

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.,
Petitioner,

v.

ORCKIT CORPORATION,
Patent Owner.

Case IPR2023-00402
Patent 8,830,821

EXHIBIT 2025

DECLARATION OF MIGUEL GOMEZ

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. QUALIFICATIONS.....	1
III. BASES OF OPINIONS	8
IV. APPLICABLE LEGAL STANDARDS.....	8
A. Ordinary Skill in the Art	8
B. Claim Construction	10
C. Obviousness (35 U.S.C. § 103).....	10
V. OVERVIEW OF THE '821 PATENT AND THE STATE OF THE ART	14
A. Technical Background of the '821 Patent.....	14
B. Petitioner's Cited References	16
1. Doshi (Ex. 1005).....	17
2. Guichard (Ex. 1006).....	22
3. Huang (Ex. 1007).....	23
4. Xu (Ex. 1025)	25
VI. LEVEL OF ORDINARY SKILL.....	27
VII. CLAIM CONSTRUCTION.....	28
A. "Entity"	28
VIII. NO COMBINATION OF THE CITED PRIOR ART REFERENCES TEACHES OR SUGGESTS EVERY LIMITATION OF THE CLAIMS.....	28

A.	The Asserted References Do Not Teach Or Suggest “Selecting An Entity Pair From Said Plurality Of Transport Entities Based At Least In Part Upon Said Overall Cost” (the “selection limitation”).....	28
B.	The Asserted References Do Not Teach Or Suggest “Digital Logic Configured To Select Said Working Entity And Said Protection Entity From Said Plurality of Transport Entity Descriptors.”	35
C.	The Asserted References Do Not Teach Or Suggest “if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entitie of said entity pair”	35
IX.	PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA COMBINE DOSHI AND GUICHARD	37
X.	PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA WOULD ADAPT SUCH A COMBINATION.	43
XI.	PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA WOULD HAVE EXPECTED THE MODIFIED DOSHI-GUICHARD COMBINATION WOULD SUCCEED.	47
A.	The Asserted References Do Not Teach Or Suggest “Logic Configured To Determine A Probability Of Concurrent Failure Of Said Working Entity And Said Protection Entity”	51
XII.	DOSHI.....	54
XIII.	XU.....	56
A.	Claims 2 and 18 Are Patentable Over Doshi	59
B.	Claim 13 is Patentable Over Doshi	61

XIV. CONCLUSION	68
------------------------------	-----------

I. INTRODUCTION

1. My name is Miguel Gomez. I have been retained as an expert witness to provide my independent opinion in regards with matters at issue in the *inter partes* review of U.S. 8,830,821 (“the ’821 Patent”) in the IPR2023-00402 proceeding. I have been retained by Orckit Corporation (“Orckit”), the Patent Owner, in the above proceedings. Petitioner in this case is Cisco Systems, Inc. (“Cisco”).

2. Unless otherwise noted, the statements made herein are based on my personal knowledge, and if called to testify about this declaration, I could and would do so competently and truthfully.

3. My CV is being submitted herewith as Exhibit 2026 and is summarized in Section II, *infra*.

4. I am not a legal expert and offer no opinions on the law. However, I have been informed by counsel of the various legal standards that apply, and I have applied those standards in arriving at my conclusions.

5. The opinions I set forth herein are my own and are based on my education, experience, training, and skill I have accumulated.

II. QUALIFICATIONS

6. I received a Bachelor of Science in Electrical Engineering degree from Yale University in 1983. I have thirty-five years of experience in the design and development of hardware and software technology used in computer systems,

communications systems, networking and storage infrastructure, and database systems. My experience includes extensive knowledge of computer operating systems, programming languages, and mobile and embedded applications. I am also highly skilled in the use of microelectronics simulation software, ASIC and FPGA coding, in hardware design languages, including Verilog and VHDL. I'm also skilled at programming for embedded systems and applications in C , C++ and C#.

7. From August 2006 to March 2009, I was VP of engineering for ActSolar. ActSolar was a solar power converter and instrumentation company where I developed the system architecture that defined the hardware and software for a solar inverter and a data gathering and control sub-system. I wrote a control program in the C programming language to manage power conversion for the solar inverter. I also managed a small team of engineers that built the software defined radio responsible for retrieving solar panel data. The data gatherer utilized the IEEE 802.15.4 low-rate wireless personal area network (WPAN) standard to retrieve power and temperature readings from each solar panel and aggregate that data at a server. This information was used to maximize the solar array output.

8. In 2006, ActSolar was sold to National Semiconductor. The transfer of technology included the software radio for data collection, the inverter and a variant of the power converter hardware that performed shading compensation.

9. From September 2004 to August 2005, I was a consultant for BridgeWave Inc. I built a 10 Gbps Ethernet test platform to facilitate manufacturing tests of BridgeWave's high speed wireless network system. This system replaced the Smart Bits Ethernet test analyzer used in the manufacturing test environment. I also built a test generator, specifically an Ethernet packet generator, which generated packets of various sizes, address spaces, payload and checksums. I built a test receiver and analyzer which detected various failure mechanisms, maintained statistics and reported them to a user interface written in Lab View. A Xilinx Vertex II FPGA was employed in the first version of the test platform and a Xilinx Spartan 3e in the second version.

10. From September 2003 to March 2004, I was a consultant for PA Consulting Group. I provided hardware evaluation services for corporate mergers and acquisitions. These services included systems and hardware reviews in the following areas: (1) Analysis of system performance to the cost of the hardware implementation; (2) Review of the VHDL and Verilog code implementations, documentation and development strategies; (3) VHDL and Verilog code management including compilation processes, simulation, test coverage, bug tracking and source code control; (4) FPGA and ASIC tool chain management; (5) Circuit board design and layout for production environments; (6) Circuit board certification testing for FCC, UL, and Environmental tests; (7) Review of production

line management including assembly and test processes; and (8) Review of hardware code for copyright or license violations.

11. From January 2003 to May 2003, I was a consultant for Santel Networks. I developed an optical duo-binary (ODB) encoder. Operating at 10.7 Giga bits per second (GBps), ODB reduces the spectral spread to $\frac{1}{2}$ of the ordinary NRZ format. This format allows high spectral efficiency over long distances by reducing nonlinear penalties due to narrower channel spacing and allows a higher bit rate or increases transmission distances. The project was implemented in a Xilinx FPGA running at 667 Mhz with a 10.7 GHz Mux/Demux feeding 32 LVDS streams to the FPGA. The ODB behavior was simulated in a C-coded simulation. The circuit board's highest clock speed was 10.7 Gigahertz and was modeled in HSPICE. Board level simulation was performed using NCVerilog. Several patents were applied for and the board was shown at the Optical Fiber Communication Conference & Exposition in March of 2003.

12. From March 2001 to December 2002, I was Director of Hardware, Content Networking Division for Extreme Networks, Inc. I managed a team of hardware and software engineers that developed and maintained a Layer 2-5 Content Addressable Switch. The switch was capable of L2, L3, L4 switching as well as L5 switching. In the L5 mode, the SLB used HTTP and HTML headers to make routing decisions. All routing decisions were made in real time (1 GBPs) by hardware logic

in FPGAs while a Power PC provided a user interface, configuration management and housekeeping diagnostics. Communications between the hardware elements and the PowerPC CPU was via a PCI bus. I was responsible for the system design as well as for managing the product manufacturing process.

13. From January 2000 to March 2001, I was Director of Hardware for Webstacks, Inc. (now Extreme Networks, Inc.). Extreme Networks and several venture founding partners commissioned Webstacks to build a Content Addressable Switch which would operate as a L2-5 server load balancer. This switch included content routing, load balancing, firewall and security capabilities utilizing Secure Sockets Layers (SSL). I was hired to be the system architect. My responsibilities included hiring the hardware and software engineers that would implement the server load balancer as well as setting up the manufacturing process. I hired 13 ASIC, FPGA and software engineers into my department and managed them for the duration of the project. Within one year, our team had developed the product and met Extreme Networks' acquisition criteria. Extreme Networks purchased Webstacks for \$68MM cash and stock.

14. From February 1997 to December 1999, I was a consultant for Philips Semiconductor. I developed Philips Semiconductors' portion of a certification environment used by Philips Semiconductor and Microsoft to verify operation of Philips' Poseidon handheld chipset running Windows CE. This included board

development, C coding, FPGA coding and simulation as well as porting to Windows CE. Later, I supported the development of Philip's version of the MIPS 4000 with communications and storage devices to emulate a system on a chip. I also designed ASIC logic and board level hardware for a 56Kbs software modem while at Philips.

15. From September 1994 to January 1997, I was President and Founder of Minden Group, Inc. The Minden Group developed and sold several types of memory adapters and a video conferencing system. The adapters were designed to expand the memory capacity of Personal Computers while saving the customer money. Both the adapters and the video conferencing system were sold through retail stores throughout the United States and Canada. Over 250,000 memory adapters were sold at stores such as Fry's, CompUSA, Computer City and Future Shop. After developing these products, I organized and managed a team of 10 to 15 employees focused on sales, marketing, finance and production. Manufacturing of the primary components was performed through contract manufacturing in Taiwan with final assembly and distribution in the United States. Distribution included order fulfillment, contracted sales channel distribution networks and end-point retail management.

16. From August 1992 to September 1994, I was Senior Hardware Manager of RAID Product Development for MTI, Inc. In this capacity, I was responsible for architectural development of the next generation RAID 1-5 systems and for

managing the product development group. I introduced several design approaches that resulted in lower product costs as well as methods for interfacing with the company's other hierarchical storage components.

17. From April 1989 to August 1992, I was Founder and President of Spectrum Analysis, Inc., a consulting services company. I founded and operated Spectrum Analysis, Inc. with three other partners. Spectrum Analysis, Inc. specialized in electrical circuit design with emphasis on FPGA and ASIC emulation and complex PCB level designs.

18. From September 1988 to March 1989, I was Applications Manager for Quickturn Systems, Inc. Quickturn developed an ASIC emulator using Xilinx FPGAs. I was hired as Quickturn's first applications engineer, and was initially responsible for developing the training and demonstration systems for Quickturn's initial customers. Later, as the company matured, I managed the customer support team. Because this ASIC emulator worked at a fraction of clock speed, I developed a methodology of slowing board designs to a fraction of the original clock speed, (on the order of 100s of KHz) while maintaining clock and memory (DRAM) integrity.

19. From July 1987 to August 1988, I was an Applications Engineer for Telestream Corporation and was responsible for demonstration system development, training and sales support for a software based communications protocol converter product. From March 1985 to July 1987 I was a design engineer for StrataCom Inc.,

an early voice and data network equipment manufacture. StrataCom developed the IPX, a packet over leased T1 transmission and management system. I was responsible for the voice compression card as well as designing a way to include data over a primarily voice solution. The hardware I developed eventually evolved into the Frame Relay protocol. Cisco bought StrataCom in 1996.

20. From June 1983 to March 1985, I was a Design Engineer for ROLM Corporation. I was responsible for the development of a multiplex/demultiplex hardware board used in the ROLM BUS 295 telephone switch. A novel elastic buffer was used to automatically determine the length of an inter-cabinet data bus, whose length varied based on the customer's requirements. This eliminated the need for manufacturing technicians to manually adjust clock skew.

21. A more detailed account of my work experience and other qualifications is listed in my Curriculum Vitae attached as Exhibit 2026 in this IPR.

III. BASES OF OPINIONS

22. In the course of conducting my analysis and forming my opinions, I have reviewed the exhibits submitted to date in this IPR, as well as the Petition, Preliminary Response and Institution Decision.

IV. APPLICABLE LEGAL STANDARDS

A. Ordinary Skill in the Art

23. My opinions in this declaration are based on the understandings of a person of ordinary skill in the art, which I understand is sometimes referred to as an “ordinary artisan” or by the acronym “POSITA” (person of ordinary skill in the art), as of the time of the invention. For purposes of this Declaration, I have been instructed to use April 7, 2006 as the time of invention. I understand that the person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention. By “relevant,” I mean relevant to the challenged claims of the ’821 Patent.

24. I understand that, in assessing the level of skill of a person of ordinary skill in the art, one should consider the type of problems encountered in the art, the prior solutions to those problems found in the prior art references, the rapidity with which innovations are made, the sophistication of the technology, the level of education of active workers in the field, and my own experience working with those of skill in the art at the time of the invention.

25. I was at the time of invention, and am currently, one of more than ordinary skill in the art through my education and research experience. As of the date of the invention, I am very familiar with the types of problems encountered in computer networking , the types of prior art solutions described in prior art references, and the rapidity at which innovations are made. Indeed, I am very familiar with people having this level of skill in the area of computer networking and

am very familiar with commonly used technologies as Ethernet and the internet protocol / transmission control protocol (TCP/IP) suite. At the time of the invention, and since that time, I have been involved in the design and development of computer networks and in that capacity have worked with and managed teams of numerous individuals that would qualify as a POSITA.

B. Claim Construction

26. I understand that claims of the patent-at-issue in this IPR are generally interpreted according to their ordinary and customary meaning taking into consideration the so-called “intrinsic evidence” of the patent consisting of (1) the claim language; (2) the specification; and (3) the prosecution history.

27. I understand that claim terms may be explicitly defined in the patent specification, or they may be implicitly defined through consistent usage in the specification. I also understand that the scope of claim terms may be limited by statements in the specification or prosecution history where the applicant clearly disavows or disclaims subject matter.

C. Obviousness (35 U.S.C. § 103)

28. I have been informed that a patent may be invalid if the claimed invention would have been obvious at the time the invention was made to a person having ordinary skill in the art. I have been informed that the following factors must

be evaluated to determine whether Petitioner has met its burden of proof that a claimed invention is obvious:

1. The scope and content of the prior art relied upon by Petitioner;
2. The difference or differences, if any, between each claim of the patent and the prior art;
3. The level of ordinary skill in the art at the time the invention of the patent was made; and
4. Any objective indicia of non-obviousness, including, for example:
 - (1) commercial success of an embodiment; (2) a long-felt need; (3) skepticism; (4) failure by others to find the solution provided by the claimed invention; (5) copying by others of the subject matter of the claim invention; (6) unexpected results of the claimed invention; (7) acceptance of others and industry praise; and (8) licensing of the patents.

29. I have been informed that a claim is not proved obvious merely by demonstrating that each of the elements was independently known in the prior art. I have been informed that many, if not all, inventions rely on building blocks already known, and claimed inventions almost of necessity will likely be combinations of what is already known.

30. I have been informed that a combination of prior art references must disclose, teach or suggest all the recited limitations of the claim for the invention to be found obvious.

31. I have been informed that it is important in the obviousness inquiry to identify whether a reason existed at the time of the invention that would have motivated a person of ordinary skill in the art in the relevant field to combine or modify the prior art references in the manner proposed by the Petitioner to arrive at the claimed invention. Put another way, a finding of obviousness should be supported by an apparent reason to combine or modify the prior art references as proposed by the Petitioner.

32. I have been informed that it is important in the obviousness inquiry that it is understood how the combination of references is supposed to work. An explanation of the operation of the combined references is often a prerequisite to showing that a person of ordinary skill in the art would have been motivated to make the proposed combination and would have had a reasonable expectation of success.

33. I have been informed that a finding of obviousness “requires the additional showing that a person of ordinary skill at the time of the invention would have selected and combined those prior art elements in the normal course of research and development to yield the claimed invention.” *Unigene Labs., Inc. v. Apotex, Inc.*, 655 F.3d 1352, 1360 (Fed. Cir. 2011). In addition, I have been informed that a

person of ordinary skill in the art should have had a reasonable expectation of success in combining the references as proposed.

34. In assessing obviousness, I have been instructed to consider both the ordinary creativity and common sense of the person of ordinary skill in the art. I also understand that it is impermissible for common sense to be applied to fill gaps in prior art that fails to teach or suggest a limitation of the claim. I also understand that the obviousness inquiry should guard against hindsight bias or hindsight reconstruction where after-the-fact reasoning is applied to combine prior art elements using the claimed invention as a template, without establishing that, as of the date of the invention, there exists a motivation to combine or apparent reason to combine the prior art as proposed.

35. In assessing obviousness, I have been instructed that, to qualify as proper prior art for an obviousness analysis, a reference must qualify as analogous art. I have been informed that a reference qualifies as analogous art with respect to the claims if it is either: (1) from the same field or endeavor as the patent; or (2) the reference is reasonably pertinent to the particular problem addressed by the invention. I have also been informed that for a reference to be reasonably pertinent, it must logically have commended itself during the ordinary course of development to an inventor's attention in considering his problem.

36. In assessing obviousness, I have also been instructed that, as noted above, secondary considerations can often be probative evidence of non-obviousness and can serve to avoid hindsight bias. I have been instructed that secondary considerations can be sufficient to overcome a prima facie case of obviousness. I have been instructed that Patent Owner can show non-obviousness if: (1) the industry praised their work; (2) their design was copied by others; (3) their design achieved a high level of commercial success; or (4) their invention had unexpected beneficial results, for example.

V. OVERVIEW OF THE '821 PATENT AND THE STATE OF THE ART

A. Technical Background of the '821 Patent

37. The general field of the '821 Patent is computer networking. The '821 Patent teaches methods and systems implementing high availability and low operating costs in a Mult-Protocol Label Switched (MPLS) network. In particular, the '821 Patent “includes a method to provide a high probability of sub-50 ms protection against multiple span or node failures...” Ex1001 ['821 Patent], 3:15-17. “In addition, the method attempts to optimize network efficiency by minimizing the cost of an active network.” *Id.*, 3: 18-20.

38. The '821 patent recognizes that pre-existing resiliency solutions are available for most networks, such as TCP/IP, but are deficient in managing high availability. “[M]ost network protocols, for example TCP/IP, have integral recovery

mechanisms, such as retransmission of dropped packets, that are intended to minimize disruption of network traffic during a switchover from active to standby equipment or connections. ... However, if the disruption is more prolonged, connections may drop.” Ex1001 [’821 Patent], 1:30-37.

39. High availability is especially important for loss intolerant networks such as fiber networks, or for network traffic such as Voice over IP or streaming Video. The criticality of network delivery in high capacity networks requires a network resilient to component failures, and in the case of a failure, the ability to recover within a defined time limit. This is (and was in the 2011 time frame) an area of extreme interest. Ex1001 [’821 Patent], 1:23-24. (“Availability of network connections has been the focus of increasing interest in the telecommunications field.”).

40. Multi-Protocol Label Switching (MPLS) provides an alternate routing mechanism to route traffic at the speeds required for these high capacity and traffic critical delivery networks. Multi-Protocol Label Switching (MPLS) simplifies routing decisions by introducing a more efficient and flexible method for forwarding packets within a network. Traditional IP routing relies on examining the IP header at each router to make forwarding decisions based on the destination IP address. MPLS, on the other hand, uses labels to simplify this process. The labels enable a router to expediently move traffic along predetermined routes. The ‘821 patent

teaches a method of ensuring that the traffic continues unabated in the event of a network change such as a router failure, a change in the network configuration or a change in the costs associated with a network span. Ex1001 ['821 Patent], 3:15-22. The '821 Patent provides examples of several MPLS linear protection scheme, including those defined in IETF draft-ietf-mpls-tp-linear-protection, IETF draft-zulr-mpls-tp-linear-protection-switching, and ITU-T G.8131 T-MPLS Linear Protection, but recognizes that “other protection switching mechanisms may be employed... .” Ex1001 ['821 Patent], 3:54-62. The '821 patent recognizes that high availability in a managed network is only one of several metrics that should be used to determine the route between the start point and the end point. In making the optimal selection of a start point to endpoint connection, the '821 patent teaches that other cost considerations should be included including route costs, determination costs and, for example, metrics from the Interior Gateway Protocol. Ex1001 ['821 Patent], 3:36-40.

B. Petitioner's Cited References

41. The Petition asserts four grounds, as shown below:

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


Pet., 14. In each ground, Doshi is the primary reference.

1. Doshi (Ex. 1005)


42. Doshi (Pat. Publ. No. US2004/0205239) is directed to a “method for determining primary and restoration paths for a new service in a mesh network.” Ex. 1005 [Doshi], Abstract. This method involves “(1) for each of a plurality of candidate primary/restoration path pairs for the new service, generating a path cost for each candidate pair, where the path cost for each restoration path is a function of the sum of the cost of links within the restoration path, and (2) selecting the primary and restoration paths for the new service from the plurality of candidate path pairs based on the path cost.” *Id.* These paths are determined between network nodes.

43. Doshi Figure 1 (below) illustrates the concept of primary (or working) LSPs and protection (or restoration) LSPs between different nodes N1 through N6. Fig. 1 illustrates five primary LSPs LSP-1 through LSP-5 and five “protection paths” LSP-6 through LSP-10. Ex. 1005 [Doshi], ¶¶[0052]-[0054]. Each primary LSP corresponds to one protection LSP that creates a path between the same nodes.

Primary LSP-1 corresponds to protection LSP-6, and so on. For example, in the figure below, primary LSP-4 is highlighted yellow and runs directly between nodes N2 and N5. LSP-4's corresponding protection node LSP-9 is highlighted green and also runs between nodes N2 and N5 but uses a different path (i.e., N2-N1-N4-N5).




44. In Doshi, two different paths can be “routed disjointly” or can “share common links” (e.g., L1-L7 in Fig. 1) along their respective routes. *Id.*, ¶[0040]. Using the example of Doshi Fig. 1, above, paths LSP-4 and LSP-9 are “routed disjointly” because they do not share any of links L1-L7. On the other hand, in the annotated example of Doshi Fig. 1 below, there are two paths between nodes N1 and N6. Path 1 traverses links L2, L5 and L7. Path 2 traverses links L2, L4 and L6. Thus, these two paths are “shared” paths because they both share link L2.



45. To identify a primary/restoration path pair, Doshi must evaluate a number of potential paths. In the example above, Path 1 and Path 2 are two paths between nodes N1 and N6, but there are other potential paths, including L1-L3-L6 or L1-L3-L4-L5-L7. In more complex networks, there could be numerous different paths to choose from. Thus, there must be a selection to determine which two paths to use as the primary and protection path pair.

46. Doshi teaches a selection process based on path cost. Ex. 1005 [Doshi], Abstract. Doshi Fig. 11 (below) illustrates this process. Fig. 11 calculates “a minimum-cost primary/restoration path pair for a new service in consideration of link utilization, utilization threshold, and sharing degree.” Ex. 1005 [Doshi], ¶[0142].



47. First, in step 1100 the bandwidth required and source and destination IDs for the new service are input. *Id.* Step 1102 determines a group of “minimum cost primary paths” using “Option=1 (i.e., no sharing considered).” *Id.*, ¶[0143]. Step 1104 determines a group of “minimum-cost restoration paths ... using Option=2 (e.g., where sharing is considered and the appropriate primary path is passed to the PathCost routine.)” *Id.* Thus, at Step 1104, a group of primary paths have already been selected and Doshi will consider shared paths for the restoration

path associated with each primary path. In Step 1106, the path costs for each primary and restoration path pair are summed together and the lowest pair is saved. *Id.*

48. Steps 1110 through 1114 calculate path cost via a second method. *Id.*, ¶[0142]. In Step 1110, cost is determined for “L disjoint lowest total-cost pairs of paths” with no sharing considered. *Id.*, ¶[0144]. In step 1112, each restoration path identified in step 1110 has its cost recalculated using Option 2 (sharing considered). The sum of the lowest pair of primary path (from the group of L primary paths) and restoration paths (from the group of L restoration paths) is stored. *Id.* In step 1118, whichever cost is lowest between the two methods is selected as the minimum-cost pair. *Id.*, ¶[0145].

49. In certain embodiments, Doshi describes selecting primary LSPs and protection LSPs that are “strictly disjoint” from each other. Ex. 1005 [Doshi], ¶[0055]. “[S]trictly disjoint means that the two paths in each primary/protection path pair share no common links or nodes other than their common ingress and egress nodes.” *Id.* However, “this need not be the case in all implementations of” Doshi and in some embodiments, “no strictly disjoint path exists between two nodes” or there may be other factors or constraints that lead to selection of a protection path “that is not strictly disjoint from the primary path (e.g., the protection path may share some common links with the primary path)”.

2. Guichard (Ex. 1006)

50. Guichard (“Definitive MPLS Network Designs”) appears to be a brief excerpt from a book titled “Definitive MPLS Network Designs.” Ex1006 [Guichard], 1. The section in the provided excerpt is titled “Reoptimizing a Traffic Engineering LSP.” *Id.*, 3. It discusses the need to identify more optimal working LSPs in a network as they become available and of establishing the new working LSP before tearing the old working LSP down *Id.*, 3-5. The excerpt of Guichard provided by Petitioner does not address finding optimal protection LSPs, as it is directed to recovery systems based on “on-the-fly” restoration, rather than selection of a protection path.

51. In a different part (identified in an excerpt provided by the Patent Owner), Guichard makes clear that “on-the-fly” restoration—as discussed in Petitioner’s excerpt of Guichard describing optimization—belongs to a different “famil[y] of network recovery mechanisms” than network protection through precomputed backup paths, to which the ’821 Patent methods and systems are directed. Ex2028 [Guichard], 96.

52. Guichard also explains that there are several different criteria and network objectives to consider in choosing between different network recovery mechanisms, including the “scope of recovery” (i.e., whether the recovery mechanism includes “link failures”), quality of service during failure (i.e., whether


the backup path “provide[s] equivalent QoS”), network overhead (i.e., the extra information or “state” required by the recovery mechanism), cost, and network stability. Ex2028 [Guichard], 96.

53. Guichard explains further that though combining recovery mechanisms may be desirable in some cases, it might also require “extra care to avoid race conditions between recovery mechanisms and double protection, which may lead to a lack of optimality in terms of backup resources.” Ex2028 [Guichard], 96. Thus, Guichard lays out a number of considerations that a POSITA would have to consider before selecting between different network recovery mechanisms as well as combining such mechanisms to create a hybrid recovery mechanism.

54. Notably, Petitioner’s expert Dr. Houh did not include or address any of these teachings from Guichard in his declaration. *See generally* Houh Declaration.

3. Huang (Ex. 1007)

55. Huang (U.S. Pat. Pub. 2003/0117950) is directed to identifying backup paths for network routing in case of “the event of a failure in the fiber.” Ex. 1007 [Huang], Abstract. Huang Fig. 5 (below) “illustrates steps performed by a head end node in a link redial method” in one embodiment of Huang. *Id.*, ¶[0024].



56. At step 502, the node receives a link failure notification. *Id.*, ¶[0048].

At step 504, the node selects a “backup bundle” (that is, a group of backup LSPs). *Id.*, ¶[0048]. At step 506, the node determines “whether the protection bandwidth [i.e., the amount of bandwidth needed for a proposed level of protection] is in use on the first outgoing link.” *Id.* If the bandwidth is already in use, then another “backup bundle” is selected. If the bandwidth is not in use, the node “activates the backup connection map (step 508)” and “mark[s] the protection bandwidth on the first outgoing link as being used (step 510).” *Id.* In step 512, “the interrupted traffic [is] bridged to the backup bundle.” *Id.* Finally, at step 514, the node “sends a Link Label Restore Request message to the first node 102C on the backup route.” *Id.* This message ultimately allows the node to “begin[] sending traffic over the backup route.” *Id.*, ¶[0052].

4. Xu (Ex. 1025)

57. Xu (U.S. Pat. Pub. 2011/0141877) is directed to a network load balancing method that “uses fixed paths and a fixed path-splitting strategy to enable ingress routers to efficiently reroute traffic after a failure.” Ex. 1025 [Abstract]. “An off-line management system computes a set of fixed paths and a set of splitting ratios for routing demand from ingress routers to egress routers, with sufficient capacity to meet demands under each failure scenario.” *Id.*, “That data is then used by the ingress router to reroute demand after observing a failure.” *Id.*

58. In particular, Xu teaches load balancing according to the following method:

At an ingress router u_d , traffic is routed by automatically balancing traffic load over a predefined set of paths P_d from the ingress router u_d to the egress router v_d according to a set of splitting ratios α_p . Each splitting ratio defines a fraction of the demands d to be transmitted over a path p of the set of paths P_d in a case of a detectable failure state o corresponding to a subset of a set of modeled failure states s .

59. Ex. 1025 [Xu], ¶[0006]. Thus, the routers of Xu transmit traffic over multiple paths p simultaneously according to “splitting ratios” that define the fraction of traffic sent over multiple paths p . If one or more paths p fail, the fraction

of traffic sent over other paths p is altered in accordance with the relevant splitting ratio.

60. Xu explains the advantages of its routing scheme as follows:

Multipath routing between each pair of ingress and egress routers is utilized. Such end-to-end routing has several benefits. First, these routes do not need to change when a failure is detected, which saves time, reduces overhead, and improves path stability. Second, end-to-end load balancing spreads the traffic in the network more effectively than local rerouting. Third, it enables faster recovery and lower protocol overhead than conventional link-state routing protocols, like Open Shortest Path First (OSPF), which not only rely on flooding link-state advertisements and re-computing shortest paths but also suffer from transient forwarding loops. A fourth benefit is the solution's simplicity; most of the functionality is incorporated in the Network-Management Software (NMS) rather than the network elements. The NMS is responsible for selecting the end-to-end paths and calculating the path-splitting parameters for each router. The solution has a minimalist control plane, used only for failure detection, which leads naturally to a simpler network where smart, expensive routers are replaced with cheaper routers and/or switches

with a limited feature set.

Id., ¶[0030]. Thus, Xu teaches its multipath routing solution has a number of benefits over local routing techniques that need to change paths when a path fails.

61. Xu teaches the detection of “failure states.” *Id.*, ¶[0005]. “The set of modeled failure states are states that cause at least one path p of the predefined set of paths p_d to fail.” *Id.*, ¶[0006]. “Each modeled failure state has at least a predetermined probability of occurrence.” *Id.* This “predetermined probability of occurrence may be sufficient to meet a predetermined network availability objective when the network is restored against the set of modeled failure states.” *Id.*, ¶[0007]. However, not all possible failure states are included in the set of modeled failure states, and only the most probable are included. *Id.*, ¶[0049] (“Of course, any subset has a non-zero probability of simultaneous failure, but because of the monstrous size of ‘all’ subsets, for practicality the set is usually restricted to the most probable.”). Thus, Xu relies on simple “heuristic” techniques to approximate optimal performance. *Id.*, ¶¶[0057], [0065].

VI. LEVEL OF ORDINARY SKILL

62. The Petition and Petitioner’s expert assert:

A Person of Ordinary Skill in The Art (“POSITA”) on June 22, 2011, would have had a working knowledge of network communications and multiprotocol label switching (“MPLS”) techniques available at the time. A POSITA would have had a bachelor’s degree in computer

science, computer engineering, electrical engineering, or equivalent training, and approximately two years of experience working in the field of network communications and been knowledgeable regarding MPLS techniques. Lack of professional experience can be remedied by additional education, and vice versa.

Pet., 12. I do not dispute Petitioner’s level of ordinary skill for purposes of this declaration.

VII. CLAIM CONSTRUCTION

A. “Entity”

63. The Petition proposes a construction of the term “entity” that “includes an MPLS LSP or a PW.” Pet., 13. For purposes of this declaration, I have not provided any opinions on this construction.

VIII. NO COMBINATION OF THE CITED PRIOR ART REFERENCES TEACHES OR SUGGESTS EVERY LIMITATION OF THE CLAIMS

A. The Asserted References Do Not Teach Or Suggest “Selecting An Entity Pair From Said Plurality Of Transport Entities Based At Least In Part Upon Said Overall Cost” (the “selection limitation”)

64. Independent claims 1 and 17 require “selecting an entity pair from said plurality of transport entities based at least in part upon said overall cost” (the claimed “selection step”). Ex. 1001 [’821 Patent], claims 1 and 17. Petitioner asserts that Doshi renders this element obvious in view of Guichard. Pet., 35-40, 65. In my opinion, Petitioner’s assertion fails because Doshi discloses a simplified selection in

which the number of available pairs does not correspond to “said plurality of transport entities.”

65. A POSITA would understand the claimed selection step—which requires selecting an *entity pair* from “*said* plurality of transport entities”—in the context of express requirements of claim 1 that precede, provide antecedent basis for, and otherwise set up the claimed selection step. Those express requirements include (i) “providing a plurality of ... transport entities between a first endpoint and a second endpoint” and (ii) “determining an overall cost *for each entity pair of said* plurality of entities.” Ex1001 [’821 Patent], claims 1 and 17. Thus, the claims require the selection is made from each of the “entity pairs of said plurality of entities”, for which the overall cost has been determined.

66. Petitioner asserts that Doshi discloses the claimed selection step. Pet. 35-36 (“Doshi discloses selecting an LSP pair from the plurality of MPLS LSP...”). In my opinion, Doshi does not disclose selecting an entity pair from an already provided plurality of entities or the pairs associated with *said* plurality; rather, Doshi teaches a simplified selection process in which an entity pair is selected from a narrow subset of the total number of candidate pairs associated with *said* plurality of entities.

67. In particular, Doshi discloses calculating two sets of entities (or paths), one set of K primary paths and one set of K restoration paths, as shown in Figure 11, steps 1102 and 1104:




FIG. 11

See also Ex. 1005 [Doshi], ¶[0143]. To provide a numerical example, consider if K is set to 10, step 1102 chooses 10 of these links to serve as potential minimum-cost primary paths. Step 1104 then chooses 10 links to serve as potential minimum-cost restoration paths.

68. At steps 1106-1108, Doshi instructs calculating the cost for K pairs of paths and selecting a pair from those K pairs; thus, for a plurality of K entities,

Doshi instructs a POSITA to create a set of candidate pairs with members. To illustrate with a numerical example, if there are 10 potential entities in each set (i.e., $K=10$), Doshi teaches a POSITA to calculate 10 candidate pairs at steps 1106-1108.

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

70. To further illustrate the difference between Doshi and the '821 patent, if the same POSITA were to ignore the primary/restoration path designation (i.e., treat pairs "AB" and "BA" as interchangeable, or $\{A, B\} = \{B,A\}$) then the same

POSITA would have 45 candidate path pairs to choose from, i.e., 4.5 times the number of candidate paths under Doshi.).


71. It would be immediately apparent from this to a POSITA that Doshi teaches a selection process that substantially curbs the size of the candidate selection set and is effectively a shortcut, compared to the method taught by the '821 patent and required by claim 1. In other words, instead of examining every possible pair of paths (as in the '821 Patent), Doshi only looks at K pairs of paths.

72. Doshi discloses the details of its selection step, which further illustrate that it does not select an entity pair from *said* plurality of entities and associated pairs. Doshi's method first calculates a group of potential primary paths, which may be a subset of the total provided paths. Thus, Doshi does not disclose selecting primary paths from a plurality of the provided primary paths.

73. Moreover, Doshi discloses calculating a set of secondary "restoration paths" that correspond to the primary paths. Even assuming the primary paths are selected from *said* plurality of transport entities, there is no disclosure in Doshi that the restoration paths are calculated from the same *said* plurality of primary paths. In other words, Doshi's set of K minimum-cost primary paths is different than Doshi's set of K minimum-cost restoration paths. Thus, Doshi does not disclose selecting a *pair* of entities from "a plurality of transport entities" because the restoration path is not selected from the same plurality of transport entities as the primary path.

74. Further, even assuming that the total number of paths K is the same as the plurality of paths provided, the pair is still not selected from the same plurality of paths because, as explained *supra*, such a selection would require determination of “each entity pair” for the total plurality of paths, a condition that Doshi explicitly does not meet.

75. Doshi likewise fails to disclose a selection from *said* plurality of entities with respect to the selection process shown at Doshi Fig. 11, at step 1116 shown below:



76. In particular, at 1116, Doshi discloses selecting an entity pair (*i.e.*, a primary and restoration path) from a group of L primary/restoration path pairs. Ex. 1005 [Doshi], ¶[0144], Fig. 11. Those L path pairs are described as the “L disjoint lowest total-cost pairs of paths” at step 1110. Because the disjoint path pairs are necessarily a subset of the total number of path pairs (since not all pairs are disjoint), this selection likewise is made from a set of paths that is less or equal than the plurality K. Accordingly, for the same reasons that Doshi’s first selection (*i.e.*, with respect to K paths and pairs) fails to disclose the claimed selection step, Doshi’s second selection likewise fails to disclose this step.

77. Lastly, Doshi fails to disclose the claimed selection step from *said* plurality of entities for the separate reason that—both as to its selection of a pair from the K and L groups—it discloses, at step 1116, selecting a pair either from the K group or from the L group, not from the full plurality of provided paths, as required by the claimed selection step. Ex1005 [Doshi], ¶ [0144]. Indeed, the claim language expressly requires selecting an entity pair from *said* plurality of entities (*i.e.*, all the pairs that can be made from that plurality), not from a subset of entity pairs that is derived from that plurality (e.g., the subsets of K and L pairs created in Doshi).

78. Accordingly, because Doshi discloses a shortcut selection that is materially different from the claimed selection step of claim 1, in my opinion, Doshi does not disclose or render obvious this limitation or claims 1 and 17.

B. The Asserted References Do Not Teach Or Suggest “Digital Logic Configured To Select Said Working Entity And Said Protection Entity From Said Plurality of Transport Entity Descriptors.”

79. Similar to independent claim 1, independent claim 14 requires “a data structure comprising a plurality of transport entity descriptors” along with “digital logic configured to *select said working entity and said protection entity from said plurality of transport entity descriptors.*” Ex. 1001 [’821 Patent], claim 14. Here, the working and protection entity are equivalent to the entity pair of claim 1 and the transport entity descriptors correspond to the transport entities of claim 1. Petitioner asserts that the same combination (Doshi in combination with Guichard) renders this limitation obvious, relying to its analysis of claim 1 in support. Pet., 78-79.

80. Given that there are no substantial differences between this limitation and the selection step of claim 1 and Petitioner asserts the same arguments, Petitioner’s challenge to claim 14 fails for the same reasons as for claims 1 and 17.

C. The Asserted References Do Not Teach Or Suggest “if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair”

81. Independent claims 1 and 17 require that “if an entity pair reselection event occurs, reselecting said entity pair from the group consisting of said entity pair and a replacement entity pair comprising at least one entity distinct from the entities of said entity pair” (the claimed “reselection”). Ex1001 [’821 Patent], claims 1 and 17.

82. Petitioner does not contend that either Guichard or Doshi (the references on which its obviousness grounds are based) disclose the claimed reselection of an entity pair. Rather, Petitioner makes the following assertions in support of its obviousness challenge based on the Guichard/Doshi combination: (i) Guichard “teaches event-driven reoptimization ... upon the occurrence of a particular event” (e.g., “each time a new link is advertised as newly operational”) and (ii) “it would have been obvious to a POSITA, in view of Guichard, to perform event-driven reoptimization, when implementing Doshi’s MPLS architecture.” Pet., 40-41.

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

path (as disclosed in Guichard) and a protection path (which is not disclosed in either Doshi or Guichard).

84. In my opinion, Petitioner has failed to show that a POSITA would have found either of these two steps obvious. First, Petitioner fails to provide a reason for which a POSITA would have combined Doshi and Guichard. Second, even assuming a POSITA would have combined Doshi with Guichard to practice primary path reoptimization, Petitioner fails to contend—or to provide a reason for which—a POSITA would adapt the Doshi/Guichard combination to practice the claimed reselection of an entity pair (i.e., “reselecting said entity pair”). Third, Petitioner has failed to provide a reason for which a POSITA would have had a reasonable expectation that such an adapted combination would succeed in the manner in which it is set forth in the claims of the ’821 patent.

IX. PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA COMBINE DOSHI AND GUICHARD

85. Petitioner asserts that a POSITA would have combined Guichard and Doshi because (i) Guichard teaches that reoptimization is “highly desirable” (Petition at 20) and (ii) a “POSITA considering Doshi would have naturally considered the teachings of Guichard” because both references “generally pertain to MPLS networks and address the problem of selecting LSPs” with a low cost. Petition at 20-21. In particular, Petitioner asserts that the “combination represents what would have been common sense to a POSITA—that it is desirable to evaluate

whether selected LSPs remain optimal and to reoptimize the selection if a better option is available.” Petitioner’s assertions are incorrect and, in any event, fail as a matter of law.

86. That both Guichard and Doshi are purportedly directed to an efficient and reliable MPLS network does not (and cannot) in my opinion provide a reason to combine these two specific references. These needs—for efficiency and reliability—are common to virtually every innovation in the field of MPLS networks. Reliability and efficiency are generic goals recognized by early MPLS specifications such as the RFC 3469, “Framework for Multi-Protocol Label Switching (MPLS)-based Recovery”, published February 2003, which states on page one: “MPLS may facilitate the convergence of network functionality on a common control and management plane. Further, a protection priority could be used as a differentiating mechanism for premium services that require high reliability, such as Virtual Leased Line services, and high priority voice and video traffic.” Ex. 2032 [RFC 3469], 1. On page three of the same document: “An important component of providing QoS, however, is the ability to transport data reliably and efficiently.” *Id.*, 3.

87. Continual commitments to these goals are expressed in later MPLS publications such as RFC 6372, “MPLS-TP Survivability Framework”, published September 2011. RFC 6372 states “[t]he functionality provided is intended to be

similar to or better than that found in established transport networks that set a high benchmark for reliability” (pg. 1) and “[t]his framework applies to general recovery schemes as well as to mechanisms that are optimized for specific topologies and are tailored to efficiently handle protection switching.” (pg. 8). Ex. 2033 [RFC 6372], 1, 8.

88. These are generic needs applicable to all networks. I understand that generic needs applicable to all references are insufficient to provide a reason to combine. Accordingly, in my opinion, this rationale would not motivate a POSITA to combine Guichard with Doshi.

89. Further, Guichard teaches reoptimization in a system that omits the protection path featured in Doshi; and Doshi omits the reoptimization technique taught in Guichard. There are any number of possible ways to implement path selection in MPLS and Doshi and Guichard teach two different techniques. There is no obvious reason for a POSITA to pick and choose certain portions of Doshi’s technique to combine with certain portions of Guichard’s technique. Stated differently, while efficiency and reliability are generic benefits to any MPLS system, the Petition provides no evidence that combining unrelated aspects of Doshi and Guichard would actually result in more efficiency and reliability, particularly given the additional overhead cost to Doshi’s system of implementing reoptimization in Doshi.

90. Moreover, Guichard contains disclosures that Petitioner chose not to include in its Petition and its expert's declaration, which discuss systems featuring precomputed protection entities like Doshi:

There are actually two families of network recovery mechanisms—protection and restoration. A protection mechanism relies on the precomputation (and signaling) of a backup (alternate) path before any failure. In contrast, a restoration mechanism requires the computation of the backup path only after the failure has occurred; in other words, the backup path is computed on the fly.

Ex1006 [Guichard], 96-97. Notwithstanding its discussion of protection systems (like the Doshi system), Guichard does not teach reoptimization in the context of protection systems like Doshi; rather, it discloses that restoration systems (like Guichard) and protection systems (like Doshi) belong to different “families of network recovery systems.” *Id.* It is in this context that a POSITA would understand Guichard's disclosure that “it is [] highly desirable to detect the existence of [network changes] and reoptimize a TE LSP along a better path when it becomes available;” and, in this context—where Guichard discloses reoptimization and protection as two different families and says nothing about combining the two. Stated differently, Guichard acknowledges the existence of protection systems like Doshi yet chooses not to extend its optimization teachings to cover precomputed protection systems, teaching instead to perform dynamic

reoptimization. Thus, in my opinion, the statement upon which Petitioner relies is insufficient to provide a reason to combine. Ex1006 [Guichard], 96-97.

91. Indeed, a POSITA reading Guichard would understand that protection systems like Doshi's and reoptimization techniques like Guichard's entail significant costs and tradeoffs that might make such a combination disadvantageous. Such tradeoffs include recovery speed versus overhead—that is, faster recover requires more overhead and thus less efficiency. For example, Cisco documentation on MPLS Recovery Failure Detection (Ex2034 [IPCisco.com], 1) agrees in that additional speed for recovery incurs additional system complexity, control plane overhead or both:

“Firstly the failure tried to detect by hardware. If the failure cannot detect by hardware, it is tried to detected by upper layers. IGP hellos can be used here. Periodic hellos sent and if no response come, the failure detected. But here there are some disadvantages. By IGP hellos, the detection time is much (almost 30 seconds) and the CPU usage is high. This time can be reduced, but this time the control plane overhead will increase.”

92. Indeed, Guichard identifies that protection systems like Doshi “are usually faster but also require more network overhead because extra state is required”:

There are actually two families of network recovery mechanisms—protection and restoration. A protection mechanism relies on the

precomputation (and signaling) of a backup (alternate) path before any failure. In contrast, a restoration mechanism requires the computation of the backup path only after the failure has occurred; in other words, the backup path is computed on the fly. ***In terms of convergence time, protection mechanisms are usually faster but also require more network overhead because extra state is required.***

Ex2028 [Guichard], 96-97. Guichard also explains that combining recovery mechanisms requires “extra care to avoid race conditions between recovery mechanisms and double protection, which may lead to a lack of optimality in terms of backup resources.” Ex2028 [Guichard], 96. This disclosure further illustrates that a POSITA would understand the combination of Guichard with Doshi to allow reoptimization of Doshi’s primary path to involve significant tradeoffs that would deter such a combination.

93. Petitioner fails to address, let alone analyze those tradeoffs to show that a POSITA would have opted to combine the two references. Instead, Petitioner’s arguments are conclusory assertions that a POSITA would have combined Doshi’s protection systems with Guichard’s reoptimization based on generic rationales. Pet., 67-69. I understand that the inquiry whether to combine references must be thorough and searching, and there is a need for specificity. Petitioner has not

provided a thorough analysis. Accordingly, the Board should reject Petitioner's grounds as to all claims that recite the reselection limitation.

X. PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA WOULD ADAPT SUCH A COMBINATION.

94. Petitioner's assertion that the Doshi/Guichard combination renders the reselection limitation obvious also fails because Petitioner has failed to provide any reason for adapting the combination to practice reselection with respect to entity pairs.

95. Claim 1 is clear that "if an *entity pair reselection event occurs*, reselecting *said entity* pair from the group consisting *of said entity pair and a replacement entity pair* comprising at least one entity distinct from the entities of *said entity pair*." Ex1001 ['821 Patent], claim 1. As discussed above, Doshi discloses the use of entity pairs (i.e., a primary entity and a protection entity) but does not disclose the reoptimization (or reselection) of either its primary or protection entities. Moreover, Doshi discloses a selection process in which the primary entity is selected from one set (with K or L members) and the corresponding protection entity is selected from a second, different set (also with K or L members). *Supra*, §§V.A.1, VIII.A. Thus, to the extent that a POSITA were to adapt Doshi to practice reoptimization (or reselection), that reoptimization would occur with respect to individual entities in a pair, rather than all pairs associated with a given plurality of entities. That is, selection in Doshi occurs on an entity basis (i.e., selecting a

primary entity or selecting a protection entity), not on an entity *pair* basis as required by the claims; thus, a POSITA would likewise understand any reselection implemented in Doshi to occur on an entity basis, rather than an entity pair basis.

96. Further, Guichard teaches reoptimization of a single (active) entity in a restoration system, not reoptimization of an entity pair, i.e., in a protection system. Thus, to the extent that a POSITA had a reason to combine Doshi and Guichard (it would not), such a POSITA would create a combination in which reoptimization would be used only with respect to the primary path, because Guichard only teaches reoptimization on the active (i.e., primary) path. Because Guichard does not teach reoptimization of a protection path, a POSITA would have to adapt the Guichard/Doshi combination beyond what either reference actually discloses to reoptimize the both the primary path and the protection path.

97. Petitioner has not provided any evidence or even contended in the Petition that a POSITA would adapt the purported Doshi/Guichard combination to practice reoptimization of both a primary and a protection path

98. In addition, Petitioner has failed to provide any reasons for which a POSITA would (i) use reoptimization in the context of a protection entity and, critically, (ii) why a POSITA would create such a system in light of the significant tradeoffs (e.g., in “network overhead”) taught by Guichard.

99. Indeed, Guichard’s teaching that protection mechanisms require substantially more “network overhead” teaches against implementing a reoptimization mechanism in such a system, in particular, because reoptimization would further increase the computational complexity of the recovery scheme and its network overhead. Petitioner’s exhibits disclose the use of heuristics (shortcuts) to reduce computational complexity (e.g., Ex1025 [Xu], Abstract, ¶[0049]), how such heuristics are actually preferable to “optimal” computational solutions that may be impractical (e.g., Ex1025 [Xu], ¶¶[0030], [0056]-[0057]), how precomputation increases network overhead (e.g., Ex2028 [Guichard] 96-97), and how a POSITA must choose carefully among different levels of protection to balance optimality with other design constraints (e.g., Ex1007 [Huang], ¶¶[0010]-[0014] (discussing different levels of protection and associated costs and tradeoffs)).

100. To explain further, heuristics aim to find approximate solutions within a reasonable amount of time, even if the problems are notoriously hard to solve optimally. Ex2029 [Softjour], 1. Xu discloses that its approach to using splitting ratios to optimally balance traffic across routers that already are carrying traffic is a very difficult calculation. Xu discloses that “Several of the optimization problems are NP-hard [Xu: 0057] where NP hard stands for non-deterministic polynomial time hard and defines a degree of difficulty of a problem. Rather than solve the problem directly, Xu reduces the degree of difficulty by finding an approximate

solution thereby reducing the computational requirements. This contrasts with the invention of the '821 patent where calculations to balance traffic are not necessary. Instead the protection entities in the '821 are not active until designated as active. The protection entity can easily and immediately switch into an active state as there are no conflicts with pre-existing traffic.

101. Pre-computation does require some additional network overhead. Some reasons include: 1. Increased network signaling to establish and maintain the paths. 2. Depending on the dynamic nature of the network, the pre-computed paths may need periodic path re-computations. Maintenance, changing costs, human error and actual failures change the nature of a network in a dynamic manner. These changes must be accounted for and managed resulting in additional communication overhead. 3. Path monitoring and management. Pre-computed paths may require monitoring to ensure their readiness for failover.

102. Combining protection, restoration and other methods of resiliency increases the computational overhead as well as the network communication overhead necessary to ensure all network routers are properly configured, and then monitored, for the various employed protection schemes. The computational complexity can escalate very quickly in such scenarios as the interactions between the protections mechanisms can grow exponentially. *See, e.g.,* Ex2028 [Guichard], 96 (combining recovery mechanisms “requires extra care to avoid race conditions

between recovery mechanisms and double protection” and “may lead to a lack of optimality in terms of backup resources.”).

103. Petitioner and its expert do not address these issues. Instead, Petitioner and its expert have simply assumed that a POSITA would do so on the basis of a generic need for optimality (which Petitioner’s art even suggests is impractical and teaches against). Petitioner’s conclusory contentions about such a combination fail as a matter of law, and, for this additional reason, the Board should reject Petitioner’s obviousness challenge.

XI. PETITIONER HAS FAILED TO PROVIDE A REASON FOR WHICH A POSITA WOULD HAVE EXPECTED THE MODIFIED DOSHI-GUICHARD COMBINATION WOULD SUCCEED.

104. Cisco contends that a POSITA would have a reasonable expectation of success in combining Doshi and Guichard on the basis of four assertions: (i) “Doshi and Guichard both address the same technology;” (ii) “Guichard provides significant implementation details;” (iii) “skilled artisans [] had reoptimized similar systems;” and (iv) “Doshi already sets forth the calculation and selection techniques that would be utilized during the reoptimization evaluation.” Pet., 22-23. I do not agree with Petitioner’s contentions. As a threshold matter, as also explained *supra*, Petitioner has not provided a reason for adapting the Guichard/Doshi combination to reoptimize both a primary and a protection path as required by the claims.

105. I understand that the fact that Doshi and Guichard address the same technology is not sufficient by itself to supply a POSITA with a reasonable expectation of success. As explained above, Guichard is directed to a restoration system that differs substantially from Doshi's protection system and even classifies the two systems in different "families of network recovery mechanisms." Neither reference discloses any rationale that might have led a POSITA to conclude that combining disparate elements from these two families would work successfully or be an improvement over existing technologies; and Petitioner's rationales lack detail. Pet., 22-23. Without a rationale as to why Doshi's system would work with Guichard's, a POSITA would not have had a reasonable expectation of success in combining the two.

106. Petitioner argues that Guichard provides "significant implementation details." Pet., 22-23. I disagree. Though Guichard gives some high level guidance on the types of tools that can be used for reoptimization of a primary path, the excerpt of Guichard that Petitioner provides does not discuss how to reoptimize both a protection and a primary path, the scenarios in which such a reoptimization would even be appropriate, the relevant design trade-offs, or any other detail that would lead a POSITA to form a reasonable expectation of success. Ex1006 [Guichard], 3-5. Guichard's reoptimization is limited to a primary path; it does not explain how or why a protection path would also need reoptimization. If anything, Guichard teaches

against reoptimization in a system with precomputed protection paths by disclosing the increased “network overhead” of such a system, which would only be further increased if it were applied to both a primary and a protection path (i.e., a path pair), as required by the claims of the ’821 patent.

107. In my opinion, Petitioner is wrong that “skilled artisans had reoptimized similar systems.” Pet., 69. Petitioner’s expert admitted that all but one of the references cited by Petitioner fail to disclose reoptimization of a protection path. Ex2027 [Houh Dep. Tr.], 90:20-24, 92:17-23, 95:25-96:7, 96:12-25, 130:21-131:5. Thus, these systems are not “similar” in that they do not disclose reoptimization of a protection path or an “entity pair” as required by the claims. The only reference that mentions reoptimization in the context of a protection system is the Zamfir reference. Ex1034 [Zamfir]; Pet., 23; Ex1003 [Houh Decl.], ¶65. Zamfir is a very different protection system than that of the ’821 patent. Zamfir does not describe a protection system based on entity pairs between two end points. Instead, Zamfir describes a dynamic and virtual protection system by creating paths on-the-fly. Ex1034 [Zamfir]. Petitioner’s assertion concerning “similar systems” likewise fails.

108. In my opinion, Petitioner is incorrect that “Doshi already sets forth the calculation and selection techniques that would be utilized during the reoptimization evaluation.” Pet., 23. Doshi says nothing about reoptimization. That Doshi contains some techniques that might work without reoptimization would not provide

a POSITA with a reasonable expectation that the same techniques would work in a different context.

109. In particular, MPLS complexity varies according to the scale and features implemented in a network. In a modern network the complexity can be high as a network may span over a large number of nodes, a large geographical area and a large number of disparate physical networks. Implementing an MPLS network is likewise complex and involves planning, device configuration, testing, and protocol deployment. For example, Cisco provides extensive documentation on configuring Multiprotocol Label Switching for their routers such as their online book “Configure a Basic MPLS VPN Network”. *See* Ex2030 [MPLS Configuration Guide]. The fact that companies like Cisco provided such extensive documentation (including writing entire books), along with numerous examples and access to tech support to their customers illustrates the complexity of implementing MPLS.

110. Planning includes defining the services to be provided and should account for the topology of the network, the devices that make up the current network and extensions to the existing network that need to be acquired. In addition, MPLS allows for distributed computational engines which may require modifications to the software or script development.

111. Further, the MPLS network must be tested before deployment in a simulated or emulated environment. The proposed network must be extensively

tested to ensure each device is correctly configured, responsive to normal traffic as well as failures and other environmental changes.

112. Doshi’s techniques are designed to work in a system that employs shortcuts and heuristics to reduce computational complexity, including by using disjoint paths to avoid determining a probability of failure (*See* §VIII.D, *infra*), and by limiting the number of path pairs available for selection (*see* §VIII.A, *supra*). Performing the same computations repeatedly in response to state changes—as taught by Guichard—would increase the computational complexity of the Doshi system. In other words, each reoptimization performed in Doshi would require additional CPU resources. Given the teachings in Doshi and in other art (e.g., Xu, Guichard) concerning the desirability of reducing network overhead and simplifying network calculations, a POSITA would not understand Doshi’s disclosed techniques as a reason for which combining Doshi with Guichard and adapting the resulting combination would enjoy a reasonable expectation of success.

113. Accordingly, in my opinion, it would not be obvious to a POSITA to combine Doshi and Guichard in the manner suggested by Petitioner.

A. The Asserted References Do Not Teach Or Suggest “Logic Configured To Determine A Probability Of Concurrent Failure Of Said Working Entity And Said Protection Entity”

114. Independent claim 14 requires “logic configured to determine a probability of concurrent failure of said working entity and said protection entity”

(the “probability logic”). Ex. 1001 [’821 Patent], claim 14. Petitioner asserts that Doshi and Xu render this limitation obvious. Petitioner’s assertion fails, as neither Doshi or Xu disclose determining a probability, let alone the probability of concurrent failure of a working and protection entity, nor do they provide a reason for adapting either of the two references in light of the other to practice the claimed probability logic.

115. The claimed logic is expressly “configured to determine a probability of concurrent failure” of two entities (the working entity and protection entity). Ex1001 [’821 Patent], claim 14. The ’821 Patent teaches that it is desirable, in high availability MPLS networks, to provision primary and protection paths that are not likely to fail at the same time, to minimize service disruptions. *Id.*, 1:61-66 (“in network environments where span or node failure likelihood is relatively high, existing MPLS-TP linear protection mechanisms may not meet operator availability requirements since they may not protect against one or several failures affecting both the working entity and the protection entity.”). To accomplish this purpose, the ’821 teaches configuring logic to determine that probability and to factor that probability into the cost of the selected entity pairs. For instance, the ’821 describes the use of a probability matrix, which is used to accomplish this determination. Ex1001 [’821 Patent], 3:40-47 (“The concurrent failure probability matrix 150 examines the set of entities 120 and calculates the probability of concurrent failure for each entity

pair.”). The probability matrix is further evidence—beyond what is already evident from the claim language—that the determination of a probability of concurrent failure requires a logical process.

116. To explain further, a logical process refers to the ordered steps or activities that are performed in a systematic manner. Determining a probability of concurrent failures includes the steps of collecting data from the network, storing the data, and using computational hardware for computing the matrix of probabilities. The collected data would include status from the network routers including the state of the power sub-system, the strength of the transmitted signals, the strength of the received signals, the link bit error rates and other diagnostics that would indicate the health of each router. Status from each router is collected by the monitoring system via system messages. The monitoring system in turn includes a second set of logical processes to calculate the probability of concurrent failures using a combination of physical hardware, software or software scripts. As one example, such scripts could use the health status information from each router, and relate that to probabilities of failure based on past behavior or default configuration values.

117. In contrast, neither Doshi nor Xu disclose determining such a probability; rather, as described below, both Doshi and Xu describe different heuristics (or shortcuts) for selecting entity pairs that are less susceptible to concurrent failure. The Xu and Doshi approaches do not equate to determining a

probability; rather, they avoid determining a probability, the opposite of what is required by the '821 patent.

XII. DOSHI

118. Petitioner asserts that “[s]ince Doshi explains that the ‘probability of occurrence’ of two such failures is ‘very insignificant,’ a POSITA would have understood that determining that two paths are disjoint with no nodes or links in common is also a determination that the paths have a ‘very insignificant’ probability of failing concurrently.” Pet., 81 (*citing* Doshi ¶[0040]). In particular, Petitioner identifies Paragraph 40 of Doshi as a disclosure of this limitation:

Normally the extra capacity is sized for recovering from only a single failure at a given time in the network. This is due to the fact that the probability of occurrence of more than one failure during the time interval required to repair a failure is very insignificant, and therefore most of the time it is sufficient to cover just a single network failure.

Ex. 1005 [Doshi], ¶[0040]. Contrary to Petitioner’s assertion, this disclosure from Doshi teaches against determining a probability of failure.

119. Consistent with other heuristics (or shortcuts) that it adopts to reduce the complexity associated with MPLS path protection, Doshi is prompting a POSITA to ignore the probability of failure and instead opt for protection based on the fact that the protection path is disjoint from the primary path. Ex. 1005 [Doshi], ¶[0040]. According to Doshi, the probability of concurrent failure when two paths

are disjoint is “very insignificant;” therefore, when the pair selected is disjoint, the probability of concurrent failure need not be determined. But Doshi does not disclose determining a probability, as expressly required by the claim.

120. That Doshi is silent as to determining the probability of concurrent failure is further evident from the lack of any structure in Doshi for determining such a probability. In contrast, claim 14 recites “digital logic ... comprising: logic configured to determine a probability of concurrent failure”. Ex. 1001 [’821 Patent], claim 14. The specification provides, as one example of such a logic, a probability matrix, the role of which is to determine the probability of concurrent failure of the entity pairs associated with the claimed plurality of entities. Ex. 1001 [’821 Patent], 3:40-47. Further, the ’821 Patent discloses that the probability of failure is a factor in the claimed cost determination for the entity pairs. Ex. 1001 [’821 Patent], 3:40-47 (“[a]n overall cost logic block 140 uses results from a concurrent failure probability matrix... to select the working entity 122 and the protection entity 124...”). In contrast, Doshi discloses a system in which the number of pairs available for selection is restricted (to K pairs as discussed *supra*) and a disjoint candidate pair from an even smaller subset of pairs is compared to the lowest cost candidate pair to arrive at an eventual selection.

121. Accordingly, far from disclosing the determination of a probability of concurrent failure, Doshi teaches against such a determination by adopting a shortcut

in its protection system. Thus, independent claim 14 and claims 15 and 16 that depend therefrom are patentable over Doshi.

XIII. XU

122. Petitioner also asserts that Xu discloses determining a probability of failure or renders this limitation obvious. Pet., 81-82. In particular, Petitioner asserts that “[t]o inform path selection, Xu models failure states that would cause at least one path to fail” and therefore “[a] POSITA would have understood that at least some of the modeled failure states would be states in which the failure affects two paths simultaneously—corresponding to a state of concurrent failure.” Pet., 81. Essentially, Petitioner asserts that Xu discloses a scenario in which a primary and a protection entity fail concurrently. For reasons explained below, this first assertion is wrong. Petitioner also asserts that “Xu explains that each failure state has a ‘predetermined probability of occurrence’” from which a POSITA would understand that said “probability has been ‘determine[d].’” Pet., 82 (quoting Ex. 1025 [Xu], ¶[0006]. For reasons explained below, this second assertion, too, is wrong.

123. Petitioner’s first assertion is wrong because the claim do not require determining a probability of concurrent failure in the abstract; rather, they require determining the probability of concurrent failure of “said working entity and said protection entity.” Ex1001 [’821 Patent], claim 14. Unlike the claims, Xu’s “failure states” do not relate to a working entity and protection entity as required by claim

14. Xu does not even teach the concept of a “protection entity;” rather, Xu describes a conceptually and fundamentally different architecture for “balancing traffic load over a predefined set of paths P_d .” (e.g., multipath routing). Ex. 1025 [Xu], ¶¶[0006], [0030].

124. In particular, Xu discloses that each available path p in the set of paths P_d may simultaneously carry a fraction of the traffic between two routers according to a “splitting ratio.” *Id.*, ¶[0006]. If one or more paths fail, Xu uses a different splitting ratio to determine how much traffic to transmit over each remaining path p . *Id.*, ¶[0008]. Because each path p is constantly transmitting if it is available, each path is by default a “working entity,” *i.e.*, it is constantly transmitting and receiving traffic, and there is therefore no designated “protection entity” as required by the claims here. *See id.*, ¶¶[0030], [0035]. Indeed, because all entities are actively transmitting and receiving, Xu does not describe the concept of a “working entity” at all, let alone of a “protection entity.”

125. This difference between claim 14 and Xu is captured by an analogy comparing water drain systems. A typical bathroom sink has a main drain at the bottom and an overflow drain near the top in case the bottom drain is stopped. This is similar to claim 14, where there is a working entity (the main drain) and a protection entity (the overflow drain) for use only if the working entity (the main drain) fails. Xu does not work like this. Instead, Xu is similar to storm drains, where

a street has multiple storm drains and each storm drain is designed to carry a portion of the runoff water during a rainstorm. If one storm drain is blocked (or fails), the other drains are designed with excess capacity so they simply redistribute the extra water and each carry slightly more water. In this architecture, there is no concept of at least a “protection entity” or of “a probability of concurrent failure” for a “protection entity.” Thus, Xu is not determining “a probability of concurrent failure of said working entity and said protection entity” as required by claim 14.

126. But even setting aside the problems with Petitioner’s first assertion, Petitioner’s second assertion (concerning the determination of a probability) is also wrong. Like Doshi, Xu discloses a heuristic (or shortcut) that is used to *avoid* determining a probability of failure. Xu teaches a “predetermined probability of occurrence” that is a predetermined threshold probability of failure (used to inform the appropriate splitting ratios in a static system (Ex. 1025 [Xu], ¶[0049]), not the actual probability of individual failure states. In particular, Xu discloses a “set of failure scenarios S” comprised of the links and nodes that are most likely to fail (*id.*, ¶[0049]), not to independently determine the probability that two specific entities will concurrently fail as required by claim 14.

127. Petitioner’s argument that the predetermined probabilities fed into the Xu system must have been determined is simply incorrect and does not make logical sense. The claims require logic configured to determine a probability—not using a

predetermined probability, which is a prebaked heuristic. Indeed, Xu emphasizes the use of heuristics (shortcuts) in connection with implementing network protection “because of the monstrous size of ‘all’ subsets” that would need to be calculated if a POSITA were to opt for determining a probability on the fly. Ex. 1025 [Xu], ¶[0049]. Moreover, Xu does not explain how or whether its system determines probabilities for individual failure states, let alone provide any reason to a POSITA that those probabilities were actually determined (as opposed to, for example, extracted from historic data or other assumptions).

128. Thus, independent claim 14 and claims 15 and 16 that depend therefrom are patentable over Doshi and Xu.

A. Claims 2 and 18 Are Patentable Over Doshi


129. Claim 2 requires “[t]he method of claim 1, wherein said step of selecting an entity pair further comprises: selecting a working entity from said entity pair; and selecting a protection entity from said entity pair.” Ex1001 [’821 Patent], claim 2. Claim 18 is similar but also includes “selecting an active entity from the set consisting of said working entity and said protection entity.” Ex1001 [’821 Patent], claim 18. The Petition relies on Doshi alone for these limitations. Pet., 46-47, 65. In my opinion, Doshi does not disclose this limitation.

130. The language of claim 2 requires both working and protection entity are chosen “from said entity pair.” *Id.* “Said entity pair” refers to the entity pair chosen

in claim 1: “selecting an entity pair from said plurality of transport entities.” *Id.*, claim 1. Thus, claim 2 requires a pre-existing “entity pair” (i.e., 2) of transport entities. *Id.*, claim 2. From this pre-existing “entity pair” (i.e., a set of 2 entities), claim 2 requires that one transport entity is selected as a working entity and one is selected as a protection entity. *Id.*

131. Doshi does something different. First, Doshi selects “a set of K minimum cost primary paths” out of all available paths. Ex1005 [Doshi], ¶[0143], Fig. 11. Note there may be more than K paths available in the network; Doshi is simply selecting an arbitrary number K of paths with minimum costs. Then, Doshi selects “a set of K minimum-cost restoration paths” out of all available paths. *Id.* Next, Doshi selects a primary path (e.g., a working entity) from the “set of K minimum cost primary paths” and a restoration path (e.g., a protection entity) from the “set of K minimum-cost restoration paths.” *Id.* Thus, Doshi does not select a working entity from an entity pair or a protection entity from an existing entity pair as required by claim 2. Instead, Doshi selects a working entity from a group of potential working entities and a protection entity from a group of potential protection entities. This is a fundamentally different process than claim 2 of the ’821 Patent.

132. The difference between the techniques taught by ’821 Patent claim 2 and Doshi is shown in the illustration below:



133. In claim 2, the entity pair is selected before the working and protection entity are selected from the entity pair; In Doshi, the entity pair is selected only after the potential working entities and protection entities are selected. Thus, claims 2 and 18 are patentable over Doshi.

B. Claim 13 is Patentable Over Doshi

134. Claim 13 requires “[t]he method of claim 1, further comprising the step of: if said entity pair reselection results in both working and protection entities being replaced, sequentially replacing said working entity and said protection entity.”

Ex1001 [’821 Patent], claim 13. The Petition relies on Doshi in view of Guichard for this limitation. Pet., 59-63. The Petition provides no explanation or evidence that either Doshi or Guichard discloses “sequentially replacing said working entity and said protection entity.” *Id.* Thus, the Petition effectively concedes that neither Doshi nor Guichard disclose this limitation.

135. Instead, Petitioner argues, without evidentiary support, that “[i]t would have been obvious for the LSPs of the originally selected pair (e.g., a primary LSP and a protection LSP) to be sequentially replaced as the ongoing reoptimization efforts identify lower cost LSPs.” Pet., 60. Thus, Petitioner argues a POSITA would import the completely missing limitation “sequentially replacing said working entity and said protection entity” from their background knowledge.

136. I understand adding a limitation that is completely missing from the prior art should generally only be done when the missing limitation in question is unusually simple and the technology particularly straightforward and only where the Petitioner provides reasoned analysis and evidentiary support.

137. In my opinion, this limitation is not unusually simple or the technology particularly straightforward. To the contrary, entity pair reselection for path selection in MPLS technology is a very complex subject. Multiple references in this IPR alone (including Doshi, Guichard, Huang and Xu) demonstrate that there are numerous ways to perform MPLS path selection, none of which are obviously preferred by a POSITA to the others.

138. The ‘821 patent teaches high availability where no working link is left unguarded or is unguarded for the minimum amount of time. The transition from a working entity to a protection entity without losing traffic is a critical aspect of ensuring network reliability and continuity in MPLS. Other publications related to

MPLS such as RFC4090 “Fast Reroute Extensions to RSVP-TE for LSP Tunnels” and The RFC 6445 “Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base for Fast Reroute” includes methods different than those defined by the ‘821 patent including MPLS Fast Reroute (FRR). Exs2035, 2036 [RFC4090], [RFC6445]. Fast Reroute methodologies promote local detours using label stacking as opposed to an entity pair and overall cost determination from a set of entity pairs.

139. Issues that arise from pre-computed protection systems include stale protection decisions, label stacking and the distribution of packets already in the network, the time delay for messages to pass between the control plane and the data plane, complex protection schemes that include more than one type of protection, and processes of graceful restarting and restoring the failed entities. Given the range of issues that must be addressed, there is not necessarily one “best” way to create a successful transfer of the active state between a working entity and a protection entity.

140. The ‘821 patent takes the approach of carefully guarding the traffic by transferring the active state of the working link to the inactive link upon a failure of the working link in a step-wise or sequential fashion. Taking steps to transfer control that are performed in a sequential fashion minimizes the chances of traffic loss. “Entity pair reselection may not impact traffic as long as the newly selected

pair does not change the currently active entity. To ensure that the replaced entity is not active during entity reselection, if entity reselection results in both working and protection entities being replaced, the working entity and the protection entity may be sequentially replaced...” Ex1001 [‘821 Patent], 4:66-5:4.

141. Further, as the Petition implicitly concedes, none of these references teach or suggest sequentially replacing both a working entity and a protection entity as required by claim 13. In my opinion, the Petition does not provide reasoned analysis or evidentiary support supporting its allegation that “sequentially replacing said working entity and said protection entity” in the context of claims 1 and 13 of the ’821 Patent is part of the background knowledge a POSITA would inherently apply to its Doshi and Guichard combination. Petitioner does not cite any evidence showing that “sequentially replacing said working entity and said protection entity” in the MPLS network transport technology of claim 1 would be within the background knowledge of a POSITA. Pet., 59-63, Ex1003 [Houh Decl.], ¶¶174-181. Nor does Petitioner provide a rationale for why a POSITA would want to combine this limitation with its hypothetical combination of Doshi and Guichard.

142. In my opinion, a POSITA’s background knowledge would not include “sequentially replacing said working entity and said protection entity” in the context of claims 1 and 13. Even with a basic understanding of MPLS, a POSITA would not simply conclude that this limitation should be implemented in the hypothetical

combination of Doshi and Guichard. The method taught by claims 1 and 13 were not at the time a common technique for transmitting data over MPLS, and considerations such as overhead, reliability and up time had to be taken into account for the particular system being implemented. Because of these factors, a POSITA could not know whether adding claim 13 to the method of claim 1 would be desirable or even possible. None of the prior art even discloses reoptimization of both a working and protection entity. Thus, the concept of replacing a working entity and protection entity, let alone sequentially, was not common knowledge for someone with a basic understanding of MPLS. Instead, a POSITA would be directed by the published RFC specifications to utilize other approaches such as Fast Reroute found in RFC4090 and others. Fast reroute promotes local detours using stacked labels: “When a failure occurs along the protected LSP, the PLR redirects traffic onto the local detour.” Ex. 4090 [RFC4090], 7. This is different than what the ’821 patent discloses and claims.

143. The Petition starts with the hindsight assumptions that the missing limitation is part of the background knowledge of a POSITA that a POSITA would naturally want to combine with its hypothetical combination of Doshi and Guichard. Based on these assumptions, the Petition speculates that a POSITA would start with “Doshi’s modified and annotated [by Petitioner] Figure 1” and then, in view of Guichard, choose to implement the undisclosed limitation. Pet., 61-62.

144. In my opinion, Petitioner’s speculation is incorrect. First, Petitioner’s analysis starts with Petitioner’s hypothetical annotations to Doshi Fig. 1 rather than Doshi itself. Pet., 61 (showing Doshi Fig. 1 as “modified and annotated” by Petitioner. As discussed above, Petitioner’s combination of Doshi and Guichard is flawed, and the combination does not teach claim 1. *Supra*, §VIII.C. Further, Petitioner’s *annotations* are not reasoned analysis or support for how a POSITA would understand the prior art. Thus, Petitioner’s argument starts with unsupported assumptions about how a POSITA would interpret Doshi Fig. 1 before it even addresses the missing limitation.

145. Next, Petitioner speculates that “[w]hen both LSPs are determined to be suboptimal, a POSITA would have found it obvious to replace the LSPs sequentially to minimize interruption in service.” Pet., 62. But Petitioner points to no evidence that “both LSPs” in Doshi in view of Guichard would ever be “determined to be suboptimal.” As discussed above, Doshi does not discuss reoptimization and Guichard only discusses reoptimizing a single link (*supra*, §VIII.C), thus Petitioner’s hypothetical combination lacks reoptimization of two LSPs, let alone any disclosure that both LSPs are “determined to be suboptimal.”

146. Petitioner states that “a POSITA would have appreciated the desirability of maintaining a working and protection path throughout the replacement process.” Pet., 62. Petitioner and its expert provide no evidence in

support of this statement. I disagree that maintaining a working and a protection path through reoptimization was known by a POSITA at the time to be desirable, let alone to reoptimize in a sequential manner. Further, even if true, there are any number of other techniques by which a POSITA could choose to maintain a working and protection path that do not involve “sequentially replacing said working entity and said protection entity.” As shown earlier, Fast Reroute as described in the publication RFC4090 is one alternative option that entails local detours as opposed to the teachings of the ‘821 patent. Ex2035 [RFC4090], 7. Petitioner provides no explanation or evidence explaining why a Petitioner would choose to implement the specific technique taught by claim 13.

147. Petitioner cites its expert declaration in support of its argument. Pet., 59-63. But the expert declaration mostly contains the same assertions as the Petition, which are incorrect for the reasons discussed above. Ex1003 [Houh Decl.], ¶¶174-180. Dr. Houh does not cite any evidence in support of his opinions. *Id.* As discussed above, I disagree with Dr. Houh’s conclusions.

148. Petitioner has not shown that this limitation is unusually simple or the technology particularly straightforward, or provide reasoned analysis or evidentiary support that the missing limitation is within the background knowledge of a POSITA. Indeed, as I explain throughout my declaration, the technology at issue is far from simple or straightforward. It is in this context that Petitioner provides an

explanation that rests hindsight, with several layers of assumptions that contradict the evidence of record. Accordingly, Petitioner has shown that claim 13 is unpatentable.

XIV. CONCLUSION

149. I hereby declare that all statements made herein of my own knowledge are true and all statements made herein on information and belief were and are believed by me to be true, and that all statements herein were and are 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 and that any such willful false statements may jeopardize the validity of the application or any patents issued thereon.

Respectfully submitted,



Miguel Gomez

Dated: December 15, 2023