

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

Juniper Networks Inc.,
Petitioner,

v.

Portsmouth Network Corporation
Patent Owner.

Case No. IPR2024-00893
U.S. Patent No. 6,917,986

DECLARATION OF JAMES OLIVIER, Ph.D.

TABLE OF CONTENTS

I.	INTRODUCTION	1
II.	PROFESSIONAL BACKGROUND.....	1
III.	COMPENSATION	5
IV.	OVERVIEW OF TASK & BASIS FOR OPINIONS	6
V.	MATERIALS CONSIDERED	7
VI.	LEGAL CONSIDERATIONS	7
VII.	LEVEL OF ORDINARY SKILL IN THE ART	10
VIII.	The '986 Patent.....	11
IX.	OVERVIEW OF PRIOR ART REFERENCES.....	15
	A. U.S. Patent No. 6,032,194 (Gai)	15
	B. Automatic Fault Detection and Recovery in Real Time Switched Ethernet Networks (IEEE 1999)	16
	C. U.S. Patent No. 6,330,229 (Jain).....	20
X.	A POSITA WOULD NOT HAVE COMBINED GAI AND IEEE 1999	20
XI.	A POSITA WOULD NOT HAVE COMBINED GAI, IEEE 1999, AND JAIN.....	22
XII.	GAI IN VIEW OF IEEE 1999 DOES NOT DISCLOSE EACH CLAIM FEATURE	24
	A. Independent Claims 1 and 12.....	24
	B. Independent Claim 13	29
	C. Independent Claim 24	30
	D. Dependent Claims 6 and 18	31
	E. Dependent Claims 11 and 23	33

Declaration of James Olivier, Ph.D.

XIII. GAI IN VIEW OF IEEE 1999 AND JAIN DOES NOT DISCLOSE EACH CLAIM FEATURE	34
A. Dependent Claims 4, 5, 16, and 17	34
XIV. CONCLUSION.....	34

I, Dr. James L. Olivier, declare under penalty of perjury as follows:

I. INTRODUCTION

1. My name is James L. Olivier, Ph.D. I have been retained on behalf of Patent Owner, Portsmouth Network Corporation, as an independent expert and have been asked to provide my expert opinions regarding U.S. Patent No. 6,917,986. EX1001 (“the ’986 Patent”).

II. PROFESSIONAL BACKGROUND

2. My qualifications for forming the opinions given in this declaration are summarized here and are addressed more fully in my curriculum vitae, which is submitted as Exhibit 2003. My curriculum vitae also includes a list of my publications and technical consulting experience.

3. I am a recent Fulbright award winner in the areas of computer science and blockchain technologies, and I currently serve as an adjunct professor at Southern Methodist University.

4. I possess the knowledge, skills, experience, training, and education to form an expert opinion and testimony in this case. I am an expert in the field of networking technologies and have been an expert in this field since before 2002. I have designed and developed access systems such as those described in the ’986 Patent while working at Marconi Communications and teach MAC bridging and

Declaration of James Olivier, Ph.D.

Virtual Private LAN Services like those systems described in the '986 patent in my Multiprotocol Label Switching class.

5. I have a Ph.D. from The Ohio State University in Electrical Engineering with minors in Discrete Mathematics, Computer Science, and Microelectronics. My Ph.D. dissertation was based on coding theory and was titled "Concurrent Error Detection in Arithmetic Processors using GAN Codes," in which I developed new codes for use in Arithmetic Processors such as microprocessors. I have published papers on coding theory and continued my work in coding theory for arithmetic processors while at the General Motor Research Laboratory.

6. I have extensive experience in network systems design and development and have specialized in networking product development. I have developed and designed equipment for networks since my time at AT&T Bell Laboratories in the 1990s, where I worked on AT&T 5ESS, Network Control Point, and Autoplex Series base stations as a Member of the Technical Staff. These systems made extensive use of a counter-rotating ring network topology known as the Common Network Interface, much like the network described in the '986 Patent. It was there where I first began my work with telecommunication standards bodies as a contributing member of the Asynchronous Transfer Mode ("ATM") ATM Forum. Later at DSC Communications, I was the Senior Manager of the ATM systems engineering group developing ATM packet switches for a new generation base

station for Motorola's use in their Centralized Base Station Controller. While at DSC, I was also their corporate representative to the ATM Forum, participating in bi-monthly standards body development. I was one of the first contributors to the wireless standards at the ATM Forum.

7. At Samsung Telecommunications America, I was a Principal Engineer for wireless broadband services over Universal Mobile Telecommunications System ("UMTS"). UMTS is a third-generation ("3G") broadband standard developed by the 3rd Generation Partnership Project ("3GPP"). At Samsung, I worked on designing their next-generation cellular switch, a UMTS mobile switching center ("MSC"). While at Samsung, I was also their corporate representative to the International Telecommunication Union ("ITU"), the United Nations agency responsible for standardizing information and communication technologies. The ITU, under its International Mobile Telecommunications-2000 effort, was part of the 3GPP standards body. It was there that I participated in the development of standards for advanced wireless networks. While at Samsung, I was responsible for the 'services portion' of the UMTS switch. This subsystem provided Internet service, email, and the like to the subscribers.

8. At Marconi, I worked on several systems for the access market, such as Digital Subscriber Line ("DSL") modems and Digital Subscriber Line Access Multiplexers ("DSLAMs"), along with the design of point-to-point wireless systems.

Declaration of James Olivier, Ph.D.

At Navini Networks, I was responsible for layer 2 and layer 3 network protocols for their Wideband Code Division Multiplexed Access (“WCDMA”) wireless base stations and broadband modems. These layers were responsible for packet transmissions for various services over the WCDMA air interface. I was also responsible for the development of a Single Sign On (“SSO”) system which allowed users to gain information relating to their wireless subscription and delete or add services.

9. Currently I am an Adjunct Professor in the Network Engineering Graduate Program at Southern Methodist University’s School of Electrical Engineering, where I teach MPLS networked-enabled applications. Among the MPLS applications I teach is the use of MPLS in advanced network design, including content delivery networks, (CDNs) service provider network design and Intranet vs Internet network designs. In my class I also teach SONET/SDH rings, ADMs, and bandwidth models. I am also the owner of Olivier Consulting, where I provide consulting services for advanced network/product design along with Intellectual Property consulting. This year I received a three-year Fulbright specialist award in the area of computer science for my work at the Hunt Institute for Engineering and Humanity, which is part of the Lyle School of Engineering at Southern Methodist University in Dallas, Texas. There, I served as the Program Lead for Transformational Technology. Among the transformational technologies I

investigate at the Hunt Institute is the Internet of Things (“IoT”), which makes extensive use of Fourth-Generation and Fifth-Generation networks.

10. I am a co-inventor of U.S. patent No. 8,334,775 issued Dec. 18, 2012 and entitled “RFID-Based Asset Security and Tracking System, Apparatus and Method.” This invention relates to a Radio-Frequency Identification (RFID)-based Global Positioning System (GPS) tracking system to provide control and security of assets. This system integrates an RFID-based detection system with a conventional GPS tracking system. The GPS tracking system includes GPS receivers connected over a network to a centrally based GPS monitoring system.

11. Because of my background, training, and experience, I am qualified as an expert to explain the background of the technology encompassed by the ’986 Patent as well as to offer opinions related to the validity of the claims of the ’986 Patent in view of the prior art references discussed herein.

12. If called upon to do so, I could and would testify truthfully regarding the points stated below in this declaration.

III. COMPENSATION

13. I am being compensated for my independent analysis as an expert with respect to this *inter partes* review proceeding. My work on this matter is being billed at my customary rate of \$685.00 per hour. Also, I am being reimbursed for reasonable expenses I incur in relation to my services. My compensation is not

Declaration of James Olivier, Ph.D.

contingent in any way on the content of my analysis, the opinions that I offer, or the outcome of this proceeding. I have no financial interest in Petitioner or Patent Owner, and I have no financial interest in the '986 Patent.

IV. OVERVIEW OF TASK & BASIS FOR OPINIONS

14. I have been asked to review the '986 Patent in light of the materials cited and discussed herein and my knowledge and experience in this field during the relevant time frame, which I understand to be January 7, 2002. I have been asked to provide opinions related to certain issues from the perspective of a person of ordinary skill, having knowledge of the relevant art, January 7, 2002 and—except where otherwise noted—the opinions stated in this declaration are from that perspective. I describe the qualifications and abilities of such a person in Section “Level of Ordinary Skill in the Art” below. I have also been asked to consider whether certain combinations of references have rendered obvious the inventions described in the claims of the '986 Patent.

15. My findings, as explained below, are based on my study, education, training, and experience as well as items that I reviewed to prepare my opinions, including the documents listed in Mr. McNair’s declaration and any other items I reference in my below analysis.

16. I understand that other issues may arise that require further explanation, and I will provide that explanation if appropriate. As a result, I respectfully reserve

the right to update and supplement this declaration and the information and opinions provided herein.

V. MATERIALS CONSIDERED

17. In preparing this declaration and forming my opinions, I have considered: the claims, specification, and prosecution history of the '986 Patent; the Petition and declaration by Mr. McNair (including materials cited in each); and all materials cited and discussed in this declaration. As part of my analysis, I have also considered my own knowledge and experience, including my work and experience in the field.

VI. LEGAL CONSIDERATIONS

18. I am not an attorney. My understanding of the law is based on information provided by counsel for Patent Owner.

19. It is my understanding that assessing the validity/patentability of a U.S. patent claim based on a prior art analysis requires two essential steps. First, one must analyze the terms of the patent claims to understand what meaning one of ordinary skill in the art would give the terms. Second, one may then assess validity by comparing a patent claim to the "prior art." I understand that the teaching of the prior art is viewed through the eyes of a person of ordinary skill in the art at the time the invention was made. My opinion as to what constitutes a relevant person of ordinary skill in the art is set forth below.

20. I understand that a claim is invalid as obvious under 35 U.S.C. § 103 if the differences between the subject matter sought to be patented and the prior art are such that the subject matter of the claim as a whole would have been obvious at the time of the invention to a person having ordinary skill in the art. I have been informed that the following factors are used to determine whether or not the claimed subject matter would have been obvious: (i) the scope and content of the prior art; (ii) the differences between the prior art and the claimed invention; (iii) the level of ordinary skill in the field of the invention; and (iv) any relevant objective considerations of non-obviousness.

21. I also understand that a party asserting obviousness based on a combination of prior art references must demonstrate that one of ordinary skill in the art would have been motivated to combine the teachings of those references to achieve the claimed invention and that the skilled artisan would have had a reasonable expectation of success in doing so. It is my understanding that it is not enough to show that one skilled in the art could combine elements of multiple references, but instead there must be some reason that would have prompted a person of ordinary skill in the art to combine the elements in the way the claimed invention does. I understand that there must be some reasoned explanation as to why one of ordinary skill in the art would combine the references.

22. I understand that a combination of references would not have been obvious if the alleged modification(s) to be made to the reference(s) would have been inconsistent with the **reference's stated goals** or would have rendered the **combination inoperable for its intended purpose**. I further understand that for something to have been obvious, the party asserting obviousness must explain why a POSITA would have selected components for combination in the manner claimed.

23. I also understand that a party asserting obviousness based on a single prior art reference must nonetheless show a motivation to modify the reference.

24. I have been informed that a prior art reference must be considered in its entirety, i.e. as a whole, including portions that would lead away from the claimed invention.

25. It is my further understanding that obviousness cannot be based on a hindsight combination of components selected from prior art references. I also understand that an invention would not have been obvious simply because all of the elements of the invention may have been known separately in the prior art.

26. My opinions address what would have been logical to, and within the skill level of, a person of ordinary skill in the art ("POSITA") at the time of the invention, given the state of the relevant art, the knowledge and skill that a POSITA would have, the teachings of the references discussed below, and how a POSITA would have understood those teachings. My opinions also address whether a

POSITA would have had a reasonable expectation of the modified systems discussed below successfully yielding the results discussed below. I understand that a POSITA is a hypothetical person. I understand that a POSITA thinks along the lines of conventional wisdom. I understand that a POSITA is also a person of ordinary creativity, not an automaton (see Section “Level of Ordinary Skill in the Art,” below, for qualifications of a POSITA).

27. I understand that a claim term is given the meaning that the term would have to a POSITA at the time of the invention, which generally is the ordinary and customary meaning of the term. I further understand that the ordinary and customary meaning of a term may be evidenced by a variety of sources, including the words of the claims themselves, the specification, the prosecution history, and extrinsic evidence concerning relevant scientific principles, the meaning of technical terms, and the state of the art. I also understand that a patentee can act as a lexicographer by clearly expressing an intent to define a claim term.

VII. LEVEL OF ORDINARY SKILL IN THE ART

28. Counsel for Patent Owner has informed me that a 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.

29. For the purpose of this Declaration, and based on my opinions herein, the time of invention for the '986 Patent is January 7, 2002. However, I understand

that the conception dates of the inventions may have been earlier than that date. Any earlier or later conception date around January 7, 2002 would not alter my opinions herein.

30. Mr. McNair opined that a POSITA “would have had a bachelor’s degree in electrical engineering or a related field, plus at least two years of industry experience or graduate-level academic experience with computer networking.” EX1003 (McNair Decl.) at ¶ 111.

31. I agree with Mr. McNair’s opinion on the qualifications of a POSITA, and have applied this definition throughout my analysis.


VIII. The ’986 Patent

32. The 986 Patent is directed to a novel system to provide rapid reconfiguration of a network in response to network topology changes. As described in the Patent, prior art systems “provide[d] multiple, redundant paths between switches in the network” EX1001 (’986 Patent) at 2:7–9. However, existing techniques created “loops that disrupt the proper functioning of the learning process” and also caused serious network congestion, such as broadcast storms. *Id.* at 2:23–29. Additionally, existing techniques to remove network loops were “very slow in adapting to topology changes that may occur when a network link fails or is subsequently restored.” *Id.* at 2:65–67. To address these issues, the ’986 Patent teaches a system that uses redundant ports and links to connect switches together so

that data transmissions may continue in case a particular port or link fails. For example, when a switch detects a failure in one of its ports, it may activate a backup port and begin transmitting data through the link connected to the backup port. The '986 patent discloses techniques and methods to detect and respond to a network failure without needing to modify existing standards or use special command.

33. The independent claims of the '986 Patent map out this process, which comprises: (i) coupling an edge switch to one or more access switches via at least first and second redundant links, (ii) placing a first port of the edge switch in an active state and a second port of the edge switch in a blocking state so that communication traffic is conveyed over the first port, (iii) placing the second port in an active state and the first port in a blocking state in response to a failure of a first link associated with the first port, and (iv) sending dummy traffic from the edge switch over the second link to the one or more access switches so as to cause each switch to modify its respective database to account for the second port being activated. *See, e.g.*, EX1001 at claim 1.

34. For illustrative purposes, an annotated version of '986 patent Figure 2A is provided below.




EX1001 at Fig. 2A (annotated).

35. The annotated version of Figure 2A shows an edge switch (red) coupled to access switch A and access switch B (both in blue) over a first redundant link 30 and a second redundant link 32 (both in orange). Port P3 of the edge switch (green) is initially placed into an active state while port P4 (purple) is placed in a blocking state. Therefore, traffic from the edge switch may be transmitted from the edge switch to the access switches via port P3 and the first redundant link 30, but cannot be transmitted to access switch B directly via second redundant link 32 (red X). The '986 patent recognizes that access switches and edge switches are common types of

Media Access Control (MAC) bridges. As a result, the '986 patent refers to switches and bridges interchangeably. See, e.g., EX1001, 1:26–30.

36. In response to a failure associated with the first link, port P3 is placed in a blocking state while port P4 is placed in the active state to provide an alternate route for network traffic. Dummy traffic is then sent from the edge switch over the second redundant link to access switch B and access switch A so that one or more of the access switches can modify their respective databases, as shown below.



EX1001 at Fig. 2A (annotated).

IX. OVERVIEW OF PRIOR ART REFERENCES

A. U.S. Patent No. 6,032,194 (Gai)

37. Gai is titled “Method and Apparatus for Rapidly Reconfiguring Computer Networks,” and was issued on Feb. 29, 2000. Gai is directed to “rapidly reconfigur[ing] a computer network having a plurality of devices executing the spanning tree algorithm” EX1005 (Gai) at Abstract.

38. Gai discloses transmitting dummy multicast messages toward the root through a newly-active port, each of the messages containing the source address of an entity coupled to the switch or is downstream of the switch. See EX1005 at Abstract (“First, one or more devices are configured and arranged so that **one port, providing connectivity to the root, is in the forwarding state** and the **remaining ports, providing connectivity to the root, are in the blocked state**. Next, one or more of the blocked ports are designated as back-up ports. Upon detection of a failure at the active forwarding port, **one of the back-up ports immediately transitions from blocked to forwarding**, thereby becoming the new active port for the device.”) (emphasis added); see also *id.* at 5:40–60 (“In accordance with the method, one or more devices are configured and arranged so that **one trunk port** is in the forwarding state and **other trunk ports are in the blocked state**. Additionally, one or more of the blocked ports are designated as back-up ports. Upon detection of a failure at the active forwarding port, the state of one of the **back-up**

ports immediately **transitions** from blocked to **forwarding**. ...Upon transition to the new forwarding port, the device begins transmitting ‘dummy’ multicast messages through the new port. These dummy multicast messages carry the **source address** of each entity that is directly coupled to the device with the new active port or downstream thereof (**relative to the root**) and are received by other devices in the computer network.”) (emphasis added).

B. Automatic Fault Detection and Recovery in Real Time Switched Ethernet Networks (IEEE 1999)

39. IEEE 1999 is a paper authored by Srinidhi Varadarajan and Tzi-cker Chiueh and published in 1999 during the IEEE INFOCOM '99 Conference on Computer Communications.

40. IEEE 1999 discloses the *EtheReal* switch architecture including fault detection and recovery mechanisms supported by the *EtheReal* architecture. See EX1009 (IEEE 1999) at Abstract. IEEE 1999’s fault detection recovery mechanism begins with detecting a change in the network topology. See *id.* at 166. In the interior of the network, the fault would be independently detected by switches both upstream as well as downstream of the fault. *Id.* (“In the interior of the network, the fault would be independently detected by switches both upstream as well as downstream of the fault.”). After detection, these switches perform the following functions. : (1) discard all best effort traffic, (2) terminate all connections going through the faulty link, (3) rebuild the spanning tree, and (4) reestablish the

connections that were torn down. *See id.* IEEE 1999 discloses that the forwarding of best effort traffic results in a broadcast “when the route to the destination of a packet is unknown.” *Id.* Because broadcast at the link layer requires a spanning tree to avoid network loops, IEEE 1999 avoids this issue by discarding all best effort traffic altogether. *See id.*

41. IEEE 1999 also discloses that real time connection oriented traffic “is never broadcast in the network.” *Id.* Therefore, once a real time connection is set up, traffic “can follow the same route through the network without causing network loops, independent of spanning tree reconfiguration, as long as there is no change detected along its connection path.” *Id.* IEEE 1999 further discloses that real time connection oriented traffic is forwarded over a link “independent of whether the link belongs to the current spanning tree, or is blocked” and that real time connection oriented data can be forwarded even during the spanning tree rebuild process. *Id.*

42. To terminate the connections through a failed link, the switches first terminate all real time connections on the link and then inform the source of these connections (e.g. the host computer which initiated the connection) that the network is unstable. *Id.* at 166 (“The fault recovery mechanism then terminates all real time connections going through the offending link and **informs the source of the connection** that the network is **unstable**”) (emphasis added). An example of the

switches and their connections to the host computers in the *EtheReal* architecture is provided below:




Figure 1: *EtheReal Network Architecture*

Id. at Figure 1.

To terminate these connections going through a faulty link or switch, tear down messages are sent to the sources of the connections. *See id.* at 167.

43. Once the teardown messages are received and the connections are torn down, the fault recovery mechanism begins rebuilding the spanning tree. *See id.* at 166. The reconfiguration of the spanning tree is handled in three phases: (1) a startup phase (SP), a (2) startup completion phase (SCP), and (3) a topology discovery complete phase (TDC). *See id.* at 165, 167. In the SP phase of the spanning tree reconfiguration process, a switch in the network transmits a message through all its ports to its neighbors to join the spanning tree. *See id.* at 165. A node receiving the

message can accept the offer to join the spanning tree and then ask its own neighbors to join, or the node can reject the offer. *See id.* If a node has received rejections from all its neighbors after a reply time interval, the node assumes that it is a leaf node and initiates the SCP phase by sending a SCP message to its parent. *See id.* The parent node waits to receive a SCP message from all its children before sending its own SCP message to its own parent further up the spanning tree. *See id.* This process repeats until the root of the tree detects that all of its children have completed the SP phase. *See id.* The root signals the completion of the SCP phase and initiates the TDC phase by sending a TDC message to all of its children down the tree. *See id.* The TDC message is an indication to a node that it can safely forward data through its links. *See id.*

44. After the spanning tree is rebuilt, hosts that tore down their connections initiate new connections. *Id.* at 166 (“At the end of the spanning tree rebuild phase, all the hosts connected to the network are informed that the network is stable. Hosts that had their connections torn down due to network failure now initiate a new connection to their destination.”). When a real-time application attempts to set up a connection, it sends the reservation request to the Real Time Communication Daemon (RTCD) on the same host. *Id.* at 162. The RTCD is responsible for setting up the connections. *Id.*

C. U.S. Patent No. 6,330,229 (Jain)

45. Jain is titled “Spanning Tree with Rapid Forwarding Database Updates,” and was issued on Dec. 11, 2001. Jain is directed to “management of forwarding databases in the case of link failure on bridges.” EX1006 (Jain) at Abstract. Jain discloses limiting the propagation of topology change notifications to only those parts of the network which are affected by the change and triggering partial flushing of learned MAC addresses in forwarding databases for sets of addresses associated with the ports affected by the topology change instead of complete flushing. *See id.* at 2:66–3:8.

X. A POSITA WOULD NOT HAVE COMBINED GAI AND IEEE 1999

46. In my opinion, a POSITA would not have been motivated to combine Gai’s teachings with IEEE 1999 to send dummy traffic downstream or to send dummy traffic to modify a database responsive to the second port being in the active state, as recited in the independent challenged claims of the ’986 patent.

47. Mr. McNair appears to acknowledge that Gai does not disclose sending dummy traffic downstream and instead proposes modifying Gai with IEEE 1999 to address the deficiency. *See e.g., id.* at ¶¶ 137, 173. In my opinion, Mr. McNair’s reasoning for such modifications are unsupported and flawed. For example, Mr. McNair states that “POSAs would have had reason (i.e., faster recovery from a link failure as explained in IEEE 1999) to modify Gai’s technique by sending Gai’s

dummy traffic downstream after a downstream link failure (as the ‘986 patent describes), in addition to sending dummy traffic upstream after an upstream link failure (as Gai describes).” *Id.* at ¶ 137. Mr. McNair acknowledges that IEEE 1999 describes its messages as “teardown messages” and not dummy messages. *Id.* at 138. However, Mr. McNair asserts that the function of the teardown messages “is the same as Gai’s ‘dummy messages.’” *Id.* Mr. McNair’s comparison is incorrect. As previously mentioned in section IX.B, IEEE 1999’s fault detection recovery mechanism detects a fault in a link and performs the following functions: (1) discard all best effort traffic, (2) terminate all connections going through the faulty link, (3) rebuild the spanning tree, and (4) reestablish the connections that were torn down. *See supra* section IX.B; EX1009 at 166. IEEE 1999’s teardown messages are transmitted to host computers to tear down connections going across the faulty link. They are not received by switches in the network. In addition, these messages are NOT used to learn MAC address, but to instead to inform host computers that the network is unstable and that their particular connections have been torn down. Also, this happens before the spanning tree is reconfigured and before each switch begins forwarding traffic along new routes. *See* EX1009 at 166, 167. This is not how Gai describes dummy messages. Gai discloses that dummy multicast messages are transmitted “[u]pon transition to the new forwarding port” and are sent so that “other devices within the network learn to utilize the new forwarding port, rather than the

failed port, when directing messages to these entities.” EX1005 at 5:54–6:5. Therefore, Gai is clear that its dummy multicast messages are transmitted as part of the reconfiguration process and only transmitted once a device has already successfully transitioned to a new forwarding port. As such, Gai’s dummy messages are sent at a different time and serve a different purpose than IEEE 1999’s teardown messages.

48. Therefore, a POSITA would understand that IEEE 1999’s transmission of teardown messages differs from Gai’s transmission of dummy messages. Because the purpose and timing of the messages differ in both IEEE 1999 and Gai, a POSITA would not have been motivated to apply IEEE 1999’s transmission of teardown messages to Gai’s system as proposed by Mr. McNair.

XI. A POSITA WOULD NOT HAVE COMBINED GAI, IEEE 1999, AND JAIN

49. In my opinion, a POSITA would not have been motivated to combine Gai’s teachings with IEEE 1999 and Jain to update the database of the edge switch, as recited in the dependent challenged claims of the ’986 patent. As previously mentioned in section X, a POSITA would not have been motivated to combine Gai and IEEE 1999. *See supra* section X. Furthermore, it is my opinion that a POSITA would not have been motivated to incorporate Jain’s techniques to Gai’s system for the reasons below.

50. Mr. McNair asserts that “Gai expresses concern that messages may be misdirected or lost” and that “POAs would have been motivated to modify Gai by incorporating Jain’s technique of updating a filtering database responsive to a link failure for the benefit of reducing the possibility of misdirected/lost data messages during the self-healing process.” EX1003 at ¶ 214. Mr. McNair specifies that the modifications to Gai “would be implemented by remapping entries from the failed port to a backup port” and that the update “would be implemented before receiving communication traffic on the backup port to reduce the possibility of the backup port misdirecting or losing received messages using outdated database information.” *Id.* But Gai already describes a solution to this problem by sending dummy multicast messages. EX1005 at 12:62–66 (“Switch 114, via rapid reconfiguration entity 234, preferably informs the other devices of its new forwarding port by transmitting dummy multicast packets through the new port, as indicated at block 348.”).

51. Mr. McNair’s proposed modifications are therefore unnecessary and taught away by Gai. While Gai does mention the risk of messages being lost or misdirected, Gai presents its own solution to avoid that outcome. For example, Gai discloses that when access switch 114 transitions from initial port number three to back-up root port number four, “the identity of the new root port (e.g., port number four) must be propagated to the other devices, such as access switches 115-116 and backbone switches 120-125 to prevent messages from being lost or misdirected.”

EX1005 at 12:59–62. Gai further discloses that when a dummy multicast message 400 is received at a new source port, “this new location information is entered by backbone switch 124 into its filtering database, replacing the previous information that was stored therein.” *Id.* at 13:34–38. By doing so, Gai’s dummy multicast message 400 “effectively apprises backbone switch 124 of the change in forwarding ports that occurred at access switch 114.” *Id.* at 13:43–45. Gai’s transmission of dummy multicast messages records the location information of the new root port and propagates that information to other devices in the system so that filtering databases can be updated. As a result, Gai already provides its own solution so that messages are not lost or misdirected when a switch transitions to a new port.

52. Therefore, a POSITA would not have been motivated to incorporate Jain’s techniques to Gai’s system when Gai already provides a solution for the issues identified by Mr. McNair.

XII. GAI IN VIEW OF IEEE 1999 DOES NOT DISCLOSE EACH CLAIM FEATURE

A. Independent Claims 1 and 12

53. In my opinion, Mr. McNair fails to show how Gai in view of IEEE 1999 renders obvious each limitation of claim 1 and claim 12.

- 1. Sending dummy traffic “from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective**

database responsive to the second port being in the active state” (claim limitation [1E])

54. In my opinion, Gai in view of IEEE 1999 does not disclose or render obvious sending dummy traffic “from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state” (claim limitation [1E]). Mr. McNair relies on Gai as allegedly disclosing sending dummy traffic upstream to update filtering databases on each device that receives the dummy multicast message. *See, e.g.*, EX1003 at ¶¶ 165, 172. However, Mr. McNair appears to acknowledge that Gai does not disclose sending dummy multicast messages from an edge switch to one or more access switches to modify the respective databases of devices located further downstream. *See id.* at ¶ 173. Instead, Mr. McNair simply argues that a POSA “would have been motivated, by IEEE 1999, to modify Gai to send dummy traffic downstream” to achieve such a result. *Id.* It is my opinion that Mr. McNair’s arguments fail.

55. Mr. McNair relies on IEEE 1999 to disclose the transmission of self-healing messages from an edge switch to one or more access switches and specifically identifies IEEE 1999’s teardown messages as the appropriate message to modify filtering databases. *See id.* at ¶ 138. It is not clear why Mr. McNair has renamed the tear down messages as “self-healing” as this term is not used in IEEE 1999. However, as previously mentioned in section IX.B, IEEE 1999’s teardown

messages are sent to the sources of the connections (e.g. the hosts) and are not sent to modify filtering databases of any switches or reconfigure the spanning tree. *See supra* section IX.B. IEEE 1999 teaches that teardown messages are sent to host computers to tear down existing connections. *See id.*

56. IEEE 1999 is clear that its fault detection recovery mechanism detects a fault in a link and performs the following functions: (1) discard all best effort traffic, (2) terminate all connections going through the faulty link, (3) rebuild the spanning tree, and (4) reestablish the connections that were torn down. EX1009 at 166. As the name suggests and as outlined in IEEE1999, the teardown messages are specifically transmitted to tear down specific connections at the source of these connections, before the spanning tree is reconfigured and before each switch begins forwarding traffic along new routes. *See supra*, section IX.B; *see also* EX1009 at 166, 167. A POSITA would therefore understand that the teardown messages play no role in modifying or updating filtering databases in the switches.

57. Therefore, it is my opinion that Mr. McNair fails to establish that the combination of Gai and IEEE 1999 teaches or renders obvious “sending dummy traffic from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state.”

2. Sending dummy traffic “from the first bridge over the second link downstream to the second bridge, so as to cause the second

bridge to modify its database responsive to the second port being in the active state.” (claim limitation [12E])

58. Mr. McNair relies only on his arguments for claim limitation [1E] to support his discussion of claim limitation [12E]. As previously mentioned in section XII.A.1, Mr. McNair’s arguments for claim limitation [1E] fail. *See supra* section XII.A.1. Additionally, Mr. McNair’s arguments fail for the reasons set forth below.

59. The ’986 patent explicitly references and distinguishes Gai as part of the ’986 patent’s discussion of the background of the invention. *See* EX1001, 3:15–41. The ’986 patent discloses that “[Gai’s] method assumes that the switches in the network, including both core and access switches, have multiple trunk ports providing connectivity to the STP root via trunk switches ... ***The switch then transmits dummy multicast messages toward the root (upstream)*** through the newly-active trunk port, each containing the source address of a network entity that is coupled to the switch itself or is downstream of the switch.” EX1001, 3:18–33 (emphasis added). Thus, the ’986 patent expressly distinguishes Gai from the claims by stating that Gai transmits dummy multicast messages towards the root upstream and that the source address comes from an entity downstream of the switch. *Id.*

60. In other words, the ’986 patent distinguishes Gai from the claimed invention by stating that Gai’s method is limited to only transmitting dummy multicast messages upstream and that Gai does not disclose sending any transmissions downstream. In my opinion, as claim limitation [12E] explicitly

requires downstream transmission, the '986 patent's differentiation establishes that Gai does not disclose this limitation and that Gai does not render the '986 patent's claims obvious.

61. Not only does the '986 patent explain that Gai is limited to upstream transmissions, but Gai's own disclosure confirms this. See EX1005, Abstract ("First, one or more devices are configured and arranged so that **one port, providing connectivity to the root, is in the forwarding state and the remaining ports, providing connectivity to the root,** are in the blocked state. Next, one or more of the blocked ports are designated as back-up ports. Upon detection of a failure at the active forwarding port, **one of the back-up ports immediately transitions from blocked to forwarding,** thereby becoming the new active port for the device.") (emphasis added); 6:42–50 ("Upon failure of the current root port, this command additionally configures the access switch to immediately transition one of its blocked trunk ports to the forwarding state and to also begin *transmitting dummy multicast messages through the new port*, as mentioned above. Upon detection of a new or repaired link or device representing a **better path toward the root**, this command additionally configures the access switch to transition to the new path without suffering a loss of connectivity.") (emphasis added).

62. Additionally, in my opinion, IEEE 1999 does not cure Gai's deficiency. Mr. McNair states that "IEEE 1999 discloses sending such messages both upstream

and downstream” and that “IEEE 1999 describes the messages as ‘teardown messages.’” EX1003 at ¶ 138. However, as previously mentioned in sections IX.B and XII.A.1, IEEE 1999’s teardown messages are sent to the source of these connections to terminate them. They are not received by or operated by any switches which may pass the message through to modify its database. *See supra*, sections IX.B, XII.A.1. Therefore, even assuming that IEEE 1999’s teardown messages are sent downstream, they are not sent downstream “to cause the second bridge to modify its database responsive to the second port being in the active state.”

63. Therefore, it is my opinion that Mr. McNair fails to establish that the combination of Gai and IEEE 1999 teaches or renders obvious “sending dummy traffic from the first bridge over the second link downstream to the second bridge, so as to cause the second bridge to modify its database responsive to the second port being in the active state.”

B. Independent Claim 13

64. In my opinion, Mr. McNair fails to show how Gai in view of IEEE 1999 renders obvious each limitation of claim 13.

65. In my opinion, Mr. McNair’s analysis for claim 13 fails for the same reasons discussed in section XII.A regarding claim 1, as claim 13 is substantially similar to claim 1. That is, limitation [13C] is substantially similar to limitation [1E]. A comparison of the limitations [1E] and [13C] is provided below.

[1E] sending dummy traffic from the edge switch over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective database responsive to the second port being in the active state	[13C] a protocol processor, adapted to [...] send dummy traffic through the second access port over the second link to the one or more access switches, so as to cause each of the one or more access switches to modify its respective access switch database responsive to the second access port being in the active state
---	---

66. Mr. McNair relies on his analysis of the limitations for claim 1 for purposes of analyzing the limitations of claim 13. For example, Mr. McNair expressly cites to his analysis of limitation [1E] for purposes of limitation [13C]. See EX1003 at ¶¶ 196, 197. However, in my opinion, Mr. McNair’s argument does not cure the deficiencies of his analysis in connection with claim 1. Thus, claim 13 is non-obvious for the same reasons as claim 1.

67. Therefore, it is my opinion that Gai in view of IEEE 1999 does not render claim 13 obvious.

C. Independent Claim 24

68. In my opinion, Mr. McNair fails to show how Gai in view of IEEE 1999 renders obvious each limitation of claim 24.

69. In my opinion, Mr. McNair’s analysis for claim 24 fails for the same reasons discussed in section XII.A.2 regarding claim 12, as claim 24 is substantially similar to claim 12. That is, limitation [24C] is substantially similar to limitation [12E]. A comparison of the limitations [12E] and [24C] is provided below

[12E] sending dummy traffic from the first bridge over the second link downstream to the second bridge, so as to cause the second bridge to modify its respective database responsive to the second port being in the active state	[24C] the protocol processor being further adapted [...] to send dummy traffic over the second link downstream to the further bridge device, so as to cause the further bridge device to modify its database responsive to the second downstream port being in the active state
--	---

70. Mr. McNair expressly relies on his analysis of limitations [1B]-[1E] for purposes of limitation [24C], which is also repeated in his analysis for claim 12. EX1003 at ¶¶ 209, 210, 193. Mr. McNair does not present any new arguments in his analysis of limitation [24C]. Thus, claim 24 is non-obvious for the same reasons as claim 12 and claim 1.

71. Therefore, it is my opinion that Gai in view of IEEE 1999 does not render claim 24 obvious.

D. Dependent Claims 6 and 18

72. In my opinion, Mr. McNair fails to show how Gai renders obvious each limitation of claim 6 and claim 18. For reference, a comparison of dependent claims 6 and 18 is provided below.

[6] A method according to claim 1, wherein sending the dummy traffic comprises sending data frames originating at the edge switch while having source addresses corresponding to network addresses of other entities in the network that are accessible to the access switches via the edge switch.	[18] A device according to claim 13, wherein the dummy traffic comprises data frames originating at the edge switch while having source addresses corresponding to network addresses of other entities in the network that are accessible to the access switches via the edge switch.
---	---

73. In my opinion, Gai does not teach “wherein [sending] the dummy traffic comprises [sending] data frames originating at the edge switch while having source addresses corresponding to network addresses of other entities in the network that are accessible to the access switches via the edge switch.” Mr. McNair states that “it would have been obvious for backbone switch 122 (i.e., the edge switch) to send dummy multicast messages and thereby originate them.” EX1003 at ¶ 179. However, no portion of Gai actually discloses this scenario, and Mr. McNair does not cite to any passage in Gai to support his argument. *Id.* In fact, Gai discloses the transition at access switch 114 from an initial root port number three to back-up root port number four and the dummy messages associated with the transmissions from the access switch. *See, e.g.*, EX1005 at 12:56–13:33.

74. Because Mr. McNair cannot cite to any passage in Gai which discloses sending data frames originating at the **edge switch** while having source addresses corresponding to other network entities accessible to the access switches via the edge switch, it is my opinion that Gai does not disclose dependent claims 6 or 18. IEEE 1999 does not cure the deficiencies in Gai for the same reasons as previously discussed in sections XII.A.1–2. *See supra*, sections XII.A.1–2.

75. Therefore, it is my opinion that Gai in view of IEEE 1999 fails to teach or render obvious limitations 6 and 18 and Mr. McNair has failed to establish that claims 6 and 18 are rendered obvious.

E. Dependent Claims 11 and 23

76. In my opinion, Mr. McNair fails to show how Gai in view of IEEE 1999 renders obvious each limitation of claim 11 and claim 23. For reference, a comparison of dependent claims 11 and 23 is provided below.

[11] A method according to claim 1, wherein sending the dummy traffic comprises sending data frames downstream from the edge switch toward the access switches.	[23] A device according to claim 13, wherein the dummy traffic comprises data frames sent downstream from the edge switch toward the access switches.
---	---

77. Mr. McNair relies on his arguments in connection with claim 1 for his analysis of dependent claims 11 and 23. See EX1003 at ¶¶ 188, 206. Gai and IEEE 1999 fail to disclose “wherein [sending the dummy traffic comprises sending data frames downstream]/[the dummy traffic comprises data frames sent downstream] from the edge switch toward the access switches” for the same reasons as discussed in sections XII.A.1–2. See *supra* sections XII.A.1–2.

78. Therefore, it is my opinion that Gai in view of IEEE 1999 fails to teach or render obvious dependent claim 11 and 23 and thus fails to render obvious claims 11 and 23.

XIII. GAI IN VIEW OF IEEE 1999 AND JAIN DOES NOT DISCLOSE EACH CLAIM FEATURE

A. Dependent Claims 4, 5, 16, and 17

79. As Ground 2 only challenges dependent claims, the fact that Gai in view of IEEE 1999 fails to render obvious the claims in Ground 1 equally applies to the claims in Ground 2.

80. Mr. McNair also argues that a POSA would have been motivated to combine Gai, IEEE 1999, and Jain to disclose dependent claims 4, 5, 16, and 17. *See* EX1003 at ¶¶ 215–220. As previously discussed in section XI, a POSITA would not have been motivated to combine Gai’s teachings with IEEE 1999 and Jain to update the database of the edge switch. *See supra*, section XI.

81. Therefore, it is my opinion that Mr. McNair has failed to establish that claims 4, 5, 16, and 17 are rendered obvious.

XIV. CONCLUSION

82. This declaration reflects my opinions as currently held. However, my analysis may continue such that I may consider additional information that may result in additional observations.

83. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false

Declaration of James Olivier, Ph.D.

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.

March 4, 2025

Date

/s/ James Olivier

James Olivier, Ph.D.