

UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC,
Petitioner,

v.

WSOU INVESTMENTS d/b/a BRAZOS LICENSING AND
DEVELOPMENT,
Patent Owner.

Case IPR2025-00429
Patent 9,450,884

EXPERT DECLARATION OF DR. DOUGLAS CHRISSAN, PH.D.

I, Dr. Douglas A. Chrissan, declare as follows:

I. INTRODUCTION

1. I have been retained by WSOU Investments, LLC, d/b/a Brazos Licensing and Development (“Brazos”), as an expert in connection with the above-captioned *Inter Partes Review* to offer opinions relating to the 9,450,884 (“the ’884 Patent”).

2. I have been asked to provide my opinions on various technical aspects of the ’884 Patent and other references listed below. I submit this declaration in support of Brazos’ Patent Owner Preliminary Response.

3. I am being compensated as an independent consultant in this matter at the rate of \$300 per hour for analysis of documents and preparation of any declaration or report. This compensation is not dependent on my opinions or testimony, or the outcome of this litigation.

II. QUALIFICATION AND EXPERIENCE

4. The following sections recite the most relevant aspects of my educational, professional, and technical background. My curriculum vitae, which describes my qualifications in more detail, is attached to this report as Exhibit 1.

5. I am presently a technical consultant in the areas of communications systems, multimedia systems, computer systems, and digital signal processing.

6. I earned a B.S. and M.S. in Electrical Engineering from the University

of Southern California in 1988 and 1990, respectively, and a Ph.D. in Electrical Engineering from Stanford University in 1998.

7. From 1993 to 1997, I was a full-time doctoral student at Stanford University. My Ph.D. dissertation and related publications were in the fields of statistical signal processing and communication systems, and more specifically in the area of impulsive noise modeling for low-frequency communication systems. In 1995, I was the instructor for the graduate Statistical Signal Processing class (EE278) in the Electrical Engineering department at Stanford University. Prior to teaching this class, I was a teaching assistant for ten different classes in signal processing and radio frequency electronics at Stanford.

8. I was a Masters Fellow and Member of the Technical Staff at Hughes Aircraft Company in El Segundo, California, from 1988–1993. While at Hughes Aircraft, I designed and developed wireless communication systems for commercial and military spacecraft, including for the MILSTAR satellite program.

9. Between 1992 and 1993, while at Hughes Aircraft Company, I designed and built a state-of-the-art, 800 megabit-per-second (Mbps) telecommunications modem for the NASA Lewis Research Center.

10. From 1997–2003, I worked at 8x8, Inc., starting as a DSP software engineer in 1997, becoming a manager in 1998, a director in 1999, and Vice President of Engineering in 2000 (managing a team of approx. 60 engineers in the

company's microelectronics group). I played a key role in developing several semiconductor products used worldwide in multimedia and communications devices, mainly for video conferencing systems and Internet Protocol ("IP") telephones. Some of these semiconductor products were in production more than ten years.

11. From 2003–2007, I was a Systems Architect and Engineering Program Manager at Texas Instruments in the Digital Subscriber Line ("DSL") product business unit. At Texas Instruments, I was directly involved in the architecture, design, development and production of DSL modem and router products (also known as "customer premises equipment", or "CPE") products. My work included architecting a multicarrier DSL semiconductor and software product with routing and switching functionality, and managing all aspects of its development from inception to production. I also gained further experience with IEEE 802.11 (WLAN/WiFi) and WiMAX systems.

12. I have developed, and managed the development of, several successful semiconductor, software and systems products in the communications and multimedia fields. These products are listed in the attached curriculum vitae.

III. MATERIALS CONSIDERED

13. In reaching the opinions expressed in this declaration and conclusions reached, I have reviewed the following documents:

- Petition for *Inter Partes Review* of U.S. Patent No. 9,450,884
- The '884 Patent (Ex. 1001)
- The File History of the '884 Patent (Ex. 1002)
- Declaration of Paul Min Ph.D. In Support of Petition for *Inter Partes Review* (Ex. 1003)
- U.S. Patent No. 7,453,804 (Ex. 1005, “Feroz”)
- U.S. Patent No. 7,406,522 (Ex. 1006, “Riddle”)
- U.S. Patent Publication No. 2007/0248009 (Ex. 1008, “Petersen”)
- U.S. Patent Publication No. 2009/0097495 (Ex. 1009, “Palacharla”)

IV. APPLICABLE LEGAL STANDARDS

14. My opinions in this declaration are based on the understandings of a person of ordinary skill in the art—sometimes referred to as an “ordinary artisan” or by the acronyms “POSA”, “POSITA” (person of ordinary skill in the art) or “PHOSITA” (person having ordinary skill in the art)—as of the time of the invention, which I understand to be the filing date (June 11, 2015) of the utility application of the '884 Patent. I understand that the person of ordinary skill in the art is a hypothetical person who is presumed to have known the art relevant to the challenged claims of the '884 Patent at the time of the invention.

15. 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.

16. Petitioner submits that a POSITA at the time of the '884 patent was a person who “would have had a working knowledge of the computer networking art that is pertinent to the '884 Patent, including techniques for managing and allocating bandwidth within network devices. A [POSA] would have had a master’s degree in computer science, computer engineering, electrical engineering, or an equivalent, and two years of professional experience relating to packet-based computer networking. Lack of professional experience can be remedied by additional education, and vice versa.” Pet. at 11–12. I agree. I have used this level of skill for a POSA for my analysis herein and have employed it for the purposes of the subject IPR proceedings.

17. I was at the time of invention, and am, one of more than ordinary skill in the art through my education, research experience and work experience. As of the date of the invention, I was familiar with the types of problems encountered in computer networking systems and the rapidity at which innovations are made. I was and am familiar with people having this level of skill in the area of computer networking systems.

V. OVERVIEW OF THE '884 PATENT

18. The '884 Patent is titled “Software Defined Networking Based Congestion Control.” It was issued by the USPTO to inventors Jae Hyun Hwang and Thierry Klein on September 20, 2016 and assigned to Alcatel Lucent. The '884 Patent's Abstract describes:

Methods and systems of adjusting bandwidth allocation by a network element in a communications network includes monitoring a data flow traversing a target port, determining a bandwidth allocation for the target port, determining a fair-share bandwidth allocation for the target port, and adjusting the bandwidth allocation for the target port based on the fair-share bandwidth allocation. The bandwidth allocation for the target port is a bandwidth that is currently allocated for the data flow. The fair-share bandwidth allocation is a proportional allocation of a total bandwidth of the network element.

Ex. 1001 at Abstract.

19. The '884 Patent is directed to a technique for “adjusting bandwidth allocation by a network element in a communications network.” Ex. 1001 at 1:46–47. The inventors recognized that existing solutions to network congestion, e.g., Explicit Congestion Notification (ECN), “may not be scalable in terms of the number of concurrent data flows.” *Id.* at 1:35–36. To solve this network congestion problem, the inventors introduced a specific technique of “adjusting the bandwidth

allocation for the target port based on the fair-share bandwidth allocation.” *Id.* at 1:59–60. The claimed invention achieves reduced packet losses while supporting Quality of Service (QoS) awareness. *Id.* at 7:54–59.

VI. OPINIONS RELATED TO THE PATENTABILITY OF THE ’884 PATENT

20. I understand that the Petition includes three grounds for unpatentability of the ’884 Patent: (1) obviousness in light of Feroz in combination with Riddle; (2) obviousness in light of Petersen; and (3) obviousness in light of Petersen in combination with Palacharla. Pet. at 12.

21. I understand that the claims the Petition challenges as unpatentable are independent claims 1 and 17. Both claims require a “network switching element,” meaning a device that operates at least as a network switch.¹ Claim 1 repeatedly recites “network switching element”, and claim 17 recites both “edge switch” and “network switching element”.

22. For ground 1, petitioners have mapped the claimed switches to Feroz’s “application traffic management device” or a combination of Feroz’s “application traffic management device” and “router 22.” Pet. at 24. For ground 2, petitioners

¹ Such a network switching element may perform other functions, but at least includes the ability to operate as a network switch.

have mapped the claimed switches to Petersen’s “network switching system.” Pet. at 64.


23. In the computer network context, a network switch (or just “switch”) is a device that connects computers on a network and forwards received data frames (i.e., “packets”)² from any given input port to one or more specific output ports. *See, e.g.,* Ex. 2022 (Computer Networks, a systems approach, Larry Peterson, Fifth Edition 2012), 169–74. Typically, each physical port is bi-directional (i.e., the electrical or optical cable assembly that is attached to it communicates data in both directions) and includes a designated input port associated with a designated output (target) port. *See* Ex. 2022, pp. 169–74, 189–203, *passim*.

24. If a switch receives a unicast packet with a destination address that is already in that switch’s forwarding information base (i.e., forwarding tables), it will transmit that unicast packet on the output port known to be associated with the recipient device having that destination address. *Id.* For broadcast, unknown unicast and multicast packets (“BUM” traffic), the switch will transmit copies of the received packet on multiple output ports. *Id.* Switches have a “broadcast domain,” which is the set of output ports to which an unknown packet (i.e., for which the intended destination is unknown) will be “flooded” (i.e., retransmitted on all output

² In the well-known Open Systems Interconnection (OSI) model, a layer 2 “data frame” may encapsulate a layer 3 “packet,” but I use the terms “data frames” and “packets” interchangeably herein unless otherwise noted.

ports in that set). The broadcast domain may be every output port except that associated with the incoming packet. In a Virtual Local Area Network (VLAN) configuration, it may be a subset of output ports. *Id.*

25. In Feroz, the described application traffic management device is not a network switch and does not have the functionality of a network switch—not even the simplest functionality of a network switch described in the preceding paragraphs. Instead, it is an application traffic management device that is inserted within a single communications connection to monitor traffic flows. Ex. 1005, 5:64–6:3. By its name (“application”), it operates on application flows, which are at a higher layer (e.g., a higher Open Systems Interconnection (OSI) layer) than a conventional layer 2 switch or even a layer 3 network switch. Additionally, the application traffic management device 130 is depicted to have only two ports, one connecting to a router 22, and the other connecting to a network of connected computing devices 42, 44. *See* Ex. 1005, 5:49-53, FIG. 1 (reproduced below).



26. According to the Feroz specification, the application traffic management device 130 classifies and manages the application traffic traversing access link 21, which in fact is not directly connected to the application traffic management device but instead connects to router 22. Ex. 1005, 5:62–67, Fig. 1.

27. In Feroz’s Figs. 1 and 2, the application traffic management device has two “hardware communication interfaces,” as the specification calls them, that connect to the other network components. Ex. 1005, 6:26–29. From this depiction

and description, it does not appear that the application traffic management device is forwarding packets to destination ports as a conventional network switch with multiple output ports would do.

28. Feroz mentions that “application management device 130 can include additional network interfaces, beyond network interfaces 71 and 72, to support additional links or other functionality.” Ex. 1005, 6:42–45. However, “supporting additional links” refers to additional access links 21 to additional network cloud service providers, as these are the “links” that Feroz contemplates. *See* Ex. 1005, 6:41–51. Furthermore, if Feroz had contemplated its device being a switch or router, it would not need to depict a separate router 22 in Fig. 1. Feroz does not suggest in any way that if the application traffic management device has more than one output port, it would operate as a switch to transfer data frames between those output ports. *See generally* Ex. 1005, 6:26-50. Therefore, in light of both this passage and Feroz as a whole, Feroz does not convey to a POSA that the application data management device is a network switching element.

29. Claims 1 and 17 both recite “target ports” that belong to the switch, i.e., “network switching element including a target port” (claim 1) and “edge switch including a target port” (claim 17). The claims describe monitoring, determining and adjusting steps that occur at the target port of these switches. *See* Ex. 1001, 22:1–6, 24:55–64.

30. The Petition in Ground 1 maps Feroz's partitions to "bandwidth allocation for the target port." *See, e.g.,* Pet., 35-46. However, Feroz's partitions are not "bandwidth allocation for the target port." Feroz teaches that partitions correspond to traffic class: "According to one implementation of the present invention, the user configuration interface implemented by application traffic management device 130 allows the user to configure a bandwidth partition scheme. Such as the hierarchical configurations illustrated in FIGS. 9A thru 9F. In one implementation, the partition configuration can directly map, on a one to-one basis, to the hierarchical traffic class configuration used to classify network traffic, where all traffic gets matched to a leaf node in the traffic class tree." Ex. 1005, 13:66-14:7. The hierarchical partitioning of traffic classes does not correspond to bandwidth allocation for the target port.

31. Petitioner states: "As part of its partition-based bandwidth allocation scheme, Feroz explains that a partition may correspond to all traffic passing in one direction through one network interface, for example, all traffic at network interface 71 that is destined for network interface 72—in other words, all traffic inbound to network interface 71" Pet., 39. This scenario also fails to show that partitions correspond to the bandwidth allocation of the target port. In the scenario where the "root node may represent the access link in both the inbound and outbound directions, while the child partitions from the root can correspond to traffic

encountered at a given network interface (such as interface 71) and destined for egress from a second network interface (such as interface 72),” this only allocates the total bandwidth of the link in the inbound and outbound directions. Taking Feroz’s example of a T1 line with 1.5 Mbps, Feroz teaches bandwidth allocation of this T1 line among outbound and inbound traffic. This fails to teach bandwidth allocation for the target port itself.

32. For ground 2 (Petersen + obviousness) and ground 3 (Petersen + Palacharla + obviousness), where the Petition discusses monitoring, determining bandwidth allocations or adjusting the target ports, it identifies “virtual output queues” or “VOQs” as the target ports. Pet. at 66, 70, 72, 76.


33. However, the ports of a network switch are where devices physically and logically connect to the network through the switch. Switches are frequently described as having a certain number of these physical ports. Petersen clearly points out such ports using indices 110 and 112 (Fig. 1), 212 and 214 (Fig. 2), 412, 414 and 416 (Fig. 4). Ex. 1008 at Figs. 1, 2, 4. A “target port” as recited in the claims might be, for example, an output port on such a switch, but not a virtual output queue. Additionally, the ’884 Patent describes congestion occurring at the one of the switch ports, distinguishing it from a virtual output queue. Ex. 1001 at 11:13–32.

34. The virtual queues are not the target ports recited in claims 1 and 17 or described in the ’884 Patent. They are not physical connections to the network, but

instead are virtual queues implemented as memory buffers in a computing device. Thus, a POSA would not consider these VOQs to be the target ports recited in the claims.

I declare under the penalty of perjury of the laws of the United States that the foregoing statements made in this declaration are based on my own knowledge and that all opinions given are my own.

Dated: April 30, 2025



Douglas A. Chrissan, Ph.D.