. For example, where embodiments of the present invention are implemented using predominantly special-purpose hardware (e.g., application specific integrated circuits, special-purpose processors, programmable logic devices, or the like) or specific (e.g., vectored) interrupts, a link between a network element and a downstream portion of a communications network may be disabled within a period of time substantially less than or equal to 50 milliseconds of detecting a failure of another link between the network element and an upstream portion of the communications network.”), 4:10-62 (“As the link state or link failure information is propagated downstream, downstream network elements with the capability to switch to alternate (e.g., redundant) links may do so, preserving higher-level communication channels between the downstream and upstream portions of the communications networks and providing seamless failover. Within the present description, the term “downstream” is intended to indicate in a direction from a network's core to a network's edge or towards a network's edge, the term “upstream” by contrast is intended to indicate in a direction from a network's edge to a network's core or towards a network's core. Further within the present description, the term “endstation” is intended to indicate a network element (e.g., a file/data server, personal computer, or other data

processing system) residing at the edge of a communications network, the term “switch” is intended to indicate a data link layer network element (e.g., an Ethernet switch), the term “link” is intended to indicate a data link layer connection and may include one or more logical sub-links (e.g., an “Ether-Channel” or “Port-Channel”), and the term “failure” is intended to indicate the loss of a data link layer link. FIG. 2 illustrates a data processing system including a primary switch network element according to an embodiment of the present invention. As previously described with respect to FIG. 1, data processing system 200 of the illustrated embodiment of FIG. 2 includes an upstream communications network portion 102, a primary switch 104 a, a secondary switch 104 b, and a number of redundantly-linked network elements (e.g., multi-homed endstations 106 a, 106 b . . . 106 n). Upstream communications network portion 102, (e.g., a network core, wide, metropolitan, or local area network, or the like) including any of a number of network elements is coupled to primary switch 104 a using a primary upstream link 108 a and to secondary switch 104 b using a secondary upstream link 108 b. Multi-homed endstations 106 a, 106 b . . . 106 n are each similarly coupled to primary switch 104 a via one of a plurality of primary downstream links 110 a-110 n (e.g., via a primary network interface card or host bus adapter) and to secondary switch 104 b via one of a plurality of secondary downstream links 112 a-112 n (e.g., via a secondary network interface card or host bus adapter) as shown. While links 108,

110, and 112 have been illustrated as direct connections between network elements (e.g., upstream communications network portion 102, primary and secondary switches 104, and endstations 106) in alternative embodiments, any of links 108, 110, and 112 may be wireless (e.g., using electro-magnetic, optical, infrared and/or acoustic signals or transmission media). For example, one or more of primary downstream links 110 and secondary downstream links are implemented in one embodiment of the present invention using a wireless local area network (e.g., IEEE 802.11x standard) communication protocol.”), 5:30-45 (“According to one embodiment, configuration interface 202 is configured to receive data from a user specifying various configuration parameters such as whether or not link failure propagation is to be enabled or not, and if enabled, how link failure propagation is to be performed (e.g., automatically or on demand, for one or more individually specified ports, for all ports associated with one or more specified virtual networks (e.g., virtual local area or storage area networks), or for all ports of a switch). In another embodiment of the present invention, configuration parameters additionally include data specifying what action is to be taken when a previously failed link (or a new link) becomes operational. For example, one or more associated endstations may alternatively “fail back” to a previous primary link which becomes operational again following a failure or the endstation(s) may continue to use a “failed over” secondary link or set of links.”), 6:6-36 (“While the

illustrated embodiment of FIG. 2 includes a single primary upstream link 108 a and secondary upstream link 108 b, in alternative embodiments of the present invention, multiple links between upstream communications network portion 102 and primary switch 104 a and/or secondary switch 104 b may be used. According to one embodiment, two or more primary upstream links between primary switch 104a and upstream communications network portion 102 are provided. Such primary upstream links may operate simultaneously or in a fail-over configuration with one another. Consequently, a failure of one or more of the primary upstream links need not cause any of downstream links 110 to be immediately disabled. For example, a downstream link 110 may be disabled only upon the failure of all primary upstream links, upon the failure of a predetermined or defined number or proportion of primary upstream links, or upon available bandwidth between primary switch 104a and upstream communications network portion 102 falling below a predetermined or defined threshold level due to the failure of one or more primary upstream links. Once a port and/or link (e.g., one or more of downstream links 110) is disabled, an associated multi-homed endstation may failover to a corresponding secondary downstream link 112 by activating a network interface associated with the secondary downstream link and deactivating a network interface associated with the failed primary downstream link. Data is then transmitted between the multi-homed endstation which has failed-over and

upstream communications network portion 102 via an associated secondary downstream link 112, a secondary switch 104b, and a secondary upstream link 108b.”), 7:34-55, FIGS. 2-4.



EX1005, FIG. 4 (annotated).

- ii. [1a] provisioning different first and second instances of a **multi-point-to-multi-point (MP-MP)** communication service over respective **first** and **second alternative sets of links** that connect a **plurality of endpoints** in the network,

65. Mitchell renders obvious claim limitation [1a].

66. Mitchell provides services for multipoint-to-multipoint communication. For example, Mitchell discloses **multi-homed endstations 414a-414n** connected to each other through “redundantly-linked network elements.”

EX1005, 1:25-29 (“[R]edundantly-linked network elements are said to be ‘multi-


homed’.”), 7:36-49. Data is transmitted between these *endstations*. See EX1005, 5:46-50 (“[D]ata is transmitted between one or more of *multi-homed endstations 106a-106n* and upstream communications network portion 102 using primary downstream links 110a-110n, primary switch 104a, and primary upstream link 108a.”), 6:26-36.

67. Mitchell gives multiple examples of communication services, including virtual storage area networks (VSANs), virtual local area networks (VLANs), wireless local area network (LAN) services using IEEE 802.11x, and Ethernet based services. EX1005, 4:58-62 (“For example, one or more of primary downstream links 110 and secondary downstream links are implemented in one embodiment of the present invention using a wireless local area network (e.g., IEEE 802.11x standard) communication protocol.”), 5:35-37 (“...for all ports associated with one or more specified virtual networks (e.g., virtual local area or storage area networks)...”), 5:59-65 (“In another embodiment, only those downstream links or ports which are individually predetermined or identified using configuration interface 202 are disabled, and in yet another embodiment, only those downstream links or ports associated with a virtual network such as a VLAN or VSAN which was predetermined or identified using configuration interface 202 are disabled.”), 6:37-44 (“Using one embodiment of the present invention, a network element (e.g., a datalink layer-capable Ethernet switch) is configured to

track the state of a virtual network (e.g., a VLAN or VSAN) on all identified upstream interfaces (e.g., links, ports, interface cards, or the like).”), 7:34-49 (“FIG. 4 illustrates a data processing system including a primary Ethernet switch network element according to an embodiment of the present invention. In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404a and 404 b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414 n similarly coupled to primary Ethernet switch 412 a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown.”).

68. As shown in Figure 4 below, Mitchell describes “a plurality of *multi-homed endstations 414 a-414 n*.” EX1005, 7:36-49. *Multi-homed endstations 414a-414n* are “network elements residing at the edge or terminating point of a communications network” that are redundantly-linked. EX1005, 1:25-29. The *endstations 414a-414n* transmit data between them using the redundant links. EX1005, 5:46-50 (“In operation, data is transmitted between one or more of multi-

homed endstations 106a-106n and upstream communications network portion 102 using primary downstream links 110a-110n, primary switch 104a, and primary upstream link 108a.”), 6:26-36 (“Once a port and/or link (e.g., one or more of downstream links 110) is disabled, an associated multi-homed endstation may failover to a corresponding secondary downstream link 112 by activating a network interface associated with the secondary downstream link and deactivating a network interface associated with the failed primary downstream link. Data is then transmitted between the multi-homed endstation which has failed-over and upstream communications network portion 102 via an associated secondary downstream link 112, a secondary switch 104b, and a secondary upstream link 108b.”).



EX1005, FIG. 4 (annotated).

69. As shown in Figure 4 above, Mitchell describes two different instances of a communication service operating over two sets of links (“*group of links*” and “*set of links*”). For example, Figure 4 of Mitchell shows a first communication service instance that includes the *group of links* (including *links* 416a-416n), core router 404a, network switch 410a, primary ethernet switch 412a, endstations 414a-414n, and network core 402 (including the network elements within the network core 402). EX1005, FIG. 4. Similarly, Figure 4 shows a second communication service instance that includes the *set of links* (including *links* 418a-418n), core router 404b, network switch 410b, primary ethernet switch 412b, endstations 414a-414n, and network core 402 (including the network elements within the network core 402). EX1005, FIG. 4.

70. Thus, a POSITA would have understood that Mitchell describes that the different core routers (e.g., core routers 404a and 404b), network switches (e.g., network switches 410a and 410b), and primary ethernet switches (e.g., primary ethernet switches 412a and 412b) are part of two communication service instances that share an upstream network (e.g., network core 402). Thus, a POSITA would have recognized network core 402 share a number of network elements between its first and second communication service instances.

71. Contextually, a POSITA would have also recognized that the network core 402 in Figure 4 of Mitchell should be understood with Figure 1, where

Mitchell discloses that its upstream network 102 may include “any of a number of network elements.” EX1005, 1:36-38. A POSITA, when reading Mitchell, would have understood upstream network 102 of Figure 1 is analogous to network core 402 of Figure 4. Similarly, such a POSITA would have understood that Mitchell’s “network elements” may also include devices, such as a “switch, router, hub, gateway or other redistribution point.” EX1005, 1:20-25. As a result, a POSITA would have immediately recognized that the network core 402 further includes network elements, such as a router, switch, etc.

72. Mitchell discloses that the *first set of links* includes *links 416a-416n* and other *links* between core router 404a, network switch 410a, and primary Ethernet switch. This “*group of links*,” are coupled to core router 404a, network switch 410a, and primary Ethernet switch, and make up the first instance of a multipoint-to-multipoint communication service that traverses each network element. EX1005, 3:35-43, 5:46-55 (discussing transmission of data across the *multi-homed endstations*), 7:36-49.

73. Mitchell further discloses that the *second set of links* include the *links 418a-418n* and other *links disposed in* between core router 404b, network switch 410b, and secondary Ethernet switch. This “*set of links*,” along with core router 404b, network switch 410b, and secondary Ethernet switch, and make up a second instance of a multipoint-to-multipoint communication service that traverses each of


the various network elements. EX1005, 2:24-33 (“Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained.”), 3:35-43 (“According to one embodiment of the present invention, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected and a second link (e.g., a group of links which are logically related to the first link) between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 5:46-55 (“In operation, data is transmitted between one or more of multi-homed endstations 106a-106n and upstream communications network portion 102 using primary downstream links 110a-110n, primary switch 104a, and primary upstream link 108a. Following a failure of primary upstream link 108a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary

switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n.”), 7:34-49 (“FIG. 4 illustrates a data processing system including a primary Ethernet switch network element according to an embodiment of the present invention. In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404a and 404b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414n similarly coupled to primary Ethernet switch 412a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown.”).

iii. [1b] each of the sets traversing a plurality of network nodes, which provide physical layer resources for operating the links;

74. Mitchell teaches claim limitation [1b].

75. As shown below in Figure 4, each of Mitchell’s sets of links traverses multiple network nodes. EX1005, FIG. 4.



EX1005, FIG. 4 (annotated).

76. For example, the *group of links* traverse core router 404a, network switch 410a, and primary Ethernet switch 412a. EX1005, 7:36-49. The *group of links* traversing core router 404a, network switch 410a, and primary Ethernet switch 412a are associated with the first communication service instance. EX1005, 7:36-49. In addition, the *set of links* traverse core router 404b, network switch 410b, and secondary Ethernet switch 412b. EX1005, 7:36-49. The *set of links* traversing core router 404b, network switch 410b, and secondary Ethernet switch 412b are associated with the second communication service instance. EX1005, 7:34-49. (“FIG. 4 illustrates a data processing system including a primary Ethernet switch network element according to an embodiment of the present invention. In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented

including a core network 402, primary and secondary core routers (404a and 404b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414n similarly coupled to primary Ethernet switch 412a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown.”), FIG. 4.

77. The *group of links* and the *set of links* described in Mitchell provide physical layer resources for operating the links. A failure of a link may be caused by the failure of physical resources, such as the physical link hardware, network interface card, or switch. EX1005, 5:50-52 (“Following a failure of primary upstream link 108 a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104 a)...”). Additionally, Mitchell describes deactivating the physical layer resources, including ports, following a link failure. The ports on the nodes are used for operating the links. EX1005, 5:56-6:5 (“According to the illustrated embodiment of FIG. 2, all downstream links or ports of primary switch 104 a are disabled in response to the detection of a failure of primary upstream link 108 a. In another embodiment, only those downstream links

or ports which are individually predetermined or identified using configuration interface 202 are disabled, and in yet another embodiment, only those downstream links or ports associated with a virtual network such as a VLAN or VSAN which was predetermined or identified using configuration interface 202 are disabled. To disable one or more links according to embodiments of the present invention, any of a number of techniques may be implemented. For example, according to one embodiment of the present invention, a physical layer protocol circuit or “PHY” associated with one or more ports and/or links is disabled using an administrative command, causing all ports and/or links associated with the PHY to become disabled.”).

iv. [1c] providing the communication service to the endpoints over the first set of links using the first instance;

78. Mitchell teaches claim limitation [1c].

79. As I explained above in Section IX.A.1.ii [1a], Mitchell’s group of links (*e.g.*, the **group of links**) provide the communication service to the endpoints (*e.g.*, **multi-homed endstations 414a-414n**) using the first instance. The first instance is made up of the **group of links** and the network elements in the network core 402, core router 404a, network switch 410a, and primary Ethernet switch 412a. EX1005, 7:36-49 (“In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and

secondary core routers (404a and 404b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414n similarly coupled to primary Ethernet switch 412a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown.”), 4:15-62 (“Within the present description, the term ‘downstream’ is intended to indicate in a direction from a network's core to a network's edge or towards a network's edge, the term ‘upstream’ by contrast is intended to indicate in a direction from a network's edge to a network's core or towards a network's core. Further within the present description, the term ‘endstation’ is intended to indicate a network element (e.g., a file/data server, personal computer, or other data processing system) residing at the edge of a communications network, the term ‘switch’ is intended to indicate a data link layer network element (e.g., an Ethernet switch), the term ‘link’ is intended to indicate a data link layer connection and may include one or more logical sub-links (e.g., an ‘Ether-Channel’ or ‘Port-Channel’), and the term ‘failure’ is intended to indicate the loss of a data link layer link. FIG. 2 illustrates a data processing system including a primary switch network element according to

an embodiment of the present invention. As previously described with respect to FIG. 1, data processing system 200 of the illustrated embodiment of FIG. 2 includes an upstream communications network portion 102, a primary switch 104 a, a secondary switch 104 b, and a number of redundantly-linked network elements (e.g., multi-homed endstations 106 a, 106 b . . . 106 n). Upstream communications network portion 102, (e.g., a network core, wide, metropolitan, or local area network, or the like) including any of a number of network elements is coupled to primary switch 104a using a primary upstream link 108 a and to secondary switch 104b using a secondary upstream link 108b. Multi-homed endstations 106a, 106b . . . 106n are each similarly coupled to primary switch 104a via one of a plurality of primary downstream links 110a-110n (e.g., via a primary network interface card or host bus adapter) and to secondary switch 104b via one of a plurality of secondary downstream links 112a-112n (e.g., via a secondary network interface card or host bus adapter) as shown. While links 108, 110, and 112 have been illustrated as direct connections between network elements (e.g., upstream communications network portion 102, primary and secondary switches 104, and endstations 106) in alternative embodiments, any of links 108, 110, and 112 may be wireless (e.g., using electro-magnetic, optical, infrared and/or acoustic signals or transmission media). For example, one or more of primary downstream links 110 and secondary downstream links are implemented in one embodiment of the present invention

using a wireless local area network (e.g., IEEE 802.11x standard) communication protocol.”), 5:33-65 (“...how link failure propagation is to be performed (e.g., automatically or on demand, for one or more individually specified ports, for all ports associated with one or more specified virtual networks (e.g., virtual local area or storage area networks), or for all ports of a switch). In another embodiment of the present invention, configuration parameters additionally include data specifying what action is to be taken when a previously failed link (or a new link) becomes operational. For example, one or more associated endstations may alternatively “fail back” to a previous primary link which becomes operational again following a failure or the endstation(s) may continue to use a “failed over” secondary link or set of links. In operation, data is transmitted between one or more of multi-homed endstations 106a-106n and upstream communications network portion 102 using primary downstream links 110a-110n, primary switch 104a, and primary upstream link 108a. Following a failure of primary upstream link 108a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n. According to the illustrated embodiment of FIG. 2, all downstream links or ports of primary switch 104a are disabled in response to the detection of a failure of primary upstream link 108a. In another

embodiment, only those downstream links or ports which are individually predetermined or identified using configuration interface 202 are disabled, and in yet another embodiment, only those downstream links or ports associated with a virtual network such as a VLAN or VSAN which was predetermined or identified using configuration interface 202 are disabled.”), 6:26-44 (“Once a port and/or link (e.g., one or more of downstream links 110) is disabled, an associated multi-homed endstation may failover to a corresponding secondary downstream link 112 by activating a network interface associated with the secondary downstream link and deactivating a network interface associated with the failed primary downstream link. Data is then transmitted between the multi-homed endstation which has failed-over and upstream communications network portion 102 via an associated secondary downstream link 112, a secondary switch 104b, and a secondary upstream link 108b. Using one embodiment of the present invention, a network element (e.g., a datalink layer-capable Ethernet switch) is configured to track the state of a virtual network (e.g., a VLAN or VSAN) on all identified upstream interfaces (e.g., links, ports, interface cards, or the like). In the described embodiment, a virtual network can be available on multiple (but not necessarily all) upstream interfaces and can be associated with one or more downstream link(s).”)

v. [1d] upon detecting a **failure** in the **first set of links** by
a **detecting node** of the plurality of network nodes:


80. Mitchell teaches claim limitation [1d].

81. Mitchell discloses detecting a failure in the first set of links.

Specifically, Mitchell discloses “a method is provided in which a **failure** of *a first link* between a network element and an upstream portion of a communications network *is detected*.” EX1005, Abstract; *see also* EX1005, 7:49-53 (“Using one or more embodiments of the present invention, a failure such as that illustrated on a link between primary network switch 410a and core router 404a may be propagated downstream to one or more of endstations 414.”).

82. Mitchell describes *network switch 410a* detecting a failure on the link between core router 404a and *network switch 410a*, as shown below in Figure 4. EX1005, Abstract (“According to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link or group of links between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 7:49-53 (“Using one or more embodiments of the present invention, a failure such as that


illustrated on a link between primary network switch 410 a and core router 404 a may be propagated downstream to one or more of endstations 414.”).



EX1005, FIG. 4 (annotated).


83. Additionally, Mitchell teaches that a failure can occur and be detected in any of the links in the group of links. When *link 108a* fails, primary switch 104a detects the **failure**, as shown below in Figure 2. EX1005, 5:14-29 (“In the illustrated embodiment of FIG. 2, primary switch 104 a includes a configuration interface 202 and a link failure propagation module 204. While configuration interface 202 and link failure propagation module 204 have been illustrated in FIG. 2 as included within primary switch 104 a, in an alternative embodiment of the present invention one or more of configuration interface 202 and link failure propagation module 204 may be stored at or executed from a switch or other

network element within data processing system 200 which is directly or indirectly coupled to one or more of endstations 106 a-106 n. In one embodiment, configuration interface 202 is used to provide a user interface for configuring link failure propagation module 204 and the operation thereof and link failure propagation module 204 is used to propagate link failure and/or link state as further described herein.”), 5:50-55 (“Following a **failure** of primary upstream *link 108 a* (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n.”).



EX1005, FIG. 2 (annotated).

84. Accordingly, based on the combination of the embodiments shown in Figures 2 and 4 of Mitchell, a POSITA would have readily understood that just as *primary switch 104a* in Figure 2 is used to detect a **failure**, it would have been possible for the *core router 404a* in Figure 4 to detect a **failure** in the *link* between the *core router 404a* and the network core 402.



EX1005, FIG. 4 (annotated).

85. Figures 2 and 4 of Mitchell show two examples where link failures are detected at different link locations. EX1005, 5:16-24 (“While configuration interface 202 and link failure propagation module 204 have been illustrated in FIG. 2 as included within primary switch 104a, in an alternative embodiment of the present invention one or more of configuration interface 202 and link failure propagation module 204 may [sic] be stored at or executed from a switch or other

network element within data processing system 200 which is directly or indirectly coupled to one or more of endstations 106a-106n.”). In Figure 2, the failed link is between the upstream network and the first network element. EX1005, FIG. 2. In Figure 4, the failed link is between two network elements. EX1005, FIG. 4.

86. Specifically, Mitchell states that a link failure occurs upstream from the detecting node. EX1005, 3:35-43 (“According to one embodiment of the present invention, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected and a second link (e.g., a group of links which are logically related to the first link) between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 2:43-47 (“The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting.”), 3:9-12 (“On the contrary, it is intended to cover such alternatives, modifications, and equivalents as can be reasonably included within the scope of the invention as defined by the appended claims.”).


87. A POSITA would have recognized that any network element in Figures 2 and 4 can detect an upstream link failure. EX1005, 2:33-42 (“According

to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network using one or more alternate links.”), 3:35-43 (“According to one embodiment of the present invention, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected and a second link (e.g., a group of links which are logically related to the first link) between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 5:46-55 (“In operation, data is transmitted between one or more of multi-homed endstations 106a-106n and upstream communications network portion 102 using primary downstream links 110a-110n, primary switch 104a, and primary upstream link 108a. Following a failure of primary upstream link 108a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n.”), 7:3-

15 (“If a determination is made that link failure propagation is enabled on a system associated with the illustrated process embodiment, a loop is entered in which a failure of any upstream link of the associated network element is detected (process block 306). Once an upstream link failure has been detected, a determination is made whether link failure propagation is to be performed on-demand or automatically (process block 308). Where link failure propagation is to be performed on-demand, a determination is made whether or not the current link failure is to be propagated (process block 310) (e.g., in response to a response to a user prompt or lack thereof, or based upon additional processing or analysis of system or environmental attributes).”)

88. A POSITA therefore would have understood that it would be possible, for example, for *core router 404a* to detect a **failure** in the link between the *core router 404a* and network core 402 (as shown in Figure 4 below), based on a combination of Figures 2 and 4. EX1005, 7:3-15 (“If a determination is made that link failure propagation is enabled on a system associated with the illustrated process embodiment, a loop is entered in which ***a failure of any upstream link of the associated network element*** [e.g., *core router 404a*] is detected (process block 306).”), 1:19-25 (“For example, redundant network links or connections are frequently used to couple a single network element or node, (e.g., a client, server, or other host data processing system, a communications network appliance, or a

switch, router, hub, gateway or other redistribution point) to one or more communications networks or portions thereof (e.g., network segments).”).



EX1005, FIG. 4 (annotated).

89. Hence, a POSITA would have made such a modification based on the different embodiments of Mitchell because it would have provided a more flexible communications network that can readily respond to failure at any link location. And a POSITA would have had a reasonable expectation of making the modification because Mitchell itself teaches that the failure propagation module can be carried out from any one of the network elements. *See* EX1005, 5:16-24.

vi. [1e] propagating **failure** information by the **detecting node** to **each node of the other nodes** of the plurality of network nodes traversed by **the first set of links**;

[1f] for **each node of the other nodes** of the plurality of network nodes traversed by **the first set of links**;

[1g] receiving the **failure** information; and


90. Mitchell teaches each of claim limitations [1e], [1f], and [1g].

91. Mitchell describes a link failure propagation method where the **detecting node** detects a failure in a link and propagates failure information to the nodes downstream from the failed link. EX1005, 4:10-15 (“As the link state or link failure information is propagated downstream, downstream network elements with the capability to switch to alternate (e.g., redundant) links may do so, preserving higher-level communication channels between the downstream and upstream portions of the communications networks and providing seamless failover.”), 2:23-33 (“Disclosed is a method and system for fast link failover. Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with

the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained.”)

92. The *failure* propagation is carried out by a link failure propagation module, shown as stored in *primary switch 104a* in Figure 2 below. EX1005, 5:50-55 (“Following a failure of primary upstream link 108a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n.”),

FIG. 2.



EX1005, FIG. 2 (annotated).

93. Though illustrated as part of *primary switch 104a* in Figure 2, Mitchell explains that any network element may store the link failure propagation


module. EX1005, 5:13-29 (“In the illustrated embodiment of FIG. 2, primary switch 104 a includes a configuration interface 202 and a link failure propagation module 204. While configuration interface 202 and link failure propagation module 204 have been illustrated in FIG. 2 as included within primary switch 104 a, in an alternative embodiment of the present invention one or more of configuration interface 202 and link failure propagation module 204 may be stored at or executed from a switch or other network element within data processing system 200 which is directly or indirectly coupled to one or more of endstations 106 a-106 n. In one embodiment, configuration interface 202 is used to provide a user interface for configuring link failure propagation module 204 and the operation thereof and link failure propagation module 204 is used to propagate link failure and/or link state as further described herein.”).

94. Therefore, a POSITA would have understood that the link failure propagation module could be stored in *core router 404a*, which would allow *core router 404a* to detect a **failure** in the *link* between *core router 404a* and the network core 402 (as shown in Figure 4 below) and propagate the **failure** information to the downstream network elements traversed by the group of links (e.g., *network switch 410a* and *primary Ethernet switch 412a*). EX1005, 7:49-53 (“Using one or more embodiments of the present invention, a failure such as that

illustrated on a link between primary network switch 410 a and core router 404 a may be propagated downstream to one or more of endstations 414.”),

95. Further, Mitchell discloses “[a]s the link state or link failure information is propagated downstream, downstream network elements with the capability to switch to alternate (e.g., redundant) links may do so, preserving higher-level communication channels between the downstream and upstream portions of the communications networks and providing seamless failover.”

EX1005, 4:10-15. Mitchell explains that “[w]here link failure propagation is to be performed automatically (or on-demand in a particular instance) one or more downstream links of a network element associated with the illustrated process embodiment are disabled (process block 312).” EX1005, 7:25-29; *see also* EX1005, FIG. 4.



EX1005, FIG. 4 (annotated).

96. During the failure propagation described in Mitchell, the links downstream of the failed link (e.g., the *link* between *core router 404a* and the network core 402) in the *group of links* are deactivated. For the node to deactivate the links, the *failure* information must be received by the downstream nodes (e.g., *network switch 410a* and *primary Ethernet switch 412a*). EX1005, 7:25-29 (“Where link failure propagation is to be performed automatically (or on-demand in a particular instance) one or more downstream links of a network element associated with the illustrated process embodiment are disabled (process block 312).”), Abstract (“According to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link or group of links between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 2:24-33 (“Disclosed is a method and system for fast link failover. Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with


the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained.”), 7:36-53 (“In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404a and 404b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414n similarly coupled to primary Ethernet switch 412a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown. Using one or more embodiments of the present invention, a failure such as that illustrated on a link between primary network switch 410a and core router 404a may be propagated downstream to one or more of endstations 414.”).

vii. [1h] deactivating a physical layer of the first set of links connected thereto, thereby causing a loss of connectivity in the first set of links; and

97. Mitchell renders obvious claim limitation [1h].


98. A POSITA would have understood that the physical layer of the links downstream of the network elements traversed by the group of links (e.g., *core*

router 404a, *network switch 410a*, and *primary Ethernet switch 412a*) are *deactivated*, causing a loss of connectivity in the *group of links*.



EX1005, FIG. 4 (annotated).

99. Although not explicitly described with respect to Figure 4, Mitchell describes that the links downstream from the failed primary link 108a are disabled after propagation of the failure, as shown below in Figure 2. EX1005, 5:55-67 (“According to the illustrated embodiment of FIG. 2, *all downstream links or ports of primary switch 104a are disabled in response to the detection of a failure of primary upstream link 108a.*”), FIG. 2.



EX1005, FIG. 2 (annotated).

100. A POSITA would have understood that this disclosure applies to the links downstream of the failed link in Figure 4 as well. EX1005, 5:67-6:5 (“For example, according to one embodiment of the present invention, a physical layer protocol circuit or ‘PHY’ associated with one or more ports and/or links is disabled using an administrative command, causing all ports and/or links associated with the PHY to become disabled.”), 7:3-33 (“If a determination is made that link failure propagation is enabled on a system associated with the illustrated process embodiment, a loop is entered in which a failure of any upstream link of the associated network element is detected (process block 306). Once an upstream link failure has been detected, a determination is made whether link failure propagation is to be performed on-demand or automatically (process

block 308). Where link failure propagation is to be performed on-demand, a determination is made whether or not the current link failure is to be propagated (process block 310) (e.g., in response to a response to a user prompt or lack thereof, or based upon additional processing or analysis of system or environmental attributes). According to one embodiment of the present invention, a determination of whether or not a particular link failure is to be propagated on-demand is made by determining (e.g., using a timer) whether any user intervention has occurred within a predetermined amount of time from the occurrence of the link failure. If no user input or intervention is received within the predetermined amount of time in the described embodiment, a default behavior or action (e.g., propagation of the link failure) is performed. Otherwise, a behavior or action specified by the user input or intervention is performed. Where link failure propagation is to be performed automatically (or on-demand in a particular instance) one or more downstream links of a network element associated with the illustrated process embodiment are *disabled* (process block 312). Thereafter, or following a determination that a particular link failure instance is not to be propagated on-demand, the aforementioned loop is re-entered in which a failure of any upstream link is detected (process block 306).”), 6:26-31 (“Once a port and/or link (e.g., one or more of downstream links 110) is disabled, an associated multi-homed endstation may failover to a corresponding secondary downstream link 112

by activating a network interface associated with the secondary downstream link and deactivating a network interface associated with the failed *primary downstream link*.”), FIGS. 3-4.

101. In other words, it causes a loss of connectivity in the “*group of links*.”

Hence, Mitchell renders obvious limitation [1h].

viii. [1i] responsively to sensing the loss of connectivity, resuming the communication service over the second instance by automatically transferring communication among the endpoints to the second set of links.

102. Mitchell teaches claim limitation [1i].


103. Responsively to the deactivation of the physical layer, Mitchell discloses a failover to the second *set of links*, resuming the communication service. EX1005, 6:26-36 (“Once a port and/or link (e.g., one or more of downstream links 110) is disabled, an associated *multi-homed endstation* may failover to a corresponding *secondary downstream link 112* by activating a network interface associated with the *secondary downstream link* and deactivating a network interface associated with the failed primary downstream link. Data is then transmitted between the multi-homed endstation which has failed-over and upstream communications network portion 102 via an associated secondary downstream link 112, a secondary switch 104b, and a secondary upstream link 108b.”), 2:23-42 (“Disclosed is a method and system for fast link failover. Using

one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to *failover to alternate, redundant links* such that the state of one or more connections or communications channels with the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained. According to one embodiment, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected, and a second link between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network using one or more alternate links.”), 4:10-15 (“As the link state or link failure information is propagated downstream, downstream network elements with *the capability to switch to alternate (e.g., redundant) links* may do so, preserving higher-level communication channels between the downstream and upstream portions of the communications networks and providing seamless failover.”).

104. Mitchell explains that after failing over to the *secondary downstream links*, the communication service resumes between the *endpoints* (i.e., *endstations*

414a-414n) over the *set of links* (e.g., *links 418a-418n*). EX1005, 2:23-33

(“Disclosed is a method and system for fast link failover. Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated, thereby enabling downstream network elements not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of one or more connections or communications channels with the upstream portion(s) of a communications network may be preserved and the connection(s)/channel(s) may be maintained.”)



EX1005, FIG. 4 (annotated).

105. Additionally, Mitchell discloses that the failover to the *second set of links* can be automatic. EX1005, 5:30-38 (“According to one embodiment,

configuration interface 202 is configured to receive data from a user specifying various configuration parameters such as whether or not link failure propagation is to be enabled or not, and if enabled, how link failure propagation is to be performed (e.g., automatically or on demand, for one or more individually specified ports, for all ports associated with one or more specified virtual networks (e.g., virtual local area or storage area networks), or for all ports of a switch).”), 7:3-10 (“If a determination is made that link failure propagation is enabled on a system associated with the illustrated process embodiment, a loop is entered in which a failure of any upstream link of the associated network element is detected (process block 306). Once an upstream link failure has been detected, a determination is made whether link failure propagation is to be performed on-demand or automatically (process block 308).”), 7:25-29 (“Where link failure propagation is to be performed automatically (or on-demand in a particular instance) one or more downstream links of a network element associated with the illustrated process embodiment are disabled (process block 312).”), FIG. 3.

2. Independent Claim 16

106. As I explained above in Paragraphs 45-46, the Applicant claimed during prosecution that support for original claims 23 and 24, which correspond to issued claims 16 and 17, could be found in block 76 of Figure 3 and in the below text of the specification.

Embodiments of the present invention provide improved methods for rapidly initiating the diversion of traffic from
25 the failed primary topology to the backup topology. In these
embodiments, when a network node detects a local failure or
other loss of connectivity in the primary topology, the node
propagates this information to the other nodes. Upon receiving
the propagated failure information, each node deactivates the
30 physical layers of the primary set links that are connected to
the node. As a result, the endpoints rapidly detect a local
loss of connectivity and cause upper-layer protocols to perform
a topology change, i.e., a diversion of traffic to the backup
set of links.

EX1004, 13.

107. As I explained in Section VII.B, Applicant amended claims 23-24
(now issued claims 16-17) to obviate an indefiniteness rejection in the first Office
Action. *See e.g.*, EX1004, 160-164. The then-Applicant noted that “[c]laim[] ... 24
... [is] unchanged,” EX1004, 163. But, in the very same Office Action Response
(on the very next page), Applicant contradicted itself, stating: “Applicant *is*
amending claim 24 by deleting the phrase ‘another network node’ and reciting ‘a
first network node’, ‘a second network node’, and a ‘third network node’ to
distinctly point out the subject matter claimed.” EX1004, 164.

108. In support of the amendments that would eventually become issued
claims 16 and 17, Applicant pointed to portions of the specification cited below.

The claimed features recited in claims 1, 13, and 23
[corresponding to issued claim 16] as “propagating ...

failure information to each of the other nodes of the plurality of network nodes”; and “upon receiving the second message, to cause a loss of connectivity in the first set of links” are disclosed in paragraph [0030] and Figure 3, block 76;

EX1004, 163.

109. Applicant further argued that all of the cited prior art references:

fail to show or suggest at least the following features in amended independent claims 1, 13, and 23 [corresponding to issued claim 16]: propagating failure information by a detecting node to each node of the other nodes of the plurality of network nodes traversed by a first set of links; and receiving the failure information at a node and deactivating a physical layer of the first set of links connected thereto, thereby causing a loss of connectivity in the first set of links.

EX1004, 165.

110. Next, Applicant amended claim 24 again, which is now issued claim 17, in an attempt to obviate another 35 U.S.C. §112 indefiniteness rejection.

EX1004, 250 (“The Office Action states the limitation ‘the message’ in line 4 of claim 24 lacks sufficient antecedent basis. The Applicants have amended claim 24 to distinguish a third message from the first message and second message of base claim 23.”).

111. Therefore, in my analysis below, I regard the limitations of claim 16 as similar to those of claim 1. The “first message” recited in [16e], “second message” recited in [16f], and “third message” recited in [17] are treated as the failure information propagated by the detecting node to each of the other network nodes. The “first network node” recited in [17] is treated as the detecting node. “All other nodes” recited in [16e], “another node” recited in [16f], “that node” recited in [16g], “second network node” recited in [17], and “third network node” recited in [17] are treated as the network nodes that receive the propagated failure information from the detecting node.

i. [16pre] A communication network serving a plurality of endpoints, the network comprising:

112. To the extent the preamble may be interpreted to be a claim limitation, Mitchell teaches the preamble.

113. As I explained above in Sections IX.A.1.i, IX.A.1.ii, IX.A.1.iii, and IX.A.1.iv [[1pre], [1a], [1b], and 1[c]], Mitchell teaches a communication network serving *multi-homed endstations 414a-414n*.

ii. [16a] a plurality of links, comprising at least first and second alternative sets of links that interconnect the endpoints via the network; and

114. Mitchell teaches claim limitation [16a].

115. As I explained above in Sections IX.A.1.ii, IX.A.1.iii, and IX.A.1.iv [[1a], [1b], and [1c]], Mitchell teaches first and second alternative sets of links (the *group of links* and the *set of links*) that interconnect the *multi-homed endstations 414a-414n*.

iii. [16b] network nodes, which are traversed by the *first* and *second* sets of links and are arranged: to provide physical layer resources for operating the links;

116. Mitchell teaches claim limitation [16b].

117. As I explained above in Section IX.A.1.iii [1b], Mithcell teaches network nodes that are traversed by the *group of links* and the *set of links* and provide physical layer resources for operating the links.

iv. [16c] to provide a communication service to the *endpoints* over the *first set of links* using a first instance of the communication service;

118. Mitchell teaches claim limitation [16c].

119. As I explained above in Sections IX.A.1.ii, IX.A.1.iii, amd IX.A.1.iv [[1a], [1b], and [1c]], Mitchell teaches that the communication service is provided to the *multi-homed endstations 414a-414n* over the *group of links* using the first instance.


v. [16d] to detect a local *failure* in the *first set of links*;

120. Mitchell teaches claim limitation [16d].

121. As I explained above in Section IX.A.1.v [1d], Mitchell discloses that the *core router 404a* detects a *failure* in the *link* between *core router 404a* and the network core 402.

122. The *failure* in the *link* between *core router 404a* and the network core 402 is a local *failure*, because the *failure* occurs in a *link* adjacent to the detecting node (e.g., *core router 404a*). Additionally, Mitchell explains that the *failure* is detected towards the network core 402 relative to the *core router 404a*. EX1005, 3:35-43 (“According to one embodiment of the present invention, a method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected and a second link (e.g., a group of links which are logically related to the first link) between the network element and a downstream portion of the communications network is responsively disabled to maintain a communications channel between the downstream and upstream portions of the communications network.”), 2:43-47 (“The foregoing is a summary and thus contains, by necessity, simplifications, generalizations and omissions of detail; consequently, those skilled in the art will appreciate that the summary is illustrative only and is not intended to be in any way limiting.”), 3:9-12 (“On the contrary, it is intended to cover such alternatives, modifications, and equivalents as can be reasonably included within the scope of the invention as

defined by the appended claims.”). The network core 402 is locally upstream of the *core router 404a*, as shown in Figure 4 below. EX1005, FIG. 4.



EX1005, FIG. 4 (annotated).

vi. [16e] to propagate a first message notifying all other nodes in the network that are traversed by the **first set of links** upon detecting the local **failure**; and

[16f] to receive a second message from another node indicating a remote **failure** in the **first set of links**; and


123. Mitchell teaches each of claim limitations [16e] and [16f].

124. As I explained above in Section IX.A.1.vi [[1e], [1f], and [1g]],

Mitchell teaches the propagation of failure information following a **failure** in the **group of links** to all other nodes traversed by the **group of links**.

125. After detecting the local failure, the *core router 404a* can propagate failure information downstream to network switch 410a, which in turn sends the

failure information to primary ethernet switch 412a, where the failure information message indicates the remote failure occurred at the *link* between *core router 404a* and network core 402.



EX1005, FIG. 4 (annotated).

- vii. [16g] responsively to receiving the second message indicating the remote **failure**, to **deactivate** the physical layer resources of each node to links in the **first set of links** connected to that node upon receiving the second message, to cause a loss of connectivity in the **first set of links** and

126. Mitchell teaches claim limitation [16g].

127. As I explained above in Sections IX.A.1.vi and IX.A.1.vii [[1e], [1f], and [1g] and [1h]], Mitchell teaches that after the primary ethernet switch 412a receives the failure information that indicates a remote failure at the *link* between

core router 404a and the network core 402, the physical layer resources associated with the *group of links* are *deactivated* upon receipt of a failure message, which causes a loss of connectivity in the *group of links*.


viii. [16h] thus cause an automatic transfer of communication among the *endpoints* to the *second set of links*, and to resume the communication service over the *second set of links* using a second instance of the communication service different from the first instance.

128. Mitchell teaches claim limitation [16h].

129. As I explained above in Sections IX.A.1.vii and IX.A.1.viii [[1h] and [1i]], Mitchell teaches an automatic transfer of communication among *endstations 414a-414n* to the *set of links*. Specifically, the failover causes an automatic transfer to the *set of links* such that the collection of network elements in the network core 402, core router 404b, network switch 410b, secondary ethernet switch 412b coupled together by the *set of links* are now used instead of the first instance of network elements in the network core 402, core router 404a, network switch 410a, primary ethernet switch 412a coupled together by the *group of links*.
EX1005, 7:36-49.

130. As shown in Figure 4 below, the *set of links* provides the communication service to the *endstations 414a-414n* using the second instance (e.g., *set of links* and core router 404b, network switch 410b, and secondary

ethernet switch 412b). The second instance is different from first instance (e.g., *group of links*, core router 404a, network switch 410a, and primary Ethernet switch 412a) used by the *group of links*. EX1005, 7:36-49 (“In the illustrated embodiment of FIG. 4, a data processing system 400 is implemented including a core network 402, primary and secondary core routers (404a and 404b), primary and secondary network switches (410a and 410b), and primary and secondary Ethernet switches (412a and 412b) coupled together using a plurality of links and further including a plurality of multi-homed endstations 414a-414n similarly coupled to primary Ethernet switch 412a via one of a plurality of primary downstream links 416a-416n (e.g., via a primary network interface card or host bus adapter) and to secondary Ethernet switch 412b via one of a plurality of secondary downstream links 418a-418n (e.g., via a secondary network interface card or host bus adapter) as shown.”).



EX1005, FIG. 4 (annotated).

3. Dependent Claim 17

- i. [17] The network according to claim 16, wherein a **first network node** is arranged to detect the **failure** and to send a third message indicating the **failure** to a **second network node**, and wherein a **third network node**, which provides the physical layer resources to links in the first set, is arranged to **deactivate** the physical layer resources responsively to accepting the third message.

131. Mitchell teaches claim limitation [17].

132. As I explained above in Sections IX.A.1.v [1d], IX.A.1.vi [1e]-[1g], and IX.A.1.vii [1h], Mitchell teaches this claim limitation.

4. Dependent Claims 7, 13, 19, and 23

- i. [7] The method according to claim 1, wherein the one or more of the links in **the first set of links** comprise wire links, and wherein said deactivating the physical layer resources comprises disabling respective carrier signals of the wire links.

[13] The method according to claim 1, wherein a link in **the first set of links** includes a wire link, and wherein deactivating the physical layer includes disabling a carrier signal of the wire link.

[19] The network according to claim 16, wherein the one or more of the links in **the first set of links** comprise wire links, and wherein the network nodes are arranged to deactivate the physical layer resources by disabling respective carrier signals of the wire links.

[23] The communication network of claim 16, wherein a link in **the first set of links** includes a wire link, and wherein deactivating the physical layer resources includes disabling a carrier signal of the wire link.

133. Mitchell teaches the additional limitations of claims [7], [13], [19], [23].

134. As I explained in Sections IX.A.1 [1] and IX.A.2 [16], Mitchell renders obvious claims 1 and 16.

135. Mitchell teaches that “**links 108, 110, and 112** have been illustrated as *direct connections* between network elements.” EX1005, 4:45-62.

136. A POSITA would have understood that wired links are a direct connection. *See, e.g.,* EX1013, 5:56-67 (“[C]ommunication media includes wired media, such as a wired network or direct-wired connection....”).

137. For example, Mitchell notes that direct connections (e.g., wired connections) are distinguished from wireless connections. *See, e.g.,* EX1005, 4:45-62 (describing wireless connections as “e.g., using electro-magnetic, optical, infrared and/or acoustic signals or transmission media.”); *see also* EX1005, 1:46-60. Consequently, a POSITA would have recognized that the direct connections forming the “*group of links*” are also wired connections.

138. Further, a POSITA would have known that wired connections transmit carrier signals. For example, it was well-known that wired connections can carry carrier signals/waves to send and receive data across communication networks. *See, e.g.,* EX1013, 5:56-67 (“Communications media typically includes ... a modulated data signal, such as a carrier wave or other transport mechanism.... [C]ommunication media includes wired media, such as a wired network or direct-wired connection....”). Thus, a POSITA would have recognized that when Mitchell disables “a physical layer protocol circuit or ‘PHY’ associated with one or more ports and/or *links* ... ***causing all ports and/or links associated with the PHY to become disabled,***” the carrier signals are disabled. *See, e.g.,* EX1005, 5:56-6:5; *see also* EX1014, 2:56-3:14 (describing it is well-known and well-understood that

disabling a wired connection (e.g., a telephone line) may be performed in response detecting an inactive line state associated with a carrier signal).

5. Dependent Claims 9, 15, and 25

- i. **[9] The method according to claim 1, wherein the one or more of the links in the first set of links comprise a logical link, and wherein said deactivating the physical layer resources comprises releasing an association of the logical link with the first instance of the communication service.**

[15] The method according to claim 1, wherein a link in the first set of links includes a logical link, and wherein deactivating the physical layer includes releasing an association of the logical link with the first instance of the communication service.

[25] The network according to claim 16, wherein a link in the first set of links includes a logical link, and wherein deactivating the physical layer resources includes releasing an association of the logical link with the first instance of the communication service.

139. Mitchell teaches the additional limitations of claims [9], [15], and [25].

140. As I explained in Sections IX.A.1 [1] and IX.A.2 [16], Mitchell renders obvious claims 1 and 16.

141. First, the '637 Patent describes an example of a logical link that includes a virtual LAN (VLAN). *See, e.g.*, EX1001, 9:53-55 ("Remote failure propagation can be carried out in such cases by locally shutting down the logical link (in the present example the VLAN)....").

142. Mitchell discloses that logical links may be included among the *group of links* and *set of links*. Mitchell expressly describes that a “link” may include a “logical” link or links. *See* EX1005, 4:26-29 (“[T]he term ‘link’ is intended to indicate a data link layer connection and may include one or more *logical sub-links* (e.g., an ‘Ether-Channel’ or ‘Port-Channel’)...”). Mitchell further discloses that its links may be implemented in a VLAN (e.g., as a logical VLAN link). EX1005, 6:37-41 (“Using one embodiment of the present invention, a network element (e.g., a datalink layer-capable Ethernet switch) is configured to track the state of a virtual network (e.g., a *VLAN* or VSAN) on all identified upstream interfaces (e.g., links, ports, interface cards, or the like).”).

143. In addition, Mitchell describes that, in some examples, disabling a logical VLAN link associated with the “*group of links*” may involve deactivating physical layer resources by releasing an association with the *group of links* (that are coupled with the network elements core router 404a, network switch 410a, primary ethernet switch 412a). *See, e.g.*, EX1005, 5:62-65 (“[I]n yet another embodiment, only those downstream links or ports associated with a virtual network such as a *VLAN* or VSAN which was predetermined or identified using configuration interface 202 are disabled.”), 3:35-43 (“[A] method is provided in which a failure of a first link between a network element and an upstream portion of a communications network is detected and a second link (e.g., *a group of links*

which are logically related to the first link) between the network element and a downstream portion of the communications network is responsively *disabled* to maintain a communications channel between the downstream and upstream portions of the communications network.”).

6. Dependent Claims 10 and 21

- i. [10] The method according to claim 1, wherein a link of the **first set of links** is a critical link, and wherein said deactivating the physical layer is performed only when the **failure** occurs in the critical link.

144. Mitchell teaches the additional limitation of claim [10].


145. As I explained in Sections IX.A.1 [1] and IX.A.2 [16], Mitchell renders obvious claims 1 and 16.

146. It should be noted that the '637 Patent describes that a “critical link is defined as a link whose failure triggers a switch-over to the backup topology.” EX1001, 7:46-56. In fact, the '637 Patent goes further, contrasting critical links with non-critical links. For instance, the '637 Patent describes that, unlike a failure in a critical link, a “[f]ailure of a non-critical link is considered tolerable, and does not initiate such a change.” EX1001, 7:54-56.

147. Similarly, Mitchell also discloses that its communications network includes both critical and non-critical links. For example, as explained in Sections IX.A.1.v[1d], IX.A.1.vi [1e]-[1g], IX.A.1.vii [1h], and IX.A.1.viii [1i], Mitchell

discloses that the critical link may be the *link* between *core router 404a* and network core 402, for example, because when Mitchell detects a **failure** in that *link*, an automatic failover to the secondary *set of links* is triggered. EX1005, 5:30-38, 5:46-6:5, 6:26-36 (“Once a port and/or link ... is disabled, an associated multi-homed endstation may failover to a corresponding secondary downstream link ... by activating a network interface associated with the secondary downstream link and deactivating a network interface associated with the failed primary downstream link.”), 7:7-10, 7:25-55, FIGS. 2, 4.

148. As shown below, Figure 4 of Mitchell depicts a **failure** in a critical link.



EX1005, FIG. 4 (annotated).

149. In addition, Mitchell also uses non-critical links. For instance, Mitchell discloses that certain failures (e.g., in non-critical links) may not result in deactivating a physical layer resource or trigger any failover to a backup topology (e.g., secondary *set of links*). EX1005, 6:6-25 (“Consequently, ***a failure of one or more of the primary upstream links need not cause any of downstream links 110 to be immediately disabled.*** For example, a downstream link 110 may be disabled only upon the failure of all primary upstream links, upon the failure of a predetermined or defined number or proportion of primary upstream links, or upon available bandwidth between primary switch 104 a and upstream communications network portion 102 falling below a predetermined or defined threshold level due to the failure or one or more primary upstream links.”), 6:26-36, 7:10-15, FIGS. 2, 4.

150. Therefore, Mitchell discloses the added limitations of claims [10] and 21.

7. Dependent Claim 11

- i. [11] The method according to claim 1, wherein said provisioning different first and second instances of a multipoint-to-multipoint communication service over respective *first* and *second* alternative sets of links comprises sharing respective allocations of network resources reserved for communicating over the *first* and *second* sets of links.

151. Mitchell teaches the additional limitation of claim [11].

152. As I explained in Section IX.A.1 [1], Mitchell renders obvious claim 1.

153. The '637 Patent describes allocating shared network resources. EX1001, 10:39-47 (“Since in most cases traffic actually flows in only one of the two topologies at any given time, it is possible to share bandwidth and other resource allocations between the two topologies.”). Otherwise stated, the '637 Patent describes sharing resource allocations between topologies, such as VPLS1 and VPLS2, is possible because traffic flows may be present in either of VPLS1 or VPLS2 at any given time.


154. Further, Mitchell describes that a failover (e.g., to secondary *set of links*) may occur, while still “preserving higher-level communication channels between the downstream and upstream portions of the communications networks [to] provid[e] a seamless failover.” EX1005, 4:10-15. Thus Mitchell discloses the same seamless failover as it fails over from the primary *group of links* to the secondary *set of links* and shares allocations of the network elements inside the network core 402, which is just like the shared resource allocations described in the '637 Patent (e.g., when seamlessly failing over from *VPLS1* to *VPLS2*). See EX1005, 1:46-60, 4:31-44, 6:6-36.

155. I note that the '637 Patent erroneously points to U.S. Patent Application No. 11/305,485 instead of U.S. Patent Application No. 11/305,486, as

being incorporated by reference. For example, '637 Patent states that: "U.S. patent application Ser. No. 11/305,485, entitled 'Resource Sharing among Network Tunnels,' filed Dec. 15, 2005, which is assigned to the assignee of the present patent application." EX1001, 10:48-53. However, this appears to be a typographical error because it is U.S. Patent Application No. 11/305,486 that is in fact titled "Resource Sharing among Network Tunnels." *See* EX1015, Title. Indeed, U.S. Patent Application No. 11/305,486 was filed December 15, 2005, was originally assigned to Orckit Corrigen Ltd., and is now U.S. Patent No. 7,463,580 to Agmon *et al.* ("Agmon") (EX1015).

156. However, the '637 Patent's incorporation of Agmon by reference indicates that a network resource may be any resource associated with a network element. Specifically, Agmon explains that the "term 'resource' means *any resource* associated with a network element." EX1015, 5:12-21.

157. Returning to Mitchell, annotated figure 4 of Mitchell is shown below.



EX1005, FIG. 4 (annotated).

158. As shown, Mitchell provides network connections to its endstations using various network elements across a network that may include a network core, LAN, WAN, and/or MAN. *See, e.g., EX1005, 1:46-60 (“In operation, **data is transmitted between multi-homed endstations 106 and upstream communications network portion 102** using primary links 110a-110n, primary switch 104a, and primary link 108a. Following a failure ... **[d]ata is then transmitted between the multi-homed endstation which has failed over and upstream communications network portion 102** via an associated secondary link 112, a secondary switch 104b, and a secondary link 108b.”), 5:46-50, 6:26-36, 7:36-49.*

159. In one example, Mitchell shares network resource allocations by sharing resources in various network elements across an upstream network (e.g.,

the network elements contained in network core 402 between the primary *group of links* and the secondary *set of links*). Therefore, Mitchell discloses sharing respective allocations of network resources through network elements in an upstream network.

B. Ground 2: Mitchell and Bhate Render Obvious Claims 2, 6, 12, 18, and 22

1. Dependent Claim 2

i. [2] The method according to claim 1, wherein the network comprises a ring network.

160. Mitchell and Bhate render obvious the additional limitation of claim [2].

161. Mitchell does not expressly state that the network may include a ring network. However, ring networks were well-known network topologies. Further, it would have been obvious to an ordinarily skilled artisan to modify Mitchell's system to include a ring network.

162. For example, Bhate discloses a ring network. Specifically, Bhate discloses a unidirectional path switched ring (UPSR) ring network. EX1006, 4:33-64. Bhate's UPSR ring network provides communication traffic "in two different directions along a working path and a protection path between a source node and a destination node." EX1006, 4:33-64. As shown below, Bhate's Figure 3 includes a UPSR ring network 300.




FIG.3

EX1006, FIG. 3. Bhate's UPSR 300 includes four network nodes 305 (e.g., Node 1, Node 2, Node 3, and Node 4) that are disposed along the ring network 300.

EX1006, 4:33-48.

163. Nodes 305 are described in Bhate as functioning in a similar manner to the network elements described in Mitchell. For example, Mitchell's specification refers to "network elements" and "nodes" interchangeably. *See, e.g.,*

EX1005, 1:19-25. Further, Mitchell discloses that examples of a suitable node may include switches and routers. EX1005, 1:19-25 (“For example, redundant network links or connections are frequently used to couple a single **network element or node**, (e.g., a client, server, or other host data processing system, a communications network appliance, or a **switch, router**, hub, gateway or other redistribution point) to one or more communications networks or portions thereof (e.g., network segments).”).


164. Further, each of Mitchell’s network elements/nodes and Bhate’s nodes perform similar functions. For example, detecting a link failure is a function that is performed by both Mitchell’s network elements/nodes and Bhate’s nodes. *See e.g.*, EX1005, 5:50-55 (“Following a failure of primary upstream link 108a (e.g., due to failure of the physical link hardware, a network interface card, or primary switch 104a) link failure propagation module 204 of primary switch 104a is used to detect the failure and responsively disable one or more of primary downstream links 110a-110n.”); EX1006, 2:59-3:2 (“As indicated in FIG. 2B, responsive to detecting a fault (e.g., a loss of signal) at the line side of an intervening passthrough or regeneration node 215, a corresponding laser transmitter used to propagate the signal further downstream is turned off 220.”).

165. In addition, Mitchell and Bhate each describe that when a failure is detected, their respective network elements/nodes perform the function of

propagating failure information to one or more downstream nodes. *See, e.g.*, EX1005, 4:10-20 (“As the link state or link failure information is propagated downstream, downstream network elements with the capability to switch to alternate (e.g., redundant) links may do so, preserving higher-level communication channels between the downstream and upstream portions of the communications networks and providing seamless failover.”), 5:50-55; EX1006, 2:59-3:2, 4:25-32 (“As an illustrative example, if a node is configured to provide SONET service, conventional SONET AIS alarms for loss of signal (LOS), loss of frame (LOF), BER, or signal fail may be transmitted downstream in response to a fault whereas for a service lacking an AIS protocol the node may turn off a laser transmitter to propagate fault information towards the destination node.”).


166. As a result of these similarities, a POSITA would have readily that Mitchell and Bhate’s respective network elements/nodes are similar devices that perform similar functions. Specifically, Mitchell and Bhate’s respective network elements/nodes provide network communications and connectivity, detect failures (e.g., in an upstream connection), and propagate failure information to other downstream network elements/nodes in response to the detected failures. Because of these similarities, a POSITA would have recognized that Mitchell’s network elements, including core routers 404, could be substituted for Bhate’s nodes 335 (nodes 1-4).

167. To illustrate this combination, I have reproduced annotated Figure 4 of Mitchell and Figure 3 of Bhate below.



EX1005, FIG. 4 (annotated); EX1006, FIG. 3.

168. It would have been obvious to substitute Mitchell's core routers 404 for Bhate's nodes 335 (nodes 1-4), placing the core routers 404 onto the UPSR ring network 300, as shown below.



169. In making the combination above, a POSITA would have had a reasonable expectation of success. Disposing Mitchell's core routers 404 onto the USPR ring network 300 as a replacement for Bhate's nodes 335 would have merely required substituting one network element for another network element that provides nearly identical functionality. Consequently, a POSITA would have known that making such a substitution would only require using a well-known

network element instead of using another well-known network element to achieve a predictable result of providing network communications and connectivity over a ring topology.

170. Further, a POSITA would have readily recognized additional reasons for using Mitchell's network elements/nodes with Bhate. For example, a POSITA would have readily recognized that Mitchell and Bhate both disclose forwarding/propagating failure information. *See, e.g.*, EX1005, 2:23-33 ("Using one or more embodiments of the present invention, network connectivity (e.g., data link layer connectivity) information is propagated..."), 4:10-15 ("As the link state or link failure information is propagated downstream..."); EX1006, Abstract ("[T]he the fault information is forwarded to the destination node...."), 2:26-32 ("The present invention generally comprises ... forwarding fault information..."). Bhate similarly detects a link failure and forwards failure information over a ring network. *See, e.g.*, EX1005, 2:23-33 ("Using one or more embodiments of the present invention, **network connectivity** (e.g., **data link layer connectivity**) **information is propagated**, thereby enabling **downstream network elements** not immediately adjacent to the site of a link failure or directly coupled to a network element experiencing link failure to failover to alternate, redundant links such that the state of **one or more connections or communications channels** with the upstream portion(s) **of a communications network** may be preserved and **the**

connection(s)/channel(s) may be maintained.”), 4:10-15 (“As the link state or *link failure information is propagated downstream, downstream network elements* with the capability to switch to alternate (e.g., redundant) links may do so, preserving *higher-level communication channels* between the downstream and upstream portions of *the communications networks* and providing seamless failover.”), 4:33-39 (“As previously described with respect to FIG. 1, *data processing system 200* of the illustrated embodiment of FIG. 2 includes *an upstream communications network portion 102*, a primary switch 104a, a secondary switch 104b, and a number of redundantly-linked network elements (e.g., multi-homed endstations 106a, 106b ... 106n).); EX1006, Abstract, 2:36-3:2, 3:20-35, 4:33-64, FIGS. 2A-2D, 7. As a result, a POSITA would have further understood that, because Mitchell and Bhate both relate to link failure recovery, their respective systems’ shared functionalities would have been combinable.

171. A POSITA would have readily recognized that Mitchell and Bhate both provide redundancies across their respective network topologies. *See, e.g.*, EX1005, 4:15-30, 4:58-62, 5:35-37, 5:59-65, 6:37-44, 7:34-49; EX1006, 1:15-37. Thus, a POSITA would have known that adding Mitchell’s fast link failover technique to Bhate’s redundant ring networks would further improve an overall network resiliency through the combined redundancies.

172. Such a POSITA would have also been motivated to include Mitchell's core routers 404 in Bhate's USPR ring network 300 because there were additional benefits that were readily recognizable in doing so. As I explained in Section IX.A.1 [1], Mitchell's multipoint network seeks to improve data accessibility and reliability. *See e.g.*, EX1005, 1:11-18. Thus, a POSITA would have recognized that placing Mitchell's network elements/nodes on Bhate's USPR ring network 300 would have the added benefit of improving data accessibility and reliability across a communications network.

173. Indeed, it was well-known and well-understood that ring networks are frequently utilized to improve data accessibility and reliability. For instance, it was well-known that bandwidth utilization efficiencies can be improved with ring networks. *See, e.g.*, EX1012, [0010]-[0012]. Further, a POSITA would have known that using ring networks can improve data accessibility and reliability, for example, by increasing network traffic efficiencies and decreasing network congestion across a network. EX1012, [0040]. Thus, a POSITA would have been motivated to place Mitchell's core routers onto Bhate's USPR ring network 300 to yield the added benefits of improvements to the communication network's data accessibility and reliability.

174. Therefore, Mitchell in combination with Bhate renders obvious the added limitation of claim [2].

2. Dependent Claims 6, 12, 18, And 22

- i. [6] The method according to claim 1, wherein the one or more of the links in the **first set of links** comprise optical links, and wherein said deactivating the physical layer resources comprises disabling respective laser sources of the optical links.

[12] The method according to claim 1, wherein a link in **the first set of links** includes an optical link, and wherein deactivating the physical layer includes disabling a laser source of the optical link.

[18] The network according to claim 16, wherein the one or more of the links in **the first set of links** comprise optical links, and wherein the network nodes comprise respective laser sources for driving the optical links and are arranged to deactivate the physical layer resources by disabling the laser sources.

[22] The communication network of claim 16, wherein a link in **the first set of links** includes an optical link, and wherein deactivating the physical layer resources includes disabling a laser source of the optical link.

175. Mitchell in combination with Bhate render obvious the additional limitations of claims [6], [12], [18], and [22].

176. For example, Mitchell discloses the **group of inks** may be **optical links**. EX1005, 4:51-58 (“[A]ny of **links 108, 110, and 112** may be wireless (e.g., using electro-magnetic, **optical**, infrared and/or acoustic signals or transmission media).”).

177. Mitchell does not explicitly disclose its optical links include laser sources. However, laser sources for driving optical links was well-known. Further, disabling optical links using their laser sources was also well-known.

178. Bhate discloses an optical transmitter that “includes a laser to generate optical signals at a line wavelength.” EX1006, 5:21-23. Bhate uses these optical signals to transmit network communications. *See* EX1006, 5:24-41. Consequently, a POSITA would have known that *optical links* are driven by a laser source.

179. In addition, Bhate describes *optical links* may be deactivated, for example, by “turning off a laser transmitted used to propagate the signal towards its destination.” EX1006, 2:36-42 (“For example, in response to a node detecting a fault condition at a tributary interface or a line interface, the node turns off a corresponding laser transmitter used to propagate the signal towards its destination.”), 2:50-3:2, 3:19-35, Abstract; *see also* EX1006, 4:33-57, 5:42-52 (discussing controllers to shut down a laser transmitter 520), FIG. 5, 550.

180. Indeed, Bhate explains that to do so, a “controller 550 is configured to detect faults and turn off an appropriate laser transmitter to propagate the fault information downstream, e.g., towards the destination node.” EX1006, 5:42-44. According to Bhate, controller 550 drives the signal (*e.g.*, shuts off a laser of transmitter 520) in response to determining a loss of signal (LOS) condition.

181. A POSITA would have been motivated to combine Mitchell with Bhate. Further, such a POSITA would have known that a combination of Mitchell and Bhate (“Mitchell/Bhate”) would have had a reasonable expectation of success. Specifically, a POSITA would have understood Mitchell/Bhate would have had a reasonable expectation of success because deactivating an optical link by disabling an optical laser (e.g., by driving a laser of an optical source) would merely involve using a known technique of turning a laser on and off (i.e., driving or disabling) the laser of an optical link. *See, e.g.*, EX1005, 4:51-58; EX1006, Abstract, 2:26-35.

182. Therefore, Mitchell in combination with Bhate renders obvious the added limitations of claims [2], [12], [18], and [22].

C. Ground 3: Mitchell And Raahemi Render Obvious Claims 3 And 4

1. Dependent Claim 3


- i. **[3] The method according to claim 1, wherein the communication service comprises a virtual private local area network service (VPLS), and wherein the first and second instances comprise first and second VPLS instances.**

183. Mitchell in combination with Raahemi render obvious the additional limitation of claim [3].

184. Mitchell is silent as to whether its communication service includes a virtual private local area network service (VPLS). But a POSITA would have understood that using a VPLS as a communication service was well-known.

185. Indeed, a POSITA would have known that as early as 2005, VPLS communications services were “already being used by service providers to aggregate services for delivery to residential and enterprise customers.” EX1007, [0002]. Further, Raahemi explains there were well-known benefits associated with using VPLS communications services, including that a VPLS “allows multiple sites to be connected in a single bridged domain over a provider managed IP/MPLS network.” EX1007, [0003]. And in doing so, a POSITA would have also known “[a]ll customer sites in VPLS instance appear to be on the same LAN regardless of their location.” EX1007, [0003].

186. As shown below, Figure 2 of Raahemi depicts an exemplary VPLS network.




EX1007, FIG. 2.

187. Raahemi describes that Figure 2 depicts a multimedia broadcast service. EX1007, [0026] ("FIG. 2 shows an example of a Steiner Tree-based

VPLS network used for Digital TV broadcast services.”). Further, Raahemi discloses that the multimedia broadcast service shown in Figure 2 (above) includes a “network [that] is comprised of a video source attached to a multicast router, the Provider Edge (PE) routers, the core routers (P), the Customer Equipment (CE) or aggregation devices, and the Set-top Boxes (STB) that are attached to customer TV sets.” EX1007, [0026].

188. Further, as shown below, Figure 2 shows that Raahemi’s network includes two different examples of VPLS instances that are disjoint from one another (e.g., a *working tree* on the left and a *backup tree* on the right). For example, Raahemi discloses that “[i]n the specific application of this invention, that is *broadcasting digital multimedia over VPLS*, load balancing on the two trees (*the working tree and the backup tree*) during the normal operation is discussed.” EX1007, [0046]. Raahemi further describes that “[d]uring the service initialization stage the Network Management System (NMS) pre-computes 2 disjoint Steiner Trees as shown in FIG. 2.” EX1007, [0049].




EX1007, FIG. 2 (annotated).

189. Raahemi describes that the benefits of having two disjoint VPLS instances is that, in the case of a fault, the backup VPLS can continue to service the

destination PE routers. *See, e.g.*, EX1007, Abstract (“Destination PE routers in the VPLS network are connected to the disjoint trees so that they can be serviced by either tree in the case of a fault.”), [0046] (“In the present invention Load Balanced Disjoint Steiner Trees with Split Horizon and Dual connectivity at the PEs are provided. These are provided by methods and architecture to protect the minimum cost (Steiner) tree-based VPLS from the failure of the Branching node in particular, and many other possible types of failures in general, thus providing a resilient architecture for broadcasting multimedia services.”), [0047], [0048] (“The Network Management System (NMS) pre-computes two disjoint Steiner trees taking advantage of dual connectivity at the receiver PE’s.”), [0049] (“During the service initialization stage the Network Management System (NMS) pre-computes 2 disjoint Steiner Trees as shown in FIG. 2.”).

190. Combining Mitchell with Raahemi (Mitchell/Raahemi) would have been obvious to a POSITA. For example, it would have been obvious to a POSITA to modify Mitchell’s fast link failover system to include the VPLS instances, as described Raahemi and discussed below, because doing so would enable customers with multiple different physical sites to “appear to be on the same LAN regardless of their location.” EX1007, [0003].


191. I have reproduced annotated versions of Mitchell’s Figure 4 and Raahemi’s Figure 2, which can be combined as described below.



EX1005, FIG. 4 (annotated); EX1007, FIG. 2 (annotated).

192. Mitchell's communication service can be included in a system that is similar to Raahemi's VPLS service. For example, as shown below, Raahemi's PE routers at the edge of the core network shown in Figure 2 of Raahemi by using Mitchell's core routers, which provide similar functionality as the PE routers described in Raahemi.

MITCHELL/RAAHEMI



193. A POSITA would have recognized that, just like the PE routers shown as being located at the edge of the communication network in Figure 2 of Raahemi above, Mitchell similarly discloses that network elements, such as routers, may be used to provide redundant network links at the edge of a communication network.

EX1005, 1:19-29 (“For example, redundant network links or connections are frequently used to couple a single *network element or node*, (e.g., a client, server, or other host data processing system, a communications network appliance, or a *switch, router*, hub, gateway or other redistribution point) to one or more communications networks or portions thereof (e.g., network segments).”). As a result, a POSITA would have readily recognized that replacing Raahemi’s PE routers at the edge of its communication network with Mitchell’s core routers would have merely amounted to a substitution of a well-known prior art element (e.g., Raahemi’s PE routers) for another (e.g., Mitchell’s core routers 404) according to known network communication methods to yield the predictable results of a communication service that provides two different VPLS instances with additional, similar redundancies that allow a fast link failover to occur.

194. Further, a POSITA would have been motivated to modify Mitchell’s fast link failover system to incorporate Raahemi’s disjoint VPLS service to take advantage of both of the failover methodologies described in Mitchell and Raahemi, respectively. A POSITA would have understood that the combined Mitchell/Raahemi VPLS communication service would have been desirable to further improve the failover systems described in both Mitchell and Raahemi without any loss or degradation of these desirable failover techniques.

195. In addition, a POSITA would have expected that the combined Mitchell/Raahemi VPLS communication service would have yielded the predictable result of a system that provides communication services with fast link failover capabilities with even faster recovery times. A POSITA would have understood that Mitchell's fast link failover system is described as being capable of improving link failure recovery times significantly (e.g., reducing recoveries from multiple minutes to under 50 milliseconds). EX1005, 2:7-12 ("The described techniques suffer from a number of significant shortcomings however. To account for ordinary communications network congestion and to avoid falsely declaring a link failure, the threshold time periods (and resultant latency) associated with the described techniques *are typically multiple seconds or more.*"), 3:52-63 ("Using embodiments of the present invention, a link failure may be propagated within a bounded or predetermined period of time. For example, where embodiments of the present invention are implemented using predominantly special-purpose hardware (e.g., application specific integrated circuits, special-purpose processors, programmable logic devices, or the like) or specific (e.g., vectored) interrupts, a link between a network element and a downstream portion of a communications network *may be disabled within a period of time substantially less than or equal to 50 milliseconds of detecting a failure of another link* between the network element and an upstream portion of the communications network."). Similarly,

Raahemi describes outperform existing systems with a recovery time of less than 50 milliseconds. EX1007, [0058] (“Also, *a pre-planned recovery scheme, as proposed herein, outperforms dynamic failure recovery such in speed.* Dynamic failure recovery is usually slow and does not provide an equate restoration time of *less than 50 mSec.*”).

196. Consequently, a POSITA would have readily recognized that the combination of Mitchell and Raahemi would have led to a resulting system that maintains its desired fast link failover with the added benefit of providing a redundant VPLS service over the network. And the POSITA would have understood that, in doing so, the combination of Mitchell/Raahemi would have had a reasonable expectation of success.

197. Therefore, Mitchell in combination with Raahemi renders obvious the added limitation of claim [3].

2. Dependent Claim 4

i. [4] The method according to claim 1, wherein the communication service comprises a bridging service.

198. Mitchell in combination with Raahemi renders obvious the additional limitation of claim [4].

199. As an initial matter, I note that '637 Patent incorporates U.S. Publication 2006/0109802 (“Zelig”) (EX1011) by reference. *See, e.g.*, EX1001,

1:46-57. And Zelig describes VPLS services, explaining that VPLS provides communication services that is sometimes also referred to as a transparent LAN service (“TLS”). EX1011, [0007]. Zelig further explains that VPLS is used to provide a “bridge-like service.” EX1011, [0007]. Thus, a POSITA would have understood that VPLS communication services are just one well-known type of a bridging service.

200. In addition and as discussed in Section IX.C.1.i, it would have been obvious to set up a VPLS communication service over Mitchell’s fast link failover system combined with Raahemi’s resilient VPLS service.

201. For example, Raahemi describes that a VPLS service “is a layer 2 multipoint Virtual Private Network (VPN) that allows multiple sites *to be connected in a single bridged domain* over a provider managed IP/MPLS network.” EX1007, [0003]. Additionally, Raahemi describes using *split horizon bridging*, for example, to provide rules for implementing its VPLS trees and preventing undesirable loops from forming when doing so. EX1007, Abstract (“In the event of a fault, the services on the faulty tree are switched to the other tree using split horizon bridging.”), [0006] (“To prevent forward loops, the Split Horizon rule is used. In the VPLS context, this rule basically implies that a PE must never send a packet on a PW if that packet has been received from a PW. This ensures that traffic cannot form a loop over the backbone network using PWs.

The fact that there is always a full mesh of PWs between the PE devices ensures that every destination within the VPLS will be reached by a broadcast packet.”), [0014] (“The limitations of the prior art are overcome by the present invention which proposes a scheme for implementing Split Horizon (a characteristic of VPLS-enabled edge routers) together with dual connectivity at the receiver PE's to facilitate the switch-over process.”), [0046] (“In the following description the details as how to switch over traffic from failed tree to backup tree are provided. There is also a description of the process for implementing Split Horizon together with dual connectivity at the receiver PE's to facilitate the switch over process.”), [0047] (“An alternate solution is provided herein based on computing a disjoint minimum-cost (Steiner) tree together with employing split horizon at the PE's as shown in FIG. 2.”), [0048], [0050] (“To attach receiver CE's to the working tree, split horizon rules are used at the receiver PE's.”), [0057] (“The combination of Steiner tree, with split horizon and dual connectivity at the PE's results in a unique architecture presented in this invention which supports resiliency in tree-based VPLS broadcast systems.”).

202. Therefore, Mitchell in combination with Raahemi renders obvious the added limitation of claim [4].

D. Ground 4: Mitchell And Huai Render Obvious Claim 5

1. Dependent Claim 5

- i. **[5] The method according to claim 1, wherein said propagating failure information comprises advertising the failure using an Opaque Link State Advertisement (LSA) message of an Open Shortest Path First (OSPF) protocol.**

203. Mitchell in combination with Huai render obvious the additional limitation of claim [5].

204. While Mitchell is silent about using an Opaque LSA message of an OSPF [Open Shortest Path First] protocol, it would have been obvious to do so because doing so was well-known in the prior art.

205. By way of example, Huai discloses that its network uses the OSPF protocol. EX1008, 1:66-2:6 (“*It is well known* for routers of messages (in a packet-switched network) *to communicate* with one another *in accordance with the Open Shortest Path First (‘OSPF’) protocol*, which is governed by a standard of the Internet Engineering Task Force (‘IETF’). IETF document RFC 1583 (‘IETF RFC 1583’), shown in Appendix G, describes OSPF Version 2 and is incorporated herein by reference in its entirety.”). Huai describes that OSPF messages may include information about the location of network elements and their connections, which is called the network topology. EX1008, 2:8-11 (“In OSPF, messages containing information about the location of various routers and interconnections

among the routers (also called ‘network topology’) are sent by the routers to one another.”). Huai further describes that certain OSPF messages that are called LSAs [link state advertisements] may be used to maintain and update OSPF databases (e.g., that maintain information about a given network topology). *See, e.g.*, EX1008, 2:8-16 (“.... Each router maintains and updates a database of network topology information retrieved from such messages, which are also called ‘Link State Advertisements’ (‘LSA’). Each router uses the network topology information to determine the shortest path from itself to all other routers in the network.”).

206. In addition, Huai discloses an ability to propagate failure information, using an OSPF flooding mechanism, to other network elements in the network. EX1008, Abstract (“Information relating to the communications links are automatically propagated in the network using a protocol used by routers in communicating with other routers.”), 2:32-3:2 (“This invention relates to a method and associated apparatus for automatically propagating circuit information in a network.... Information relating to a link coupling two circuit switches is automatically propagated on the network using a protocol which is ordinarily used by routers(“packet routing protocol”) in communicating with other routers.... Information which is automatically propagated using the packet routing protocol is used to create and maintain a table describing each operational link connecting two circuit switches.... In another embodiment, a router and a circuit switch are in a

single network element, the network element being part of a network....

Information relating to the links coupling the circuit switches is automatically propagated using a packet routing protocol. Such information is also used to create and maintain a table which is included in all network elements..... In another embodiment, the links connecting the circuit switches are Synchronous Optical Network (“SONET”) links and the packet routing protocol is OSPF. In one embodiment, the OSPF flooding mechanism is used to automatically propagate circuit information (as opposed to packet network type of information). Such circuit information includes a description of the interfaces coupled to a link and other information relating to the link.”).

207. Further, Huai discloses that the OSPF flooding mechanism to update OSPF databases, for example using link state updates that are transmitted as opaque LSA’s. EX1008, 13:7-39. Huai describes that these opaque LSAs are sent using the OSPF flooding technique when a link failure occurs, for example, by propagating the failure using the flooding mechanism. EX1008, 13:7-39. In one example, Huai discloses that an updated packet 800, which is shown in Figure 8 of Huai (below), may be an OSPF link state update packet 800. EX1008, 8:44-52. Huai further discloses that the OSPF link state update packet 800 also includes an opaque LSA. EX1008, 8:44-52.

208. I have reproduced Figure 8 of Huai below.




FIG. 8

EX1008, FIG. 8.

209. Figure 8 of Huai (above) shows that the update packet 800 includes standard LSA header 810 and standard opaque LSA header 830. EX1008, 8:53-9:16 (“FIG. 8 shows a packet 800 in accordance with an embodiment of the invention. Packet 800 includes words 810, words 830, field 840, and words 850 (i.e. words 850A, 850B . . .). Words 810 compose a standard OSPF Link State Update packet header (p. 179, IETF RFC 1583).... The number of link

advertisements following words 810 is indicated in field 811, referred to in IETF RFC 1583 as “#advertisements” (p. 179, IETF RFC 1583).... Words 830, shown in FIG. 8, compose a standard Opaque LSA header as described in IETF RFC 2370 (pp. 12-14, IETF RFC 2370). Words 830 include fields 831-836. Field 831 (LSA type) is set to 0xA to indicate that packet 800 contains an Opaque LSA. In one embodiment, field 832 (Opaque Type) and field 833 (Opaque ID) are set to 0xCE and 0x0010CF, respectively, to indicate that packet 800 contains information in accordance with the present invention....”).

210. Thus, it would have been obvious to a POSITA to combine Mitchell and Huai. In doing so a POSITA would have recognized that using the Opaque LSA message of an OSPF protocol would merely involve adding a well-known messaging technique to Mitchell’s fast link failover system. Further, such a POSITA would have understood that using Huai’s Opaque LSA message of an OSPF protocol in Mitchell’s system would have merely amounted combining known prior art elements according to known methods to yield the predictable result of notifying network elements that are interconnected of changes in the shared network topology.

211. In addition, a POSITA would have known that propagating failure information (e.g., circuit information) using opaque LSAs according to the OSPF protocol would merely be using one known messaging technique in a similar

network in a similar way to yield predictable results. Further, the POSITA would have understood that using opaque LSAs to propagate failure information, as described in Huai, in Mitchell's fast link failover system would only require applying a known communication protocol to a known network element to yield predictable results that allow a network to propagate failure information with a well-known messaging protocol.

212. Therefore, a POSITA would have been able to make the combination with a reasonable expectation of success.

E. Ground 5: Mitchell And Yamauchi Render Obvious Claims 8, 14, 20, And 24

1. Dependent Claims 8, 14, 20, And 24

- i. [8] The method according to claim 1, wherein the one or more of the links in **the first set of links** use a physical layer auto-negotiation protocol, and wherein said deactivating the physical layer resources comprises introducing an auto-negotiation error into the auto-negotiation protocol.

[14] The method according to claim 1, wherein a link in **the first set of links** includes a link which uses a physical layer auto-negotiation protocol, and wherein deactivating the physical layer includes introducing an auto-negotiation error into the auto-negotiation protocol.

[20] The network according to claim 16, wherein the one or more of the links in the first set of links use a physical layer auto-negotiation protocol, and wherein the network nodes are arranged to deactivate the physical layer resources by introducing an auto-negotiation error into the auto-negotiation protocol.

[24] The communication network of claim 16, wherein a link in the first set of links includes a link which uses a physical layer auto-negotiation protocol, and wherein deactivating the physical layer resources includes introducing an auto-negotiation error into the auto-negotiation protocol.

213. Mitchell in combination with Yamauchi render obvious the additional limitations of claims [8], [14], [20], and [24].

214. Mitchell does not explicitly teach deactivating the physical layer by introducing an auto-negotiation error in an auto-negotiation protocol. But it would have been obvious to do so because auto-negotiation protocol techniques were well-known.

215. For example, Yamauchi discloses an auto-negotiation monitoring system. EX1009, Abstract. Specifically, Yamauchi discloses detecting a link failure in an auto-negotiation monitoring system. EX1009, Abstract. Yamauchi explains that, in response to detecting a link failure is detected, its auto-negotiation monitoring system brings “a link of the port down and outputs an auto-negotiation code.” EX1009, Abstract. Yamauchi further describes that a downstream node notifies another network element “that ‘[t]he switch detects a failure and the link is

down.’, as an alarm,” for example, after the auto-negotiation code is received by a downstream network element. Further, Yamauchi discloses that the downstream network element can use an alarm to notify another network element further downstream. EX1009, Abstract (“The switch detects a failure and the link is down.’, as an alarm.”); *see also* EX1009, [0029] (“The 1000BASE-X auto-negotiation is performed by exchanging the auto-negotiation configuration code (/C/Ordered set) between the two network apparatuses. Where the auto-negotiation is finished, linkup is established between the two network apparatuses so that data transmission and reception can be performed therebetween. In a link_ok (linkup) state, the network apparatuses do not transmit the configuration code (/C/Ordered set).”), [0031] (“More specifically, where a failure occurs in an optical fiber on the reception side of one of the switches, the switch brings the link of its port down according to the definition described in Documents 1 and 2, and transmits the auto-negotiation configuration code (/C/Ordered set).”), [0037] (“As defined in Documents 1 and 2, the LAN switch performs LAN-frame transmission (transmission of /D/Code group, /S/Ordered set, and /T/Ordered set), or transmission of a code indicating the idle state (transmission of /I/Ordered set), where the port of the LAN switch is in the link_ok state. Where a failure is detected, the LAN switch brings the link of its port down and transmits the auto-negotiation configuration code (/C/Ordered set).”), [0046] (“The code

discriminator 111 performs code analysis for the input signal and transmits a notification to the alarm unit 113 upon receiving the auto-negotiation configuration code (/C/Ordered set).”), [0047] (“The code discriminator 112 performs code analysis for an output signal and transmits a notification to the alarm unit 113, where the auto-negotiation configuration code (/C/ordered set) is transmitted thereto.”), [0055] (“Where the failure occurs in the optical fiber 102 (a1 in FIG. 4), the switch 3 detects optical-signal shutoff or optical-signal deterioration on the reception side of the port 31 (a2 in FIG. 4), brings the link of the port 31 down (a3 in FIG. 4), and outputs an auto-negotiation configuration code (/C/Ordered set), as an LAN signal (a4 in FIG. 4).”).

216. As a result, it would have been obvious to combine Mitchell’s fast link failover system with Yamauchi’s auto-negotiation protocol monitoring system because using Yamauchi’s auto-negotiation protocol was a well-known technique. Indeed, a POSITA would have understood that modifying Mitchell’s system with Yamauchi’s system would merely involve deactivating a physical layer resource using an auto-negotiation error. In doing so, the POSITA would have known that such a deactivation would merely amount to combining known prior art elements using known methods to yield the predictable result of disabling a physical layer resource.

217. Further, a POSITA would have been motivated to combine Mitchell and Yamauchi because such a modification would have simply required using a single well-known implementation technique to deactivate a physical layer resource. A POSITA would have readily recognized that introducing such an auto-negotiation error would merely involve using a known auto-negotiation protocol technique to obtain predictable results. The POSITA would have further understood that introducing an auto-negotiation error using Yamauchi's auto-negotiation protocol would have just required applying a known communication protocol (e.g., an auto-negotiation protocol) to a known network element that was ready for improvement. In doing so, the POSITA would have known that the network element would have been improved through automation (e.g., by disabling the physical layer resource) to yield the predictable result.

218. Therefore, a POSITA would have been able to make the combination with a reasonable expectation of success.

F. Ground 6: Mitchell And McGee Render Obvious Claims 9, 15, And 25

1. Dependent Claims 9, 15, And 25

- i. [9] The method according to claim 1, wherein the one or more of the links in the first set of links comprise a logical link, and wherein said deactivating the physical layer resources comprises releasing an association of the logical link with the first instance of the communication service.

[15] The method according to claim 1, wherein a link in the **first set of links** includes a logical link, and wherein **deactivating** the physical layer includes releasing an association of the logical link with the first instance of the communication service.

[25] The network according to claim 16, wherein a link in the **first set of links** includes a logical link, and wherein **deactivating** the physical layer resources includes releasing an association of the logical link with the first instance of the communication service.

219. Mitchell in combination with McGee render obvious the additional limitations of claims [9], [15], and [25].

220. If Patent Owner contends that Mitchell does not expressly disclose “releasing” an association of a logical link responsive to deactivating a physical layer resource, doing so was well-known in the art. Indeed, McGee discloses releasing an association of a logical link responsive to deactivating a physical layer resource.

221. For instance, the VLAN logical links that are described by Mitchell are similarly disclosed in McGee. EX1005, 5:59-65 (“[I]n another embodiment, only those downstream links or ports associated with a virtual network *such as a VLAN* or VSAN which was predetermined or identified using configuration interface 202 are disabled.”); EX1010, [0006] (“LANs may also be *segmented into logical sub-networks called virtual LANs (VLANs)*, ...”), [0039], [0045] (“In the

network of FIG. 2B, however, ***the network is segmented into virtual LANs or VLANs***”), [0046], FIG. 2B.

222. In addition, McGee describes that logical VLAN links can be released when deactivating a physical link in response to detecting a link failure. *See, e.g.,* EX1010, [0109] (“Those of skill in the art will recognize an additional benefit to an embodiment of ***system 100 that includes support for VLANs such as that illustrated in FIG. 2B....*** Whenever application 600 either assigns to or recaptures resources from a team, the application 600 can instruct the teaming driver to ***either bind to or release from the team of drivers the driver associated with that resource.*** In addition, ***the resource must be either coupled to or decoupled from the VLANs to which that team is coupled.***”). Restated, McGee provides additional implementation details that describe releasing logical VLAN channels, for example, by decoupling the VLAN association with the physical layer link (*e.g.,* when a physical link is deactivated after a link failure).

223. It would have been obvious to a POSITA to combine Mitchell’s fast link failover system with McGee’s resource allocation application 600. More specifically, a POSITA would have recognized that would have been obvious to try to modify Mitchell with McGee because there are a finite number of ways to release an association of a logical link, and the application 600 would provide a mechanism to release the association with the logical VLAN link.

224. A POSITA would have been further motivated to modify Mitchell's fast link failover system with McGee's resource allocation application 600 to increase network efficiencies by preventing the system from underutilizing network resources. EX1010, [0018], [0079] ("Based on the current network usage, as well as the status of the available and allocated resources, an embodiment of the application 600 dynamically adjusts resource allocation...."), [0111] ("Load balancing teams can be managed to receive additional resources from a pool when saturated and to give up underutilized resources to the pool. And given that this additional release of logical link can be carried out using simple allocation application, a POSITA would have been able to make the combination with a reasonable expectation of success.

IX. Conclusion

225. It is my conclusion that the challenged claims 1-25 of the '637 patent are unpatentable.

X. Availability for Cross-Examination

226. In signing this declaration, I recognize that the declaration will be filed as evidence in a case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross examination in the case and that cross examination will take place within the United States. If cross examination is required of me, I will appear for cross

examination within the United States during the time allotted for cross examination.


XI. Right to Supplement

227. I reserve the right to supplement my opinions in the future to respond to any arguments that the Patent Owner may raise and to take into account new information as it becomes available to me.

XII. Jurat

228. 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 further that these statements were made with the 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.

DATE: May 30, 2024



Dr. Henry H. Houh