

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

DELL TECHNOLOGIES INC.,
DELL INC.,
and
CISCO SYSTEMS, INC.,
Petitioners

v.

CORRIGENT CORPORATION,
Patent Owner.

Case No. IPR2023-00370

U.S. Patent No. 7,593,400

**DECLARATION OF DR. NICHOLAS BAMBOS IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 7,593,400**

TABLE OF CONTENTS

	<u>Page</u>
I. INTRODUCTION	1
II. QUALIFICATIONS AND EXPERIENCE.....	2
III. PERSON OF ORDINARY SKILL IN THE ART	5
IV. APPLIED LEGAL PRINCIPLES	6
A. Prior Art.....	6
B. Anticipation	7
C. Obviousness.....	8
D. Claim Construction Standard	10
V. TECHNOLOGY BACKGROUND.....	11
A. Overview of Computer Networks	11
B. Switches, Learning MAC Tables, Spanning Trees	14
C. Link Bundling and Link Aggregation	18
VI. BACKGROUND OF THE '400 PATENT	20
A. Summary of the '400 Patent.....	20
B. File History of the '400 Patent	28
VII. MATERIALS AND PRIOR ART CONSIDERED	29
VIII. INVALIDITY ANALYSIS OF THE CHALLENGED CLAIMS	32
A. Overview of Prior Art References.....	33
1. Smith	33
2. Ishimori	38
3. Edsall.....	41

4.	Zelig	42
5.	802.1Q-1998	44
B.	Motivations to Combine.....	44
1.	Smith and Ishimori.....	44
2.	Smith, Ishimori, and Edsall.....	46
3.	Smith, Ishimori, and Zelig	48
4.	Smith, Ishimori, Zelig, and 802.1Q-1998.....	49
C.	Specific Grounds of Invalidity	50
1.	Ground 1: Claims 1, 3, 6, 11, 13, and 16 would have been obvious over Smith.....	50
a.	Claim 1[pre]: A method for communication, comprising:	50
b.	Claim 1[a]: configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;.....	51
c.	Claim 1[b]: configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined Member Line Cards;	55
d.	Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;	60
e.	Claim 1[d]: receiving a data packet on an ingress port of said network node from a MAC source address,	

	said data packet specifying a MAC destination address on said Layer 2 data network;	65
f.	Claim 1[e]: conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;	69
g.	Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;	72
h.	Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and	75
i.	Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.	79
j.	Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.	83
k.	Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via switching core of the network node to at least the second line card.	87
l.	Claim 11[pre] 11: A node for network communication, comprising:	91

- m. Claim 11[a]: a switching core;.....91
- n. Claim 11[b]: a plurality of ports;.....91
- o. Claim 11[c]: a plurality of member line cards
conjoined in a link aggregation (LAG) group of
parallel physical links between two endpoints in a
Layer 2 data network joined together into a single
logical link, having a plurality of LAG ports to
forward packets through said switching core so that
the node operates as a virtual media access control
(MAC) bridge in said Layer 2 data network, said
plurality of member line cards including at least first
and second line cards, each line card having
respective ports and having a respective forwarding
database (FDB) to hold records associating MAC
addresses with said respective ports of said line
cards;.....92
- p. Claim 11[d]: wherein said line cards are arranged so
that upon receiving a data packet on an ingress line
card from a MAC source address, said data packet
specifying a MAC destination address, said ingress
line card conveys said data packet via said switching
core to at least said first line card for transmission to
said MAC destination address, whereupon said first
line card checks said MAC source address of said
data packet against records in said FDB of said first
line card, and if said FDB database of said first line
card does not contain a record of an association of
said MAC source address with said ingress port, adds
said record to the FDB of said first line card and
sends a message to at least said second line card
informing said second line card of said association,
and arranged, when said MAC destination address
does not appear in said FDB, to flood said data
packet via one and only one of said LAG ports. 92
- q. Claim 13: The node according to claim 11, wherein
responsively to the message, the second line card
adds the record of the association to the MAC

	database of the second line card if the record does not already exist in the FDB of the second line card.....	93
r.	Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.	93
2.	Ground 2: Claims 1–3 and 11–13 would have been obvious over Smith in view of Ishimori	93
a.	Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;	94
b.	Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;	97
c.	Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and	99
d.	Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.	101
e.	Claim 2: The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC address and the respective ports.	106
f.	Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the	

record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card. 109

- g. Claim 11[c]: a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards; 112
- h. Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports. 112
- i. Claim 12: The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the

	second line card of new associations between the MAC addresses and the respective ports.....	113
j.	Claim 13: The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.....	113
3.	Ground 3: Claims 1, 4–7, 10–11, 14–17, and 20 would have been obvious over Smith in view of Ishimori in further view of Edsall.....	114
a.	Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and	114
b.	Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.	115
c.	Claim 4: The method according to claim 3 and comprising marking the records in the respective FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.....	116
d.	Claim 5[a]: The method according to claim 4, and comprising: associating a respective aging time with each of the records;.....	117
e.	Claim 5[b]: refreshing the records in the FDB responsively to further packets transmitted by the line cards; and	121

- f. Claim 5[c]: removing the records from the respective FDB if the records are not refreshed within the respective aging time. 121
- g. Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via switching core of the network node to at least the second line card. 122
- h. Claim 7: The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed.122
- i. Claim 10[a]: The method according to claim 1, and comprising: conveying a further data packet, received from a further MAC source address, to the second line card for transmission over the network; 124
- j. Claim 10[b]: checking the further MAC source address against the records in the FDB of the second line card; and.....126
- k. Claim 10[c] responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record. 127
- l. Claim 11[a]: a switching core;.....128
- m. Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to

said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports. 128

- n. Claim 14: The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards. 129
- o. Claim 15: The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time. 129
- p. Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card. 130
- q. Claim 17: The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet,

- which indicates to at least the second line card to indicate that the record has been changed. 130
- r. Claim 20: The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record. 130
4. Ground 4: Claims 8–9 and 18–19 would have been obvious over Smith in view of Ishimori in further view of Zelig 131
- a. Claim 8: The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances. 131
- b. Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.....134
- c. Claim 18: The node according to claim 11, wherein at least some of the line cards are configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records

	associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.....	135
d.	Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.	135
5.	Ground 5: Claims 9 and 19 would have been obvious over Smith in view of Ishimori, Zelig, and 802.1Q-1998	136
a.	Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.....	136
b.	Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.	136
IX.	ADDITIONAL REMARKS.....	136

TABLE OF APPENDICES

Document	Description
Appendix A	<i>Curriculum Vitae</i>
Appendix B	Text of Challenged Claims

TABLE OF EXHIBITS

Exhibit	Description
1001	U.S. Patent No. 7,593,400 (“the ’400 patent”)
1002	Copy of Prosecution History of the ’400 patent
1004	U.S. Patent Application Publication No. 2005/0198371 (“Smith”)
1005	Certified English Translation of Japanese Patent Application No. 2005086668 (“Ishimori”)
1006	U.S. Patent No. 6,735,198 (“Edsall”)
1007	U.S. Patent Application Publication No. 2004/0133619 (“Zelig”)
1008	IEEE Standard 802.1Q, Virtual Bridged Local Area Networks, 1998 Edition (“802.1Q-1998”)
1015	IEEE 802.1D, Media Access Control (MAC) Bridges, 2004 Edition
1016	IEEE Standard 802.1Q, Virtual Bridged Local Area Networks, 2005 Edition
1017	IEEE 802.3 Standard for Local and metropolitan area networks: Specific requirements, 2002 Edition
1018	Kompella et al., <i>Virtual Private LAN Service</i> , IETF (December 2005)
1019	Lasserre et al., <i>Virtual Private LAN Services over MPLS</i> , IETF (November 2005)
1020	Martini et al., <i>Encapsulation Methods for Transport of Ethernet Over MPLS Networks</i> , IETF (November 2005)
1021	U.S. Patent No. 6,917,986
1022	U.S. Patent No. 7,974,223

I, Dr. Nicholas Bambos, hereby declare as follows:

I. INTRODUCTION

1. I have been retained on behalf of Dell Technologies Inc., Dell Inc., and Cisco Systems, Inc. (collectively, “Petitioners”) as an independent expert in this Petition for *inter partes* review by the U.S. Patent and Trademark Office. As part of my engagement, I have been asked to provide analysis and expert opinions on the following topics: (1) the disclosure of the U.S. Patent No. 7,593,400 (“the ’400 Patent”); and (b) the invalidity of claims 1–20.

2. In summary, it is my opinion that claims 1–20 of the ’400 patent (“the Challenged Claims”) are invalid in view of the prior art I discuss in this Declaration. In particular, it is my opinion that the prior art discussed herein discloses, teaches, or suggests the invention claimed in the Challenged Claims, rendering the Challenged Claims invalid as obvious over the prior art. The particular references and/or combinations that invalidate the Challenged Claims, as well as the reasons for my opinions and bases thereof, are set forth in detail below.

3. I am being compensated for my work on this case at my standard consulting rate. I am also being reimbursed for my expenses that I may incur. My compensation is in no way contingent upon the results of my study, the substance of my testimony, or the outcome of this case.

4. Based upon the facts and information available to me to date, I detail my opinions relevant to this petition for *inter partes* review. I am prepared to provide expert testimony on the opinions formed herein if asked about those issues by the Board or others.

5. Additionally, I may discuss my own work, teachings, and knowledge of the state of the art in the relevant time period. I may rely on handbooks, textbooks, technical literature, and the like to demonstrate the state of the art in the relevant period and evolution of the technology.

6. Throughout this declaration, I refer to specific line numbers or page numbers of the relevant prior art or other technical documents. The citations are intended to be exemplary as opposed to being exhaustive and are not intended to convey that the citations are the only source of evidence to support the propositions for which they are cited.

II. QUALIFICATIONS AND EXPERIENCE

7. I have knowledge of these technologies based upon education, training, or experience, of the matters set forth herein.

8. I am an R. Weiland Professor of Engineering at Stanford University, having a joint appointment in the Department of Electrical Engineering and the Department of Management Science & Engineering. I have also served as the

Fortinet Founders Department Chairman of the Management Science & Engineering Department from January 2016 to August 2020.

9. Before joining Stanford as an Associate Professor in 1996, I was since 1989 an Assistant Professor and later a tenured Associate Professor in the Electrical Engineering department of the University of California at Los Angeles (UCLA).

10. I received my Ph.D. from the University of California at Berkeley (1989) in Electrical Engineering and Computer Sciences (EECS). Also, from U.C. Berkeley, I received a M.S. in EECS (1987) and a M.A. in Mathematics (1989). I graduated in Electrical Engineering from the National Technical University of Athens-Greece (1984) with first class honors. While at U.C. Berkeley, I had been a U.C. Regents Fellow, a David Gale Fellow, an Earl Anthony Fellow, and an EECS Departmental Scholar.

11. At Stanford University, I head the Network Architecture and Performance Engineering research group, working on high-performance design of computer systems and networks. I have researched various high-performance design aspects of wireless and wireline networking (including switching) and computing technologies and have published numerous papers on such topics.

12. From 1999 to 2005, I was the Director of the Stanford Networking Research Center project. I have held the Cisco Systems Faculty Development Chair (1999-2003) in computer networking at Stanford University and have won the IBM

Faculty Development Award (2002) for research in performance engineering of computer systems and networks. I have also been the recipient of the National Young Investigator Award of the National Science Foundation (1992).

13. I have served as editor of various research journals (including the “Computer Networks” research journal), as technical reviewer for numerous networking and computing research journals, and on various technical panels for the National Science Foundation.

14. For over 30 years, I have done research in and taught networking and computing technology concepts and design principles (at Stanford since 1996 and at UCLA during 1989-96). I have graduated over 35 Ph.D. students who have later been in leadership positions in academia and the information technology industry (e.g., Stanford University, California Institute of Technology, Columbia University, New York University, University of Michigan; Cisco, IBM Labs, Qualcomm, ST Micro, Google, Intel, Nokia, MITRE, Sun Labs, Facebook, Twitter, etc.).

15. I have technical expertise in system architecture and high-performance engineering of computer networks/systems, have published over 200 peer-reviewed research papers in this field, have given numerous technical talks in this field world-wide, and have supervised numerous doctoral theses and research projects in this field at Stanford and UCLA. I am a named inventor on eight patents and have served as a technical expert witness in numerous patent litigation cases. My full

qualifications and experience are set forth in my Curriculum Vitae, which is attached as Exhibit A to this declaration.

III. PERSON OF ORDINARY SKILL IN THE ART

16. In evaluating the prior art references and other material, I have used the perspective of a person of ordinary skill in the art (“POSA”) to which the patent is related as of the time of the patent’s priority date. I am informed by Counsel that a POSA is presumed to be aware of pertinent prior art and the conventional wisdom in the art and is a person of ordinary creativity. I have applied this standard in my analysis throughout my declaration.

17. The ’400 patent is entitled “MAC Address Learning in a Distributed Bridge” and is directed to the field of communication networks. ’400 patent at 1:6–7. A POSA in the 2006 timeframe, which I am informed by Counsel is the earliest claimed effective filing date of the ’400 patent, would have had at least 1) a bachelor’s degree in electrical engineering, computer engineering/science, or a related field, and 2) either (a) a master’s degree in electrical engineering, computer engineering/science, or a related field, or (b) two or more years of work or research experience in networking and computing; that is, more education can compensate for less professional experience.

18. I meet the criteria described above and I am a person with at least ordinary skill in the art pertaining to the ’400 patent. I would have qualified as such

a person by at least 2006, which I understand is the earliest priority date of the '400 patent.

IV. APPLIED LEGAL PRINCIPLES

19. I am informed by Counsel for the Petitioners about the following legal standards, which I have applied throughout my analysis.¹

A. Prior Art

20. I am informed by Counsel that a patent, published patent application, or other publication, must first qualify as prior art before it can be used to invalidate a patent claim.

21. I am further informed by Counsel that a U.S. or foreign patent qualifies as prior art to a patent claim if the date of application, issuance, and/or publication of the prior art patent is prior to the purported date of invention of the patent claim.

¹ I am informed by Counsel that the patent laws were amended by the America Invents Act (AIA), but that the earlier statutory requirements still apply to pre-AIA patents. I have been informed by Counsel that the '400 patent is a pre-AIA patent, so the pre-AIA requirements control. Unless otherwise stated, my understanding of the law about patent invalidity as set forth in this Declaration relates to the pre-AIA requirements.

22. I am additionally informed by Counsel that a printed publication, such as a technical publication, a magazine article, or newsgroup post, qualifies as prior art to a patent claim so long as the date of publication is prior to the purported date of invention of the patent claim.

23. I am moreover informed by Counsel that a U.S. patent qualifies as prior art to a patent claim if the application for the prior art patent was filed in the United States before the purported date of invention of the patent claim.

24. I am further informed by Counsel that, if a U.S. Patent incorporates another patent or document by reference, it is considered to be part of that same reference.

B. Anticipation

25. I am informed by Counsel that, once the claims of a patent have been construed, the second step in determining anticipation of a patent claim requires a comparison of the construed claim language to the prior art on a limitation-by-limitation basis.

26. I am further informed by Counsel that a prior art reference anticipates a claim, and thus renders the claim invalid, if all elements of the claimed process, system, or device are disclosed in the prior art reference, either explicitly or inherently.

27. I am also informed by Counsel that, under Section 102 of the Patent Act, claims of a patent are invalid for lack of novelty if they are anticipated by a single prior art reference. I am further informed by Counsel that a claimed invention is invalid as anticipated and hence lacks novelty if all of the limitations of the claim as construed by the Court are present in a prior art reference, or any document that the reference incorporates by reference. However, I am informed by Counsel that a limitation of the claim need not be shown directly in a reference so long as a POSA would have understood the limitation to be inherent, or necessarily present in the reference.

C. Obviousness

28. I am informed by Counsel that an invention may also be obvious, even though the invention is not identically disclosed or described in a single prior art reference, which is required for anticipation. This is true if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the relevant art.

29. I am further informed by Counsel the Supreme Court has urged caution in granting a patent based on the combination of elements found in the prior art. This is because a patent for a combination that only unites old elements with no change in their respective functions withdraws what already is known from the public,

giving it exclusively to the patentee. The combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results.

30. I am also informed by Counsel that when a work is available in one field of endeavor, design incentives and other market forces can prompt variations of it, either in the same field or a different one. If a person of ordinary skill can implement a predictable variation, the invention is likely obvious. For the same reason, if a technique has been used to improve one device, and a person of ordinary skill in the art would recognize that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill. Put another way, I am informed by Counsel that a court must ask whether the improvement is more than the predictable use of prior art elements according to their established functions. I am further informed by Counsel that this is a flexible standard and it is not confined to any particular formalistic conception.

31. I am also additionally informed by Counsel that, when conducting an obviousness inquiry, the analysis should consider the scope and content of the prior art, differences between the prior art and the claims of the challenged patent claims, and the level of ordinary skill in the art. Secondary considerations of nonobviousness can also be considered, which include commercial success enjoyed by devices practicing the patented invention, industry praise for the patented invention, copying

by others, and the existence of a long-felt but unsatisfied need for the invention. I am also informed by Counsel that there must be a nexus between these factors and the claims themselves for these secondary considerations to impact the analysis of obviousness.

32. I am informed by Counsel that there generally must be a motivation to combine the references in order to find obviousness. This motivation or reason to combine can arise from the references themselves, changing market dynamics, or from simple common sense.

33. I am further informed by Counsel that a single reference can render a patent claim obvious, even if that reference does not fully anticipate the claim. To determine whether there was an apparent reason to combine the known elements in the way a patent claims, it will often be necessary to look to interrelated teachings of multiple patents; to the effects of demands known to the design community or present in the marketplace; and to the background knowledge possessed by a person having ordinary skill in the art.

D. Claim Construction Standard

34. I am informed by Counsel that, during *inter partes* review, the same standard is used as during claim construction in district court civil actions. I am further informed by Counsel that, in the district court context and in this IPR, the words of a claim are generally given their ordinary and customary meaning. This is

the meaning the term would have to a person of ordinary skill in the art in question at the time of the invention, which generally means the time the application was filed.

35. I am further informed by Counsel that the specification informs the proper interpretation of a claim. This is because a person of ordinary skill in the art is deemed to have read the claim term not only in the context of the particular claim in which the dispute term appears, but in the context of the entire patent, including the specification. I am also informed by Counsel that the prosecution history and extrinsic evidence concerning relevant scientific principles, the meaning of technical terms, and the state of the art are also considered in ascertaining the meaning of claim terms.

V. TECHNOLOGY BACKGROUND

36. To put in context my opinions on the validity of the claims of the '400 patent, in this section I provide some background on computer networks. All of the concepts discussed in this section were well-known and widely used well before the May 19, 2006 priority date for the '400 patent.

A. Overview of Computer Networks

37. Computer communication networks are comprised of communication links connected to ports of communication nodes (*e.g.*, switches, routers, computer hosts) forming a graph of (communication) nodes and (communication link) edges.

Information communication between nodes occurs in data units called packets (or frames in some cases), which are finite sequences of 0/1 bits representing information in digital form. In general, each packet includes a destination address, that is, the identifier (ID) of the final/destination node that it wants to reach. Each intermediate node on the path of a packet (i.e., the sequence of links and nodes it traverses) from its source (i.e., the node that injected it into the network) to its destination receives the incoming packet via an input port from an upstream node and retransmits it via an output port to a downstream node, which typically depends on the destination address of the packet.

38. A rough analogy is that of viewing the packet as a “car,” the communication network as a “freeway system,” the communication nodes as “freeway intersections” and the communication links as the “freeway segments” connecting the intersections. When the “car” reaches a “freeway intersection” on a “freeway segment,” the freeway signs direct (route) the car onto another “freeway segment” towards another “freeway intersection” and so on, until it reaches its destination and exits on a freeway ramp.

39. The path of a packet through a network, from its source node to its destination node, should have no loops under normal network operation. Clearly, a packet going around in loops and visiting any node more than once creates many problems, including 1) increasing the delay to deliver the packet to its destination

(since it has to lose time to traverse a loop of links and nodes to end up back at the node it entered the loop from) and 2) misusing the communication bandwidth of the nodes and links on the loop and, hence, causing unnecessary congestion on the network and delaying also other packets that are traveling through the network. Therefore, it is important that the link/node path traversed by each packet on the network be loop-free under normal network operation. The rough analogy of the communication network to the “freeway system” also makes it clear why we should avoid loops; a “car” should not loop around and visit the same “freeway segment” and/or “freeway intersection” twice at any point while traveling towards its destination under normal “freeway” conditions.

40. A tree is a graph where there is only one path (i.e., sequence of consecutive links/edges) to go from any node to any other one, as opposed to having more than one distinct path to go from some node to some other node. (i.e., a tree is a loop-free graph). Given any general graph (e.g., even with loops), a spanning tree of the graph is a tree (i.e., loop-free) subgraph that includes all nodes of the original graph, but possibly fewer links in order to break any loops that may exist in the original graph and reduce the original graph to a tree, hence, allowing no loops.

41. Spanning tree algorithms are used to construct spanning trees of communication network graphs in order to route packets on loop-free paths. Such algorithms are embedded in the Spanning Tree Protocol (STP) which runs on

computer communication networks in order to detect and remove loops, so that packets can be forwarded to their destinations on loop-free paths.

B. Switches, Learning MAC Tables, Spanning Trees

42. The Open Systems Interconnection (OSI) reference model is a layered modular abstraction of computer communication networks, especially with respect to network protocols (*i.e.*, operational schemes or mechanisms).




Figure 43–1 — Architectural positioning of Link Aggregation sublayer

(Ex. 1017, Section 3, Fig. 43-1).

43. Layer 1 (L1) is the Physical layer, layer 2 (L2) is the Data Link layer and layer 3 (L3) is the Network layer; on top of those layers others are also defined. A sub-layer of the Data Link layer L2 is the Media Access Control (MAC) one. For example, a widely used computer communication (family of) protocol(s) generally

conforming to the L1, L2, L3 layers of the OSI model is the Ethernet (family of) protocol(s), which has been standardized by IEEE under the name 802.3.

44. Each network port/interface device on a network node has a unique hardware identifier (ID), which is called “MAC address” and is generally inserted into the device by its manufacturer. MAC addresses are used by L2 switches (also called bridges historically) to switch packets at network nodes from input ports to output ports. L2/MAC switches/bridges can interconnect separate physical LANs into a unified logical network.

45. L2 switching nodes, including in the context of Ethernet IEEE 802.3 technology, operate at the L2/MAC layer to allow transfer of data packets called “frames” from switch input ports to output ports. Each frame includes the MAC address of its source node and the MAC address of its destination node. Each node maintains a MAC table, which specifies at which output port the switch should output an incoming/received frame, based on its MAC destination address. That is, when the switch receives a frame, it “looks” at the frame’s MAC destination address and “consults” its MAC table to find the port from where it should next transmit the frame to continue on the path to its destination. Thus, each switch’s MAC table maps the incoming frames’ MAC destination addresses to the switch’s output ports, in order to transmit each received frame on to the downstream node on its path to its destination.

46. The MAC tables are “learned” dynamically and distributedly. Indeed, when an incoming frame is received at a port of the switch, the latter “looks” at the source MAC address of the frame, and infers and “learns” that this MAC address is reachable through that port, and “keeps” that information in its MAC table. Thus, the MAC table is (backward) “learned” by the switch, in the sense that if the switch sees an incoming frame with a certain source MAC address at one of its ports, then it infers and “learns” that this MAC address is reachable from that port (as a destination MAC address now). Therefore, next time the switch receives a frame with that MAC address as a destination MAC address, it will output the frame on that port.

47. When a switch is first connected to the network in a “plug-and-play” manner, its MAC table is empty and the switch “knows nothing” – it has to autonomously “learn as it goes.” As mentioned before, when it receives an incoming frame at a port it “learns” that the source MAC address of the frame is also reachable as a destination MAC address from that port and “records” this information in its MAC table (if it hasn’t done so already with a previously received such frame). When the switch receives a frame at one of its ports, the switch looks at the destination MAC address of the frame and consults its MAC table: 1) If it finds an entry that this destination MAC address is reachable from one of its ports, then it transmits the frame on that port. 2) If it does not find such an entry, the switch

transmits the frame on all its ports (floods it to its ports), except on the one where it received the frame from. As the switch operates, its MAC table is populated by more and more information entries mapping destination MAC addresses to switch ports from where these destination MAC addresses are reachable. After a while the MAC table of the switch has been built up enough to give it visibility into the network to forward frames to their destinations efficiently.

48. Of course, in general there may be multiple redundant links and paths to get from a source to a destination port on a network of switches for at least purposes of increasing reliability and robustness to component failures. Such a topology may include path loops which can cause frames to go in circles. Besides the detrimental effects mentioned above, allowing frames to go in loops is additionally detrimental to network “stability” in various ways.

49. For example, suppose that switch A is connected to switch B via two links, the first one connecting port A1 of A to port B1 of B, and the second connecting port A2 of A to port B2 of B. When a frame with a previously unobserved (by A and B) destination MAC address arrives to switch A on a port A3, it is output on both ports A1 and A2 to B1 and B2 respectively. Since the destination MAC address is also previously unobserved by switch B, when the frame is received by B1 it will be output on B2 (to A2) and also the (same) frame received on B2 will be output on B1 (to A1), and so on and so forth, which could cause an infinite

retransmission loop of the frame. To prevent formation of path loops, the switch network runs a Spanning Tree Protocol, which identifies redundant links and blocks transmission on them, leaving a single path to go from any source to any destination on the switch network.

50. Component failures, like port/link failures or whole switch/node failures, or even planned connection of new components to the network, or additionally movement (disconnection and reconnection) of existing components to other parts of the network cause reconfiguration of the whole network. Therefore, 1) the Spanning Tree Protocol will have to be rerun on the switches to identify a new spanning tree to forward frames from their sources to their destinations and 2) at least the MAC table entries corresponding to the topology changes will have to be purged or flushed and then “relearned” on the switches, as the old ones do not map to the new network topology any more.

C. Link Bundling and Link Aggregation

51. Multiple links (and their corresponding ports) connecting switches can be “bundled together,” for example, into a Link Aggregation Group in some cases for various reasons, including increased bandwidth, traffic load balancing, scalability, increased reliability, etc. These Link Aggregation Groups are treated as a single logical link. For example, the IEEE 802.3 standard (as shown in Fig. 43-1 of Ex. 1017, Section 3, reproduced above) shows a Link Aggregation Sublayer of Layer

2 (*i.e.*, the Data Link layer of the OSI model), which can combine a number of individual physical links into a single logical link and present a single MAC interface to the MAC client.




Figure 43-2—Link Aggregation sublayer block diagram

52. The architecture of this sublayer is shown in Fig. 43-2 of Ex. 1017 (IEEE 802.3 Part 3: Carrier sense multiple access with collision detection

(CSMA/CD) access method and physical layer specifications), Section 3 reproduced above showing the Frame Collection/Distribution modules and the aggregation control module, including the Link Aggregation Control Protocol (LACP) module.

53. “Link bundles” or LAGs across multiple physical switches would provide redundancy by eliminating a single point of failure: if one of the link/switches in the “bundle” or LAG went down, the other links/switches could still operate.

VI. BACKGROUND OF THE '400 PATENT

A. Summary of the '400 Patent

54. The '400 patent relates to “communication networks, and ... distributed bridging systems.” Ex. 1001 at 1:6–9. “Local Area Networks (LANs) connect computing systems together at the Layer 2 level. The term “Layer 2” refers to the second layer in the protocol stack defined by the well-known Open Systems Interface (OSI) model, also known as the logical link, data link, or Media Access Control (MAC) layer.” *Id.* at 1:13–17. The '400 patent incorporates by reference the IEEE 802.1D standard and describes that “MAC bridges that implement the 802.1D standard allow MAC devices attached to physically separated LANs to appear to each other as if they were attached to a single LAN.” *Id.* at 1:27–30. The specification also describes that “Recently, various means have been proposed and

developed for transporting Layer-2 packets, such as Ethernet frames, over high-speed, high-performance Layer-3 packet networks.” *Id.* at 1:47–51.

55. At the time of the invention, “a number of authors have described methods for creating a virtual private LAN service (VPLS), which links different LANs together over an IP network.” *Id.* at 2:9–12. Further, the specification describes that “Link aggregation (LAG) is a technique by which a group of parallel physical links between two endpoints in a data network can be joined together into a single logical link (referred to as the “LAG group”).” *Id.* at 2:37–40. “Traffic transmitted between the endpoints is distributed among the physical links in a manner that is transparent to the clients that send and receive the traffic.” *Id.* at 2:40–42. The ’400 patent states that “For ethernet networks, link aggregation is defined by Clause 43 of IEEE Standard 802.3, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications (2002 Edition), which is incorporated herein by reference.” *Id.* at 2:43–47.


56. Figure 1 of the ’400 patent (reproduced below) shows that “terminal **22** is connected to a LAN, such as an Ethernet LAN, while terminal **24** is connected to a wide area network (WAN) **28**, such as the Internet or another Layer 3 network. The VPLS, however, permits the users of terminals **22** and **24** to communicate with one another as though they were connected to the same LAN domain.” *Id.* at 5:57–62. Further, the ’400 patent states, “The specific configuration of LAN **26** and WAN **28**

is shown in FIG. 1 purely by way of illustration, and the principles of the present invention may be applied in substantially any network configuration that supports the provisioning of Layer 2 virtual private networks.” *Id.* at 6:2–7. The ’400 patent further describes Figure 1 as follows:

In the exemplary configuration shown in FIG. 1, a network node 30 links LAN 26 and WAN 28. Node 30 comprises multiple line cards 32, linked by a switching core 34. Line cards 32 have ports 36, which connect to other nodes in LAN 26 and WAN 28 (and possibly in other networks, as well). Typically, each line card comprises multiple ports, although only a few ports are shown in FIG. 1. In the description that follows, ports 36 are assumed to be Ethernet ports, for the sake of simplicity of explanation. Alternatively, some or all of the line cards may comprise ports of other types, and may connect to other types of networks, such as Internet Protocol (IP) networks. For example, in an alternative embodiment (not shown in the figures), WAN 28 comprises a Resilient Packet Ring (RPR) network, and some of line cards 32 thus comprise RPR interfaces. Features of a network node that may be used to connect an Ethernet network to a RPR network are described, for example, in U.S. patent application Ser. No. 10/993,882, filed Nov. 19, 2004, which is assigned to the assignee of the present patent application

and whose disclosure is incorporated herein by reference. Additionally or alternatively, line cards 32 may connect to tunnels, such as Multi-Protocol Label-Switching (MPLS) tunnels, through WAN 28 via appropriate label-switched routers in the WAN.

Id. at 6:8–30.



57. As stated above, a network node 30 “comprises multiple line cards 32, linked by a switching core 34,” and each line card has ports 36, which connect to other nodes in LAN 26 and WAN 28. *Id.* at 6:8–12. “MAC bridges maintain a forwarding database (FDB) to map destination MAC addresses of the packets they receive to bridge ports. The bridge builds the forwarding database by means of a learning process, in which it associates the source MAC address of each incoming packet with the port on which the packet was received. When the bridge receives an incoming packet whose destination address is not found in the database, it floods

(i.e., broadcasts) the packet through all its available ports, except the one through which the packet arrived. Other MAC bridges that do not recognize the destination address will further flood the packet to all the relevant ports. Through the flooding mechanism, the packet will eventually traverse all interconnected bridges at least once, and will ultimately reach its destination.” *Id.* at 1:33–46. The ’400 patent states:

Each of the line cards may typically serve as both ingress and egress for data packets and has a respective MAC forwarding database (FDB) that is shared by the ingress and egress functions. When an ingress line card receives an incoming data packet over the VPN on one of its ports, it consults the FDB in order to choose the line card and port through which the packet should be forwarded based on the MAC destination address (or floods the packet through the ports in the VPN when the MAC destination address does not appear in the FDB).

Id. at 3:7–17. As I discussed above, this method of flooding was well known at the time of the ’400 patent. The ’400 patent allegedly identified a need for better MAC address learning at the bridges within a VLAN to mitigate the need to flood (or broadcast) messages when the destination MAC address was unknown, as follows:

When the forwarding destination of a packet is a link aggregation group (LAG), LAG member selection (i.e., selection of the link over which

the packet is to be forwarded) is typically performed on the ingress line card. In the absence of the synchronization method described above, other members in the LAG may not receive such packets for transmission, so that the FDB of the corresponding line cards will not be updated. When these line cards receive incoming packets, the result may be constant flooding, since the FDB is incomplete. The synchronization mechanism described herein avoids this problem by **updating the FDB in all line cards in the LAG group (or across the entire VPN instance)** within the node. Typically, when the transmitting line card transmits the data packet via a port that belongs to a LAG group, the synchronization message sent by the line card identifies the VPN instance and the incoming port. The **other line cards in the same LAG group (as well as all the other line cards serving this VPN instance)** can use this information to learn the MAC address association **even when these other line cards have not yet received packets from the MAC address in question.**

Id. at 3:34–53 (emphasis added).

58. Referring to Figure 2, the '400 patent describes:

Upon receiving an incoming packet from switch **40**, port **36** passes the packet to ingress path **54**. Packet processor **52** identifies the VPLS

(typically by a lookup and classification process based on certain packet header fields), extracts the other key parameters from the incoming packet (including the MAC destination address (DA), and optionally, the VLAN 30 identifier), and uses the key to query database **58**. If the record is found, the packet processor adds a tag to the packet indicating the egress port through which the packet should be forwarded, as well as the ingress port through which the packet was received. If the output interface indicated by the record is a LAG group, the packet processor selects one of the physical ports in the LAG group (using a hash function, for example), and tags the packet for transmission via the selected port. The packet processor then passes the tagged packet to switching core **34**, which conveys the packet to egress path **56** of the appropriate port.

When packet processor **52** receives a packet on ingress path **54** for whose key there is no a corresponding record in database **58**, however, it tags the packet for flooding. In this case, switching core **34** will pass the packet for transmission via all the ports (other than the ingress port through which the packet was received) that are used by this VPLS instance. For each LAG group serving the VPLS

instance, however, the flooded packet is transmitted via only one port in the group.

Id. at 7:25–49. That is, upon receiving an incoming packet, “[i]f the record is found, the packet processor adds a tag to the packet indicating the egress port through which the packet should be forwarded.” *Id.* at 7:31–34. If there is “no . . . corresponding record in the database . . . however, it tags the packet for flooding.” *Id.* at 7:42–44. While the “switching core 34 will pass the packet for transmission via all ports . . . that are used by this VPLS instance,” “[f]or each LAG group serving the VPLS instance, however, the flooded packet is transmitted via only one port in the group.” *Id.* at 7:47–49.

59. Synchronization messages (“SYNC”) are sent at regular intervals “to report each SELF entry that it has created in the FDB 58 to the other line cards 32 in node 30.” *Id.* at 8:17–22. The MAC table differentiates between SELF entries, which are entries that are learned by the packet processor on the line card and SYNC entries. *Id.* The ’400 patent uses an aging mechanism that is applied to the “MAC database 58 in order to remove records that are no longer in effect and free space for new records.” *Id.* at 9:4–6. A record is removed from the database “if a predetermined aging time elapses following the timestamp without a further packet having been received with the same key.” *Id.* at 9:9–11. If, on the other hand, the current packet matches a record in the FDB, “the packet processor refreshes the

timestamp of the record,” and “forwards the packet to the appropriate output port.”
Id. at 9:29–31.

B. File History of the '400 Patent

60. As I state below, in forming my opinions, I have also reviewed the file history of the '400 patent.

61. The '400 patent's application was filed on May 19, 2006. Ex. 1001, [22]. The Examiner rejected certain pending claims as anticipated by Edsall, or obvious over Edsall in view of another reference. Ex. 1002, 81–85. During an interview between the applicants and the Examiner, it was discussed that if the limitation “link aggregation group having a plurality of ports” in claim 1 were amended to recite the claim term as defined by the specification (at 2:37–40), then the amended claim would overcome Edsall. *Id.* at 112. The applicant therefore amended the claim limitation to add “a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link.” *Id.* at 114. The applicant made no other amendments to overcome the rejections over Edsall. The Examiner allowed the claims, and the patent issued on September 22, 2009. *Id.* at 128, 160.

62. Based on my review of the record, I understand that the claims were amended to add “a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link”

to overcome the rejection by Edsall, which corresponds to the language found in claims 1[b] and 11[c]. My opinions do not rely on Edsall for disclosure of this particular claim language that was added to claims 1[b] and 11[c]. Instead, it is my opinion that claim 1[b] and 11[c] are disclosed by Smith, which I describe in more detail below.

VII. MATERIALS AND PRIOR ART CONSIDERED

63. I have considered information from various sources in forming my opinions. I have drawn on my decades of experience in this field. I have employed methods and analyses of a type reasonably relied upon by experts in my field in forming opinions or inferences on the subject. Additionally, in preparing this Declaration, I have relied upon '400 patent and its file history, the exhibits cited to below, and the additional exhibits listed at the beginning of this Declaration.

64. In this Declaration, I provide five prior art references with exemplary citations identifying the relevant features related to the claim language of the Challenged Claims.² I have also described combinations of these prior art references that, when combined, render the Challenged Claims of the '400 patent obvious. It is my opinion that the Challenged Claims of the '400 patent are rendered obvious in light of the prior art specifically discussed in this Declaration.

² The complete text of the Challenged Claims is set forth in Appendix B hereto.

65. The citations presented in the body of this declaration correspond to the teachings disclosed in exemplary items of prior art that identify: (1) the problem(s) confronted by one of ordinary skill in the art; (2) an express suggestion to make one or more of the combinations; (3) an implicit suggestion to make one or more of the combinations; or (4) the knowledge of those skilled in the art as to the applicable field of technology at the relevant timeframe.

66. It is my opinion that a person of ordinary skill would consider at least the prior art references below:

Prior Art Reference	Date of Publication	Exhibit No.
U.S. Patent App. Pub. No. 2005/0198371 (“Smith”)	September 8, 2005	Ex. 1004
Japanese Patent App. No. 2005086668 (“Ishimori”)	March 31, 2005	Ex. 1005
U.S. Patent No. 6,735,198 (“Edsall”)	May 11, 2004	Ex. 1006
U.S. Patent App. Pub. No. 2004/0133619 (“Zelig”)	July 8, 2004	Ex. 1007
IEEE 802.1Q-1998 (“802.1Q-1998”)	March 8, 1999	Ex. 1008

67. Specifically, it is my opinion that the claims of the ’400 patent would have been obvious as follows:

'400 Patent Claims	Basis of Invalidity
1, 3, 6, 11, 13, 16	Obvious over Smith
1–3, 11–13	Obvious over Smith in view of Ishimori
1, 4–7, 10, 11, 14–17, 20	Obvious over Smith in view of Ishimori in further view of Edsall
8–9, 18–19	Obvious over Smith in view of Ishimori in further view of Zelig
9, 19	Obvious over Smith in view of Ishimori in further view of Zelig and 802.1Q-1998

68. I am also not aware of any secondary considerations of non-obviousness that affect my conclusions regarding the obviousness of these claims.

69. Other references and knowledge described earlier in this declaration are relevant to show the state of the art (*i.e.*, what a person of ordinary skill would have known when using the references identified above) or to use in combination with the references above to render the Challenged Claims invalid. I have considered also the following, which are also either cited or incorporated by reference in the '400 patent:

- a. IEEE 802.1D, Media Access Control (MAC) Bridges, 2004 Edition (Ex. 1015)
- b. IEEE Standard 802.1Q, Virtual Bridged Local Area Networks, 2005 Edition (Ex. 1016)

- c. IEEE 802.3 Standard for Local and metropolitan area networks: Specific requirements, 2002 Edition (Ex. 1017)
- d. Kompella et al., *Virtual Private LAN Service*, IETF (December 2005) (Ex. 1018)
- e. Lasserre et al., *Virtual Private LAN Services over MPLS*, IETF (November 2005) (Ex. 1019)
- f. Martini et al., *Encapsulation Methods for Transport of Ethernet Over MPLS Networks*, IETF (November 2005) (Ex. 1020)
- g. U.S. Patent No. 6,917,986 (Ex. 1021)
- h. U.S. Patent No. 7,974,223 (Ex. 1022)

VIII. INVALIDITY ANALYSIS OF THE CHALLENGED CLAIMS

70. Below is a detailed analysis of the bases of invalidity identified above.

71. I note that claim 1 in the '400 patent is a method claim and claims 2 through 10 depend from claim 1; claim 11 is an apparatus claim and claims 12 through 20 depend from claim 11. I am informed by Counsel that a method claim is a claim that describes a series of acts or steps for performing a desired function or accomplishing an intended result. I am informed by Counsel that an apparatus claim is a claim that defines the invention by its physical or structural components and their relationships. Because of the two claim types, there exist slight differences in the claim language (*e.g.*, “sending” versus “send”). Additionally, the apparatus

claim 11 recites a switching core, which is not found in independent claim 1. However, besides these differences, the differences found in the other claim limitations are immaterial for the purposes of my conclusions regarding unpatentability and do not affect my analysis. For my analysis, I have grouped claim limitations where they are substantively similar or where it made sense to discuss them together.

A. Overview of Prior Art References

1. Smith

72. Smith is directed to a virtual network device comprising interface bundles, which are managed as a single logical interface. Smith discloses:

A virtual network device includes several different virtual network device sub-units, which collectively operate as a single logical network device. An interface bundle includes interfaces in more than one of the different virtual network device sub-units included in the virtual network device. The interface bundle is coupled to a virtual link bundle, which connects the virtual network device to another device. The interface bundle is managed as a single logical interface.

Ex. 1004, Abstract; *see also id.* ¶ 34. Smith teaches that “[t]he communication links are configured to be managed as a single link. When the first network device sends a packet to the virtual network device via the virtual link bundle, the first network

device selects one of the communication links on which to send the packet. Each packet sent between the virtual network device and the first network device is sent via only a one of the communication links.” *Id.* ¶ 9.

73. As depicted in Figure 3, virtual network device 202 is coupled to other network devices 120(1)–120(3). *Id.* ¶ 44. The virtual network device consists of virtual network device sub-units 122(1) and 122(2), which includes several line cards 304(1)–304(4). *Id.* ¶ 46. “In virtual network device sub-unit 122(1), line card 304(1) includes forwarding engine 314(1) and interfaces 320(5), 320(7), and 320(9). Interface 320(7) is coupled to network device 120(3). Interface 320(9) is also coupled to network device 120(1).” *Id.* ¶¶ 46–47. “Line card 304(4) includes forwarding engine 314(4) and interfaces 320(12), 320(14), and 320(16). Interfaces 320(12) and 320(16) are respectively coupled to satellite network devices 120(3) and 120(1).” *Id.* ¶ 48. “Interfaces 320(13), 320(9), and 320(16), which are each coupled to network device 120(1) by virtual link bundle 250(1), form an interface bundle (e.g., an EtherChannel (TM) port bundle).” *Id.* ¶ 53. “Similarly, interfaces 320(11) and 320(8) form another interface bundle that is coupled to network device 120(2) by virtual link bundle 250(2). Interfaces 320(7) and 320(12) form a third interface bundle that is coupled to network device 120(3) by virtual link bundle 250(3).” *Id.*

74. The interface bundle is described by Smith as follows:

One way to avoid the complexity of having several independent redundant links is to operate those links as single logical transmission path, such as that provided using **a link bundling technique like EtherChannel (TM) or link aggregation (defined in IEEE 802.3)**. For example, an EtherChannel (TM) port bundle can be formed from several ports on a switch, each of which is coupled to a respective link in a group of links coupling that switch to another switch. Once an EtherChannel (TM) port bundle is formed, the port bundle can be managed as a single bridge port by routing protocols such as spanning tree, thus simplifying management of the redundant links.

Id. ¶ 6 (emphasis added).

Various embodiments of methods and systems for implementing interface bundles in virtual network devices are disclosed. A virtual network device includes several different virtual network device sub-units, which collectively operate as a single logical network device. An interface bundle includes interfaces in more than one of the different virtual network device sub-units included in the virtual network device.

Id. ¶ 8.

75. Additionally, Smith teaches that:

The redundant links coupling each of network devices 120(1) and 120(2) to virtual network device 202 can be operated as a **single logical link, referred to herein as a virtual link bundle**. Network device 120(1) operates the two links coupling network device 120(1) to virtual network device 202 as a virtual link bundle 250(1). In such an embodiment, each interface in network device 120(1) that is coupled to one of the links is included in an interface bundle, which corresponds to virtual link bundle 250(1). Network device 120(2) similarly operates the two links coupling network device 120(2) to virtual network device 202 as virtual link bundle 250(2). In some embodiments, **virtual link bundles** 250(1) and 250(2) are each operated as an EtherChannel (TM) or as an **aggregated link (as described in IEEE 802.3)**.

Id. ¶ 36 (emphasis added).

76. When a packet is received at an uplink interface, the virtual network device sub-unit can “learn that the sending device’s MAC (Media Access Control) address is “behind” uplink interface 320(13) by associating the MAC address with the logical identifier of uplink interface 320(13).” *Id.* ¶ 54. The sub-unit then informs each forwarding engine within the virtual device of this association (between the MAC address and the logical identifier of the uplink interface). *Id.* As

Smith teaches, “packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier.” *Id.*




FIG. 3


77. In order to keep its MAC tables updated, Smith uses “MAC notification frames are used to keep the content of the L2 tables in virtual network device sub-unit 122(1) synchronized with the content of the L2 tables in virtual network device sub-unit 122(2) and vice versa.” *Id.* ¶ 62. If a “forwarding table already includes an entry associating the destination address with a port of one of the network devices,” that forwarding engine will generate “a MAC notification identifying this association, which is distributed to any other forwarding engines within virtual

network device sub-unit 122(2).” *Id.* ¶ 63. “If there is no hit in the forwarding table, as determined at 407, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation, the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at 409.” *Id.* ¶ 66. But to take advantage of the virtual link bundles, Smith further specifies that as a result of the virtual link bundling, the data packet is “sent via only [] one of the communication links.” *Id.* ¶ 9.

2. Ishimori

78. Ishimori is directed to a packet forwarding method to “prevent repeated flooding” in systems using “path learning” and “link aggregation.” Ex. 1005, Abstract. Ishimori discloses a system with node groups 60, 70 in between two terminals A10 and B20. *See* Ex. 1005, Fig. 9. The node groups 60, 70 have communication cards #0, #1, #2 and #3, #4, #5, respectively. *Id.* ¶ 15. Ishimori teaches a method where the source MAC address of a received packet is learned by storing the source MAC address “in a buffer [MAC table] had by each node.” *Id.* ¶ 2. Thus, the destination MAC address for a packet is searched for “in the MAC table of the local device.” *Id.* ¶ 4. If the result is found in the buffer (MAC table),

then the packet is transmitted according to the forwarding path that was learned. *Id.* Thus, Ishimori teaches that “as long as a learning result relating to the destination MAC address (DA) of the receive packet exists in the buffer, there is no need for flooding and traffic in the connected LAN can be kept to a minimum, since the forwarding path to apply to the corresponding packet can be decided unambiguously by using this information.” *Id.* ¶ 6.



79. However, if the MAC address is not found, Ishimori acknowledges that in previous approaches, the packet would be flooded to all nodes. *Id.* ¶ 2. Ishimori recognized, however, the inefficiency in this approach because the learning results on the communication cards in the same trunk are not leveraged and “flooding is constantly performed.” *Id.* ¶ 19. Ishimori then teaches a solution to this flooding

problem. Specifically, Ishimori teaches, by leveraging “link aggregation,” which bundles a plurality of ports to function as one virtual port, that “one representative port is selected from among the large number of ports.” *Id.* ¶ 13. And the packet is therefore transmitted “using this representative port.” *Id.* Thus, “a packet received from the terminal 10 and addressed to terminal 20 is received at card #0” of node 60. *Id.* ¶ 15. “One representative port is selected . . . and this packet is further transmitted.” *Id.* Ishimori also teaches that for “each communication card #0 to #5 of the node groups 60, 70, as described in conjunction with FIGS. 1 to 4, information relating to the path leading to the local device is learned in association with the source address of the packet.” *Id.* ¶ 16. Ishimori also teaches that a “learn packet” is generated at a “predetermined timing.” *Id.* ¶ 25. Ishimori further teaches that the “learn packet” that informs the plurality of line cards regarding new associations “is transmitted to all nodes having the corresponding trunk.” *Id.* ¶ 33.

80. Ishimori also teaches an aging process. When the buffers of the cards of each node learn a MAC address, the “hit bit” corresponding to the learned MAC address is set to “1.” *Id.* ¶ 9. Then, after the first aging cycle, the hit bit is set to “0,” which indicates that the MAC address should be deleted from the buffer on the next cycle. If the MAC address is learned again at the next aging process, the hit bit is again set to “1.” *Id.*

3. Edsall

81. Edsall is generally directed to techniques for updating and synchronizing “forwarding tables contained on line cards that are interconnected by a switch fabric of a distributed network switch.” Ex. 1006 at Abstract. Edsall teaches that the “forwarding table” has an L2 portion that is “used to execute forwarding decision operations for frames forwarded among ports of the line cards.” *Id.* at 5:66–6:4. Similar to Smith and Ishimori, in Edsall, “[i]f the frame is received at the ingress card for the first time, this ingress forwarding engine also ‘learns’ a source MAC address of the frame.” *Id.* at 6:26–31. This involves “creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.* at 6:31–34. “The ingress forwarding engine then performs a flood-to-fabric (FF) operation on the frame by asserting all bits in the POE mask field of the fabric frame. The asserted POE bits instruct the switch fabric to switch (“flood”) copies of the fabric frame through its port interfaces to all (egress) line cards of the network switch.” *Id.* at 6:34–39; *see also id.* at 18:44–47.

82. Edsall further teaches a “novel MN frame is provided to complement the FF operation.” *Id.* at 6:46–50. The MN (MAC notification) frame “involves use of a primary input (PI) indicator,” which “denotes a primary input MAC address that is directly attached to a port of the line card associated with the forwarding table containing this entry.” *Id.* at 6:50–56. The “PI indicator is asserted for a forwarding

table entry having a MAC address that is learned from a frame sourced through one of the ports of the line card, as opposed to being learned through the switch fabric.” *Id.* at 6:56–60; *see also id.* at 18:56–19:5. The frame additionally includes a POE (port-of-exit) field that “includes a plurality of bits, one for each port interface of the switch fabric.” *Id.* at 6:24–25. The POE bit instructs the switch which port interfaces on which line cards should receive the MN frame. *Id.* at 9:47–50. Edsall also teaches an aging process “directed to aging of entries in the forwarding tables of the distributed switch.” *Id.* at 17:41–42. “[A] MAC address entry that has not been refreshed as a source within a specified period of time is removed from the L2 portion of the forwarding table in connection with a conventional aging policy executed by the microprocessor on the line card.” *Id.* at 17:43–47; *see also id.* at 17:41–18:34.

4. Zelig

83. Zelig is directed to a data communication network that “includes a plurality of primary virtual bridges, interconnected by primary virtual connections.” Ex. 1007, Abstract. Zelig discloses that “MAC bridges that implement the 802.1D standard allow MAC devices attached to physically separated LANs to appear to each other as if they were attached to a single LAN.” *Id.* ¶ 3. A MAC bridge “includes two or more MAC devices that interconnect the bridge ports to respective LANs.” *Id.* The “MAC bridges maintain a database to map destination MAC

address of the packets they receive to bridge ports.” *Id.* ¶ 4. “The bridge builds the database by means of a learning process, in which it associates the source MAC address of each incoming packet with the port on which the packet was received.” *Id.* Zelig provides “improved mechanisms for protection from failure in virtual private networks (VPNs).” *Id.* ¶ 18. Zelig thus discloses a data communication network including “a plurality of primary virtual bridges” and “one or more backup virtual bridges” that are “arranged to transmit the packets using a virtual private LAN service (VPLS).” *Id.* ¶¶ 21–24. “[E]ach of the primary and backup virtual bridges is adapted to maintain a respective media access control (MAC) table, and to forward the data packets in accordance with entries in the MAC table.” *Id.* ¶ 27. The VPN 20 includes multiple primary virtual bridges. *Id.* ¶ 42.




FIG. 1

5. 802.1Q-1998

84. The 802.1Q-1998 standard defines the architecture for Virtual Bridged LANs, its services, protocols, and algorithms. Ex. 1008, Abstract. It was part of the effort of the IEEE to standardize virtual LAN services in bridged LANs. The standard discloses a twelve-bit VLAN identifier field that is used to “uniquely identify the VLAN to which the frame belongs.” *Id.* § 9.3.2.3. This standard would have been known to a POSA at the time of the ’400 patent.

B. Motivations to Combine

1. Smith and Ishimori

85. A POSA would have been motivated to combine the teachings of Smith and Ishimori, because both references disclose link bundling and path learning techniques in communication networks and sending packets through one port of the link bundle. As explained above, Smith discloses that “[e]ach packet sent between the virtual network device and the first network device is sent via only a one of the communication links.” Ex. 1004 ¶ 9. Smith further teaches that “[f]or interfaces (e.g., ports or uplinks) included in interface bundles, the virtual network device sub-unit selects one egress interface per interface bundle via which to send the packet.” *Id.* ¶ 66.

86. Similarly, Ishimori also discloses flooding via one representative port. Ex. 1005 ¶ 13. Ishimori recognized the shortcomings in the approach taken by others

of flooding all ports because the learning results on the communication cards in the same trunk are not leveraged and “flooding is constantly performed.” *Id.* ¶ 19. Ishimori explains that using “one representative port is selected from among the large number of ports had by this trunk according to a predetermined computation, and actual packet transmission is performed using this representative port.” *Id.* ¶ 13. “This configuration prevents the application of a different path in the same trunk to the same packet.” *Id.* ¶ 14.

87. A POSA looking to solve the shortcomings of flooding all the ports would have therefore looked to Smith and Ishimori. Smith teaches that selecting one egress interface on a virtual network bundle. Ex. 1004 ¶ 66. Ishimori teaches that using “one representative port” prevents constant flooding. A POSA looking to implement this solution would have therefore been motivated to combine their teachings because flooding over only one link in a bundle would prevent sending similar packets unnecessarily over multiple links. A POSA would have understood that flooding would also not occur over the ingress port.

88. Additionally, Ishimori teaches a “learn packet” that is generated at a “predetermined timing.” Ex. 1005 ¶ 25. If one of skill in the art were looking to limit flooding over unnecessary links, one would have been motivated to implement Ishimori’s learn packet generated at a predetermined timing into Smith’s MAC notification method. Furthermore, one of skill in the art would have had a reasonable

expectation of success in implementing this combination because it would have required a simple substitution of Ishimori's structures for Smith's relevant structures. Both Smith and Ishimori leverage link bundling/aggregation techniques in the packet forwarding methods they disclose. Thus, it would have been within the knowledge of a POSA implementing such a method to combine the teachings of Smith and Ishimori and have a reasonable expectation of success.

2. Smith, Ishimori, and Edsall

89. A POSA would have been motivated to combine Smith, Ishimori, and Edsall. All three references disclose methods of using MAC forwarding tables in a distributed network switch. Ex. 1004 ¶ 54; Ex. 1005 ¶ 9; Ex. 1006 at Abstract. It would have been obvious to a POSA to modify the MAC forwarding tables of Smith to implement Edsall's forwarding engine. Specifically, a POSA could implement the MAC tables of Smith to implement how Edsall's forwarding engine learns the source MAC address by "creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch." Ex. 1006 at 18:39–44. A POSA could further implement Edsall's "flood-to-fabric (FF) operations" by implementing the MAC tables to include a POE bits that determines which port interfaces on which line cards should receive the MN frames. *Id.* at 18:47–50. Finally, it would have been obvious to a POSA to implement the MAC tables to include the PI indicator as taught in Edsall in order to keep track of whether

a MAC address was learned from a local line card or a different line card. It would have also been obvious to a POSA to further implement the MAC forwarding tables of Smith-Ishimori to use the POE field in order to determine which ports should receive the MN frame.

90. Furthermore, creating entries in a MAC forwarding table and adding a PI indicator field and POE field to the MAC tables would have been an implementation that a POSA would have known how to make with a reasonable expectation of success.

91. Ishimori additionally teaches an aging process. Ishimori's aging process uses a "hit bit." Ex. 1005 ¶ 9. When a MAC address is first learned, the hit bit is set to 1. *Id.* Then, "at the first aging process, this hit bit is cleared," meaning the "hit bit" it is set to 0. *Id.* The learning result is not yet deleted. *Id.* If on the next aging cycle, the MAC address is learned again, the hit bit will be set again to 1. *Id.* Otherwise, the entry will be deleted. *Id.* Therefore, the "hit bit" keeps track of how long a MAC address has been in the MAC address table without having been refreshed. *Id.*

92. A POSA would have understood that MAC addresses may become stale due to changes in the network and would have been motivated to look at solutions for making sure the MAC table is up to date. Although Smith does not disclose aging the entries in the MAC table, it would have been obvious to a POSA to look

for a type of aging parameter could have been used and Ishimori describes precisely that. Moreover, because the size of a MAC table is limited, a POSA would have known that older entries have to be deleted when excessively aged to make space for newer entries. A POSA looking to implement this solution would have therefore been motivated to combine the teachings of Smith and Ishimori and it would have been an easy application of Ishimori's methods to Smith's system. For example, the aging process could be added to the MAC address table described in Smith using known techniques and a simple addition of that information to a table. *Id.*

93. Furthermore, one of skill in the art would have had a reasonable expectation of success in implementing this combination because it would have required a simple implementation of Ishimori's aging process to Smith-Edsall. Smith, Ishimori, and Edsall already leverage link aggregation techniques in the packet forwarding methods they disclose. Thus, it would have been within the knowledge of a POSA to simply implement the additional aging process taught by Ishimori to Smith-Edsall, and a POSA would have been able to do so with a reasonable expectation of success. *Id.*

3. Smith, Ishimori, and Zelig

94. A POSA would have been motivated to combine Smith-Ishimori with Zelig. All three references disclose methods of using MAC bridges in a data communication network. Ex. 1004 ¶ 54; Ex. 1005 ¶ 9; Ex. 1007 at Abstract. It

would have been obvious to a POSA to implement the MAC bridges in Smith-Ishimori to maintain a separate MAC table for each VPN MAC bridge as taught in Zelig.




FIG. 1

95. It would have also been obvious to a POSA to configure each virtual MAC bridge to serve a respective VPN, as shown above in Figure 1. Furthermore, implementing the MAC bridges as taught in Zelig to the Smith-Ishimori data communication network would have been an implementation that a POSA would have known how to make with a reasonable expectation of success.

4. Smith, Ishimori, Zelig, and 802.1Q-1998

96. A POSA would have been motivated to combine Smith-Ishimori-Zelig with 802.1Q-1998. All four references pertain to MAC bridges in a data

communication network. Ex. 1004 ¶ 54; Ex. 1005 ¶ 9; Ex. 1007 at Abstract; Ex. 1008 at Abstract. Moreover, Zelig expressly cites to the 802.1Q standard. Ex. 1007 ¶ 12. A POSA would have therefore been motivated to look to the teachings of 802.1Q-1998 in combination with Smith-Ishimori-Zelig to implement the messages from Smith-Ishimori-Zelig with a VLAN identifier in order to identify the VPLS instance. Furthermore, modifying the Smith-Ishimori-Zelig message to add the VLAN identifier would have been a modification that a POSA would have known how to make with a reasonable expectation of success.

C. Specific Grounds of Invalidity

1. Ground 1: Claims 1, 3, 6, 11, 13, and 16 would have been obvious over Smith

97. I have analyzed claims 1, 3, 6, 11, 13, and 16 and in my opinion they would have been obvious over Smith. I provide a detailed analysis of each claim limitation below.

a. Claim 1[pre]: A method for communication, comprising:

98. In my opinion, Smith discloses claim 1[pre]. Smith teaches a system that includes “several communication links.” Ex. 1004 ¶ 9. As stated in the background, computer communication networks are comprised of communication links (and nodes), and these communication links result in a method for communication. Indeed, Smith’s disclosure “relates to networking.” *Id.* ¶ 2.

- b. Claim 1[a]: configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;**

99. In my opinion, Smith discloses claim 1[a]. In figure 3 of Smith, there is a **virtual network device 202**, which is a network node. *Id.* ¶ 36. The virtual network device includes **several line cards**, such as 304(1) and 304(4), which may correspond to the first and second line cards. *Id.* ¶ 46. “In virtual network device sub-unit 122(1), line card 304(1) includes forwarding engine 314(1) and interfaces 320(5), 320(7), and 320(9).” *Id.* ¶ 47. “Line card 304(4) includes forwarding engine 314(4) and interfaces 320(12), 320(14), and 320(16).” *Id.* ¶ 48. The interfaces on line cards 304(1) and 304(4) are a plurality of ports. An “interface” and “port” are synonymous and Smith treats them interchangeably, and thus, discloses the first and second ports. *Id.* ¶ 63 (stating “port *or* uplink interface”). Therefore, it is my opinion that Smith’s virtual network device, which is a network node, is configured to have a first and second line card with respective first and second ports.




FIG. 3

100. Smith further teaches virtual link bundles, which may be link aggregation (LAG) groups and may provide Layer 2 forwarding, and are “managed as a single link.” *Id.* ¶¶ 6, 9, 30. Indeed, Smith discloses that:

- a. “One way to avoid the complexity of having several independent redundant links is to operate those links as single logical transmission path, such as that provided using a **link bundling technique like** EtherChannel (™) or **link aggregation (defined in IEEE 802.3).**” *Id.* ¶ 6 (emphasis added).

- b. “In some embodiments, virtual link bundles **250(1)** and **250(2)** are each operated as an EtherChannel (TM) or as an **aggregated link (as described in IEEE 802.3)**.” *Id.* ¶ 35. (emphasis added).
- c. “[...] the non-satellite network devices provide **L2 (Layer 2)** and **L3 (Layer 3) forwarding** and routing [...]” and “[...] the satellite network devices simply forward all packets to non-satellite network devices for **L2 forwarding** and L3 routing.” *Id.* ¶ 30 (emphasis added).
- d. “For a given virtual link bundle, that virtual link bundle can be managed (e.g., with respect to control protocols such as **L2 protocols**) in a central location.” *Id.* ¶ 57 (emphasis added).
- e. “In some embodiments, **MAC** notification frames are used to keep the content of the **L2 tables** in virtual network device sub-unit **122(1)** synchronized with the content of the **L2 tables** in virtual network device sub-unit **122(2)** and vice versa.” *Id.* ¶ 62 (emphasis added).
- f. “The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from

a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device sub-unit **122(1)** can learn that the sending device's **MAC (Media Access Control)** address is "behind" uplink interface **320(13)** by associating the MAC address with the logical identifier of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** can inform each forwarding engine in virtual network device sub-unit **122(1)** as well as each forwarding engine in virtual network device sub-unit **122(2)** of this association. Based on the association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a **packet addressed to that MAC address can be forwarded** via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added).

- g. “The communication links are configured to be **managed as a single link**.” *Id.* ¶ 9 (emphasis added).

h. “Network devices **120(1)-120(3)** each operate their multiple uplinks to virtual network device **202** as a **single logical uplink**. Additionally, in some embodiments, each network device **120(1)-120(3)** can operate as if that network device is coupled to a **single distribution-layer device**, virtual network device **202**, instead of operating as if that network device were coupled to two independent distribution-layer network devices.” *Id.* ¶ 44 (emphasis added).

101. The virtual network device 202 “route[s] and forward[s] packets to and from network devices 120(1)–120(3)” by associating the MAC address of a received data packet with the logical identifier of the uplink interface. *Id.* ¶ 54. Therefore, it is my opinion that Smith’s virtual network device 202 operates as a distributed MAC bridge in Layer 2 of the data network.

102. In conclusion, it is my opinion that Smith teaches configuring a network node (the virtual network device) having at least a first and second line card (the line cards 304(1) and 304(4)) with a plurality of ports (the interfaces on the line cards) that operates as a distributed MAC bridge in Layer 2 of the data network.

c. **Claim 1[b]: configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined Member Line Cards;**

103. In my opinion, Smith discloses claim 1[b]. Smith teaches virtual link bundles, which may be link aggregation (LAG) groups and may provide Layer 2 forwarding, and are “managed as a single link.” *Id.* ¶¶ 6, 9, 30. Indeed, Smith discloses that:

- a. “One way to avoid the complexity of having several independent redundant links is to operate those links as single logical transmission path, such as that provided using a **link bundling technique like** EtherChannel (TM) or **link aggregation (defined in IEEE 802.3).**” *Id.* ¶ 6 (emphasis added).
- b. “In some embodiments, virtual link bundles **250(1)** and **250(2)** are each operated as an EtherChannel (TM) or as an **aggregated link (as described in IEEE 802.3).**” *Id.* ¶ 36 (emphasis added).
- c. “[T]he non-satellite network devices provide **L2 (Layer 2)** and L3 (Layer 3) **forwarding** and routing” and “the satellite network devices simply forward all packets to non-satellite network devices for **L2 forwarding** and L3 routing.” *Id.* ¶ 30 (emphasis added).
- d. “For a given virtual link bundle, that virtual link bundle can be managed (e.g., with respect to control protocols such as **L2 protocols**) in a central location.” *Id.* ¶ 57 (emphasis added).

- e. “In some embodiments, **MAC** notification frames are used to keep the content of the **L2 tables** in virtual network device sub-unit **122(1)** synchronized with the content of the **L2 tables** in virtual network device sub-unit **122(2)** and vice versa.” *Id.* ¶ 62 (emphasis added).
- f. “The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device sub-unit **122(1)** can learn that the sending device's **MAC (Media Access Control)** address is "behind" uplink interface **320(13)** by associating the MAC address with the logical identifier of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** can inform each forwarding engine in virtual network device sub-unit **122(1)** as well as each forwarding engine in virtual network device sub-unit **122(2)** of this association. Based on the association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case,

uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a **packet addressed to that MAC address can be forwarded** via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**. *Id.* ¶ 54 (emphasis added).

- g. “The communication links are configured to be **managed as a single link**.” *Id.* ¶ 9 (emphasis added).
- h. “Network devices **120(1)-120(3)** each operate their multiple uplinks to virtual network device **202** as a **single logical uplink**. Additionally, in some embodiments, each network device **120(1)-120(3)** can operate as if that network device is coupled to a **single distribution-layer device**, virtual network device **202**, instead of operating as if that network device were coupled to two independent distribution-layer network devices.” *Id.* ¶ 44 (emphasis added).

104. Smith discloses that **network device 120(2)** or one of the endpoints “is coupled to **virtual network device 202**” or a second endpoint “by **virtual link bundle 250(2)**” as shown in the annotated figure below. *Id.* ¶ 44. The virtual link bundle 250(2) consists of **two uplinks**, which is a plurality of LAG ports. *Id.* ¶ 51.




FIG. 3

105. Smith further specifies that the virtual network devices “provide L2 (Layer 2) . . . forwarding and routing.” *Id.* ¶ 30. Continuing with my above example, line cards 304(1) and 304(4) may correspond to the first and second line card respectively. The other two line cards 304(2) and 304(3) may make up the conjoined member line cards, because all four line cards are part of the virtual network device sub-units 122(1) and 122(2) that “can coordinate their behavior such that they appear to be a single virtual device.” *Id.* ¶ 50. In my opinion, Smith therefore discloses configuring a link aggregation (LAG) group of parallel physical links (the virtual link bundle 250(2)) between two endpoints (virtual network device 202 and network

device 120(2)) in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports (interfaces 320(8) and 320(11)) and a plurality of conjoined Member Line Cards (line cards 304(2) and 304(3)).

d. Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;

106. In my opinion, Smith discloses claim 1[c]. The member line cards in Smith include **interfaces**, which are ports. Ex. 1004 ¶ 47. And the forwarding tables in Smith correspond to forwarding databases. Indeed, Smith discloses:

- a. “Line card **304(3)** includes forwarding engine **314(3)**, **interfaces 320(11) and 320(13)**, and **port 320(15)**. *Id.* ¶ 47 (emphasis added).
- b. “identifier of the identified **port or uplink interface . . .**” *Id.* ¶ 65 (emphasis added).
- c. “If the packet is received from a local uplink **interface or port** (i.e., if the packet is not received from another virtual network device sub-unit within the virtual network device), as determined at **405**, the virtual network device sub-unit attempts to forward the packet to its destination address. For example, the virtual network device sub-unit can provide the destination address to a **forwarding table** in order to determine which **logical identifier, if any, is associated with that destination address**. If there is no hit

in the **forwarding table**, as determined at **407**, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation, the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at **409**. For **interfaces (e.g., ports or uplinks)** included in interface bundles, the virtual network device sub-unit selects one egress interface per interface bundle via which to send the packet. If the packet was received by the virtual network device sub-unit via an interface bundle, all interfaces in that interface bundle are excluded from sending the packet.” *Id.* ¶ 66 (emphasis added).

- d. “If the packet’s destination address hits in the **forwarding table**, as determined at **407**, the virtual network device sub-unit uses the logical identifier returned by the **forwarding table** to select the interface(s) to which the packet should be sent. If the **forwarding table** does not identify an interface bundle, as determined at **411**, the packet is sent via the identified port(s) and/or uplink

interface(s), as indicated at **413**. If the **forwarding table** does identify an interface bundle, the virtual network device sub-unit sends the packet via one local interface included within the identified interface bundle, as shown at **415** (if the forwarding table identifies other non-interface-bundle interfaces, the packet is sent via those interfaces as well).” *Id.* ¶ 67 (emphasis added).

- e. “The virtual network device sub-unit determines whether that sub-unit has already learned the logical identifier associated with the packet's destination device. In this example, this is performed by providing the destination address to a **forwarding table**, as shown a **417**. If there is not a hit in the **forwarding table** (i.e., if no association has already been learned for the destination address), the virtual network device sub-unit floods the packet on the incoming VLAN.” *Id.* ¶ 69 (emphasis added).
- f. “If there is a hit in the forwarding table, and if the forwarding table does not identify an interface bundle at **421**, the packet is sent via the identified port and/or uplink interfaces, as indicated at **423**. If instead the forwarding table does identify an interface bundle, the packet is not sent via that interface bundle.” *Id.* ¶ 70 (emphasis added).

- g. “When virtual network device sub-unit **122(1)** looks up the destination address of the packet in a **lookup table**, the lookup table returns the logical identifier that identifies local uplink interfaces **320(9)** and **320(13)**. The packet is then forwarded to uplink interface **320(13)** (e.g., selected based on load-sharing considerations).” *Id.* ¶ 61 (emphasis added).
- h. “For example, control unit **310(2)** can use this information to set up or modify **lookup tables on line cards 304(2) and/or 304(4).**” *Id.* ¶ 57 (emphasis added).
- i. “The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device sub-unit **122(1)** can learn that the sending device's MAC (Media Access Control) address is "behind" uplink interface **320(13)** by associating the MAC address with the logical identifier of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** can **inform each forwarding engine** in virtual network device sub-

unit **122(1)** as well as **each forwarding engine** in virtual network device sub-unit **122(2)** of this association. Based on the association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a packet addressed to that MAC address can be forwarded via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added).




FIG. 3

107. When a packet is received on a particular uplink interface, the virtual network device, which is a network node as described with respect to claim 1[a], learns the sending device's MAC address by "associating the MAC address with the logical identifier of [the] uplink interface" or the ports of said plurality of ports. *Id.* ¶ 54. "The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device 202 to route and forward packets to and from network devices 120(1)-120(3)." *Id.* Smith teaches that this information may be used to "set up or modify lookup tables," which maps to a forwarding database (FDB). *Id.* ¶¶ 57, 61. Therefore, it is my opinion that Smith teaches that the lookup table or forwarding table or FDB may store records associating MAC addresses with the logical identifier of the uplink interface of said network node.

e. Claim 1[d]: receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;

108. In my opinion, Smith discloses claim 1[d]. Smith teaches:

- a. "The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a **packet from a sending device (e.g., a client coupled to network device**

120(1)) is received via uplink interface 320(13), virtual network device sub-unit 122(1) can learn that the sending device's MAC (Media Access Control) address is "behind" uplink interface 320(13) by associating the MAC address with the logical identifier of uplink interface 320(13). Virtual network device sub-unit 122(1) can inform each forwarding engine in virtual network device sub-unit 122(1) as well as each forwarding engine in virtual network device sub-unit 122(2) of this association. Based on the association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case, uplink interfaces 320(9) (in virtual network device sub-unit 122(1)) and 320(16) (in virtual network device sub-unit 122(2)) also have the same logical identifier as uplink interface 320(13), a packet addressed to that MAC address can be forwarded via any of uplink interfaces 320(9), 320(13), and 320(16).” *Id.* ¶ 54 (emphasis added)

- b. “If the packet is **received from a local uplink** interface or port (i.e., if the packet is not received from another virtual network device sub-unit within the virtual network device), as determined at **405**, the virtual network device sub-unit attempts to **forward the**

packet to its destination address. For example, the virtual network device sub-unit can provide the **destination address** to a forwarding table in order to determine which logical identifier, if any, is associated with that **destination address**. If there is no hit in the forwarding table, as determined at **407**, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation, the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at **409**. For **interfaces (e.g., ports or uplinks)** included in interface bundles, the virtual network device sub-unit selects one egress interface per interface bundle via which to send the packet. If the packet was received by the virtual network device sub-unit via an interface bundle, all interfaces in that interface bundle are excluded from sending the packet.” *Id.* ¶ 66 (emphasis added).

109. Smith discloses that the uplink interface receives a data packet with the “sending device’s MAC address.” *Id.* ¶ 54. The uplink interface is an ingress port

and the sending device's MAC address is the MAC source address. For example, if the network device 120(1) is the MAC source address, the ingress port would be interface 320(9) from a data packet received from network device 120(1).




FIG. 3

110. It is therefore my opinion that Smith teaches receiving a data packet on an ingress port (interface 320(9)) of said network node (virtual network device 202) from a MAC source address (network device 120(1)).

111. Smith further discloses that this data packet has a “destination logical identifier.” *Id.* ¶¶ 60, 62. The MAC destination address could be network device 120(3), for example. Smith specifies that these devices provide Layer 2 forwarding. *Id.* ¶ 30. The data packet’s MAC destination address (network device 120(3)) is on

the Layer 2 of the data network. In my opinion, Smith teaches receiving a packet on an ingress port (for example, interface 320(9)) of said network node from a MAC source address (for example, network device 120(1)), said data packet specifying a MAC destination address (for example, network device 120(3)).

- f. **Claim 1[e]: conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;**

112. In my opinion, Smith discloses claim 1[e]. Smith discloses:

- a. “The association between a packet and a particular logical identifier can be used by forwarding engines within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a **packet from a sending device (e.g., a client coupled to network device 120(1)) is received via uplink interface 320(13), virtual network device sub-unit 122(1) can learn that the sending device's MAC (Media Access Control) address is "behind" uplink interface 320(13) by associating the MAC address with the logical identifier of uplink interface 320(13). Virtual network device sub-unit 122(1) can inform each forwarding engine in virtual network device sub-unit 122(1) as well as each forwarding engine in virtual network**

device sub-unit **122(2)** of this association. Based on the association, **packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier**. Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a packet addressed to that MAC address can be forwarded via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added).

- b. “If the packet is **received from a local uplink** interface or port (i.e., if the packet is not received from another virtual network device sub-unit within the virtual network device), as determined at **405**, the virtual network device sub-unit attempts to **forward the packet to its destination address**. For example, the virtual network device sub-unit can provide the **destination address** to a forwarding table in order to determine which logical identifier, if any, is associated with that **destination address**. If there is no hit in the forwarding table, as determined at **407**, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation,

the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at **409**. For **interfaces (e.g., ports or uplinks)** included in interface bundles, the virtual network device sub-unit selects one **egress interface** per interface bundle via which to send the packet. If the packet was received by the virtual network device sub-unit via an interface bundle, all interfaces in that interface bundle are excluded from sending the packet.” *Id.* ¶ 66 (emphasis added).

113. Smith teaches that the data packet is forwarded based on “the association between a packet and a particular logical identifier.” *Id.* ¶ 54. Smith additionally teaches that its system “favor[s] local interfaces.” *Id.* ¶ 56. Again, referring to Figure 3, in order to transmit the data packet to the MAC destination address (network device 120(3)), the data packet would be conveyed from interface 320(9) to interface 320(7) (said first port), because the system favors local interfaces. Interface 320(7), which is on line card 304(1) (said first line card), would then transmit the data packet to the MAC destination address (port of network device 120(3)).




FIG. 3

- g. **Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;**

114. In my opinion, Smith discloses claim 1[f]. Smith teaches:

- a. “If the packet is received from a local uplink interface or port (i.e., if the packet is not received from another virtual network device sub-unit within the virtual network device), as determined at 405, the virtual network device sub-unit attempts to forward the packet to its destination address. For example, the virtual network device sub-unit can provide the destination address to a forwarding table in order to determine which logical identifier, if any, is associated

with that destination address. **If there is no hit in the forwarding table**, as determined at 407, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation, the virtual network device sub-unit **floods the packet to all egress ports** and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at 409. For interfaces (e.g., ports or uplinks) **included in interface bundles**, the virtual network device sub-unit selects **one egress interface per interface bundle** via which to send the packet. If the packet was received by the virtual network device sub-unit via an interface bundle, all interfaces in that interface bundle are excluded from sending the packet.” *Id.* ¶ 66 (emphasis added).

- b. “If the packet’s destination address hits in the forwarding table, as determined at 407, the virtual network device sub-unit uses the logical identifier returned by the forwarding table to select the interface(s) to which the packet should be sent. If the forwarding table does not identify an interface bundle, as determined at 411,

the packet is sent via the identified port(s) and/or uplink interface(s), as indicated at 413. If the forwarding table does identify an interface bundle, the virtual network device sub-unit sends the packet via **one local interface** included within the identified interface bundle, as shown at 415 (if the forwarding table identifies other non-interface-bundle interfaces, the packet is sent via those interfaces as well).” *Id.* ¶ 67 (emphasis added).

- c. “The virtual network device sub-unit determines whether that sub-unit has already learned the logical identifier associated with the packet's destination device. In this example, this is performed by providing the destination address to a forwarding table, as shown a **417. If there is not a hit in the forwarding table** (i.e., if no association has already been learned for the destination address), the virtual network device sub-unit **floods the packet on the incoming VLAN.**” *Id.* ¶ 69 (emphasis added).

115. Smith teaches that the virtual network device sub-unit 122(1) looks up the destination address in a “lookup table” (FDB). *Id.* ¶ 61. Smith discloses sharing entries of the FDB with all the member line cards. *Id.* ¶ 63. “If a forwarding engine within virtual network device sub-unit 122(2) already knows that the destination address is behind a particular uplink interface or port . . . , that forwarding engine

generates a MAC notification identifying this association, which is distributed to any other forwarding engines within virtual network device sub-unit 122(2).” *Id.* It is my understanding therefore that one way to read “said FDB” in the claim is that it could be referring to the FDBs associated with any of the member line cards. Alternatively, “said FDB” could refer to the FDB on the line card associated with the ingress port, which would be line card 304(1) or the first line card in my example above. Although there is no line card or FDB associated with the ingress port or first line card recited in the claim, it would have been obvious to a POSA that the first line card could have an FDB.

116. If there is no hit in the forwarding table, “the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN.” *Id.* ¶¶ 63, 66. Smith discloses that “flooding” is “via one and only one LAG port.” In particular, Smith teaches that when the network device sends a packet “via the virtual link bundle,” it “selects one of the communication links on which to send the packet.” *Id.* ¶ 9. Thus, it my opinion that Smith discloses, or at least renders obvious, sending the data packet via one and only one LAG port.

h. Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and

117. In my opinion, Smith discloses claim 1[g]. Smith teaches:

- a. “If the packet is received from a local uplink interface or port (i.e., if the packet is not received from another virtual network device sub-unit within the virtual network device), as determined at **405**, the virtual network device sub-unit attempts to forward the packet to its destination address. For example, the virtual network device sub-unit can **provide the destination address to a forwarding table** in order to determine which **logical identifier, if any, is associated with that destination address**. If there is no hit in the forwarding table, as determined at **407**, the virtual network device sub-unit has not yet learned the logical identifier of the interface(s) in front of the destination device(s). In this situation, the virtual network device sub-unit floods the packet to all egress ports and/or uplink interfaces in the incoming VLAN (Virtual Local Area Network) (the incoming VLAN is the VLAN that includes the device identified by the packet's source address), excluding the interface that the packet arrived on, as shown at **409**. For interfaces (e.g., ports or uplinks) included in interface bundles, the virtual network device sub-unit selects one egress interface per interface bundle via which to send the packet. If the packet was received by the virtual network device sub-unit via an interface bundle, all

interfaces in that interface bundle are excluded from sending the packet.” *Id.* ¶ 66 (emphasis added).

- b. “If the packet’s destination address **hits in the forwarding table**, as determined at **407**, the virtual network device sub-unit uses the **logical identifier returned by the forwarding table** to select the interface(s) to which the packet should be sent. If the forwarding table does not identify an interface bundle, as determined at **411**, the packet is sent via the identified port(s) and/or uplink interface(s), as indicated at **413**. If the forwarding table does identify an interface bundle, the virtual network device sub-unit sends the packet via one local interface included within the identified interface bundle, as shown at **415** (if the forwarding table identifies other non-interface-bundle interfaces, the packet is sent via those interfaces as well).” *Id.* ¶ 67 (emphasis added).
- c. “The virtual network device sub-unit determines whether that sub-unit has already learned the logical identifier associated with the packet's destination device. In this example, this is performed by **providing the destination address to a forwarding table**, as shown a **417**. If there is not a hit in the forwarding table (i.e., if no association has already been learned for the destination address),

the virtual network device sub-unit floods the packet on the incoming VLAN.” *Id.* ¶ 69 (emphasis added).

- d. “If there is a **hit in the forwarding table**, and if the forwarding table does not identify an interface bundle at **421**, the packet is sent via the identified port and/or uplink interfaces, as indicated at **423**. If instead the forwarding table does identify an interface bundle, the packet is not sent via that interface bundle.” *Id.* ¶ 70 (emphasis added).
- e. “When virtual network device sub-unit **122(1) looks up the destination address of the packet in a lookup table**, the lookup table returns the logical identifier that identifies local uplink interfaces **320(9)** and **320(13)**. The packet is then forwarded to uplink interface **320(13)** (e.g., selected based on load-sharing considerations).” *Id.* ¶ 61 (emphasis added).

118. Smith teaches that “[b]ased on the source address of the packet and which port or uplink interface received the packet, the virtual network device sub-unit learns the source identifier of the sending device, as indicated at 403.” *Id.*

¶ 65. These identifiers are stored in a lookup table on a virtual network device. *Id.*

¶ 61. It would have therefore been obvious to check the MAC source address in

the records of the FDB of said first line card (line card 304(1)) to determine if it exists in the FDB.

- i. **Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.**

119. In my opinion, Smith discloses claim 1[h]. Smith teaches:

- a. “The association between a packet and a particular logical identifier can be used by **forwarding engines** within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device sub-unit **122(1)** can **learn** that the sending device's MAC (Media Access Control) address is "behind" uplink interface **320(13)** by **associating the MAC address with the logical identifier** of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** **can inform each forwarding engine** in virtual network device sub-unit **122(1)** as well as **each forwarding engine** in virtual network device sub-unit **122(2)** of this association. Based on the

association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a packet addressed to that MAC address can be forwarded via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added).

- b. “In some embodiments, **MAC notification frames** are used to keep the **content of the L2 tables** in virtual network device sub-unit **122(1)** **synchronized** with the content of the L2 tables in virtual network device sub-unit **122(2)** and vice versa. Whenever a MAC notification that involves a port behind a virtual link bundle or an uplink interface included in an uplink interface bundle is generated **within a virtual network device sub-unit** (e.g., **such a notification** can be generated by **one line card** in order to **update an L2 table on another line card**), a **copy of the MAC notification is sent via** to virtual network device link **360**.

Similarly, if a virtual network device sub-unit determines that a packet should be flooded, the virtual network device sub-unit will

send a copy of that packet via virtual network device link **360**, ensuring that the virtual network device sub-unit will receive a copy of any MAC notification response generated by a forwarding engine in the peer virtual network device sub-unit.” *Id.* ¶ 62 (emphasis added).

- c. “By way of example, assume that virtual network device sub-unit **122(1)** floods a packet because the forwarding engine(s) included in virtual network device sub-unit **122(1)** do not know which port or uplink interface is associated with the packet's destination address. As part of flooding the packet, virtual network device sub-unit **122(1)** sends a copy of the packet to virtual network device sub-unit **122(2)** via virtual switch link **360**. If a forwarding engine within virtual network device sub-unit **122(2)** already knows that the destination address is behind a particular uplink interface or port (e.g., if a **forwarding table already includes an entry associating the destination address with a port** of one of network devices **120**), that forwarding engine **generates a MAC notification** identifying this association, which is **distributed to any other forwarding engines** within virtual network device sub-unit **122(2)**. Since the packet was originally received via virtual

network device link **360**, virtual network device sub-unit **122(2)** also sends a copy of the MAC notification back via virtual network device link **360**. This MAC notification can then be distributed among the forwarding engines included in virtual network device sub-unit **122(1)**. After being updated based on the MAC notification, the forwarding engines in virtual network device sub-unit **122(1)** now know the location of the device identified by the destination address. Accordingly, subsequently- received packets addressed to that device will not be flooded.” *Id.* ¶ 63 (emphasis added).

120. A POSA would have understood that if the record of an association between the MAC source address with said ingress port did not exist in the lookup table of, for example, line card 304(1) (said first line card), it would create a new entry and add it to the lookup table of line card 304(1). Smith teaches that the network device sub-unit 122(2) sends a MAC notification (a message) to update the forwarding engines when it learns of a new association. *Id.* ¶ 63. “After being updated based on the MAC notification, the forwarding engines in the virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* I understand this to mean that the MAC notification,

which is a message, is therefore sent to each member line card (304(2) through 304(4)).

- j. **Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.**

121. In my opinion, Smith discloses claim 3. Smith teaches:

- a. “The association between a packet and a particular logical identifier can be used by **forwarding engines** within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device sub-unit **122(1)** can **learn** that the sending device's MAC (Media Access Control) address is "behind" uplink interface **320(13)** by **associating the MAC address with the logical identifier** of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** **can inform each forwarding engine** in virtual network device sub-unit **122(1)** as well as **each forwarding engine** in virtual network device sub-unit **122(2)** of this association. Based on the association, packets addressed to that MAC address will be sent

from an uplink interface having the associated logical identifier.

Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a packet addressed to that MAC address can be forwarded via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added)

- b. “In some embodiments, **MAC notification frames** are used to keep the **content of the L2 tables** in virtual network device sub-unit **122(1)** **synchronized** with the content of the L2 tables in virtual network device sub-unit **122(2)** and vice versa. Whenever a MAC notification that involves a port behind a virtual link bundle or an uplink interface included in an uplink interface bundle is generated **within a virtual network device sub-unit** (e.g., such a **notification** can be generated **by one line card** in order to **update an L2 table on another line card**), a **copy of the MAC notification is sent via** to virtual network device link **360**.

Similarly, if a virtual network device sub-unit determines that a packet should be flooded, the virtual network device sub-unit will send a copy of that packet via virtual network device link **360**,

ensuring that the virtual network device sub-unit will receive a copy of any MAC notification response generated by a forwarding engine in the peer virtual network device sub-unit.” *Id.* ¶ 62 (emphasis added).

- c. “By way of example, assume that virtual network device sub-unit **122(1)** floods a packet because the forwarding engine(s) included in virtual network device sub-unit **122(1)** do not know which port or uplink interface is associated with the packet's destination address. As part of flooding the packet, virtual network device sub-unit **122(1)** sends a copy of the packet to virtual network device sub-unit **122(2)** via virtual switch link **360**. If a forwarding engine within virtual network device sub-unit **122(2)** already knows that the destination address is behind a particular uplink interface or port (e.g., if a **forwarding table already includes an entry associating the destination address with a port** of one of network devices **120**), that forwarding engine **generates a MAC notification** identifying this association, which is **distributed to any other forwarding engines** within virtual network device sub-unit **122(2)**. Since the packet was originally received via virtual network device link **360**, virtual network device sub-unit **122(2)**

also sends a copy of the MAC notification back via virtual network device link **360**. This MAC notification can then be distributed among the forwarding engines included in virtual network device sub- unit **122(1)**. After being updated based on the MAC notification, the forwarding engines in virtual network device sub-unit **122(1)** now know the location of the device identified by the destination address. Accordingly, subsequently- received packets addressed to that device will not be flooded.” *Id.* ¶ 63 (emphasis added).

122. Smith discloses that the network device sub-unit 122(2) sends a MAC notification (the message) to update the forwarding engines. *Id.* ¶ 63. Smith specifies that “[a]fter being updated based on the MAC notification, the forwarding engines in virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* I understand that this includes the other line cards in virtual network device 202, such as line card 304(4). For example, line card 304(4) (*e.g.*, the second line card) would receive the MAC notification (the message) and update its lookup tables. *Id.* ¶ 57. Thus, Smith discloses that in response to the message, the record of association is added to the FDB of the second line card if the record does not already exist in the FDB of the second line card.




FIG. 3

- k. **Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via switching core of the network node to at least the second line card.**

123. In my opinion, Smith discloses claim 6. Smith teaches:

- a. “The association between a packet and a particular logical identifier can be used by **forwarding engines** within virtual network device **202** to route and forward packets to and from network devices **120(1)-120(3)**. For example, when a packet from a sending device (e.g., a client coupled to network device **120(1)**) is received via uplink interface **320(13)**, virtual network device

sub-unit **122(1)** can **learn** that the sending device's MAC (Media Access Control) address is "behind" uplink interface **320(13)** by **associating the MAC address with the logical identifier** of uplink interface **320(13)**. Virtual network device sub-unit **122(1)** **can inform each forwarding engine** in virtual network device sub-unit **122(1)** as well as **each forwarding engine** in virtual network device sub-unit **122(2)** of this association. Based on the association, packets addressed to that MAC address will be sent from an uplink interface having the associated logical identifier. Since in this case, uplink interfaces **320(9)** (in virtual network device sub-unit **122(1)**) and **320(16)** (in virtual network device sub-unit **122(2)**) also have the same logical identifier as uplink interface **320(13)**, a packet addressed to that MAC address can be forwarded via any of uplink interfaces **320(9)**, **320(13)**, and **320(16)**.” *Id.* ¶ 54 (emphasis added).

- b. “In some embodiments, **MAC notification frames** are used to keep the **content of the L2 tables** in virtual network device sub-unit **122(1)** **synchronized** with the content of the L2 tables in virtual network device sub-unit **122(2)** and vice versa. Whenever a MAC notification that involves a port behind a virtual link bundle

or an uplink interface included in an uplink interface bundle is generated **within a virtual network device sub-unit** (e.g., **such a notification** can be generated **by one line card** in order to **update an L2 table on another line card**), a **copy of the MAC notification is sent via** to virtual network device link **360**.

Similarly, if a virtual network device sub-unit determines that a packet should be flooded, the virtual network device sub-unit will send a copy of that packet via virtual network device link **360**, ensuring that the virtual network device sub-unit will receive a copy of any MAC notification response generated by a forwarding engine in the peer virtual network device sub-unit.” *Id.* ¶ 62 (emphasis added).

- c. “By way of example, assume that virtual network device sub-unit **122(1)** floods a packet because the forwarding engine(s) included in virtual network device sub-unit **122(1)** do not know which port or uplink interface is associated with the packet's destination address. As part of flooding the packet, virtual network device sub-unit **122(1)** sends a copy of the packet to virtual network device sub-unit **122(2)** via virtual switch link **360**. If a forwarding engine within virtual network device sub-unit **122(2)** already knows that

the destination address is behind a particular uplink interface or port (e.g., if a **forwarding table already includes an entry associating the destination address with a port** of one of network devices **120**), that forwarding engine **generates a MAC notification** identifying this association, which is **distributed to any other forwarding engines** within virtual network device sub-unit **122(2)**. Since the packet was originally received via virtual network device link **360**, virtual network device sub-unit **122(2)** also sends a copy of the MAC notification back via virtual network device link **360**. This MAC notification can then be distributed among the forwarding engines included in virtual network device sub-unit **122(1)**. After being updated based on the MAC notification, the forwarding engines in virtual network device sub-unit **122(1)** now know the location of the device identified by the destination address. Accordingly, subsequently- received packets addressed to that device will not be flooded.” *Id.* ¶ 63 (emphasis added).

124. The line cards communicate with each other in order to transfer various messages between them. For example, note that “such a notification can be generated by one line card in order to update an L2 table on another line card.” *Id.* ¶ 62.

Therefore, the line cards must communicate, allowing for sending messages from one line card to another. This, internal to device 202, switching mechanism allowing the line cards to communicate with each other could be the switching core.

125. Smith teaches that the “MAC notification frames are used to keep the content of the L2 tables in virtual network device sub-unit 122(1) synchronized with the content of the L2 tables in virtual network device sub-unit 122(2) and vice versa.” *Id.* Thus, it would have been obvious to a POSA to implement the MAC notification frame as a synchronization packet and transmit it to at least the second line card via the switching core.

l. Claim 11[pre] 11: A node for network communication, comprising:

126. In my opinion, Smith discloses claim 11[pre]. Smith teaches that its system contains communications links that are coupled to network devices. *Id.* ¶ 9. A POSA would understand that a network device is a node.

m. Claim 11[a]: a switching core;

127. In my opinion, Smith discloses claim 11[a]. It would be within the knowledge of a POSA that the line cards in Smith would be communicating with each other, as discussed before. *See supra* ¶ 124.

n. Claim 11[b]: a plurality of ports;

128. In my opinion, Smith discloses claim 11[b]. The network devices in Smith contain several ports (a plurality of ports). Ex. 1004 ¶ 6. “For example, an

EtherChannel (TM) port bundle can be formed *from several ports on a switch*, each of which is coupled to a respective link in a group of links coupling that switch to another switch.” *Id.*

- o. Claim 11[c]: a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards;**

129. In my opinion, Smith discloses claim 11[c]. Claim 11[c] combines the elements of claims 1[a], 1[b], 1[c], and 11[a]. Therefore, for the same reasons that the Smith discloses claims 1[a], 1[b], 1[c], and 11[a], Smith also discloses claim 11[c]. *See supra* ¶¶ 99–107, 127.

- p. Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds**

said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.

130. In my opinion, Smith discloses claim 11[d]. This limitation is a combination of claims 1[d], 1[e], 1[f], 1[g], and 1[h]. Therefore, for the same reasons that Smith discloses claims 1[d], 1[e], 1[f], 1[g], and 1[h], Smith also discloses claim 11[d]. *See supra* ¶¶ 108–120.

- q. Claim 13: The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.**

131. In my opinion, Smith discloses claim 13. *See supra* ¶¶ 121–122.

- r. Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.**

132. In my opinion, Smith discloses claim 16. *See supra* ¶¶ 123–125.

2. Ground 2: Claims 1–3 and 11–13 would have been obvious over Smith in view of Ishimori

133. I have analyzed claims 1–3 and 11–13 and conclude that they would have been obvious over Smith in combination with Ishimori. I provide a detailed analysis of each claim limitation below.

- a. **Claim 1[c]: providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;**

134. In my opinion, the combination of Smith and Ishimori render obvious claim 1[c]. As explained above, when a packet is received on a particular uplink interface, the virtual network device learns the sending device's MAC address by "associating the MAC address with the logical identifier of [the] uplink interface." Ex. 1004 ¶ 54. Smith teaches that the information learned about the association between a packet and a particular logical identifier may be used to "set up or modify lookup tables" (e.g., the forwarding database (FDB)). *Id.* ¶¶ 54, 57, 61, 71.

135. To the extent Smith alone does not render obvious claim 1[c], Smith in combination with Ishimori does. Indeed, Ishimori's disclosures include:

- a. "As a packet forwarding method via a network, there is a method of learning a source MAC address of a receive IP packet together with a receive path thereof and, when an IP packet whose destination address is the same MAC address is received, using the corresponding learning result to decide a transmission path thereof. In this situation, when forwarding, at a node that relays IP packets, an IP packet having an unlearned MAC address as a destination address, a transmission path thereof cannot be specified. As such, in this situation, a broadcast transmission to all nodes (that is,

“flooding”) is performed. Then, at this time, path **learning** is carried out **by storing** the source MAC address (SA) had by this IP packet and a port that received such in a buffer had by each node.”

Ex. 1005 ¶ 2 (emphasis added).

- b. “FIG. 1 and FIG. 2 illustrate this. Note that in the example described below that is illustrated in FIG. 1 and the like, a configuration is supposed wherein each node—50-0, 50-1, 50-2—has one communication card—#0, #1, #2—respectively. When node 50-0 receives an IP packet from terminal A 10 at port #0 of card #0 (steps S1, P1), in a **MAC table** had **by this card**, “A”, which is a source MAC address (SA) thereof, is learned by being **stored** in a predetermined buffer in association with information on the receiving card #0 and port #0 at this time (path information) (step S2). Note that in this situation, when a learning result for the same MAC address already exists in the **MAC table**, this information is overwritten.” *Id.* ¶ 3 (emphasis added).
- c. “Next, for a destination address (DA; here, MAC address “B”) had by this receive IP packet, it is searched whether there is a learning result in the MAC table of the local device (step S3). Then, when the search result is a mishit (“No”), this IP packet is broadcast

(flooded) to all nodes (in this example, node 50-1 and node 50-2) (steps S5, P2). Meanwhile, when, as a result of the search at step S3, a learning result relating to the corresponding destination MAC address exists in the buffer (“Yes”), this IP packet is transmitted according to a forwarding path—that is, a card number and port number—included in this learning result (step S4). In this situation, there is no need for flooding.” *Id.* ¶ 4 (emphasis added).

- d. “Note that the same learning operations are also performed when this IP packet sent from node 50-0 reaches card #2 of node 50-2, which the IP packet passes through on its way to a corresponding destination terminal B 20. As a result, as illustrated in FIG. 1, source path information at this time—that is, the information on card #0, port #0 of node 50-0—is stored as a learning result in association with the MAC address “A” of the source terminal 10 **in the buffer (MAC table)** of this card #2.” *Id.* ¶ 5 (emphasis added).

136. Ishimori teaches that each of its communication cards has a “MAC table.” *Id.* ¶ 2. Referring to Figure 1, Ishimori explains that when a packet from terminal A is received at port #0 of card #0, the path information for address “A” is learned by storing the information in the MAC table “had by this card” (*e.g.*, card

#0). *Id.* ¶ 3. Ishimori therefore teaches that its line cards maintain a MAC table to store path information of the MAC address and the receiving card and port. *Id.* In other words, the MAC table holds records associating MAC addresses with ports of said plurality of ports of said network node. It would have been obvious to substitute the line cards from Ishimori into Smith and a POSA would have had a reasonable expectation of success in making this simple substitution.

b. **Claim 1[f]: if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;**

137. In my opinion, to the extent Smith does not explicitly teach the limitation of sending the packet to one and only one LAG port, Ishimori expressly discloses flooding via “one representative port” (*e.g.*, via one and only one LAG port).

a. “Next, a **link aggregation** function is described. This link **aggregation function is a function of using a plurality of ports by deeming such to be one virtual high-speed port.** This function enables a band improvement effect to be substantially obtained. The **virtual port in this situation is referred to as a trunk.** In using this trunk (referred to as “trunking”), **when, as above, flooding is to be performed when an address is not learned, one representative port is selected** from among the large

number of ports had by this trunk according to a predetermined computation, and **actual packet transmission is performed using this representative port.**” *Id.* ¶ 13 (emphasis added).

- b. “Then, at a node that receives an IP packet via the trunk, when carrying out path learning of a source address of this packet, instead of learning the port number at the time of receipt, information on the trunk including such is stored as the learning result. Afterward, in receiving an IP packet having a destination that passes through this trunk, one path is decided from among the paths constituting the trunk according to predetermined computation methods, independent of each other according to hardware, of these communication cards, and this decided path is applied to the corresponding packet. This configuration **prevents the application of a different path in the same trunk to the same packet.**” *Id.* ¶ 14 (emphasis added).

138. Ishimori teaches that when “flooding is to be performed when an address is not learned, one representative port is selected from among the large number of ports had by this trunk according to a predetermined computation, and actual packet transmission is performed using this representative port.” *Id.* ¶ 13. One of ordinary skill in the art would have been motivated to combine the teachings

of Smith with Ishimori because Ishimori explains that using “one representative port . . . prevents the application of a different path in the same trunk to the same packet.” *Id.* ¶¶ 13, 14. Thus, a POSA would have looked to both Smith and Ishimori, which explain why it would be beneficial to use only one representative port, and combine the teachings of the two references as they are directed to similar methods. Furthermore, a POSA would have had a reasonable success in combining Smith with Ishimori because it would have been a simple application of Ishimori’s methods on Smith’s virtual network device. Thus, the combination of Smith and Ishimori renders obvious claim 1[f].

c. Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and

139. In my opinion, to the extent this is not disclosed by Smith, Ishimori teaches that the source MAC address is stored with the reception path information including the card number and port number. *Id.* ¶ 3. Ishimori’s disclosures include:

- a. “As a packet forwarding method via a network, there is a method of **learning a source MAC address** of a receive IP packet together with a receive path thereof and, when an IP packet whose destination address is the same MAC address is received, using the corresponding learning result to decide a transmission path thereof. In this situation, when forwarding, at a node that relays IP packets,

an IP packet having an unlearned MAC address as a destination address, a transmission path thereof cannot be specified. As such, in this situation, a broadcast transmission to all nodes (that is, “flooding”) is performed. Then, at this time, path **learning** is carried out **by storing** the source MAC address (SA) had by this IP packet and a port that received such in a buffer had by each node.” *Id.* ¶ 2 (emphasis added).

- b. “FIG. 1 and FIG. 2 illustrate this. Note that in the example described below that is illustrated in FIG. 1 and the like, a configuration is supposed wherein each node—50-0, 50-1, 50-2—has one communication card—#0, #1, #2—respectively. When node 50-0 receives an IP packet from terminal A 10 at port #0 of card #0 (steps S1, P1), in a **MAC table** had **by this card**, “A”, which is a source MAC address (SA) thereof, **is learned** by being **stored** in a predetermined buffer in association with information on the receiving card #0 and port #0 at this time (path information) (step S2). Note that in this situation, when a learning result for the same MAC address already exists in the **MAC table**, this information is overwritten.” *Id.* ¶ 3 (emphasis added).

140. It would have been obvious to a POSA to check the MAC source address against the records of the MAC table to determine whether to add a new record, or as Ishimori discloses, to overwrite the record if the MAC address already exists in the MAC table. *Id.* Thus, the combination of Smith and Ishimori renders obvious this limitation.

- d. **Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.**

141. In my opinion, to the extent Smith alone does not render obvious this limitation, in my opinion, it would have been obvious to combine Smith with Ishimori. Ishimori discloses that the source MAC address is stored with the reception path information including the card number and port number. *Id.* ¶ 3. Indeed, Ishimori's disclosures include:

- a. "FIG. 1 and FIG. 2 illustrate this. Note that in the example described below that is illustrated in FIG. 1 and the like, a configuration is supposed wherein each node—50-0, 50-1, 50-2—has one communication card—#0, #1, #2—respectively. When node 50-0 receives an IP packet from terminal A 10 at port #0 of card #0 (steps S1, P1), in a **MAC table** had **by this card**, "A",

which is a source MAC address (SA) thereof, is learned by being **stored in a predetermined buffer** in association with information on the receiving card #0 and port #0 at this time (path information) (step S2). Note that in this situation, when a learning result for the same MAC address already exists in the MAC table, this information is overwritten.” *Id.* ¶ 3 (emphasis added).

- b. “The path having this representative port leads to card #3 of the nodes 70. At this card #3 as well, no learning result relating to the terminal 20 that is the destination is had. As such, at card #3 as well, this packet is flooded (step P13). In this manner, this packet reaches the terminal 20 that is the destination through card #5. Note that during the packet forwarding from the terminal 10 to the terminal 20 illustrated in FIG. 9 and FIG. 10, at each communication card #0 to #5 of the node groups 60, 70, as described in conjunction with FIGS. 1 to 4, information relating to the path leading to the local device is learned in association with the source address of the packet.” *Id.* ¶ 16.
- c. “According to the present invention, each packet forwarding device is configured to **generate a learn packet at a predetermined timing under predetermined conditions**. That is,

even in a situation wherein different paths are selected for a coming direction and a going direction due to trunking or the like and thus only forwarding in one of these directions is performed in a certain device, by generating a learn packet at a predetermined timing, this learn packet can force, for example, the device performing forwarding in only one direction in this manner to also perform packet reception in the other direction. This causes each path to perform packet reception in both directions as appropriate, and path learning in both directions is performed reliably. As a result, performing flooding repeatedly and thus inviting increased line traffic can be prevented.” *Id.* ¶ 25 (emphasis added).

- d. “However, because an object of this learn packet is simply for effective path learning to be performed, generating this learn packet with needless frequency actually invites increased line traffic. To prevent such adverse effects, learn-packet generation conditions—that is, a generation frequency, a generation opportunity, and the like—must be appropriately determined.” *Id.* ¶ 26.
- e. “That is, it is desirable for the learn-packet generation conditions to be a situation wherein the path leading to the destination of the

receive packet includes a trunk, the learn packet being generated for the first time not at the time of the first learning but when the learning result 30 becomes a hit in a destination search. That is, a configuration is desirable that generates the learn packet when a packet is received whose destination is the learned source. Note that at this time, it is desirable to transmit the learn packet to all forwarding devices constituting this trunk. Moreover, it is desirable to also generate the learn packet each subsequent time when the learning result again becomes a hit in the destination search.” *Id.* ¶ 28.

- f. “To facilitate description, a situation is then supposed wherein the paths of cards #1, #3 are selected as the representative ports of the trunk in the direction wherein the terminal 20 receives 20 and, in contrast, the paths of cards #4, #2 are selected as the representative ports of the same trunk in the direction of transmitting to the terminal 10. In this situation, when no learn packet is generated, as described in conjunction with FIG. 12 to FIG. 14, flooding occurs repeatedly. That is, when a packet of a transmission direction toward the terminal 20 again reaches card #3, because card #3 has not received a packet of a reception direction from the terminal 20,

no corresponding learning result is had, and flooding is again performed. However, according to this embodiment of the present invention, as above, the learn packet is sent from card #5 to cards #3, #4. As such, at both card #3 and card #4, the MAC address “B” of the terminal 20 is already learned. As such, even when a packet of a transmission direction toward the terminal 20 again reaches card #3, card #3 can specify the path to the terminal 20 according to the learning result provided by the learn packet.” *Id.* ¶ 34.

142. It would have been obvious to a POSA to check the MAC source address against the records of the MAC table to determine whether to add a new record, or as Ishimori discloses, to overwrite the record if the MAC address already exists in the MAC table. *Id.* Ishimori teaches a learning process where “each packet forwarding device is configured to generate a learn packet at a predetermined timing under predetermined conditions.” *Id.* ¶ 25. Ishimori further teaches that the “learn packet” that informs the plurality of line cards regarding new associations “is transmitted to all nodes having the corresponding trunk.” *Id.* ¶ 33. Ishimori therefore discloses sending a message of the association to each member line card for said plurality of member line cards. It would have been obvious to implement the MAC notification in Smith to perform the method taught in Ishimori to first check whether the MAC address is found in the MAC table, and if not, create a new

record of the association. Both Smith and Ishimori teach that a message of this association is sent to the plurality of member line cards. Ex. 1004 ¶ 63; Ex. 1005 ¶¶ 25, 33. Thus, the combination of Smith and Ishimori renders obvious this limitation.

- e. **Claim 2: The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC address and the respective ports.**

143. In my opinion, the combination of Smith and Ishimori discloses claim 2. Smith teaches that the network device sub-unit 122(2) sends a MAC notification (a message) to update the forwarding engines, but it does not disclose doing so periodically at predefined times. Ex. 1004 ¶ 63. However, Ishimori discloses this. Indeed, Ishimori's disclosures include:


- a. "According to the present invention, each packet forwarding device is configured to generate a **learn packet** at a **predetermined timing** under predetermined conditions. That is, even in a situation wherein different paths are selected for a coming direction and a going direction due to trunking or the like and thus only forwarding in one of these directions is performed in a certain device, by generating a learn packet at a **predetermined timing**, this learn packet can force, for example, the device

performing forwarding in only one direction in this manner to also perform packet reception in the other direction. This causes each path to perform packet reception in both directions as appropriate, and path learning in both directions is performed reliably. As a result, performing flooding repeatedly and thus inviting increased line traffic can be prevented.” Ex. 1005 ¶ 25 (emphasis added).

- b. “In Figure 15, the path to the destination of the receive packet includes a trunk. Moreover, when, at these nodes, a packet from the same address as the corresponding destination address (DA) is already received and the source-address learning result from this time is had—that is, when there is a hit in the destination search (step S3 in FIG. 2)—a learn packet is transmitted **to all nodes having the corresponding trunk**. That is, in FIG. 15, a packet is first forwarded from the terminal to the terminal 20. When, at each node that this forwarding operation passes through—for example, card #5—a reply packet from the terminal 20 to the terminal 10 is afterward received, at card #5, the learn packet is transmitted to all nodes having the corresponding trunk—that is, cards #3, #4. This causes the path for the corresponding MAC address to be learned at all of these nodes. That is, in this situation, at cards #3, #4, the

path having card #5 is learned for the MAC address “B” of the terminal 20.” *Id.* ¶ 33 (emphasis added).

- c. “[Problem] An object is to provide, as a packet forwarding system that uses a path learning function applying a link aggregation function, a system that can prevent repeated flooding even when paths differ between a time of transmission and a time of reception. [Resolution Means] A configuration is adopted of periodically sending, according to predetermined conditions, a learn packet to all nodes having a trunk.” *Id.*, Abstract.
- d. Figure 15 shows “Periodically transmit learn packet”:



Id., Fig. 15.

144. Ishimori teaches that a “learn packet” (message) is generated at a “predetermined timing” (periodically at predefined times). *Id.* ¶ 25. Ishimori further teaches that the “learn packet” that informs the plurality of line cards regarding new associations “is transmitted to all nodes having the corresponding trunk.” *Id.* ¶ 33. Therefore, it would have been obvious that this must include at least the second line card. *Id.* Thus, the combination of Smith and Ishimori discloses this limitation.

f. Claim 3: The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.

145. In my opinion, the combination of Smith and Ishimori discloses claim 3. Smith discloses that the network device sub-unit 122(2) sends a MAC notification (*e.g.*, the message) to update the forwarding engines. Ex. 1004 ¶ 63. Smith specifies that “[a]fter being updated based on the MAC notification, the forwarding engines in virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* I understand that this would include the other line cards in virtual network device 202, such as line card 304(4). For example, line card 304(4) (the second line card) would receive the MAC notification (the message) and update its lookup tables. *Id.* ¶ 57. Thus, Smith discloses that in response to the message, the record of association is added to the FDB of the second line card if the record does not already exist in the FDB of the second line card.




FIG. 3

146. Ishimori's disclosures include:

- a. "As a packet forwarding method via a network, there is a method of **learning** a source MAC address of a receive IP packet together with a receive path thereof and, when an IP packet whose destination address is the same MAC address is received, using the corresponding learning result to decide a transmission path thereof. In this situation, when forwarding, at a node that relays IP packets, an IP packet having an unlearned MAC address as a destination address, a transmission path thereof cannot be specified. As such, in this situation, a broadcast transmission to all nodes (that is,

“flooding”) is performed. Then, at this time, path **learning** is carried out by **storing** the source MAC address (SA) had by this IP packet and a port that received such **in a buffer had by each node.**” Ex. 1005 ¶ 2 (emphasis added).

- b. “The path having this representative port leads to card #3 of the nodes 70. At this card #3 as well, no learning result relating to the terminal 20 that is the destination is had. As such, at card #3 as well, this packet is flooded (step P13). In this manner, this packet reaches the terminal 20 that is the destination through card #5. Note that during the packet forwarding from the terminal 10 to the terminal 20 illustrated in FIG. 9 and FIG. 10, at each communication card #0 to #5 of the node groups 60, 70, as described in conjunction with FIGS. 1 to 4, information relating to the path leading to the local device is learned in association with the source address of the packet.” *Id.* ¶ 16 (emphasis added).

147. Thus, Ishimori further teaches that the reception path information is learned, which is done by storing the record of association of the MAC address and ingress port to the MAC table. *Id.* ¶ 2. It would have been obvious to do this via the MAC notification as disclosed in Smith. Additionally, during the packet transfer, each of Ishimori’s communication cards learns information about the association.

Ishimori discloses that “a learn packet is transmitted to all nodes having the corresponding trunk,” which includes at least a second line card. *Id.* ¶ 33. Thus, in Ishimori, each of the communication cards, which includes the second line card, would add a record of the association to the FDB in response to a MAC notification as taught in Smith. The combination of Smith and Ishimori therefore renders obvious this claim.

- g. Claim 11[c]: a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards;**

148. In my opinion, the combination of Smith and Ishimori discloses claim 11[c]. Claim 11[c] combines the elements of claims 1[a], 1[b], 1[c], and 11[a]. Therefore, for the same reasons that the combination of Smith and Ishimori discloses claims 1[a], 1[b], 1[c], and 11[a], Smith also discloses claim 11[c]. *See supra* ¶¶ 99–107, 134–136, 127.

- h. Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress**

line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.

149. In my opinion, the combination of Smith and Ishimori discloses claim 11[d]. The combination of Smith and Ishimori disclose this limitation for the same reasons that the combination of Smith and Ishimori discloses claims 1[d], 1[e], 1[f], 1[g], and 1[h]. *See supra* ¶¶ 108–120, 137–142.

- i. Claim 12: The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.**

150. In my opinion, the combination of Smith and Ishimori discloses claim 12. *See supra* ¶¶ 143–144.

- j. Claim 13: The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.**

151. In my opinion, the combination of Smith and Ishimori discloses claim 13. *See supra* ¶¶ 145–147.

3. Ground 3: Claims 1, 4–7, 10–11, 14–17, and 20 would have been obvious over Smith in view of Ishimori in further view of Edsall

152. I have analyzed claims 1, 4–7, 10–11, 14–17, and 20 and conclude that they would have been obvious over Smith in view of Ishimori in further view of Edsall. I provide a detailed analysis of each claim limitation below.

a. Claim 1[g]: checking said MAC source address of the data packet against records in said FDB of said first line card; and

153. In my opinion, to the extent this is not disclosed by the combination of Smith and Ishimori, it is disclosed in further view of Edsall. Edsall teaches that the forwarding engine learns the source MAC address of a frame received at the ingress card for the first time. “If the frame is received at the ingress card for the first time, the forwarding engine also ‘learns’ a source MAC address of the frame.” Ex. 1006 at 18:39–41. It does so by “creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.* at 18:42–44. “[I]f there is not a current entry,” the forwarding engine “learn[s] the source address/index of the frame.” *Id.* at 18:54–55.

154. It would have been obvious to use Edsall’s learning methods with the Smith-Ishimori MAC tables, and there would have been a reasonable expectation of success in doing so.

- b. Claim 1[h]: if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card and said plurality of member line cards.**

155. In my opinion, to the extent the combination of Smith and Ishimori does not render obvious this limitation, it would have been obvious to combine Smith and Ishimori with Edsall. Edsall teaches that the forwarding engine learns the source MAC address of a frame received at the ingress card for the first time. *Id.* at 18:39–41. Edsall further teaches that “if there is not a current entry,” the forwarding engine “learn[s] the source address/index of the frame.” *Id.* at 18:54–55. It does so by “creating/updating an entry of the L2 forwarding table with the source MAC address and its location (index) within the switch.” *Id.* at 18:42–44. It then floods copies of the fabric frame to all the egress line cards of the network switch, which Edsall calls the “flood-to-fabric (FF) operation.” *Id.* at 18:47–50. This “forces each forwarding engine associated with each egress card to either (i) update its current L2 forwarding table entry with the newly-learned source MAC address and index of the frame or, if there is not a current entry, (ii) learn the source address/index of the frame.” *Id.* at 18:50–55.

156. It would have been obvious to add Edsall’s learning method and flood-to-fabric operation in the Smith-Ishimori path learning operations. A POSA would

have been able to implement this teaching with a reasonable expectation of success, because the combination of Smith and Ishimori already teach that a message of new associations are sent to the plurality of member line cards. Ex. 1004 ¶ 63; Ex. 1005 ¶ 16. It would have been within the knowledge of a POSA to create new entries in a FDB when an association of a MAC address and port does not yet exist. Thus, the combination of Smith, Ishimori, and Edsall renders obvious this limitation.

- c. Claim 4: The method according to claim 3 and comprising marking the records in the respective FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.**

157. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 4. Edsall discloses that the “PI indicator” is asserted when a forwarding table entry for the MAC address is learned through one of the ports (*e.g.*, data packets transmitted via a port of the line card) “as opposed to through the switch fabric” (*e.g.*, received from another of the line cards). Ex. 1006 at 6:46–64. Thus, the PI indicator is different for the MAC address learned through one of the ports, which correspond to the first type of records in the claim (“data packets transmitted via a port of the line card”) “as opposed to through the switch fabric, which corresponds to the second type of records (“messages from another of the line cards”). Ex. 1006 at 6:46–64.

158. Additionally, during prosecution, the Examiner found that Edsall disclosed this limitation and the applicant did not amend the claims based on this rejection. Ex. 1002 at 82, 120–122. The Examiner determined that Edsall also discloses that the “PI indicator is asserted for a destination MAC address entry of the forwarding table on the egress card and the DI contained in the switched fabric frame (*i.e.*, the ingress DI) is different from the DI stored in this egress forwarding table (*i.e.*, the egress DI).” Ex. 1006 at 18:56–19:5; *see also* Ex. 1002 at 82. Based on my review of the file history, it is my understanding that this disclosure from Edsall was cited by the examiner during prosecution and the applicant did not amend the claims based on this rejection. Ex. 1002 at 120–122.

d. Claim 5[a]: The method according to claim 4, and comprising: associating a respective aging time with each of the records;

159. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 5[a]. Ishimori’s disclosures include:

- a. “Furthermore, in this system, based on an approach of effectively utilizing a limited buffer storage capacity, upon learning, when there is no DA search hit for the same address in a certain period, the learning result relating to this MAC address is deleted. This operation leading to deletion is referred to as “**aging**.” Specifically, as illustrated in FIG. 5, the buffers of the cards of each node have a

“hit bit” corresponding to the learned MAC address; when this is learned, the hit bit = “1”. Then, at the **first aging process, this hit bit is cleared** (that is, made to be “0”; steps S11, S12). Here, the learning result is not yet deleted. Then, **when the hit bit is still “0” at the next aging process** (step S14, step S15, “No” at step S11), the **corresponding MAC address is deleted from the buffer** (step S13). Meanwhile, when overwriting learning is performed previous to this or when there is a hit at the DA search (that is, the destination search operation of step S3 in FIG. 2 above), the **hit bit is again set to “1”** (as a result, step S11 “Yes” and step S12). This configuration prevents an address currently being used (that is, being used for communication) from being deleted from the buffer.” Ex. 1005 ¶ 9 (emphasis added).

- b. “This aging process is further described in conjunction with FIGS. 6 to 8. Steps P1 to P6 in FIG. 7 are the same steps as steps P1 to P6 in FIG. 4 above. Now, each time an IP packet is forwarded, at steps P1, P2, the MAC address “A” is learned at card #0, and the hit bit thereof is made to be “1”. Likewise, at steps P2, P3, the MAC address “A” is learned at card #1 and card #2. Continuing in the

same vein, at steps P4, P5 and steps P5, P6, the MAC address “B” is learned at card #2 and card #0.” *Id.* ¶ 10 (emphasis added).

- c. “Afterward, as illustrated in the uppermost row in FIG. 8, the first aging process clears the respective hit bits of the MAC addresses “A”, “B” at each card #0, #1, #2 (steps S12). It is then supposed that before the second aging process, the next IP packet, likewise addressed to terminal B 20, is sent from terminal A 10 (step P7). At card #0 that receives this, the source MAC address “A” is learned by overwriting. At the same time, in the destination search, the destination MAC address “B” becomes the search target. As such, the MAC addresses “A”, “B”, which are temporarily cleared to “0” as above, have their hit bits respectively returned to “1”.” *Id.* ¶ 11 (emphasis added).
- d. “Then, this IP packet is forwarded to card #2 according to the learning result regarding the MAC address “B” that yielded a hit in the destination search (step P8). Then, in card #2 as well, the MAC addresses “A”, “B” temporarily cleared to “0” as above have their hit bits respectively returned to “1” by the same operation as above. It is then supposed that afterward, the next aging-process timing arrives without either of the IP packets of terminals A, B 10,

20 being received. Here, as illustrated in the lowermost row in FIG. 8, at card #0 and card #2, as above, the hit bits of the MAC addresses “A”, “B” are each returned to “1” at steps P7, P8. As such, these are again cleared to “0”. However, the corresponding learning result is not yet erased at this point (step S12).

Meanwhile, for card #1, no corresponding IP packet passes through during this time, and the hit bits of the MAC addresses “A”, “B” remain “0”. As such, the corresponding learning result is deleted at this point (step S13).” *Id.* ¶ 12 (emphasis added).

160. Ishimori teaches that “the buffers of the cards of each node have a ‘hit bit’ corresponding to the learned MAC address; when this is learned, the hit bit = ‘1’.” *Id.* ¶ 9. Therefore, Ishimori associates a respective aging time with each of the records. “[A]t the first aging process, this hit bit is cleared.” *Id.* “[W]hen the hit bit is still ‘0’ at the next aging process,” “the corresponding MAC address is deleted from the buffer.” *Id.* Thus, the hit bit is how Ishimori associates a respective aging time with each of the records.

161. It would have been obvious to implement a “hit bit” on the lookup tables in the line cards of Smith and Edsall (where Edsall leverages the PI Indicator) in order to associate a respective aging time with each of the records.

e. Claim 5[b]: refreshing the records in the FDB responsively to further packets transmitted by the line cards; and

162. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 5[b]. As described above, Ishimori teaches that when the hit bit is “0,” “the corresponding MAC address is deleted from the buffer.” *Id.* ¶¶ 9, 11. Ishimori then teaches that when the destination MAC address is received at card #0, the source MAC address is learned by “overwriting,” and the hit bit is “returned to ‘1’” (*e.g.*, refreshing the records in the FDB responsively to further packets transmitted by the line cards). *Id.* ¶ 11.

163. It would have been obvious to a POSA implement a “hit bit” on the lookup tables in the line cards of Smith and Edsall (where Edsall leverages the PI Indicator) and to refresh the hit bit by learning the source MAC address of the data packets transmitted by the line cards.

f. Claim 5[c]: removing the records from the respective FDB if the records are not refreshed within the respective aging time.

164. In my opinion, the combination of Smith and Ishimori discloses claim 5[c]. Ishimori teaches that if the hit bit remains “0” in the next aging process, then the corresponding MAC address is deleted. *Id.* ¶¶ 9, 12. Thus, Ishimori discloses removing the records from the FDB if the records are not refreshed within the respective aging time.

165. It would have been obvious to a POSA implement a “hit bit” on the lookup tables in the line cards of Smith and Edsall (where Edsall leverages the PI Indicator) to remove the records from the respective lookup tables in Smith-Edsall if the record has not been refreshed by the next aging process as taught by Ishimori.

- g. Claim 6: The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via switching core of the network node to at least the second line card.**

166. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 6. As I explained above, Smith discloses this limitation. *See supra* ¶¶ 123–125. Edsall further discloses the “plurality of line cards” are “interconnected by a switch fabric 550.” Ex. 1006, 8:20–27. The “switch fabric” in Edsall could be a “switching core.” Thus, a POSA would have understood that the packet is sent from the first line card via a switching core to at least the second line card.

- h. Claim 7: The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed.**

167. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 7. Smith teaches that its virtual network device includes several line cards, which include several interfaces. Ex. 1004 ¶ 46. “[W]hen updating control protocol

behavior of virtual link bundle 250(1), a user can simply access virtual network device sub-unit 122(1) (instead of accessing both virtual network device sub-units 122(1) and 122(2)). *Id.* Virtual network device sub-unit 122(1) can then automatically propagate to network device 122(2) any changes made by the user to the control protocols.” *Id.* ¶ 59. Smith teaches that “MAC notification frames are used to keep the content of the L2 tables in virtual network device sub-unit 122(1) synchronized with the content of the L2 tables in virtual network device sub-unit 122(2) and vice versa.” *Id.* ¶ 62.




FIG. 3

168. Thus, if the record in the FDB of the first line card 304(1) associates the MAC source address (the MAC address of virtual device 120(1)) with a port

different from the one of the ports on which the data packet was received (a port other than 320(9)), the MAC notification frames will notify and update the L2 tables in the virtual network device 202, which would include at least the second line card 304(4) to indicate the record has been changed. Similarly, Ishimori teaches that the routes are transmitted via a “learn packet.” Ex. 1005 ¶ 25. A POSA would understand that the “learn packet” as taught in Ishimori functions as a synchronization packet. During the packet transfer, each of the communication cards learns information about the route (*e.g.*, the second line card). *Id.* ¶¶ 28, 33.

169. Edsall further discloses that the forwarding engine generates an MN frame (*e.g.*, synchronization packet) that may get sent to the SMC (switch management card) to ensure that FwdT0 (*e.g.*, the forwarding table) is synchronized. Ex. 1006 at 17:26–38. The forwarding engine also asserts an appropriate bit of the POE field (port-of-exit field) when generating the MN frame, which is a port different from one of the ports on which data was received. *Id.* It would have been obvious to a POSA to implement the MN frame from Smith to include the POE field in order to note the port interface of the switch fabric. *Id.* at 6:24–25. This disclosure was cited by the examiner during prosecution and the applicant did not amend the claims based on this rejection. Ex. 1002 at 83, 120–122. Thus, the combination of Smith, Ishimori, and Edsall discloses this limitation.

- i. **Claim 10[a]: The method according to claim 1, and comprising: conveying a further data packet, received**

from a further MAC source address, to the second line card for transmission over the network;

170. In my opinion, Smith discloses claim 10[a]. *See also supra* ¶¶ 108–111. Smith teaches that the uplink interface receives a data packet with the “sending device’s MAC address,” which is a MAC source address. Ex. 1004 ¶ 54. Based on Figure 3, a further data packet received from a further MAC source address (such as network device 120(3)) would be conveyed to the second line card 304(4) for transmission over the network.




FIG. 3

171. To the extent Smith does not render obvious this claim, Edsall does. Edsall teaches that in each subsequent frame, the encoded address recognition logic (EARL) circuit looks up the MAC address and “sends the corresponding rewrite information over the local bus after the frame” (*e.g.*, conveying a further data packet to the second line card for transmission over the network). Ex. 1006 at 14:22–34. This was cited by the examiner during prosecution and the applicant did not amend the claims based on this rejection. Ex. 1002 at 83–84, 120–122.

j. Claim 10[b]: checking the further MAC source address against the records in the FDB of the second line card; and

172. In my opinion, Smith discloses claim 10[b]. *See also supra* ¶¶ 117–118. Smith teaches that the virtual network device sub-unit learns the source identifier of the sending device, which is the MAC source address. Ex. 1004 ¶ 65. These identifiers are stored in a lookup table (*e.g.*, FDB) on virtual network device sub-unit. *Id.* ¶ 61. The MAC source address, which is the MAC address of network device 120(3), would be checked in the records of the FDB of said second line card 304(4).

173. To the extent Smith does not disclose this limitation, Edsall does. Edsall teaches checking to see if the rewrite information matches, in other words, checking the MAC source address against the records in the FDB of the second line card). Ex. 1006 at 14:22–34 (“The destination port circuitry (or, alternatively, a

UDlink or central rewrite engine) matches the frame with the rewrite information and modifies the frame as needed by replacing, *inter alia*, the destination and source MAC addresses.”). This was cited by the examiner during prosecution and the applicant did not amend the claims based on this rejection. Ex. 1002 at 83–84, 120–122.

- k. Claim 10[c] responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record.**

174. In my opinion, the combination of Smith, Ishimori, and Edsall renders obvious claim 10[c]. *See also supra* ¶¶ 117–118. Smith teaches that that the network device sub-unit 122(2) sends a MAC notification (*e.g.*, a message) to update the forwarding engines (*e.g.*, sending a further message to inform at least the first line card) when it learns of a new association (*e.g.*, the further record). Ex. 1004 ¶ 63. “After being updated based on the MAC notification, the forwarding engines in the virtual network device sub-unit 122(1) now know the location of the device identified by the destination address.” *Id.* A POSA would understand that the MAC notification (*e.g.*, message) is therefore sent to inform at least the first member line card (*e.g.*, line card 304(1)) of the new association.

175. Edsall further teaches that the forwarding engine “modifies the frame as needed by replacing . . . the destination and source MAC addresses” (*e.g.*, adding

a further record with respect to the MAC source address to the FDB). Ex. 1006 at 14:22–34. This was cited by the examiner during prosecution and the applicant did not amend the claims based on this rejection. Ex. 1002 at 83–84, 120–122. It would have been obvious to use any of the learning methods from Smith, Ishimori, or Edsall, which all disclose updating a second line card with a further record and sending a message to inform at least the first line card of the further record.

l. Claim 11[a]: a switching core;

176. To the extent this is not disclosed by the combination of Smith and Ishimori, Edsall discloses this. The '400 patent describes the “switching core” as linking the multiple line cards. Ex. 1001 at 6:8–10. Edsall teaches that the “plurality of line cards . . . are interconnected by a switch fabric 550.” Ex. 1006 at 8:20–27. Thus, the “switching fabric” in Edsall could be a “switching core.”

- m. Claim 11[d]: wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address**

does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.

177. In my opinion, the combination of Smith, Ishimori, and Edsall renders obvious claim 11[d]. This limitation is a combination of claims 1[d], 1[e], 1[f], 1[g], and 1[h]. Therefore, the combination of Smith, Ishimori, Edsall renders obvious this limitation for the same reasons that the combination of Smith, Ishimori, and Edsall renders obvious claims 1[d], 1[e], 1[f], 1[g], and 1[h]. *See supra* ¶¶ 108–120, 137–142, 153–156.

- n. **Claim 14: The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.**

178. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 14. *See supra* ¶¶ 157–158.

- o. **Claim 15: The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time.**

179. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 15. *See supra* ¶¶ 159–165.

- p. Claim 16: The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.**

180. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 16. *See supra* ¶ 166.

- q. Claim 17: The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet, which indicates to at least the second line card to indicate that the record has been changed.**

181. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 17. *See supra* ¶¶ 167–169.

- r. Claim 20: The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record.**

182. In my opinion, the combination of Smith, Ishimori, and Edsall discloses claim 20. *See supra* ¶¶ 170–175.

4. Ground 4: Claims 8–9 and 18–19 would have been obvious over Smith in view of Ishimori in further view of Zelig

183. I have analyzed claims 8–9 and 18–19 and conclude that they would have been obvious over Smith in view of Ishimori in further view of Zelig. I provide a detailed analysis of each claim limitation below.

- a. Claim 8: The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.**

184. In my opinion, the combination of Smith, Ishimori, and Zelig discloses claim 8. Zelig’s disclosures include:

- a. “In another preferred embodiment, at least one of the backup virtual bridges is connected by a single secondary virtual connection to a selected one of the primary virtual bridges, such that upon the failure of the at least one of the primary virtual bridges with which the at least one of the backup virtual bridges is associated, the at least one of the backup virtual bridges transmits and receives the data packets over the network via the selected one of the primary virtual bridges over the single secondary virtual connection. Preferably, **each of the primary virtual bridges** is

adapted to **maintain a respective media access control (MAC) table**, and to **forward the data packets in accordance with entries in the MAC table**, and wherein each of the primary virtual bridges is adapted, upon detecting the failure of the at least one of the primary virtual bridges, to update the entries in the respective MAC table that point to the at least one of the primary virtual bridges so as to point instead to the selected one of the primary virtual bridges.” Ex. 1007 ¶ 31 (emphasis added).

- b. “FIG. 1 is a block diagram that schematically illustrates a **VPN 20** with a hierarchical **VPLS** topology, implementing a protection Scheme in accordance with a preferred embodiment of the present invention. **VPN is built around a virtual private LAN service (VPLS)**, operating within a network 22, typically an IP or MPLS network. **The VPLS is based on virtual bridges 30, 32,34, 36,38, and 40**,or VPLS-capable PEs, which are connected by PWs 70, 72, 74 and 76 through network 22. Although for clarity of illustration, network 22 includes only a small number of PES and represents only a single VPLS instance, the principles embodied in this network may be extended in a straightforward manner to larger

networks and to multiple VPLS instances.” *Id.* ¶ 42 (emphasis added).

- c. “Three primary virtual bridges 30, 32 and 34, referred to as primary core nodes, are connected with each other in a full mesh with PW connections 70. Typically, the PW connections comprise MPLS tunnels, but they may alternatively comprise virtual connections of other types, such as GRE or L2TP tunnels. Each of the primary core nodes 30, 32 and 34, is paired with a corresponding backup virtual bridge, referred to as a standby core node, 36, 38 and 40, respectively. The standby core nodes are connected in the network by redundant backup connections 72. Each standby core node has a topology identical to its corresponding primary core node. For example, standby core node 36 has the same topology image as primary core node 30. Each of the primary and standby core nodes is connected to all the other core nodes in the network except for the standby or primary core node with which it is paired. An optional connection 79 between a primary core node and its corresponding standby core node may also be included, as described hereinbelow.” *Id.* ¶ 43.

185. Zelig discloses that its VPN 20 contains multiple primary virtual bridges 30, 32, and 34. *Id.* ¶¶ 42–43. It would have been obvious to a POSA that the virtual bridges servicing VPN 20 are configured to serve that VPN. Zelig further teaches that “each of the primary virtual bridges is adapted to maintain a respective media access control (MAC) table, and to forward the data packets in accordance with entries in the MAC table” *Id.* ¶ 31. Thus, Zelig discloses that each MAC bridge maintains its own MAC table, and because the MAC bridge serves only that VPN instance, the records associating the MAC addresses with the respective ports are maintained independently for each VPN instance.

b. Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.

186. In my opinion, the combination of Smith, Ishimori, and Zelig discloses claim 9. Smith teaches that the MAC notification is “distributed to any other forwarding engines within virtual network device sub-unit 122(2).” Ex. 1004 ¶ 63. Zelig provides that “VPN 20 is built around a virtual private LAN service (VPLS), operating within a network 22.” Ex. 1007 ¶ 42.

187. It would have been obvious to a POSA to adopt the notification scheme in Smith to the architecture in Zelig so that the MAC notification, which is a message, is sent to all the line cards that serve the VPLS instances. Furthermore, the

IEEE 802.1Q standards, which is cited by Zelig, Ex. 1007 ¶ 2, disclose a VLAN identifier (VID) as a twelve-bit field that “uniquely identif[ies] the VLAN to which the frame belongs.” Ex. 1008 § 9.3.2.3; *see also* Ex. 1016 § 9.6. Thus, it would have been within the knowledge of a POSA to identify the VPLS instance in the message, such as using the VID, in order to inform all the line cards in the VPLS.

- c. **Claim 18: The node according to claim 11, wherein at least some of the line cards are configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.**

188. In my opinion, the combination of Smith, Ishimori, and Zelig discloses claim 18. *See supra* ¶¶ 184–185.

- d. **Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.**

189. In my opinion, the combination of Smith, Ishimori, and Zelig discloses claim 19. *See supra* ¶¶ 186–187.

5. Ground 5: Claims 9 and 19 would have been obvious over Smith in view of Ishimori, Zelig, and 802.1Q-1998

190. I have analyzed claims 9 and 19 and conclude that they would have been obvious over the combination of Smith, Ishimori, Zelig, and 802.1Q-1998. I provide a detailed analysis of each claim limitation below.

- a. Claim 9: The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.**

191. In my opinion, the combination of Smith, Ishimori, Zelig, and 802.1Q-1998 discloses claim 9. *See supra* ¶¶ 186–187.

- b. Claim 19: The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.**

192. In my opinion, the combination of Smith, Ishimori, Zelig, and 802.1Q-1998 discloses claim 19. *See supra* ¶ 189.

IX. ADDITIONAL REMARKS

193. I currently hold the opinions expressed in this Declaration. I reserve the right to further explain and supplement my opinions as I may acquire additional information and/or attain supplemental insights that may result in added observations.

194. I hereby declare that to the best of my knowledge all statements made are true and that all statements made on information and belief are believed to be true. I further declare that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of this proceeding.

Executed this 23rd day of December, 2022.

A handwritten signature in black ink, appearing to read "N. Bambos", written over a horizontal line.

Dr. Nicholas Bambos

Appendix A

CURRICULUM VITAE

Nicholas Bambos

R. Weiland Professor of Electrical Engineering
and of Management Science & Engineering

Stanford University

December 2022

Nick Bambos is a R. Weiland Professor in the School of Engineering at Stanford University, having a joint appointment in the Department of Electrical Engineering and the Department of Management Science & Engineering. He has also been the Fortinet Founders Department Chairman of the Management Science & Engineering Department (2016-20). He heads the Computer Systems & Networks Performance Engineering Lab at Stanford. This is comprised of PhD students and industry visitors engaged in research projects, which require solid understanding of the physical principles of computing and networking hardware and software architectures (e.g., digital design, file systems), but largely focus on system architectures and high-performance engineering of computer systems and networks. His research contributions span the areas of wireless & wireline networking, the Internet, cloud computing and data centers, high-speed networking and packet switching, multimedia streaming, information service engineering, computer systems security, etc. His methodological interests and contributions span the areas of stochastic modeling and control, online task scheduling, routing and distributed processing, network control, etc.

He has graduated over 40 Ph.D. students (including two postdocs), who have moved on to key positions in academia (Stanford, CalTech, Michigan, GaTech, NYU, USC, Columbia, etc.), the information technology industry (Cisco, IBM, Broadcom, Qualcomm, Nokia, MITRE, ST Micro, Intel, Samsung, Google, Facebook, Twitter, etc.), technology startups and the finance tech industry. From 1999 to 2005 he served as Director of the Stanford Networking Research Center project, a partnership between Stanford and information technology industries, involving tens of corporate members, faculty and doctoral students.

He received his Ph.D. in Electrical Engineering and Computer Sciences (EECS) from the University of California at Berkeley (1989), as well as the M.S. in EECS (1987) and the M.A. in Mathematics (1989) from the same University. He graduated in Electrical Engineering from the National Technical University of Athens-Greece (1984) with first class honors. Before joining Stanford as an Associate Professor in 1996, he served as Assistant (1990-95) and tenured Associate Professor (1995-96) in the Electrical Engineering Department of the University of California at Los Angeles (UCLA).

Nick Bambos has received several best research papers awards, has held the Cisco Systems Faculty Development Chair (1999-2003) in computer networking at Stanford, and has won the IBM Faculty Award (2002) for high-impact research in performance engineering of computer systems and networks, as well as the Griffin Award (1997). He has been the David Morgenthaler Faculty Scholar (1996-99) at Stanford, and has received the National Young Investigator Award (1992) from the National Science Foundation (NSF) for research in computer networks and distributed computing architectures, as well as the NSF Research Initiation Award (1990) for studies in performance modeling of computer systems. He has also been a U.C. Regents Fellow, a David Gale Fellow, and an Earl Anthony Fellow at U.C. Berkeley.

He has served as editor of various research journals, on international technical and scientific committees, and on review panels for networking and computing technologies. He has also served on the technical boards of start-up companies, as a consultant for high-technology development and management, and as an expert witness in high-profile patent litigation cases involving networking, computing and information technologies.

EDUCATION

- **Ph.D. in Electrical Engineering and Computer Sciences**, University of California at Berkeley, 1989.
- **M.A. in Mathematics**, University of California at Berkeley, 1989.
- **M.S. in Electrical Engineering and Computer Sciences**, University of California at Berkeley, 1987.
- **Diploma in Electrical Engineering - First Class Honors**, National Technical University of Athens-Greece, 1984.

PROFESSIONAL EXPERIENCE

- **R. Weiland Professor** in the School of Engineering, Stanford University, 2016 – present.
- **Fortinet Founders Department Chairman**, Department of Management Science & Engineering, Stanford University, 2016 – 2020.
- **Professor**, Department of Management Science & Engineering, and Department of Electrical Engineering, Stanford University, 2003 – present.
- **Head**, Computer Network Architecture and Performance Engineering Research Lab (NetLab), Stanford University, 1998 – present.
- **Director**, Stanford Networking Research Center (SNRC) Project Initiative, Stanford University, 1999 – 2005.
- **Associate Professor**, Department of Management Science & Engineering, and Department of Electrical Engineering, Stanford University, 1999 – 2003.
- **Associate Professor**, Department of Engineering-Economic Systems & Operations Research, and (courtesy appointment at) Department of Electrical Engineering, Stanford University, 1996 – 1999.
- **Associate Professor** (tenured), Department of Electrical Engineering, University of California at Los Angeles, 1995 – 1996.
- **Assistant Professor**, Department of Electrical Engineering, University of California at Los Angeles, 1989 – 1995.

HONORS AND AWARDS

- **2022 Best Paper Award for Telecommunications and Network Analytics** from INFORMS on for the paper “Robust Power Management via Learning and Game Design” (Z. Zhou, P. Mertikopoulos, A. Moustakas, N. Bambos, P. Glynn), INFORMS Annual Meeting, IN, October 2022.
- **Outstanding Paper Award at Healthcom** for the paper “Physiological Waveform Imputation of Missing Data using Convolutional Autoencoders” (with D. Miller, A. Ward, D. Scheinker, Andrew Shin). *Proceedings of the 2018 IEEE 20th International Conference on e-Health Networking, Applications and Services (Healthcom’18), Ostrava, Czech Republic, September 2018.*
- **Best Paper Award in Green Communication and Computing** by the TAOS (Transmission, Access & Optical Systems) Technical Committee for the paper: *Dynamic Resource Management in Virtualized Data Centers with Bursty Traffic* (with M. Valdez-Vivas, J. Apostolopoulos) presented in the IEEE International Communications Conference (ICC’14) in Sydney, Australia, June 2014.
- **“Best of Globecom” Paper Award** (top 2%) for the paper: *Dynamic Resource Management in Cloud Computing with Frictions and Congestion “Weather”* (with M. Valdez-Vivas, J. Apostolopoulos) presented in IEEE Global Communications Conference (Globecom’14) in Austin, TX, December 2014.
- **2013-14 Eugene L. Grant Teaching Award**, Stanford University.
- **Best Journal Paper Award for 2011** from the IEEE Multimedia Communications Technical Committee for the paper: *Channel, Deadline, and Distortion Aware Scheduling of Video Streams over Wireless Links* (with A. Dua, C. Chan, J. Apostolopoulos), IEEE Trans. on Wireless Communications, 9(3):1001-1011, 2010.
- **HP Faculty Award**, Hewlett-Packard Laboratories, Palo Alto, 2006
- **Cisco Systems Faculty Scholar**, School of Engineering, Stanford University, 1999 - 2003.
- **IBM Faculty Award**, IBM Corporation Research Labs, 2002.
- **Dana Adams Griffin Award**, School of Engineering, Stanford University, 1997.
- **David Morgenthaler Faculty Scholar**, School of Engineering, Stanford University, 1996 - 1999.
- **National Young Investigator Award**, National Science Foundation, 1992 - 1997.
- **Research Initiation Award**, National Science Foundation, 1990 - 1992.
- **EECS Departmental Scholar**, University of California at Berkeley, 1988 - 1989.
- **U.C. Regents Fellow**, University of California at Berkeley, 1986 - 1988.
- **Earl C. Anthony Fellow**, University of California at Berkeley, 1985 - 1986.
- **David and Sylvia Gale Fellow**, University of California at Berkeley, 1984 - 1985.
- **First Class Honors** for top graduating students in 1984 (among 150 students), Electrical Engineering Department, National Technical University of Athens-Greece, 1984.

- **Academic Excellence Awards** (annual), Hellenic State Scholarships Foundation, Greece, 1979 - 1984.
- **Academic Prizes** (annual), Hellenic Engineering Society, Greece, 1979 - 1984.
- **First Class Honors** for top students admitted to the University in 1979 (among tens of thousands of students participating in a national-level entrance exam), National Technical University of Athens-Greece, 1979.

DISTINGUISHED INVITED LECTURES AND TALKS (selected; additional talks on pg. 19)

- *Keynote Talk – Geometry of Resource Scheduling in Service Systems*. The 15th International Conference on Service Systems and Service Management (ICSSSM), Hangzhou, China, July 21, 2018.
- *Plenary Talk -- Transmitter Power Control in Wireless Networking*. The 2018 International Conference on Computing, Networking and Communications (ICNC), Maui, Hawaii, March 6, 2018.
- *Keynote Address – Power Control and Bandwidth Management in Wireless Networking*, IEEE Wireless On-Demand Network Systems and Services Conference (WONS), Jackson Hole, WY, February 2017.
- *Plenary Address – Digital Living 2030*, Shandong Chinese Academy of Sciences, Jinan, China, January 2017.
- *Keynote Address – A Risk Management View to Information Security*, IEEE Conference on Decision and Game Theory for Security, Berlin, Germany, Nov. 2010.
- *Keynote Address – Wireless Computing: From Infrastructure to Services*, IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), Athens, September 2007.
- *Invited Distinguished Lecture - Three Control Problems in Packet Switching*, Workshop on Mathematics of Networks, Ecole Normale Supérieure, Paris, France, June 2007.
- *Invited Distinguished Talk - Throughput and Power Issues in Packet Switches*, Center for Ultra-Broadband Information Networks, University of Melbourne, Australia, August 2007.
- *Keynote Talk - Opportunities and Challenges in Wireless Computing and Multimedia Networking*, Network Ventures Conference, Redwood Shores, CA, March 2005.
- *IBM Distinguished Talk - Resource Management and Load Balancing in Flexible Computing/Networking Service Platforms*, IBM Research Labs, Almaden, CA, June 2005
- *Distinguished Speaker - Power Control in Wireless Networks*, Wireless Industry Forum, Electrical Engineering Department, University of Texas at Austin, April 2004.
- *Plenary Talk - Power Control in Wireless Networks* and member of Panel of Experts on Network Technology Trends, IEEE Workshop on Local and Metropolitan Area Networks (LANMAN), pp. 69-74, CA, April 2004
- *Invited Lecture - Geometry of Packet Scheduling in Communication Switches*, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Nov. 2004.
- *Plenary Talk – Next Generation Wireless System Architectures* and member of Panel of Experts on Wireless Technologies in the IEEE International Conference in Broadband Networks (BroadNets), San

Jose, CA, October 2004.

- *Plenary Talk – Next Generation Wireless Data Networks: Control Plane Principles and Architectures.* Strategic Architecture Council of Sun Microsystems, Nov. 2003
- *International Expert invited* to serve on the International Evaluation Committee for network engineering research at INRIA (the French/European National Institute of Informatics and Automation of 900 scientists working in six campuses), 2003-2004.

UNIVERSITY AND PROFESSIONAL SERVICE

- *Academic Senate* of Stanford University (Member), 2005 – 2010
- *Committee on Research*, Stanford University (Member), 2009 – 2012.
- *Member of the Steering Board*, Conference on Decision and Game Theory for Security (GameSec), 2010-present.
- *Member of the Scientific Board*, LINCS – Laboratory on Information, Networking and Communication Sciences, Paris, France (founded by INRIA, Institut Telecom, France Telecom and Alcatel-Lucent), 2010-present.
- *Member the Editorial Boards* of the Wireless Networks Journal (2001-07), Computer Networks Journal (2000-02), Queueing Systems Theory and Applications (1999-2007).
- *Member of the Technical Program Committees* of the various networking/computing conferences (IEEE Infocom 2001-03, IEEE Globecom 2001-03, IEEE Mobicom 2002, IEEE VTC 2003, IEEE MobiWac 2005, SpasWin 2006, BroadNets 2006-07, Wireless Broadband World Forum 2006-07, GameComm 2008)
- *Technical Reviewer* for several research journals in computing and networking, as well as in queueing systems and applied probability (IEEE/ACM Transactions on Networking, IEEE Tran. on Vehicular Technology, IEEE Tran. on Automatic Control, Advances in Applied Probability, Journal of Applied Probability, Annals of Applied Probability, Operations Research, Journal of the ACM, Math of OR, IEEE Trans. on Communications)
- *Member of various Technical Panels* of the National Science Foundation (NSF) and other research and technology development agencies (NSF Information Technology Research Panel, 2002; INRIA Project Evaluation Panel, 2003; NSF CAREER Panel, 2005; NSF Foundations Panel, 2006; NSF NeTS CAREER Panel, 2006; NSF FIND Panel, 2007; NSF CDI Panel, 2008).

RESEARCH PUBLICATIONS, REPORTS, TALKS, ETC.

BOOK CHAPTERS:

1. **Power Control in Wireless Networking: Basic Principles and Core Algorithms** (with S. Gitzenis). Signal Processing for Mobile Communications Handbook (M.Ibnkahla, Editor), CRC Press, 2004.
2. **Transmitter Power Control in Wireless Computing** (with S. Gitzenis). Handbook of Mobile Computing (I. Mahgoub, Editor), CRC Press, 2004.
3. **Adaptive Batch Scheduling for Packet Switching with Delays** (with K. Ross). High-Performance Packet Switching Architectures (Editors: I. Elhanany, M. Hamdi), Springer-Verlag, pp. 66-80, 2006.
4. **Geometry of Packet Scheduling: Maximal Throughput Cone Scheduling Algorithms** (with K. Ross). High-Performance Packet Switching Architectures (Editors: I. Elhanany, M. Hamdi), Springer-Verlag, pp. 82-100, 2006.
5. **Autonomic Admission Control for Networked Information Servers** (with G. Paleologo). Telecommunications Network Design and Management, Operations Research and Computer Science Interfaces, Kluwer, vol. 23, pp. 227-244, 2003

JOURNAL PAPERS:

6. **Smart Greedy Distributed Energy Allocation: a Random Games Approach** (I. Bistritz, A. Ward, Z. Zhou, N. Bambos). *IEEE Transactions on Automatic Control* (accepted).
7. **No Weight-Regret Learning in Adversarial Bandits with Delays** (I. Bistritz, Z. Zhou, X. Chen, N. Bambos). *Journal of Machine Learning Research (JMLR)* (accepted).
8. **Distributed Stochastic Optimization with Large Delays** (Z. Zhou, P. Mertikopoulos, N. Bambos, P. Glynn and Y. Ye). *Mathematics of Operations Research*, accepted and published online in Dec. 2021 (forthcoming as regular publication).
9. **Wireless Body-Area Network Control Policies for Energy-Efficient Health Monitoring** (Y.B. David, T. Geller, I. Bistritz, I. Ben-Gal, N. Bambos, E. Khmelnitsky). *Sensors Journal*, 21/4245, pp. 1-21, June 2021.
10. **Prediction of Prolonged Opioid Use After Surgery in Adolescents: Insights From Machine Learning** (A. Ward, T. Jani, E. De Souza, D. Scheinker, N. Bambos, A. Anderson). *Journal of Anesthesia & Analgesia*, 133(2):304-313, 2021.
11. **Consensus-Based Stochastic Control for Model-Free Battery Cell Balancing** (I. Bistritz, N. Bambos). *IEEE Transactions on Control of Network Systems*, 8(3):1139-1150, Sept. 2021.

12. **Robust Power Management via Learning and Game Design** (Z. Zhou, P. Mertikopoulos, A. Moustakas, N. Bambos, P. Glynn). *Operations Research*, 69(1):331-345, 2021.
13. **One for All and All for One: Distributed Learning of Fair Allocations with Multi-Player Bandits** (I. Bistritz, T.Z. Baharav, A. Leshem, N. Bambos). *IEEE Transactions on Selected Areas in Information Theory*, 2(2):584-598, June 2021.
14. **Incidence of and Factors Associated with Prolonged and Persistent Postoperative Opioid Use in Children 0 to 18 Years of Age** (A. Ward, E. De Souza, D. Miller, E. Wang, E. Sun, N. Bambos, T. Anderson). *Journal of Anesthesia & Analgesia* 131(4):1237-1248, Oct. 2020.
15. **On the Convergence of Mirror Descent Beyond Stochastic Convex Programming** (with Z. Zhou, P. Mertikopoulos, S. Boyd, P. Glynn). *SIAM Journal of Optimization*, 30(1):687-716, 2020.
16. **Clustering Users by Their Mobility Behavioral Patterns** (with I. Ben-Gal, S. Weinstock, G. Singer). *ACM Transactions on Knowledge Discovery from Data*, 13(4):1-28, 2019.
17. **Deterministic and Stochastic Wireless Network Games: Equilibrium, Dynamics and Price of Anarchy** (with Z. Zhou, P. Glynn). *Operations Research*, 66(6):1498-1516, 2018.
18. **Infinite Time Horizon Maximum Causal Entropy Inverse Reinforcement Learning** (with Z. Zhou, M. Bloem). *IEEE Transactions on Automatic Control*, 63(9): 2787-2802, September 2018.
19. **Stochastic Models of Ground Delay Program Implementation for Prediction, Simulation and Insight** (with M. Bloem). *Journal of Aerospace Operations*, 5(1-2):85-117 (2017), April 2018.
20. **Video Streaming Schemes for Industrial IoT** (with H. Kanzaki, K. Schubert). *Journal of Reliable Intelligent Environments*, 3(4):233-241, Dec. 2017.
21. **Improving Predictions of Pediatric Surgical Durations with Supervised Learning** (with N. Master, Z. Zhou, D. Miller, D. Scheinker, P. Glynn). *Int. Journal of Data Science and Analytics*, 4:35-52, May 2017.
22. **Service Rate Control of Tandem Queues with Power Constraints** (with L. Xia, D. Miller, Z. Zhou). *IEEE Transactions on Automatic Control*, March 2017.
23. **Power Control for Packet Streaming with Head-of-Line Deadlines** (with N. Master). *Journal of Performance Evaluation*, 106:1-18, 2016.
24. **Power Optimization in Random Wireless Networks** (with A. Moustakas, P. Mertikopoulos). *IEEE Transactions on Information Theory*, 62(9):5030-5058, September 2016.
25. **Adaptive Prefetching in Wireless Computing** (with N. Master, A. Dua, D. Tsamis, J.P. Singh). *IEEE Transactions on Wireless Communications*, 15(5):3296-3310, 2016.
26. **Myopic Policies for Non-Preemptive Scheduling of Jobs with Decaying Value** (with N. Master). *Journal of Probability in Engineering and Informational Sciences*, November 2016.
27. **Maximum Weight Matching with Hysteresis in Overloaded Queues with Setups** (with C. Chan, M. Armony). *Queueing Systems*, 82:315-351, 2016,
28. **Cone Schedules for Processing Systems in Fluctuating Environments** (with K. Ross, G. Michailidis). *IEEE Transactions of Automatic Control*, 60(10):2710-2715, 2015.

29. **Personalized Monitors for Real-Time Detection of Physiological States** (with L. Chow, A. Gilman, A. Chander). International Journal of e-Health and Medical Communications, 5(4):1-19, October 2014.
30. **Air Traffic Control Area Configuration Advisories from Near-Optimal Distinct Paths** (with M. Bloem). Journal of Aerospace Information Systems, 11(11):764-784, November 2014.
31. **Ground Delay Program Analytics with Behavioral Cloning and Inverse Reinforcement Learning** (with M. Bloem). AIAA Journal of Aerospace Information Systems, 2015.
32. **Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions** (with C. Chan, V.F. Farias, G. Escobar). Operations Research (INFORMS), 60(6):1323-1341, Jan. 2013.
33. **BEST-AP: Non-Intrusive Estimation of Available Bandwidth and its Application for Dynamic Access Point Selection** (with P. Dely, L. Chow, N. Bayer, H. Einsiedler, C. Peylo). Computer Communications (Elsevier), 39:78-91, Jan 2014.
34. **A Software-Defined Networking Approach for Handover Management with Real-Time Video in WLANs** (with P. Dely, A. Kassler, L. Chow, N. Bayer, H. Einsiedler, C. Peylo, D. Mellado, M. Sanchez). Journal of Modern Transportation (Springer), 21:58-65, March 2013.
35. **Network Assisted Mobile Computing with Optimal Uplink Query Processing** (with C. Chan, J. Singh). IEEE Transactions on Mobile Computing, 60(6):1323-1341, Jan. 2013.
36. **Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions** (with C. Chan, V.F. Farias, G. Escobar). Operations Research (INFORMS), 60(6):1323-1341, Jan. 2013.
37. **Power and Delay Aware Management of Packet Switches** (with L. Mastroleon, D. O'Neill, B. Yolken). IEEE Transactions on Computers, , 61(12):1789-1799, Dec. 2012.
38. **Optimal State Surveillance under Observation Budget Constraints** (with P. Bommannavar). Int. Journal on Advances in Intelligent Systems, 4(4):57-67, Dec. 2011.
39. **A Scalable Delay-Power Control Algorithm for Bandwidth Sharing in Wireless Networks** (with F. Baccelli and N. Gast). IEEE/ACM Transactions on Networking, 19(5):1458-1471, October 2011.
40. **An Integrated Approach to Security Risk Management for IT-Intensive Organizations** (with Jeff Mounzer and Tansu Alpcan). Journal of Information Assurance and Security, 6(2):115-123, 2011.
41. **Channel, Deadline, and Distortion Aware Scheduling of Video Streams over Wireless Links** (with A. Dua, C. Chan, J. Apostolopoulos). IEEE Trans. on Wireless Communications, 9(3):1001-1011, March 2010.
42. **A Characterization of Max-Min SIR-Balanced Power Allocation with Applications** (with S. Stanczak, M. Kaliszan). ACM Wireless Networks, 16: 2335-2347, Nov. 2010.
43. **Game Based Capacity Allocation for Utility Computing Environments** (with B. Yolken). Telecommunications Systems Journal, pp. 1-17, May 2010.
44. **Scheduling Algorithms for Broadcasting Media with Multiple Distortion Measures**, (with C. W. Chan, S. Wee, J. Apostolopoulos), IEEE Trans. on Wireless Communications, 8(8) 4188-4199, 2009.
45. **Adaptive Data-Aware, Utility-Based Scheduling in Resource-Constrained Systems** (with D. Vengerov, L. Mastroleon, D. Murphy). Journal of Parallel and Distributed Computing, 70:871-879, 2009.

46. **Projective Cone Scheduling (PCS) Algorithms for Maximal Throughput Packet Switches** (with K. Ross). IEEE/ACM Transactions on Networking, 17(3):976-989, June 2009.
47. **Joint Task Migration and Power Management in Wireless Computing** (with S. Gitzenis). IEEE Transactions of Mobile Computing, 8(9): 1189-1204, Sept. 2009.
48. **Content-Aware Payout and Packet Scheduling for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos) IEEE Trans. on Multimedia, 10(5):885-895, Aug. 2008.
49. **Joint Transmitter Power Control and Mobile Cache Management in Wireless Computing** (with S. Gitzenis). IEEE Transactions of Mobile Computing, 7(4):498-512, April 2008.
50. **Transmission Power and Duration-Aware Payout Control for Packetized Media Streaming over Wireless Links** (with Y. Li). Wireless Communications and Mobile Computing, 8:309-326, Jan. 2008
51. **Downlink Wireless Packet Scheduling with Deadlines** (with A. Dua). IEEE Trans. on Mobile Computing, 6(12), pp. 1410-1425, 2007.
52. **TCP Performance Dynamics and Link-Layer Adaptation Based Optimization Methods for Wireless Networks** (with J. P. Singh, Y. Li, A. Bahai, B. Xu, G. Zimmerman). IEEE Trans. Wireless Communications, 6(5), pp. 1864-1878, 2007.
53. **Transmission Power and Duration-Aware Payout Control for Packetized Media Streaming over Wireless Links** (with Y. Li). Wireless Communications and Mobile Computing, 2008.
54. **Joint Power-Payout Control for Media Streaming over Wireless Links** (with Y. Li, A. Markopoulou, J. Apostolopoulos), IEEE Transactions in Multimedia, 2006.
55. **Optimal Processor Allocation to Differentiated Job Flows** (with G. Michailidis and K. Wasserman), Performance Evaluation, v. 63(1), pp. 1-14, Jan. 2006.
56. **Queueing Networks of Random Link Topology: Stationary Dynamics of Maximum Throughput Schedules** (with G. Michailidis), Queueing Systems Theory and Applications, v. 50(1), pp. 5-52, May 2005.
57. **A Fuzzy Reinforcement Learning Approach to Power Control in Wireless Transmitters** (with D. Vengerov and H. Berenji). IEEE Trans. SMC (Part B), v. 35-4, pp. 768-778, Aug. 2005.
58. **Queueing and Scheduling in Random Environments** (with G. Michailidis). Advances in Applied Probability, vol. 36 (1), 293-317, 2004.
59. **Scheduling Optical Traffic Bursts in Time-Domain Wavelength Interleaved Optical Networks** (with K. Ross, K. Kumaran, I. Saniee, I. Widjaja). IEEE Journal of Selected Areas in Communications, vol. 21(9), pp. 1441-1451, 2003.
60. **Queueing Dynamics and Maximal Throughput Scheduling in Switched Processing Systems** (with M. Armony). Queueing Systems Theory and Applications, vol. 44(3), pp. 209-252, 2003.
61. **On Stability of Queueing Networks with Job Deadlines** (with A. Ward). Journal of Applied Probability, vol. 40(2), pp. 293-304, 2003.
62. **On Parallel Queueing with Random Server Connectivities and Routing Constraints** (with G. Michailidis). Probability in the Engineering and Informational Sciences, vol. 16, pp. 185-203, 2002.

63. **Optimal Control of Parallel Queues with Batch Service** (with C. Xia, G. Michailidis, P. Glynn). *Probability in the Engineering and Informational Sciences*, vol. 16, pp. 289-307, 2002.
64. **Power Controlled Multiple Access (PCMA) Schemes in High Performance Wireless Packet Networks** (with S. Kandukuri). *IEEE Personal Communications*, vol. 9(3), pp.58-64, 2002.
65. **Dynamic Online Task Scheduling on Parallel Processors** (with C. Xia, G. Michailidis). *Performance Evaluation*, vol. 46, pp. 219-233, 2001.
66. **Globally Constrained Power Control Across Multiple Channels in Wireless Packet Networks** (with S. Kandukuri). *ACM Mobile Networks*, vol. 6(5), pp. 427-434, 2001. (invited paper)
67. **Channel Access Algorithms with Active Link Protection for Wireless Communication Networks with Power Control** (with S. Chen, G. Pottie). *IEEE/ACM Transactions on Networking*, vol. 8(5), pp. 583-597, 2000.
68. **On Synchronization of Poisson Processes and Queueing Networks with Service and Synchronization Nodes** (with B. Prabhakar, T. Mountford). *Advances in Applied Probability*, vol. 32(3), pp. 824-843, 2000.
69. **Power-Induced Time Division on Asynchronous Channels** (with J. Rulnick). *ACM Wireless Networks*, vol. 5(2), pp. 71-80, 1999. (invited paper)
70. **Towards Power-Sensitive Network Architectures in Wireless Communications: Concepts, Issues and Design Aspects**. *IEEE Personal Communications*, vol. 5(3), pp. 50-59, 1998. (invited paper)
71. **Mobile Power Management for Wireless Communication Networks** (with J. Rulnick). *ACM Wireless Networks*, vol. 3(1), pp. 3-14, 1997. (invited paper)
72. **On a Singular Feature of Critical G/M/1 Queues** (with B. Prabhakar). *Systems and Control Letters*, vol. 28(5), pp. 239-245, 1996.
73. **Wireless, Mobile, Multimedia Networks** (with A. Alwan, R. Bagrodia, N. Bambos, J. Cong, M. Gerla, L. Kleinrock, J. Villaseñor). *IEEE Personal Communications*, vol. 3(2), pp. 34-51, 1996.
74. **Convergence of Departures in Tandem Networks of GI/∞ Queues** (with B. Prabhakar, T. Mountford). *Probability in the Engineering and Informational Sciences*, vol. 10(4), pp. 487-500, 1996.
75. **The Supercomputer Supernet Testbed: A WDM-Based Supercomputer Interconnect** (with L. Kleinrock, M. Gerla, J. Cong, E. Gafni, L. Bergman, J. Bannister, S.P. Monacos, T. Bujewski, P-C. Hu, B. Kannan, B. Kwan, E. Leonardi, J. Peck, P. Palnati, S. Walton). *IEEE/OSA Journal of Lightwave Technology*, vol. 14(6), pp. 1388-1399, 1996.
76. **Optimal Server Allocation to Parallel Queues with Finite Capacity Buffers** (with K. Wasserman). *Probability in the Engineering and Informational Sciences*, vol. 10(2), pp. 279-285, 1996.
77. **Asymptotic Optimality of Statistical Multiplexing in Pipelined Processing** (with K. Wasserman). *Queueing Systems Theory and Applications*, vol. 21, pp. 97-123, 1995.
78. **OPTIMIC: A Scalable, Distributed, All-Optical Terabit Network** (with L. Kleinrock, J. Cong, E. Gafni, M. Gerla, L. Bergman, J. Bannister). *Journal of High Speed Networks*, vol. 4(4), pp. 407-424, 1995.

79. **On Infinite Queueing Tandems** (with B. Prabhakar). *Systems and Control Letters*, vol. 23, pp. 305-314, 1994.
80. **Optimality Aspects of Greedy Schemes in Parallel Processing of Random Graph-Structured Jobs** (with S. Chen). *Probability in the Engineering and Informational Sciences*, vol. 8, pp. 229-243, 1994.
81. **On Stationary Tandem Queueing Networks with Job Feedback** (with K. Wasserman). *Queueing Systems Theory and Applications*, vol. 15, pp. 137-164, 1994.
82. **Scheduling and Stability Aspects of a General Class of Parallel Processing Systems** (with J. Walrand). *Advances in Applied Probability*, vol. 25, pp. 176-202, 1993.
83. **On Closed Ring Queueing Networks**. *Journal of Applied Probability*, vol. 29(4), pp. 979-995, 1992.
84. **On Flows in Stochastic Marked Graphs** (with F. Baccelli & J. Walrand). *Probability in the Engineering and Informational Sciences*, vol. 5, pp. 145-157, 1991.
85. **On Stability and Performance of Parallel Processing Systems** (with J. Walrand). *Journal of the Association for Computing Machinery*, vol. 38(2), pp. 429-452, 1991.
86. **An Invariant Distribution for the G/G/1 Queueing Operator** (with J. Walrand). *Advances in Applied Probability*, vol. 22, pp. 254-256, 1990.
87. **On the Asymptotic Execution Time of Multi-Tasked Processes on Tandem Processors** (with J. Walrand). *Systems and Control Letters*, vol. 13, pp. 391-396, 1989.
88. **On State-Dependent Queues and Acyclic Queueing Networks** (with J. Walrand). *Advances in Applied Probability*, vol. 21, pp. 681-701, 1989.
89. **On Queues with Periodic Inputs** (with J. Walrand). *Journal of Applied Probability*, vol. 26, pp. 381-389, 1989.

PEER-REVIEWED CONFERENCE PAPERS:

1. **Power-Optimized Processor slowdown Control** (A.J. Mann, N. Bambos). *Proceedings of the IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2022.
2. **Learning in Games with Quantized Payoff Observations** (K. Lotidis, P. Mertikopoulos, N. Bambos). *Proceedings of the IEEE Conference on Decision and Control*, Cancun, Mexico, Dec. 2022.
3. **Framed Projective Cone Scheduling: Latency vs. Context-Switching Tradeoff in Data Centers** (E. Zeger, A.J. Mann, N. Bambos). *Proceedings of the IEEE Global Communications Conference*, Rio de Janeiro, Brazil, Dec. 2022.
4. **Queue Up your Regrets: Achieving the Dynamic Capacity Region of Multi-player Bandits** (I. Bistritz, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems*

(*NeurIPS*), 2022 (accepted).

5. **Active Testing for an Emerging Epidemic** (A.J. Mann, I. Bistriz, N. Bambos). *Proceedings of the IEEE International Conference on e-Health Networking, Applications & Services (HealthCom)*, 2022 (accepted).
6. **Online Learning for Load Balancing of Unknown Monotone Resource Allocation Games** (I. Bistriz, N. Bambos). *Proceedings of the 38th International Conference on Machine Learning (ICML)*, PMLR 139:968-979, July 2021.
7. **Controlling Epidemics via Testing** (K. Lotidis, A. Moustakas, N. Bambos). *Proceedings of the 2021 60th Conference on Decision and Control (CDC)*, pp. 2092-2097, Austin, TX, Dec. 2021.
8. **An OpenAI-OpenDSS Framework for Reinforcement Learning on Distribution-Level Microgrids** (K. Moy, C. Tae, Y. Wang, G. Henri, N. Bambos, R. Rajagopal). *Proceedings of the 2021 IEEE Power & Energy Society General Meeting (PESGM)*, pp. 1-5, Washington, DC, July 2021.
9. **Power Controlled Slowdown in Data Centers** (A. Mann, N. Bambos). *Proceedings of IEEE International Conference on Communications (ICC)*, pp. 1-6, Korea, May 2022.
10. **Cooperative Multi-Player Bandit Optimization** (I. Bistriz, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 22 pages, 2020.
11. **Distributed Distillation for On-Device Learning** (I. Bistriz, A. Mann, N. Bambos). *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 26 pages, 2020.
12. **My Fair Bandit: Distributed Learning of Max-Min Fairness with Multi-Player Bandits** (I. Bistriz, T.Z. Baharav, A. Leshem, N. Bambos). *Proceeding of the International Conference on Machine Learning (ICML)*, PMLR 119: 930-940, 10 pages, 2020.
13. **Distributed Scheduling of Charging for On-Demand Electric Vehicle Fleets** (I. Bistriz, M. Klein, N. Bambos, O. Maimon, R. Rajagopal). *Proceedings of the IFAC Workshop on Discrete Event Systems (WODES)*, 6 pages, 2020. [**Best student paper finalist award**; 1 best paper award and 3 finalist awards out of 75 papers presented at the conference.]
14. **Operationally-Informed Hospital-Wide Discharge Prediction Using Machine Learning** (by A. Ward, A. Mann, J. Vallon, G. Escobar, N. Bambos, A. Schuler). *Proceedings of the 22nd IEEE Int. Conf. on e-Health Applications & Services (HealthCom)*, 6 pages, Shenzhen, China, Dec. 2020. [**Best Paper Runner-Up Award**; 1 best paper award and 2 runner-up awards out of 100 papers presented at the conference.]
15. **Rates of, and Risk Factors for, Substance Use Disorder in Children and Young Adults Before and After Surgery** (I. Thapa, A. Ward, B. De Souza, N. Bambos, T. Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020*; extended abstract published in the *Journal of Anesthesia & Analgesia Supplement*, 130(5)S:689-690.
16. **Rates of, and Risk Factors for, Chronic Pain in Children and Young Adults Before and After Surgery** (I. Thapa, A. Ward, B. De Souza, N. Bambos, T. Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020*; extended abstract published in the *Journal of Anesthesia & Analgesia Supplement*, 130(5)S:691.

17. **Development and Validation of a Machine Learning Algorithm for Prediction of Prolonged Opioid Use After Surgery in Adolescents and Young Adults** (Trisha Jani, A. Ward, B. De Souza, N. Bambos, Thomas A Anderson). *Proceedings of the International Anesthesia Research Society 2020 Annual Meeting and International Science Symposium, San Francisco, CA, 2020*; extended abstract published in the *Journal of Anesthesia & Analgesia Supplement*, 130(5)S:712-713.
18. **Distributed Scheduling of Charging for On-Demand Electric Vehicle Fleets** (with I. Bistritz, M. Klein, O. Maimon, R. Rajagopal). Proceedings of the 2020 Workshop on Discrete-Event Systems (WODES), Rio de Janeiro, Brazil, Nov. 2020.
19. **Online EXP3 Learning in Adversarial Bandits** (with I. Bistritz, Z. Zhou, X. Chen, J. Blanchet). Proceedings of the Conference on Advances in Neural Information Processing Systems (NeurIPS), Vancouver, Canada, December 2019.
20. **Controlling Contact Network Topology to Prevent Measles Outbreaks** (with I. Bistritz, D. Kahana, I. Ben-Gal, D. Yamin). Proceedings of the 2019 IEEE Global Communications Conference (Globecom), pp. 1-6, Waikoloa, Hawaii, December 2019.
21. **Asymptotically Optimal Distributed Gateway Load-Balancing for the Internet of Things** (with I. Bistritz). Proceedings of the 2019 IEEE International Conference on the Network of the Future (NoF), pp. 98-101, Rome, Italy, October 2019.
22. **Noninvasive Identification of Hypotension Using Convolutional-Deconvolutional Networks** (with D. Miller, A. Ward, A. Shin, D. Scheinker). Proceedings of 2019 IEEE International Conference on E-Health Networking, Applications and Services (HealthCom), pp. 1-6, Bogota, Colombia, October 2019.
23. **Learning to Emulate an Expert Projective Cone Scheduler** (with A. Ward, N. Master). Proceedings of the 2019 IEEE American Control Conference (ACC), pp. 292-297, Philadelphia, PA, July 2019.
24. **Incidence of Persistent Opioid Use in Children after Surgery** (with T. Anderson, A. Ward, B. De Souza, E. Wang, D. Miller). Proceedings of the International Anesthesia Research Society 2019 Annual Meeting and International Science Symposium, May 17-20, Montreal, Quebec, Canada; published in the *Journal: Anesthesia & Analgesia Supplement* 128(5)S-354:764-765.
25. **Smart Greedy Distributed Allocation in Microgrids** (with I. Bistritz, A. Ward, Z. Zhou). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1-6, Shanghai, China, May 2019.
26. **Anesthesiologist Surgery Assignments Using Policy Learning** (with A. Ward, Z. Zhou, Ellen Wang, D. Scheinker). Proceedings of the 2019 IEEE International Conference on Communications (ICC), 6 pages, Shanghai, China, May 2019.
27. **Learning Health State Transition Probabilities via Wireless Body Area Networks** (with Tal Geller, Yair Bar David, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1-6, Shanghai, China, May 2019.
28. **The Power of Consensus: Optimal Distributed Multichannel Wireless Transmitter Power Control** (with I. Bistritz). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pp. 1- 6, Shanghai, China, May 2019.
29. **Bidding-Based Dynamic Power Pricing Scheme in Smart Grids** (with A. Ward, Z. Zhou). Proceedings of the 2019 IEEE International Conference on Computing, Networking and Communications (ICNC), pp. 729-734, Honolulu, Hawaii, February 2019.

30. **Learning in Games with Lossy Feedback** (with Z. Zhou, P. Mertikopoulos, S. Athey, P. Glynn and Y. Ye). Proceedings of 32nd Neural Information Processing Systems (NIPS) Conference, 11 pages, Montreal, Canada, December 2018.
31. **Distributed Asynchronous Stochastic Optimization under Unbounded Delays: How Slow Can You Go?** (with Z. Zhou, P. Mertikopoulos, P. Glynn and Y. Ye, Li-Jia Li, Li Fei-Fei). Proceedings of the 35th International Conference on Machine Learning (ICML), 10 pages, Stockholm, Sweden, July 2018.
32. **Power Control with Random Delays: Robust Feedback Averaging** (with A. Ward, Z. Zhou, P. Mertikopoulos). Proceedings of the 57th Conference on Decision and Control (CDC), pp. 7040-7045, Miami Beach, FL, December 2018.
33. **Robustness of Join-the-Shortest-Queue Scheduling to Communication Delay** (with S. Mehdian, Z. Zhou). Proceedings of the 2018 American Control Conference (ACC), pp. 3708-3713, Milwaukee, WI, June 2018.
34. **Optimal Health Monitoring via Wireless Body Area Networks** (with Yair Bar David, Tal Geller, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 57th Conference on Decision and Control (CDC), pp. 7040-7045, Miami Beach, FL, December 2018.
35. **Sensing-Constrained Power Control in Digital Health** (with D. Miller, Z. Zhou, N. Bambos, I. Ben-Gal). Proceedings of the 2018 American Control Conference (ACC), pp. 3708-3713, Milwaukee, WI, June 2018.
36. **Optimal Sensing for Patient Health Monitoring** (with D. Miller, Z. Zhou, I. Ben-Gal). Proceedings of the 2018 International Communications Conference (ICC), 7 pages, Kansas City, MO, May 2018.
37. **Physiological Waveform Imputation of Missing Data using Convolutional Autoencoders** (with D. Miller, A. Ward, D. Scheinker, Andrew Shin). Proceedings of the 2018 IEEE 20th International Conference on e-Health Networking, Applications and Services (Healthcom), pages 6, Ostrava, Czech Republic, September 2018.
38. **Automatic Sleep Arousal Identification from Physiological Waveforms Using Deep Learning** (with D. Miller, A. Ward, N. Bambos). Proceedings of 2018 Computing in Cardiology (CinC) Conference, pages 5, Maastricht, Netherlands, September 2018.
39. **Bidding-Based Dynamic Power Pricing Scheme in Smart Grids** (with A. Ward, Z. Zhou). Proceedings of the 2018 IEEE International Conference on Computing, Networking and Communications (ICNC), pp. 729-734, Honolulu, Hawaii, February 2019.
40. **Smart Greedy Distributed Allocation in Microgrids** (with I. Bistritz, A. Ward, Z. Zhou). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
41. **Anesthesiologist Surgery Assignments Using Policy Learning** (with A. Ward, Z. Zhou, Ellen Wang, D. Scheinker). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
42. **Learning Health State Transition Probabilities via Wireless Body Area Networks** (with Tal Geller, Yair Bar David, Evgeni Khmelnitsky, Irad Ben-Gal, A. Ward, D. Miller). Proceedings of the 2019

IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).

43. **The Power of Consensus: Optimal Distributed Multichannel Wireless Transmitter Power Control** (with I. Bistritz). Proceedings of the 2019 IEEE International Conference on Communications (ICC), pages 6, Shanghai, China, May 2019 (accepted 2018).
44. **Countering Feedback Delays in Multi-Agent Learning** (Z. Zhou, P. Mertikopoulos, N. Bambos, P. Glynn, C. Tomlin), Proceedings of Advances for Neural Information Processing Systems (NIPS), pp.1-11, Long Beach, CA, Dec. 2017.
45. **Stochastic Mirror Descent in Variationally Coherent Optimization Problems** (with Z. Zhou, P. Mertikopoulos, S. Boyd, P. Glynn), Proceedings of Advances for Neural Information Processing Systems (NIPS), pp. 1-11, Long Beach, CA, Dec. 2017.
46. **Classical Communication Network Design Via Quantum Heuristics** (with A. Ward, N. Master). Proceeding of the Conference on Decision and Control (CDC), pp. 1-6, Long Beach, CA, Dec. 2017.
47. **Mirror Descent Learning in Continuous Games** (with Z. Zhou, P. Mertikopoulos, A. Moustakas, P. Glynn). Proceeding of the Conference on Decision and Control (CDC), pp. 1-6, Long Beach, CA, Dec. 2017.
48. **Stable Power Control in Wireless Networks via Dual Averaging** (with Z. Zhou, P. Mertikopoulos, A. Moustakas, S. Mehdian, Peter Glynn). Proceedings of IEEE Global Communications Conference (GLOBECOM), pp. 1-6, Singapore, Dec. 2017.
49. **Longest-Queue-First Scheduling with Intermittent Sampling** (with S. Mehdian, Z. Zhou). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
50. **Least Action Routing: Identifying the Optimal Path in a Wireless Relay Network** (with A. Moustakas, P. Mertikopoulos, Z. Zhou). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
51. **Market-Based Dynamic Service Mode Switching in Virtualized Wireless Networks** (with M. Dimakopoulou, M. Valdez-Vivas, J. Apostolopoulos). Proceedings of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), pp. 1-6, Montreal, QC, Canada, Oct. 2017.
52. **Video Streaming Services for Industrial IoT** (with H. Kanzaki, K. Schubert). Proceedings of International Conference on Computer Communications and Networks (ICCCN), pp. 1-7, Vancouver, CA, Aug. 2017.
53. **Dynamic Control of Data Center Network and Computation Resources** (with D. Miller, Li Xia, Z. Zhou, N. Bambos). Proceedings of the International Conference on Computing, Networking and Communications (ICNC), pp. 1-6, San Jose, CA 2017.
54. **Asynchronous Best Response Dynamics for Resource Allocation Games in Utility Computing** (with K. Schubert, N. Master, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.
55. **An Infinite Dimensional Model for a Single-Server Priority Queue** (with N. Master, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.

56. **Join-the-Shortest Queue Scheduling with Delay: A Case Study** (with S. Mehdian, Z. Zhou). Proceedings of IEEE American Control Conference (ACC), 6 pages, Seattle, WA, May 2017.
57. **Cost-Sensitive Security Risk Management for Large-Scale Computing Infrastructures** (with N. Master). Proceedings of the IEEE International Conference on Computing, Networking and Communications (ICNC), 6 pages, San Jose, CA, January 2017.
58. **An Infinite Dimensional Model for a Many Server Priority Queue** (with N. Master, Z. Zhou). Proceeding of the Conference on Information Sciences and Systems (CISS), 6 pages, Johns Hopkins University, Baltimore, MD, March 2017.
59. **Wireless Channel Selection with Non-Invasive Power Probing** (with K. Schubert). Proceedings of IEEE Wireless Telecommunications Symposium, 6 pages, Chicago, IL, April 2017.
60. **Infinite Server Queueing Networks with Deadline Bases Routing** (with N. Master), IEEE Conference on Decision and Control (CDC), pp. 5360-5366, Las Vegas, NV, December 2016.
61. **Repeated Games for Power Control in Wireless Communications: Equilibrium and Regret** (with Z. Zhou, P. Glynn). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 3603-3610, Las Vegas, NV, December 2016.
62. **A Game-Theoretical Formulation of Influence Networks** (w Z. Zhou, B. Yolken, R.A. Miura-Ko). Proceeding of IEEE American Control Conference, pp. 3802-3807, Boston, MA, July 2016.
63. **Dynamics of Linear Influence Networks under Stochastic Environments** (with Z. Zhou, P. Glynn). Proceeding (Springer) of Conference on Decision and Game Theory for Security (GameSec), 6 pages, New York, NY, November 2016.
64. **A Stochastic Stability Characterization of the Foschini-Miljanic Algorithm in Random Wireless Networks** (with Z. Zhou, D. Miller, P. Glynn). Proceedings of IEEE Global Communications Conference (GLOBECOM), 6 pages, Washington, DC, December 2016.
65. **Wireless Power-Controlled TCP with Holdover** (with K. Schubert). Proceedings of IEEE International Conference on Communications (ICC), 6 pages, Kuala Lumpur, Malaysia, May 2016.
66. **Data Aggregation for Low Power Wireless Devices** (with K. Schubert). Proceedings of IEEE MILCOM, 6 pages, November 2016.
67. **Detecting Inaccurate Predictions of Pediatric Surgical Durations** (with Z. Zhou, D. Miller, N. Master, P. Glynn). IEEE International Conference on Data Science and Advanced Analytics (DSAA), pp. 452-457, Montreal, Canada, October 2016.
68. **Reliable and Efficient Performance Monitoring in Linux** (with M. Dimakopoulou, S. Eranian, N. Koziris). Proceedings of International Conference for High Performance Computing, Networking, Storage and Analysis, 11 pages, Salt Lake City, Utah, November 2016.
69. **Scalable Data Center Power Management via a Global Stress Signal** (with D. Miller, N. Master, Z. Zhou, N. Bambos). IEEE Global Communications Conference (GLOBECOM), 7 pages, 2015.
70. **Target-Rate Driven Resource Sharing in Queueing Systems** (with Z. Zhou). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 4490-4945, Osaka, Japan, December 2015.
71. **Wireless Communication Games in Fixed and Random Environments** (with Z. Zhou). Proceedings of IEEE Conference on Decision and Control (CDC), pp. 1637-1642, Osaka, Japan, December 2015.

72. **Dynamic Management of Network Risk from Epidemic Phenomena** (A. Sinha, J. Ducci). IEEE Conference on Decision and Control (CDC), 6 pages, Osaka, Japan, December 2015.
73. **Low Latency Policy Iteration via Parallel Processing and Randomization** (with N. Master). IEEE Conference on Decision and Control (CDC), pp. 1084-1091, 2015.
74. **Channel Responsive Wireless TCP** (with K. Schubert, H. Endo). IEEE International Communications Conference (ICC'14), 6 pages, London, UK, 2015.
75. **A General Model for Resource Allocation in Utility Computing** (with Z. Zhou). American Control Conference, pp. 1746-1751, Chicago, IL, 2015.
76. **Service Rate Control for Jobs with Decaying Value** (with N. Master). American Control Conference, 6 pages, Chicago, IL, 2015.
77. **Ground Delay Program Analytics with Behavioral Cloning and Inverse Reinforcement Learning** (with M. Bloem). Proceedings of AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, 6 pages, Atlanta, GA, June 2014.
78. **Infinite Time Horizon Maximum Causal Entropy Inverse Reinforcement Learning** (with M. Bloem). IEEE Conference on Decision and Control (CDC'14), pp. 4911-4916, Los Angeles, CA, December 2014.
79. **Power-Controlled Multiple Access with Queue-Dependent Backoff Threshold** (with J. Mounzer, K. Schubert, A. Goldsmith). IEEE Global Communications Conference (GlobeCom'14), pp. 4738-4744, Austin, TX, December 2014.
80. **Randomized Iterations for Low Latency Fixed Point Computation** (with N. Master). IEEE Conference on Decision and Control (CDC'14), pp. 5208-5215, Los Angeles, CA, December 2014.
81. **Distributed Smart Grid Architecture for Delay and Price Sensitive Power Management** (with N. Master, J. Mounzer), accepted to the IEEE International Communications Conference (ICC'14), pp. 3670-3675, Sydney, Australia, June 2014.
82. **Dynamic Resource Management in Virtualized Data Centers with Bursty Traffic** (with M. Valdez-Vivas, J. Apostolopoulos), accepted to the IEEE International Communications Conference (ICC'14), Sydney, Australia, June 2014.
83. **Power Control for Wireless Streaming with HOL Packet Deadlines** (with N. Master), accepted to the IEEE International Communications Conference (ICC'14), pp. 2263-2269, Sydney, Australia, June 2014.
84. **Channel Exploration for Wireless Media Streaming with Handoff and Rebuffering Control** (with L. Chow, J.P. Singh), accepted to the IEEE International Communications Conference (ICC'14), pp. 2281-2288, Sydney, Australia, June 2014.
85. **Channel Aware Rebuffering for Wireless Media Streaming with Handoff Control** (with L. Chow, B. Collins, P. Dely, A. J. Kassler, C. Peylo, H. J. Einsiedler, N. Bayer). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), pp. 4434-4439, Atlanta, Georgia, December 2013.
86. **Congestion versus Accuracy Tradeoffs in IP Traffic Classification** (with M. Valdez-Vivas). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), Atlanta, Georgia,

December 2013.

87. **Playout Buffer Responsive Wireless Streaming for Multiple Clients** (with P. Bommannavar, J. Apostolopoulos). Proceedings of the 2013 IEEE Global Communications Conference (GlobeCom'13), pp. 2276-2281, Atlanta, Georgia, December 2013.
88. **QoS-Aware Admission Control in Heterogeneous Data Centers** (with C. Delimitrou, C. Kozyrakis). Proceedings of USENIX International Conference on Autonomic Computing, San Jose, CA, June 2013.
89. **Delay-Sensitive Power Management for Packet Switches** (with M. Valdez-Vivas, D. O'Neill). Proceedings of IEEE International Communications Conference (ICC'13), Budapest, Hungary, June 2013.
90. **Latency and Energy in Quality-Driven Applications for Networked Wireless Devices** (with M. Valdez-Vivas). Proceedings of IEEE International Communications Conference (ICC'13), Budapest, Hungary, June 2013.
91. **Resource Allocation and Scheduling for Energy Efficient Tracking** (with P. Bommannavar, J. Apostolopoulos). Proceedings of IEEE International Communications Conference (ICC'13), pp. 1532-1537, Budapest, Hungary, June 2013.
92. **Deadline Aware Packet Scheduling in Switches for Multimedia Streaming Applications** (with P. Bommannavar, J. Apostolopoulos). Proceedings of IEEE International Communications Conference (ICC'13), pp. 3797-3802, Budapest, Hungary, June 2013.
93. **Power Optimization in Random Wireless Networks (with A.L. Moustakas)**. Proceedings of IEEE International Symposium on Information Theory (ISIT'13), Istanbul, Turkey, July 2013.
94. **Real-Time Physiological Stream Processing for Health Monitoring Services** (with L. Chow). Proceedings of International Conference on e-Health Networking, Applications and Services (HealthCom'13), pp. 611-616, Lisbon, Portugal, Oct. 2013.
95. **An Approach for Finding Multiple Areas of Specialization Configuration Advisories** (with M. Bloem). Proceedings of AIAA Aviation Technology, Integration, and Operations (ATIO), 6 pages, Los Angeles, CA, August 2013.
96. **Playout-Buffer Aware Hand-Off Control for Wireless Video Streaming** (with L. Chow, B. Collins, C. Peylo, H. Einsiedler, N. Bayer, P. Dely, A. Kassler), IEEE Global Communications Conference (GlobeComm), pp. 5237-5242, Los Angeles, CA, 6 pages, Dec. 2012.
97. **Resource Constrained Failure Management in Networked Computing Systems** (with P. Bommannavar). IEEE Global Communications Conference (GlobeComm), pp. 1884-1889, Los Angeles, CA, 6 pages, Dec. 2012.
98. **Approximating the Likelihood of Historical Airline Actions to Evaluate Airline Delay Cost Functions** (with M. Bloem, H. Huang). IEEE Conference on Decision and Control (CDC), Maui, HI, pp. 508-513, Dec. 2012.
99. **Power Optimization on a Network: The Effect of Randomness** (with A. Moustakas), IEEE Int. Symposium of Information Theory, pp. 606-610, June 2012.
100. **Nonlinear Cooperative Dynamics in Distributed Power Control for Wireless Networks** (with J. Mounzer). IEEE Int. Communications Conference (ICC), pp. 5042-5047, June 2012.

101. **Evaluation of Algorithms for a Miles-in-Trail Decision Support Tool** (with M. Bloem, D. Hattaway, and N. Bambos). Int. Conference on Research in Air Transportation, Berkeley, CA, 8 pages, May 2012.
102. **Power Budgeted Packet Scheduling for Wireless Multimedia** (with P. Bommannavar, J. Apostolopoulos). International Communications Conference (ICC), pp. 2108-2113, June 2012.
103. **Power Control in Random Networks: The Effect of Disorder in User Positions** (with A. Moustakas). IEEE Int. Workshop for Spatial Stochastic Models for Wireless Networks (SPASWIN), pp. 380-385, May 2012.
104. **Intrusion Detection in IT Networks with Limited Observations** (with P. Bommannavar) IEEE International Conference on Computing, Networking and Communications (ICNC), Maui, Hawaii, pp. 6, Nov. 2012.
105. **Coordinated Tactical Air Traffic and Airspace Management** (with M. Bloem) IEEE Conference on Decision and Control (CDC), pp. 5287-5292, Orlando, FL, Dec. 2011.
106. **Implementation of Fine-Grained Power/Bitrate Control for 802.11 Devices** (with Min-Wook Jeon, A. Kabani). Int. Tech. Conference on Circuits/Systems, Computers and Communications, Gyeongju, Korea, June 2011.
107. **Security Risk Management in Computing Systems with Constraints on Service Disruption** (with P. Bommannavar), IEEE International Conference on Computer Communication Networks (ICCCN), pp. 6, Maui, Hawaii, June 2011.
108. **Security Risk Management via Dynamic Games with Learning** (with Praveen Bommannavar and Tansu Alpcan. IEEE Int. Communications Conference (ICC), pp. 1-6, Kyoto, Japan, June 2011.
109. **Dynamic Control and Mitigation of Interdependent IT Security Risks** (with J. Mounzer, T. Alpcan) IEEE Int. Communications Conference (ICC), pp. 1-6, Cape Town, South Africa, May 2010 (Best Paper Award).
110. **Optimal State Surveillance under Budget Constraints** (with P. Bommannavar) Int. Conf. on Emerging Network Intelligence (Emerging'10), pp. 1-6, Florence, Italy, Oct. 2010 (Best Paper Award).
111. **Hybrid Power Control Algorithms for Streaming and Data Traffic in Wireless Networks** (with J. Mounzer) IEEE Int. Communications Conference (ICC), pp. 1-6, Cape Town, South Africa, May 2010.
112. **Integrated Security Risk Management for IT-Intensive Organizations** (with J. Mounzer, T. Alpcan) Sixth Int. Conf. on Information Assurance and Security (IAS), pp. 329-334, Atlanta, GA, Aug. 2010.
113. **Locking Dynamics and Mitigation Schemes in Distributed Power Control for Wireless Networks** (with J. Mounzer) IEEE Global Communications Conference (GlobeComm), pp. 1-6, Dec. 2010 .
114. **Fairness of Heterogeneous Queues through Cone Scheduling** (with C. Chan, M. Armony). Manufacturing and Service Operations Management Society Conference, pp. 1-3, Technion, Haifa, Israel, June 2010.

115. **Patch Scheduling for Risk Exposure Mitigation under Service Disruption Constraints** (with P. Bommannavar) 4th Int. Conf. on Emerging Security Information Systems and Technologies (SecureWare), pp. 176-181, Venice, July 2010.
116. **Admission Control for Autonomous Wireless Links with Power Constraints** (with M. Kaliszan, S. Stanczak). IEEE Int. Conf. in Acoustics Speech and Signal Processing, pp. 3138-3141, Dallas, TX, March 2010.
117. **Backlog aware scheduling for ingress memories in high-radix, single-stage switches** (with D. Tsamis, B. Yolken, W. Olesinski, H. Eberle, N. Gura). IEEE Global Com. Conference (Globecom), Honolulu, HI, pp. 1-6, Dec. 2009
118. **Throughput Loss in Task Scheduling due to Server State Uncertainty** (with C. Chan). ACM Int. Conf. in Performance Evaluation Methodologies and Tools (ValueTools), pp. 1-10, Pisa, Italy, 2009.
119. **Game Theoretic Rate Control for Mobile Devices** (with D. Tsamis, T. Alpcan). IEEE Int. Conf. on Game Theory for Networks (Gamenets), Instabul, Turkey, 2009.
120. **Dynamic Resource Modeling for Heterogeneous Wireless Networks** (with D. Tsamis, T. Alpcan, J. Singh). IEEE Int. Conf. on Communications (ICC), Dresden, Germany, pp. 1-6, 2009.
121. **Backlog aware scheduling for ingress memories in high-radix, single-stage switches** (with D. Tsamis, B. Yolken, W. Olesinski, H. Eberle, N. Gura). IEEE Global Com. Conference (Globecom), Honolulu, HI, pp. 1-6, 2009.
122. **Cost and Target-Based Scheduling for Switch Power Control** (with B. Yolken, D. Tsamis). IEEE International Conference on Communications (ICC), Dresden, Germany, 2009.
123. **Modeling Dependencies in Security Risk Management** (with T. Alpcan) ACM Int. Conf. on Risk and Security of Internet and Systems (CRiSIS), Toulouse, France, 6 pgs, 2009.
124. **A Characterization of Max-Min Power Allocation with Applications** (with Slawomir Stanczak, Michal Kaliszan, Marcin Wiczanski). IEEE Int. Symposium on Information Theory, Seoul, Korea, 6 pgs, June 2009.
125. **Maximizing Throughput of Hospital Intensive Care Units with Patient Readmissions** (with C. W. Chan, V. F. Farias, G. J. Escobar). INFORMS Manufacturing and Service Operations Management (MSOM) Conference, Boston, MA, 5 pgs, June 2009.
126. **Target-Based Power Control for Queuing Systems with Applications to Packet Switches** (with B. Yolken, D. Tsamis). IEEE Global Telecom. Conference (Globecom), New Orleans, pp. 1-6, Dec. 2008.
127. **Security Investment Games in Interdependent Organizations** (with A. Miura, A. Yolken, J. Mitchell), Proc. 46th Allerton Conference on Communication, Control and Computing, pp. 1-10, Sept. 2008.
128. **Game Based Admission Control for Wireless Systems** (with B. Yolken) International Wireless Internet Conference (WICON), pp. 1-7, Nov. 2008.
129. **Admission Control for Power-Controlled Wireless Networks under General Interference Functions** (with S. Stanczak, M. Kaliszan), 42nd Asilomar Conference on Signals, Systems and Computers, *Monterey, USA, Oct. 2008*.

130. **Fast Power-Up Active Link Protection in Autonomous Distributed Transmitter Power Control** (with S. Gitzenis), Proc. IEEE International Wireless Communication and Computing Conference (IWCMC), pp. 1-6, Chania, Greece, Aug. 2008.
131. **Game Based Capacity Allocation for Utility Computing Environments** (with B. Yolken), Proc. GameComm, pg. 1-10, Athens, Greece, Oct. 2008.
132. **Network-Assisted Wireless Computing** (with C. Chan and J. P. Singh), Proc. IEEE PIMRC, 2008.
133. **Target-Based Power Control for Queueing Systems with Applications to Packet Switches** (with B. Yolken and D. Tsamis), IEEE Globecom, New Orleans, 2008.
134. **Security Decision-Making amongst Interdependent Organizations** (with A. Miura, B. Yolken, J. Mitchell), Proc. IEEE Computer Security Foundations Symposium (CSF), Pittsburgh, PA, 2008
135. **Wireless Broadcasting to Diverse Users** (with C.W. Chan, S. Wee, J. Apostolopoulos). Proc. IEEE International Conference on Communications, Beijing, China, 2008.
136. **Receiver Based Optimization of Video Delivery over Wireless Links** (with C.W. Chan, S. Wee, J. Apostolopoulos), Proc. IEEE International Conference on Multimedia, 2008.
137. **Backlog Aware Scheduling for Large Buffered Crossbar Switches** (A. Dua, B. Yolken, W. Olesinski, H. Eberle, N. Gura). Proc. IEEE International Communications Conference, Beijing, China, May 2008.
138. **Target-Driven and Incentive-Aligned Power Control for Wireless Networks** (with B. Yolken), Proc. IEEE Globecom, Washington, DC, Nov. 2007.
139. **Optimal Scheduling of Media Packets with Multiple Distortion Measures** (with C.W. Chan, S. Wee, J. Apostolopoulos), Proc. IEEE International Conference on Multimedia, Beijing, China, June 2007.
140. **Power Aware Management of Packet Switches** (with L. Mastroleon, D. O'Neill, B. Yolken), Proc. IEEE Hot Interconnects, Stanford, CA, Aug. 2007.
141. **Dynamic Control Approach to Risk Mitigation in Computing Infrastructures** (with A. Miura-Ko). Proc. IEEE International Symposium on Information Assurance and Security, Manchester, UK, Aug. 2007.
142. **Backlog Aware Low Complexity Schedulers for Input Queued Packet Switches** (with A. Dua, W. Olesinski, H. Eberle, N. Gura). Proc. IEEE Hot Interconnects, pp. 39-46, Stanford, CA, 2007.
143. **Performance Tradeoffs in Mobile Computing: To Fetch or Not to Fetch?** (with A. Dua, J. Singh) Proc. ACM MobiWac, Crete, Greece, Oct. 2007.
144. **Buffer Management in Wireless Media Streaming** (with A. Dua). Proc. IEEE Globecom, Washington, DC, Dec. 2007
145. **Distributed Backlog-Driven Power Control in Wireless Networking** (with A. Dua). Proc. IEEE LanMan, Princeton, NJ, June 2007.
146. **Power Management of Packet Switches via Differentiated Delay Targets** (with B. Yolken). Proc. IEEE International Communications Conference, Beijing, China, May 2007.

147. **Power-Managed Packet Switching** (with A. Dua, B. Yolken). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007.
148. **Wireless Packet Scheduling with Soft Deadlines** (with A. Dua). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007
149. **SecureRank: An Efficient Algorithm for Prioritizing Vulnerabilities in Computing Systems** (with R. A. Miura-Ko). Proceedings of IEEE International Communications Conference (ICC'07), Glasgow, UK, 2007.
150. **Job Scheduling for Maximal Throughput in Autonomic Computing Systems** (with K. Ross), Proc. of International Workshop of Self-Organizing Systems (IWSOS), Passau, Germany, September 2006 (and in Lecture Notes in Computer Science 4124 Springer 2006).
151. **Capacity Maximizing Packet Scheduling Algorithms for Interconnection Networks with Finite Buffers** (with K. Cross). IEEE Global Communications Conference (GlobeCom), San Francisco, CA, 2006.
152. **Receiver-Based Optimization for Video Delivery over Wireless Links** (with C. Chan, J. Apostolopoulos, Y. Li). Proceedings of IEEE International Conference on Multimedia (ICME'06), Toronto, Ontario, July 2006.
153. **Optimal Throughput and Routing for Wireless Link Arrays** (with F. Baccelli, C. Chan). Proceedings of Infocom 2006, 12 pages, Barcelona, Spain, April 2006.
154. **Joint Power Allocation and Scheduling for Deadline Constrained Wireless Traffic** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, San Francisco, Dec. 2006.
155. **Low-Jitter Scheduling Algorithms for Deadline-Aware Packet Switches** (with A. Dua) Low-Jitter Scheduling Algorithms for Deadline-Aware Packet Switches. Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, Dec. 2006.
156. **Patching Rate Management for Controlled Service Disruption in Data Centers** (with L. Mastroleon, R. A. Miura-Ko). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'06), San Francisco, CA, San Francisco, Dec. 2006.
157. **Packet Transmission and Content-Dependent Payout for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of IEEE Multimedia Signal Processing (MMSP'05), 10 pages, Shanghai, China, Oct. 2005.
158. **Power Management in Battery-Constrained Handheld Devices**, (with A. Markopoulou, Y. Li, C. Chan) Proceedings of IEEE International Conference on Broadband Networks (BroadNets'05), pp. 441-450, Boston, MA, Oct. 2005.
159. **Dynamic Risk Mitigation for 'Self-Defending' Network Security**, SecureComm, Athens, May 2005.
160. **On the Fairness-Delay Tradeoff in Wireless Packet Scheduling** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'05), v. 5, pp. 2544-2548, St. Louis, MO, Nov. 2005.
161. **Power Control and QoS Tradeoffs for Real-Time Wireless Traffic** (with A. Dua). Proceedings of IEEE Wireless Communications and Networking Conference (WCNC'06), vol. 2, pp. 1148-1153, Las

Vegas, NV, April 2006.

162. **Downlink Scheduling of Heterogeneous Traffic** (with A. Dua). Proceedings of IEEE International Conference on Communications (ICC'06), Instabul, Turkey, June 2006.
163. **Scheduling with Soft-Deadlines for Input-Queued Switches** (with A. Dua). Proc. of the 44th Allerton Conference on Communication, Computation and Control, Allerton, IL, September 2006.
164. **Cross-Layer Multi-hop Wireless Routing for Inter-vehicle Communication** (with J.P. Singh, B. Srinivasan, D. Clawin) Proceedings of IEEE International Conference on Network Testbeds and Research Infrastructures (TridentCom'06), Barcelona, Spain, March 2006.
165. **Autonomic Power Management Schemes for Internet Servers and Data Centers** (with L. Mastroleon, C. Kozyrakis, D. Economou) Proceedings of IEEE Global Communications Conference (GlobeCom'05), v. 2, pp. 943-947, St. Louis, MO, Nov. 2005.
166. **Power-Managed Block Level File Decryption in Wireless Network Computing** (with S.Gitzenis), Proceedings of IEEE International Symposium on Modeling and Optimization in Mobile, Ad-Hoc, and Wireless Networks (WiOpt'06), Boston, MA, April 2006.
167. **On the Singular Behavior of a Queuing System with Random Connectivity** (with G. Michailidis). Proceedings of the IEEE Conference on Decision and Control (CDC'05) , pp. 5348-5353, Seville, Spain, Dec. 2005
168. **A Control Formulation of the Network Security Problem via a Risk Management Approach.** Proceedings of the 2005 IEEE International Conference on Security and Privacy for Emerging Areas in Communication Networks, 3 pages, Athens, Greece, Sept. 2005.
169. **Packet Transmission and Content-Dependent Payout for Video Streaming over Wireless** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of IEEE Multimedia Signal Processing (MMSP'05), 10 pages, Shanghai, China, Oct. 2005.
170. **Power Management in Battery-Constrained Handheld Devices**, (with A. Markopoulou, Y. Li, C. Chan) Proceedings of IEEE International Conference on Broadband Networks (BroadNets'05), pp. 441-450, Boston, MA, Oct. 2005.
171. **Joint Power/Payout Control Schemes for Media Streaming over Wireless Links** (with Y. Li, A. Markopoulou, J. Apostolopoulos). Proceedings of 14th IEEE Packet Video Workshop, 12 pages, Irvine, CA, Dec. 2004.
172. **Optimizing Timeout-Based Sleep Algorithms** (with A. Markopoulou). IEEE/ACM International Symposium on Information Processing in Sensor Networks (IPSN'05), UCLA, Los Angeles, CA, April 2005.
173. **Deadline Constrained Packet Scheduling in Wireless Networks** (with A. Dua). Proceedings of IEEE Vehicular Technology Conference VTC-F'05, v. 1, pp. 196-200, Dallas, TX, Sept. 2005.
174. **On the Fairness-Delay Tradeoff in Wireless Packet Scheduling** (with A. Dua). Proceeding of IEEE Global Telecommunications Conference (GlobeCom'05), v. 5, pp. 2544-2548, St. Louis, MO, Nov. 2005.
175. **Channel-State-Awareness Based Transmission Power Adaptation for Efficient TCP Dynamics in Wireless Networks** (with J.P. Singh, Y. Li) Proceedings of IEEE International Conference on

Communications (ICC'05), v. 5, pp. 3553-3559, Seoul, Korea, March 2005.

176. **Empirical Observations on Wireless LAN Performance in Vehicular Traffic Scenarios and Link Connectivity Based Enhancements for Multihop Routing** (with J.P. Singh, B. Shrinivasan, D. Clawin, Y. Yan). Proceedings of IEEE Wireless Communications and Networking Conference (WCNC'05), v. 3, pp. 1676-1682, New Orleans, LA, March 2005
177. **Transmission Power and Payout Control for Packetized Audio Streaming over Wireless Links**, (with Y. Li) Proceedings of IEEE International Conference on Wireless Networks, Communications and Mobile Computing, Workshop on Mobility Management and Wireless Access (MobiWac'05), v. 1, pp. 55-62, Maui, Hawaii, June 2005.
178. **Autonomic Power Management Schemes for Internet Servers and Data Centers** (with L. Mastroleon, C. Kozyrakis, D. Economou) Proceedings of IEEE Global Communications Conference (GlobeCom'05), v. 2, pp. 943-947, St. Louis, MO, Nov. 2005.
179. **Media and Data Traffic Coexistence in Power-Controlled Wireless Networks** (with S. Gitzenis). Proceedings of ACM Wireless Multimedia Networking and Performance Modeling (WMuNeP'05), pp. 78-85, Montreal, Quebec, CA, Oct. 2005.
180. **Dynamic Quality of Service Control in Packet Switch Scheduling** (with K. Ross). Proceedings of IEEE International Conference on Communications (ICC'05), v. 1, pp. 396-401, Seoul, Korea, May 2005.
181. **Packet Scheduling Across Networks of Switches** (with K. Ross) Proceedings of IEEE International Conference on Networking (ICN'05), Reunion Island, April 2005.
182. **On the Singular Behavior of a Queuing System with Random Connectivity** (with G. Michailidis). Proceedings of the IEEE Conference on Decision and Control (CDC'05), pp. 5348-5353, Seville, Spain, Dec. 2005
183. **Efficient Data Prefetching for Power-Controlled Wireless Packet Networks** (with S. Gitzenis). Proceedings of ACM/IEEE Mobile and Ubiquitous Networking Conference (MobiQuitous), pp.64-73, Boston, MA, August 22-25, 2004.
184. **Power-Controlled Packet Relays in Wireless Data Networks** (with S. Gitzenis). Proceedings of IEEE Global Communications Conference (GlobeCom), vol. 1, pp. 464-469, San Francisco, CA, December 1-5, 2003.
185. **Integrated Power Control for Circuit-Switched & Packet-Switched Traffic in Wireless Networks** (with S. Gitzenis). Proceedings of the IEEE Vehicular Technology Conference (VTC), vol. 4, pp. 2604-2609, Orlando, FL, Fall 2003.
186. **Power Control for TCP Adaptation to High-mobility Broadband Systems** (with J.P. Singh, K. Radermacher, V. Scharf-Katz). Proceedings of IEEE Vehicular Technology Conference (VTC), Los Angeles, CA, 9 pages, Fall 2004.
187. **Power-Controlled Wireless Links for Media Streaming Applications** (with Y. Li). Proceedings of the IEEE Wireless Telecommunications Symposium 2004, pp. 102-111, Pomona, CA, May 14-15, 2004
188. **Power-Controlled Media Streaming in Interference-Limited Wireless Networks** (with Y. Li). Proceedings of the 2004 IEEE International Conference in Broadband Networks, pp. 560-568, San

Jose, CA, October 25-29, 2004.

189. **Power-Controlled Wireless Links for Video Streaming Applications** (with Y. Li). Proceedings of the IEEE Vehicular Technology Conference (VTC), Los Angeles, CA, 9 pages, Sept. 26-29 2004.
190. **Local Search Scheduling Algorithms for Maximal Throughput in Packet Switches** (with K. Ross). Proceedings of the IEEE INFOCOM 2004, Hong Kong, China, 12 pages, March 7-11, 2004.
191. **Optimizing Quality of Service in Prioritized Packet Switch Scheduling** (with K. Ross). Proceedings of the IEEE International Conference on Communications (ICC), Paris, France, 5 pages, June 20-24, 2004.
192. **A Service Risk Management Approach to Capacity Protection in Optical Networks** (N. Bambos, S. Gitzenis, A. Miura, O. Gerstel, L. Paraschis). Proceedings of IEEE Workshop in Local and Metropolitan Area Networks (LANMAN'04), Mill Valley, CA, pp. 69-74, April 2004.
193. **Mobile-to-Base Task Migration in Wireless Computing** (S. Gitzenis, N. Bambos). Proceedings of IEEE Pervasive Communications Conference *PerCom* 2004, pp. 187-196, Orlando, FL, March 14-17, 2004.
194. **Distributed Power and Admission Control for Wireless Networks** (with T. Holliday, A. J. Goldsmith, P. Glynn), ISIT 2004.
195. **Dynamic Scheduling of Optical Data Bursts in Time-Domain Wavelength Interleaved Networks** (Kevin Ross, K. Kumaran, I. Saniee, I. Widjaja). Proceedings of the 11th IEEE Symposium on High-Performance Interconnects, Stanford, CA, pp. 108-113, 2003.
196. **Local/Remote Task Execution in Network Computing** (with Savvas Gitzenis). Proceedings of the International Conference on Internet Computing, IC '03, Las Vegas, NV, pp. 842-848, June 2003.
197. **Integrated Power Control for Circuit-Switched and Packet-Switched Traffic in Wireless Networks** (with S. Gitzenis). Proceedings of the IEEE Fall Vehicular Technology Conference VTC'03, Orlando, FL, vol. 4, pp. 2604-2609, Oct. 2003.
198. **Proposal and Demonstration of Link-Connectivity-Assessment-Based Enhancements to Routing in Mobile Ad-Hoc Networks** (with J. Singh, D. Clawin, B. Shrinivasan). Proceedings of the IEEE Fall Vehicular Technology Conference VTC'03, Orlando, FL, Oct. 2003.
199. **Power Management of Packet Switch Architectures with Speed Modes** (with D. O'Neill). Proceedings of the 41st Annual Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, Monticello, IL, Oct. 2003. (invited paper)
200. **Distributed Power Control for Time Varying Wireless Networks: Optimality and Convergence** (with T. Holliday, A. J. Goldsmith, P. Glynn), Proceedings of the 41st Annual Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, Monticello, IL, Oct. 2003. (invited paper)
201. **Capacity Protection/Restoration in Tree Optical Networks for Enterprise Content Distribution** (with S. Gitzenis, O. Gertsel, L. Paraschis). Proceedings of the Workshop "Protection and Restoration: from SONET/SDH to Next-Generation Networks" of IEEE Global Communications Conference (GlobeCom '03), San Francisco, CA, 2003.
202. **Power-Controlled Packet Relays in Wireless Data Networks** (with Savvas Gitzenis). Proceedings of the IEEE Global Communications International Conference (GlobeCom '03), San Francisco, CA, vol.

- 1, pp. 464-469, Dec. 2003.
203. **Learning Policies for Power Control in Wireless Networks.** (with D. Vengerov, H. Berenji). Proceedings of the IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 2390-2394, Oct. 2002.
 204. **Power Efficient MAC Scheme Using Channel Probing in Multirate Wireless Ad-Hoc Networks** (with JW. Kim). Proceedings of IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 2382-2386, Oct. 2002.
 205. **Geometry of Scheduling Algorithms for Packet Switches** (with K. Ross). Proceedings of the Allerton Conference on Communication, Control and Computing, University of University at Urbana-Champaign, Monticello, Illinois, Oct. 2002 (invited paper)
 206. **Wireless LAN Performance Under Varied Stress Conditions in Vehicular Traffic Scenarios** (with J. Singh, B. Shrinivasan, D. Clawin). Proceedings of the IEEE Vehicular Technology Conference, Vancouver, BC, Canada, pp. 743-747, Oct. 2002.
 207. **A Practical Adaptive Algorithm for Power Control in Wireless Networks** (with D. Vengerov, H. Berenji). Proceedings of the 2002 Symposium on Autonomous Intelligent Networks and Systems, University of California at Los Angeles, Los Angeles, CA, May 2002.
 208. **Power-Controlled Data Prefetching/Caching in Wireless Data Networks** (with S. Gitzenis). Proceedings of the 20th IEEE Conference on Computer Communications (INFOCOM 2002), New York, NY, vol. 3, pp. 1405-1414, June 2002.
 209. **Power Control for Multirate Wireless Networks with Groupwise Serial Multiuser Detection** (with J-W. Kim). Proceedings of the 2001 IEEE Global Communications Conference (GLOBECOM 2001), San Antonio, Texas, vol. 5, pp. 3201-3205, November 2001.
 210. **Multi-Modal Dynamic Multiple Access (MDMA) in Wireless Packet Networks** (with S. Kandukuri). Proceedings of the 20th IEEE Conference on Computer Communications (INFOCOM 2001), Anchorage, Alaska, vol. 1, pp. 199-208, April 2001.
 211. **Power Controlled Multiple Access (PCMA) in Wireless Communication Networks** (with S. Kandukuri). Proceedings of the 19th IEEE Conference on Computer Communications (INFOCOM 2000), Tel Aviv, Israel, vol. 2, pp. 386-395, March 2000.
 212. **Multi-Channel Power Control for Data Traffic in Wireless Networks** (with S. Kandukuri). Proceedings of the 1999 IEEE International Workshop on Mobile Multimedia Communications (MoMuC '99), pp. 83-92, Nov. 1999.
 213. **Queueing Networks with Interacting Service Resources** (with M. Armony). Proceedings of the 37th Allerton Conference on Communication, Control and Computing. Monticello, Illinois. Sept. 1999. (invited paper)
 214. **Routing in Networks with Random Topologies** (with K. Scott). Proceedings of the 1997 International Communications Conference (ICC '97), vol. 2, pp. 862-866, 1997.
 215. **Power Control and Time Division Multiple Access: The TDMA vs. CDMA Question** (with J. Rulnick). Proceedings of the 1997 IEEE Conference on Computer Communications (INFOCOM '97), Kobe, Japan, vol. 2, pp. 635-642, April 1997.

216. **Modeling Communication Links in Percolating Environments** (with N. Alexopoulos). Proceedings of the 1997 IEEE International Conference on Advances in Communications and Control (COMCON '97), pp. 249-261, 1997.
217. **Formation and Maintenance of Self-Organizing Wireless Networks** (K. Scott). Proceedings of the 31st Asilomar Conference on Signals, Systems and Computers, Monterey, CA, vol. 1, pp. 31-35, Nov. 1997.
218. **Algorithmic Speed-up of All Pairs Shortest paths Computation with Reconfigurable Optical Interconnection** (with S.M.P. Yip). Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '96), San Jose, CA, vol. 1, pp. 355-366, August 1996.
219. **Performance Evaluation of Power-Managed Mobile Communication Devices** (with J. Rulnick). Proceedings of the 1996 International Conference on Communications (ICC '96), Dallas, Texas, vol. 3, pp. 1477-1481, June 1996.
220. **Scalable Communication Schemes for Massively Parallel Processing Using Reconfigurable Optical Interconnects** (with S.M.P. Yip). Proceedings of the 1996 International Conference on Parallel & Distributed Systems (ICPADS '96), Tokyo, Japan, pp. 500-507, June 1996.
221. **Best Effort Bandwidth Reservation in High Speed LANs Using Wormhole Routing** (with B. Kwan, P-C. Hu, L. Kleinrock, Hong Xu, Joseph Touch). Proceeding of 5th IEEE International Conference on Computer Communications and Networks (ICCCN '96), Rockville, Maryland, 1996.
222. **Mobile Power Management for Maximum Battery Life in Wireless Communication Networks** (with J. Rulnick). Proceedings of the 1996 IEEE Conference on Computer Communications (INFOCOM '96), San Francisco, CA, vol. 2, pp. 443-450, March 1996.
223. **The Supercomputer Supernet (SSN): A High-Speed Electro-Optical Campus and Metropolitan Network** (with L. Kleinrock, J. Bannister, L. Bergman, J. Cong, E. Gafni, M. Gerla). Proceedings of the 1996 IEEE/ SPIE Conference on Optical Interconnects in Broadband Switching Architectures, San Jose, pp. 22-33, Feb. 1996.
224. **Optimal Server Allocation to Parallel Queues with Randomly Modulated Service Rates** (with G. Michailidis). Proceedings of the Conference on Information Sciences and Systems, Princeton University, Princeton, New Jersey, pp. 692-698, March 1996.
225. **Routing and Channel Assignment for Low Power Transmission in Personal Communication Systems** (with K. Scott). 1996 IEEE International Conference on Universal Personal Communications, vol. 2, pp. 498-502, 1996.
226. **On the Stationary Dynamics of Parallel Queues with Random Server Connectivities** (with G. Michailidis). Proceedings of the 34th IEEE Conference on Decision and Control (CDC '95), New Orleans, LA, vol. 4, pp. 3638-3643, Dec. 1995.
227. **The Self-Organizing Wireless Adaptive Network (SWAN) Protocol for Communication Among Mobile Users** (with K. Scott). Proceedings of the 1995 IEEE Global Telecommunications Conference (GLOBECOM '95), Singapore, vol. 1, pp. 355-359, Nov. 1995.
228. **Channel Probing for Distributed Access Control in Wireless Communication Networks** (with S. Chen, D. Mitra). Proceedings of the 1995 IEEE Global Telecommunications Conference (GLOBECOM '95), Singapore, v. 1, pp.322-326, November 1995.

229. **Radio Link Admission Algorithms for Wireless Networks with Power Control and Active Link Quality Protection** (with S. Chen, G. Pottie). Proceedings of the 1995 IEEE Conference on Computer Communications (INFOCOM '95), Boston, MA, v. 1, pp. 97-104, June 1995.
230. **Entropy Methods for High Speed Communications** (with B. Prabhakar). Proceedings of the 1995 Conference on Information Sciences and Systems, The Johns Hopkins University, Baltimore, Maryland, , pp. 448-453, March 1995.
231. **Adaptive (T1/T2)-Multiplexing Transmission Schemes for Voice/Data Networks** (with A. Nguyen, M. Sherif) Proceedings of the 1995 IEEE International Symposium on Computers and Communications (ISCC '95), pp. 430-435, June 1995.
232. **A Distributed, Mobile, Wireless Infrastructure for Multimedia Applications** (with S. Chen, M. Gerla, J. Tsai). Proceedings of the 1995 WINLAB Conference, Rutgers University, East Brunswick, NJ, March 1995.
233. **Admission Control Schemes for Wireless Communication Networks with Adjustable Transmitter Powers** (with S. Chen, G. Pottie). Proceedings of the 1994 IEEE Conference on Computer Communications (INFOCOM '94), Toronto, Canada, vol. 1, pp. 21-28, May 1994.
234. **On Distributed Power Control for Radio Networks** (with S. Chen, G. Pottie). Proceedings of the 1994 International Communications Conference (ICC '94), New Orleans, LA, vol. 3, pp. 1281-1285, March 1994.
235. **On Power Control with Active Link Quality Protection in Wireless Communication Networks** (with S. Chen and G. Pottie). Proceedings of the 1994 Conference on Information Sciences and Systems, Princeton, NJ, pp. 426-467, March 1994.
236. **The Asymptotics of Traffic Processes in Large Queueing Networks** (with T. Mountford and B. Prabhakar). Proceedings of the 1994 Allerton Conference on Communication, Control and Computing, University of Illinois at Urbana-Champaign, Urbana, Illinois, pp. 563-572, Oct. 1994.
237. **Queueing Dynamics and Throughput of Ring Networks with Spatial Reuse** (with A. Nguyen). Proceedings of the 1993 Allerton Conference on Communication, Control and Computing, University of Illinois at Urbana-Champaign, Urbana, Illinois, pp. 576-587, Oct. 1993 (invited paper).
238. **On Parallel Processing of Random Graph Structured Jobs** (with S. Chen). Proceedings of the 1993 NSF Design and Manufacturing Systems Conference, University of North Carolina at Charlotte, Charlotte, NC, pp. 3-11, Jan. 1993.
239. **Power Control Based Admission Policies in Cellular Radio Networks** (with G. Pottie). Proceedings of the 1992 IEEE Global Telecommunications Conference (GLOBECOM '92), Orlando, Florida, vol. 2, pp. 863-867, Dec. 1992.
240. **Optimal Message Flow on Ring Networks with Spatial Reuse** (with A. Nguyen). Proceedings of the 31st IEEE Conference on Decision and Control (CDC '92), Tuscon, AZ, vol. 2, pp. 2360-2363, Dec. 1992.
241. **On Power Control in High Capacity Cellular Radio Networks** (with G. Pottie). Proceedings of the 1992 WINLAB Conference, Rutgers University, East Brunswick, New Jersey, pp. 239-247, April 1992.
242. **On Parallel Processing in Tandem Queueing Networks with Feedback**. Proceedings of the 1992 NSF Design and Manufacturing Systems Conference, Georgia Institute of Technology, Atlanta, GA,

pp. 495-498, Jan. 1992.

- 243. **On Tandem Queueing Networks with Feedback** (with K. Wasserman). Proceedings of the 1991 Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, pp. 310-312, Oct. 1991.
- 244. **Queueing Theoretic Analysis of a General Class of Concurrent Processing Systems**. Proceedings of the 1991 NSF Design and Manufacturing Systems Conference, The University of Texas at Austin, Texas, pp. 1-5, January 1991.
- 245. **Maximal Throughput for Stability of a Class of Parallel Processing Systems** (with J. Walrand). Proceedings of the 29th IEEE Conference on Decision and Control (CDC '90), Honolulu, Hawaii, vol. 1, pp. 161-166, Dec. 1990.
- 246. **Flow Analysis of Stochastic Marked Graphs** (with F. Baccelli, J. Walrand). Proceedings of the 28th IEEE Conference on Decision and Control (CDC '89), Tampa, Florida, vol. 2, pp. 1528-1531, Dec. 1989.
- 247. **On Bifurcated Shortest Path Routing in Communication Networks** (with G. Stassinopoulos). Proceedings of the 1984 IEEE European Conference on Control and Communications, Patras, pp. 604-609, July 1984.

REPORTS:

- 248. **ARQ: A Multiclass Admission Control Protocol for Heterogeneous Datacenters** (with C. Delimitrou, C. Kozyrakis). Stanford Technical Multiscale Architecture & Systems Lab. Stanford Technical Report, TR-2013-002, Jan. 2013.
- 249. **Adaptive Data-Aware Utility-Based Scheduling in Resource-Constrained Systems** (with D. Vengerov, L. Mastroleon, D. Murhpy). Sun Microsystems Labs Technical Report SMLI TR-2007-164, April 2007.
- 250. **Architecting Novel Information Services and Markets**, Detecon Management Report, 2006.
- 251. **Wireless Packet Scheduling with Deadlines**. Stanford Technical Report, 2005
- 252. **Projective Processing Schedules in Queueing Structures: Applications to Packet Scheduling in Communication Networks Switches**. Stanford Technical Report. 2002.
- 253. **Network Security and Infrastructure Assurance – Architectural and Performance Engineering Issues**. Stanford Technical Report, 2004
- 254. **Noninvasive Channel Probing for Distributed Admission Control and Channel Allocation in Wireless Networks** (with J-W. Kim, D. Mitra, S. Chen). Stanford Technical Report, 2002

255. **Power-Controlled Data Prefetching and Caching in Wireless Mobile Networking** (with S. Gitzenis). Stanford Technical Report, 2002.
256. **Queueing Networks with Random Link Topologies; Stationary Dynamics of Maximal Throughput Schedules** (with G. Michailidis). Stanford Technical Report, 2001.
257. **On Stationary Multi-Processor Queues with Threshold-Based Server Activation** (with J. Rulnick). Technical Report: NetLab 1997-09-01, Stanford University, 1997.
258. **Convergence of Departures in Tandem Networks of $\cdot/GI/\infty$ Queues** (with T. Mountford, B. Prabhakar). Technical Report: HPL-BRIMS-96-15, BRIMS, HP Labs, UK, 1996.
259. **Processor Migration in Distributed Queueing: An Analysis of the two Processor Case** (with K. Scott). Technical Report UCLA-ENG-95-138, School of Engineering and Applied Science, UCLA, 1995.
260. **Multicopy Parallel Processing Networks: Throughput Analysis** (with K. Scott). Technical Report UCLA-ENG-95-137, School of Engineering and Applied Science, UCLA, 1995.
261. **Critical Phenomena in Massively Parallel Processing Networks: Throughput Analysis - Initial Results** (with A. Zarkesh). Technical Report UCLA-ENG-95-141, School of Engineering and Applied Science, UCLA, 1995.
262. **A Fully Distributed Algorithm for the Creation and Maintenance of Wireless Networks of Mobile Users** (with K. Scott). Technical Report UCLA-ENG-95-112, Department of Electrical Engineering, University of California at Los Angeles, 1995.
263. **Probing the Collective Behavior of Massively Parallel Processing Systems by the Renormalization Group Method** (with A. Zarkesh). Technical Report UCLA-ENG-95-140, School of Engineering and Applied Science, UCLA, 1995.
264. **Performance of a Switch State Dependent Timeout Scheme in a Wormhole Switching LAN** (with B. Kwan). UCLA Technical Report UCLA-ENG-95-121, School of Engineering and Applied Science, UCLA, 1995.
265. **Branching Processes in Rollback Synchronization of Parallel Processing Systems: Structural Results** (with J. Rulnick). Technical Report UCLA-ENG-95-136, School of Engineering and Applied Science, UCLA, 1995.
266. **Optimal Choice of Queue Length Thresholds for Processor Activation in the $M/M/K$ Queue** (with J.M. Rulnick). Technical Report UCLA-ENG-95-134, School of Engineering and Applied Science, UCLA, 1995.
267. **Partial Worm Retraction Routing Scheme for Wormhole Switching LANs** (with B. Kwan). UCLA Technical Report UCLA ENG-94-108, November 1994.
268. **Admission Control and Routing Protocols in Stationless Multi-Hop Radio Networks** (with K. Scott). Technical Report UCLA-ENG-94-106, School of Engineering and Applied Science, U.C. Los Angeles, 1994.
269. **Optimistic (Partial Retraction) Routing Schemes for Wormhole Switching LANs** (with B. Kwan). Technical Report UCLA-ENG-94-108, School of Engineering and Applied Science, U.C. Los Angeles, 1994.

270. **Radio Link Admission Algorithms for Wireless Networks with Power Control and Active Link Quality Protection** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-94-25. School of Engineering and Applied Science, U.C. Los Angeles, 1994.
271. **Distributed Dynamic Power Control Schemes in High Capacity Wireless Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-57, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
272. **A Localized Power and Admission Control Scheme Adapted to Varying Channel Conditions for High Capacity Wireless Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-56, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
273. **Admission Control Schemes for Guaranteed Link Quality in High Capacity Cellular Radio Networks** (with S. Chen, G. Pottie). Technical Report UCLA-ENG-93-55, School of Engineering and Applied Science, U.C. Los Angeles, 1993.
274. **On Generalized Multiserver Queues** (with J. Walrand). Technical Report UCLA-ENG-9-24-90, School of Engineering and Applied Science, U.C. Los Angeles, 1990.
275. **Recent Results on the Asymptotic Behavior of Concurrent Processing Systems Under Stationary and Ergodic Inputs** (with J. Walrand). Technical Report UCB/ERL M88/78, Electronics Research Laboratory, U.C. Berkeley, 1988.
276. **An Alternative Approach to the Asymptotic Makespan of Serially Executed Multitasked Jobs** (with J. Walrand). Technical Report UCB/ERL M88/79, Electronics Research Laboratory, U.C. Berkeley, 1988.

TALKS:

1. **Invited Seminar – A Risk Management Approach to Cyber-Security**, Operations Research Program, School of Engineering, University of Texas at Austin, March 2017.
2. **Predicting Pediatric Surgical Case Durations** (published abstract with N. Master, D. Scheinker). INFORMS International Meeting, Hawaii, 2016.
3. **Myopic Scheduling of Jobs with Decaying Value with Applications**, (abstract with N. Master, C.W. Chan) INFORMS International Meeting, Hawaii, 2016. (content essentially the same as of talk below.)
4. **Games on Influence Networks: Equilibria, Free Riding and Dynamics** (with Z. Zhou, P. Glynn). INFORMS Annual Meeting, Nashville, Tennessee, 2016.
5. **Myopic Scheduling of Jobs with Decaying Value with Applications in Patient Scheduling** (published abstract with N. Master, C.W. Chan). INFORMS Annual Meeting, Philadelphia, PA, 2015.
6. **Distinguished Invited Seminar: Power and Throughput Issues in Next-Generation Packet Switches**, National Technical University of Athens, Greece, Feb. 2009.
7. **Distinguished Invited Seminar: Wireless Multimedia Computing**, DoCoMo Research Labs, Palo Alto, Feb. 2009.

8. **Plenary Talk: Architecting Novel Information Services and Markets**, 4th International IMS Workshop, Fraunhofer Institute for Open Communication Systems – FOKUS, Berlin, Germany, Nov. 2008.
9. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, NEC Research Labs, Tokyo, July 2008.
10. **Power Management in Networking and Computing**, Hitachi Labs, Tokyo, July 2008.
11. **Power Management of Packet Switches**, Sun Labs, Menlo Park, April 2008.
12. **Power and Throughput Issues in Next-Generation Packet Switches**, UC Irvine, February 2008.
13. Invited Distinguished Lectures: 1) **Robust Maximal Throughput Schedules in Processing Systems**, 2) **Transmitter Power Control in Wireless Multimedia Networking**, Advanced Network Sciences Institute, University of California, San Diego, August 2008.
14. **Autonomous Admission (and Power) Control in Wireless Networking**, Technical University of Berlin, December 2007.
15. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Samsung Labs, Santa Clara, December 2008.
16. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Stanford OR Seminar, November 2007.
17. **Throughput and Power Issues in Packet Switching**, University of Sydney, Australia, August 2007.
18. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Bogazici University, Istanbul, Turkey, August 2007.
19. **Outsourcing in Computing and Data Storage Services**, Deutsche Telekom Laboratories, Berlin, August 2007.
20. **Dynamic Control of Packet Switches**, Sun Labs, Menlo Park, February 2007.
21. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Deutsche Telekom Laboratories, Berlin, January 2007.
22. **Network Control Issues in Video Transmission and Distribution**, Deutsche Telekom Labs, August 2006.
23. **Power Control in Wireless Networks**, IBM Research Labs, Hawthorne, July 2006.
24. **Scheduling and Power Management Issues in Data Centers and High-Performance Computing**, Sun Labs, Menlo Park, May 2006.
25. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Athens Information Technology Institute, May 2006.
26. **Power Control in Wireless Networks**, ENS, Paris.

27. **Video Streaming over Wireless Networks – Control Architectures and Performance Engineering**, Technical University of Athens, December 2005.
28. **Geometry of Resource Allocation and Load Balancing in Computing/Networking**, Athens University of Economics and Business, November 2005.
29. **Power Control in Wireless Networks**, Sprint Labs, Burlingame, December 2004.
30. **‘Geometry’ of Packet Scheduling in Communication Switches**, MIT LIDS Seminar, November 2004.
31. **Power(Rate)/Playout-Controlled Video over Wireless**, France Telecom Research Labs, Paris, June 2004.
32. **Next-Generation Wireless Data/Media Networks**, Athens Information Technology Institute, February 2004.
33. **Projective Packet Scheduling in Switching Systems**, France Telecom Research Labs, Paris, February 2004.
34. **Next-Generation Wireless Data/Media Networks**, France Telecom Research Labs, February 2004.
35. **Geometry of Packet Scheduling Algorithms for Input-Queued Communication Switches**. INFORMS Conference, San Jose, CA, Nov. 2002.
36. **Adaptive Admission Control for Multiclass Web Servers** (with G. Paleologo). INFORMS Telecommunications Conference, Boca Raton, FL, March 2002.
37. **Power-Controlled Data (Pre)Fetching and Caching in Wireless Networking** (with S. Gitzenis). INFORMS Telecommunications Conference, FL, March 2002.
38. **On Stability of Queueing Networks with Job Deadlines** (with A. Ward). INFORMS Applied Probability Conference, New York City, NY, July 2001.
39. **Throughput Maximizing Properties of Statistical Multiplexing in Stationary Ergodic Networks with Feedback** (with C. Westphal). INFORMS Applied Probability Conference, New York City, NY, July 2001.
40. **Predictive Caching Techniques in Wireless Networks** (with A. Ward). INFORMS Telecommunications Conference, Boca Raton, FL, March 2000.
41. **Dynamic Cone Policies for Skilled Based Routing in Call Centers** (with M. Armony). INFORMS Telecommunications Conference, Boca Raton, FL, March 2000.
42. **Dynamic Cone Policies for Queueing Networks with Interfering Resources** (with M. Armony). INFORMS Conference, Philadelphia, PA, November 1999.
43. **Stability and Performance Issues of a General Heterogeneous Parallel Processing System** (with G. Michailidis, K. Wasserman, S. Akin). INFORMS Applied Probability Conference, Ulm, Germany, July 1999.
44. **Bandwidth Allocation in Wireless Networks** (with G. Michailidis). INFORMS Conference, Seattle, WA, October 1998.

45. **Stability and Large Buffer Asymptotics of a General Parallel Processing System** (with G. Michailidis, K. Wasserman, S. Akin). Conference of the Canadian Operational Research Society, Windsor, Ontario, June 1999.
46. **Dynamic Cone Policies for Queueing Networks with Interfering Resources** (with M. Armony). INFORMS Conference, Philadelphia, PA, November 1999.
47. **Power Control and QoS in Wireless Communication Networks** (with M. Armony). INFORMS Conference, Cincinnati, OH, May 1999.
48. **Scheduling Algorithms for High Performance Switches** (with S. Yip). INFORMS Conference, Seattle, WA, Oct. 1998.
49. **Bandwidth Allocation in Wireless Networks** (with G. Michailidis). INFORMS Conferences., Seattle, WA, October 1998.
50. **Queueing in Interfering Resource Environments** (with M. Armony). INFORMS Conference, Tel Aviv, Israel, June 1998.
51. **Allocation of Interdependent Resources in Queueing Systems with Applications to Wireless Networks** (with M. Armony). INFORMS Telecommunications Conference, Boca Raton, FL, March 1998.
52. **On Performance and Limitations of Ultra Large Scale Digital Interconnects with Reconfigurable Topologies** (with S. Yip). IEEE/LEOS Conference on Interconnection of High Speed Digital Systems, Santa Fe, NM, May 1997.
53. **Queueing in Random Environments** (with G. Michailidis). INFORMS Applied Probability Conference, Boston, MA, June 1997.
54. **Network Flows in Random Environments** (with G. Michailidis). INFORMS Applied Probability Conference, Atlanta, GA, June 1995.
55. **On Queueing Networks with Random Link Topologies** (with G. Michailidis). INFORMS Telecommunications Conference, Boca Raton, FL, March 1995.
56. **Power Control Issues in Wireless Networks** (with S. Chen, G. Pottie). INFORMS Telecommunications Conference, Boca Raton, FL, March 1995.
57. **The Entropy and Delay of Traffic Processes in ATM Networks** (with B. Prabhakar). 1995 IEEE Information Theory Workshop on Multiple Access and Queueing, St. Louis, Missouri, April 1995.
58. **Scalable Supercomputer Network** (with L. Kleinrock, N. Bambos, L. Bergman, J. Cong, E. Gafni, M. Gerla). 1994 IEEE Gigabit Networking Workshop, Toronto, Canada, June 1994.
59. **The Supercomputer Supernet Hardware Design** (with L. Kleinrock, N. Bambos, L. Bergman, J. Cong, E. Gafni, M. Gerla). DARPA Networking Systems Meeting, Santa Fe, NM, September 1994.
60. **Dynamics of Parallel Processing: A Non-Causal Queue** (with J. Rulnick). Annual meeting of the Society of Industrial and Applied Mathematics (SIAM), San Diego, CA, July 1994.
61. **On the Synchronization of Poisson Processes and Queueing Networks with Synchronization Nodes** (with B. Prabhakar and T. Mountford). Fifth Workshop on Stochastic Analysis and Related Fields, Silivri,

Turkey, July 1994.

62. **On Flow Control in Ring Structured Local Area Networks with Spatial Reuse** (with A. Nguyen). INFORMS Telecommunications Conference, Boca Raton, Florida, March 1992.
63. **Stable Policies for Processing Systems with Precedence Constraints**. ORSA/TIMS Joint National Meeting, Anaheim, California, November 1991.
64. **On Stability of Parallel Processing Systems** (J. Walrand). SIAM Conference on Applied Probability in Science and Engineering, New Orleans, LA, March 1990.
65. **State-Dependent Queues and Queueing Networks** (with J. Walrand). International Conference on Stochastic Programming, Ann Arbor, MI, August 1989.
66. **Recent Developments in Power Control and their Impact on Next-Generation Dense Wireless Network** (invited distinguished talk). Wireless/Sensor Network Track of the 2002 IEEE Wireless Communications and Networking Conference (WCNC 2002), Orlando, Florida, March 2002.
67. **Power Control and Management Issues in Next-Generation Wireless Networks** (invited keynote talk). 2001 IEEE Symposium on Ad Hoc Wireless Networks (SAWN 2001) of the 2001 IEEE Global Communications Conference (GLOBECOM 2001), San Antonio, Texas, November 2001.
68. **Queueing Models and Problems in Wireless Network Control** (invited plenary talk). Stochastic Networks Conference, Madison, WI, June 2000.
69. **Power Controlled Multiple Access (PCMA) in Wireless Networking** (invited talk). IEEE Computer Communications Workshop, Estes Park, Colorado, November 1999.
70. **Queueing Networks with Interdependent Resources** (invited plenary talk). Oberwolfach Workshop on Applied Probability,, Mathematisches Forschungsinstitut, Oberwolfach, Germany, December 1998.
71. **Power Control Issues in Wireless Communication Networks** (invited plenary talk). Workshop on System and Control Issues in Communication Networks, National Science Foundation and Army Research Office, Airlie Center, Airlie, Virginia, August 1996.
72. **On Queueing Networks with Random Link Topologies: Bandwidth Allocation in Wireless Communication Networks** (invited plenary talk). BRIMS Workshop on Mathematical Trends in Communication Networks, Hewlett-Packard Labs, Bristol, UK, June 1996.
73. **Power Control with Active Link Protection in Wireless Networking** (invited plenary talk). EPFL Workshop on Computer Networking, EPFL, Switzerland, Summer 1996.
74. **Power Control Algorithms in High Performance Wireless Networks** (invited talk). Networking Seminar, IBM Watson Research Center, NY, July 2001.
75. **Queueing Networks in Random Environments** (invited talk). Networking Seminar, Dept. of EECS, U.C. Berkeley, March 1998.
76. **On Queueing Networks with Randomly Modulated Service Rates** (invited talk). Research Colloquium, Center for Operations Research, MIT, September 1997.
77. **High-Performance Wireless Networking**. Networking Seminar, Dept. of EECS, U.C. Berkeley, Nov. 1996

ISSUED PATENTS

Nicholas Bambos

October 22, 2018

1. **US 6,711,280 B2** 2004/2001 – Method and Apparatus for Intelligent Ranging via Image Subtraction.
2. **US 7,317,730 B1** 2008/2001 – Queueing Architecture and Load Balancing for Parallel Packet Processing in Communication Networks.
3. **US 7,634,287 B1** 2009/2001 – Power Controlled Multiple Access (PCMA) in Wireless Communication Networks.
4. **US 8,320,269 B2** 2012/2009 – Scalable Delay-Power Control Algorithm for Bandwidth Sharing in Wireless Networks.
5. **US 8,423,070 B2** 2013/2004 – Method and System to Model TCP Throughput, Assess Power Control Measures, and Compensate for Fading and Path Loss, for Highly Mobile Broadband Systems.
6. **US 8,488,500 B2** 2013/2008, **US 9,454,209 B2** 2016/2008, **EP 2,272,287 B1** 2017/2008 (European counterpart with English, French and German versions), **JP 5,997,794 B2** 2016/2008 (Japanese counterpart) – Power Management of Networked Devices – (and pending application **US 15/909,619**, filed on 2018-03-01, which is a continuation of pending application **US 15/276,573**, filed on 2016-09-26, which is a continuation of application **US 13/942,586**, filed on 2013-07-15 (now patent **US 9,457,209**), which is a continuation of application **US 12/114,721**, filed on 2008-05-02 (now patent **US 8,488,500**); and pending application **IN 6987/CHENP/2010** in India).
7. **US 9,614,958 B2** 2017/2007, **EP 2,132,921 B1** 2018/2007 (European counterpart) – Predictive Computer Network Services Provisioning for Mobile Users.

--- CONFIDENTIAL ---

Please note that most engagements identified below are governed by confidentiality agreements.

December 19, 2022

A) EXPERT WITNESS ENGAGEMENTS OF NICHOLAS BAMBOS

I, Nicholas Bambos, have had the following temporary expert witness engagements in the litigation cases listed below:

1. (2002-2003) Technical consultant and expert witness (*deposed*) for *Altima Communications, Inc.* in the case “Level One Communications, Inc. v. Altima Communications, Inc.” Case No. CIV. S-99-2488 GEB (GGH), US District Court, Eastern District of California
2. (2004-2005) Technical consultant and expert witness (*deposed*) for *Microsoft Corporation* in the litigation case “Lextron Systems, Inc. vs. Microsoft Corporation” Case No. CV 04-00588-VRW, US District Court, Northern District of California, San Francisco Division
3. (2004-2005) Technical consultant and expert witness (*not deposed*) for *Quova, Inc.* in the case “Quova, Inc. v. Digital Envoy, Inc.” Case No. C04 02674JL, US District Court, Northern District of California, San Jose Division.
4. (2004-2005) Technical consultant and expert witness (*deposed*) for *American Honda Motor Co., Inc.* in the litigation case “Donner, Inc. (Plaintiff and Counter-Defendant) v. American Honda Motor Co. Inc. and McDavid Plano-Acra, L.P. and The Beaumont Company d/b/a Classic Acura, (Defendants and Counter-Plaintiffs) v. Irah H. Donner (Counter-Defendant)” Civil Action File No. 5:03-CV-253 (DF), US District Court, Eastern District of Texas, Texarkana Division.
5. (2005-06) Technical consultant and expert witness (*deposed*) for *Pulse Engineering Inc. and Full Rise Electronics Co.* in the litigation case “Regal Electronics, Inc. v. Pulse Engineering, Inc., Fahrner-Miller Associates, Inc., Full Rise Electronics Co. Ltd., Max Lion Electronics, Inc., Bell Fuse, Inc., Stewart Connector Systems, Inc.” Civil Case No: 03-01296 JW, US District Court, Northern District of California, San Jose Division.
6. (2005-2008) Technical consultant and expert witness (*not deposed*) for *Realtek Semiconductor Corporation* in the litigation case “3COM Corporation v. D-Link Systems, Inc. and Realtek Semiconductor Corporation” Case No. CV-03-2177 VRW BZ, US District Court, Northern District of California, San Francisco Division.
7. (2006-2007) Technical consultant and expert witness (*deposed*) for *Tekelec, Inc.* in the litigation case “Bouygues Telecom, S.A. vs. Tekelec, Inc.” Case No. 4:05 CV 78-FL3, US District Court, Eastern District of North Carolina.
8. (2006-2008) Technical consultant and expert witness (*deposed*) for *Nissan North America, Inc.* in the litigation case “Mobile Micromedia Solutions LLC vs. Nissan North America, Inc.” Civil Action File No. 5:05-CV-230, US District Court, Eastern District of Texas, Texarkana Division.

9. (2006-2013) Technical consultant and expert witness (*not* deposed) for *J2 Global Communications, Inc.* and *Catch Curve, Inc.* in the litigation cases (all related) “J2 Global Communications Inc. v. CallWave, Inc.” Civil Action No. CVO4-706D8 DP (AJWx); and “J2 Global Communications Inc. v. Venali, Inc.” Civil Action No. CVO4-1172DDP (AJWx); and “Catch Curve, Inc. (Plaintiff/Counterclaim-Defendant) v. Venali, Inc. (Defendant/Counterclaimant)” Civil Action No. 05-4820 DDP (AJWx), US District Court, Central District of California; and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Protus IP Solutions, Inc.” Civil Action No. CV05-5610 DDP (AJWx); and “COA Network, Inc. v. J2 Global Communications, Inc.”; and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Captaris, Inc. and Open Text Corp.” Civil Action No. CV09-4150 DDP (AJWx); and J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Packetel, Inc.” Civil Action No. CV09-3240 DDP (AJWx); and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Easylink Services International Corporation” Civil Action No. CV 09-4189 DDP (AJWx); and “J2 Global Communications, Inc. and Advanced Messaging Technologies, Inc. v. Venali, Inc.” CV04-1172 DDP (AJWx); and the related case in Canada “J2 Global Communications, Inc. and Catch Curve Inc. v. Protus IP Solutions Inc.” Canadian Federal Court Action No. T-139-06.
10. (2007) Technical consultant and expert witness (*not* deposed) for *Netgear* in the litigation case “Serconet, Ltd. (Plaintiff) v. Netgear, Inc. (Defendant/Third-Party Plaintiff) v. Cameo Communications, Inc., Delta Networks, Inc., Hon Hai Precision Industry Co., Ltd., Sercomm Corp., and Z-Com, Inc. (Third Party Defendants)” Civil Action No. C-06-04646 PJ, US District Court, Northern District of California, San Francisco Division.
11. (2008) Technical consultant and expert witness (*not* deposed) for *Quantum Corporation* in the litigation case “Quantum Corporation v. Riverbed Technology, Inc.” Case No. C-08-0927-WHA, US District Court, Northern District of California.
12. (2008) Technical consultant and expert witness (*deposed*) for Vero Systems (Teoco) in the litigation case “Telarix, Inc. vs. Vero Systems, Inc.” Eastern District of Virginia Case No. 1-08-cv-587 (CMH)
13. (2009-2010) Technical consultant and expert witness (*not* deposed) for ActiveIdentity Corporation in the litigation case “ActiveIdentity vs. Intercede Group PLC and Intercede Ltd.” Northern District of California, Case No. C-08-4577 VRW.
14. (2010-2011) Technical consultant and expert witness (*not* deposed) for Perfect 10 Satellite Distributing Company in the litigation case “John Mezzalingua Associates, Inc. vs. Pace Electronics, Inc., Perfect 10 Antenna Company and DS Engineering, LLC” Case Number: 0.10-cv-00064-MJD-JJG.
15. (2010) Technical consultant and expert witness (*not* deposed) for Belkin International, Inc., Cisco-Linksys L.L.C., D-Link Systems, Inc., Netgear, Inc., and SMC Networks, Inc. in the litigation case “OptimumPath, L.L.C. v. Belkin Int’l, Inc., et al” Case No. 4:09-CV-1398-CW (N.D. Cal)
16. (2011) Technical consultant and expert witness (*not* deposed) for Infoblox, Inc. in the litigation case “Infoblox v. BlueCat” Case Number CV 10-01962 JVS.
17. (2011-2012) Technical consultant and expert witness (*deposed*) for Twitter, Inc. in the litigation case “Copper Notification, Inc., v. Twitter, Inc., Everbridge, Inc., Rave Wireless, Inc., Federal Signal Corp” Civil Action No. 09-865-LPS.
18. (2011-2014) Technical consultant and expert witness (*deposed*) for Barnes & Noble, Inc. in the litigation case “Barnes & Noble, Inc. v. LSI Corp.” No. 3:11-CV-2709-EMC, U.S.D.C. N.D. Cal.

19. (2011-2013) Technical consultant and expert witness (*not* deposed) for Google, Inc. in the litigation case “MasterObjects, Inc. v. Coogle, Inc.” Case No. CV 11-01054 PJH, Northern District of California, Oakland Division.
20. (2012-2013) Technical consultant and expert witness (*not* deposed) for Money Suite Company in the litigation case “The Money Suite Company v. Insurance Answer Center, LLC, et al,” USDC Central District of California – Case No. SACV 11-1847-AG (JPRx)
21. (2012-2013) Technical consultant and expert witness (*deposed*) for Brilliant Telecommunications / Juniper Networks in the litigation case “Alcatel Lucent Technologies USA Inc. v. Brilliant Telecommunications, Inc., et al,” Case No.: 1-11-CV-206910, Superior Court of California, County of Santa Clara.
22. (2012-2013) Technical consultant and expert witness (*deposed & testified at ITC trial*) for HTC Corporation / HTC America Inc. (Respondents) and for Google, Inc. (Intervener) in “The Matter of Certain Electronic Devices, Including Mobile Phones, Tablet Computers, and Components Thereof, ITC Inv. No. 337-TA-847.” (Complainants: Nokia Corp., Nokia Inc., Intellisync Corp.) Worked on three patents, two of which were dropped by Complainants and the Judge found no infringement on the third.
23. (2013) Technical consultant and expert witness (*not* deposed) for HTC Corporation / HTC America, Inc. in the litigation case “Smartphone Technologies vs. AT&T Mobility LLC, HTC Corporation, HTC America, Inc.” Case No. 6:11-CV-561 (E.D. Tex.).
24. (2013) Technical consultant and expert witness (*deposed*) for Aeritas, LLC in the litigation cases: 1) Aeritas, LLC v. Alaska Air Group, Inc., Case No. 11-967-SLR (D. Del.); 2) Aeritas, LLC v. Continental Airlines, Inc., Case No. 11-968-SLR (D. Del.); 3) Aeritas, LLC v. Delta Airlines, Inc., Case No. 11-969-SLR (D. Del.); 4) Aeritas, LLC v. United Airlines, Inc., Case No. 11-970-SLR (D. Del.); 5) Aeritas, LLC v. US Airways, Inc., Case No. 11-1267-SLR (D. Del.); and 6) Aeritas, LLC v. Virgin America, Inc., Case No. 12-1492-SLR (D. Del.).
25. (2013-2014) Technical consultant and expert witness (*deposed*) for IpLearn, LLC in the litigation case “IpLearn, LLC v. Blackboard Inc.” Case No. 1:11-CV-00876-RGA (D. Del.).
26. (2013-2015) Technical consultant and expert witness (*not* deposed) for Infoblox, Inc. in the litigation case “Versata Software, Inc., et al v. Infoblox, Inc.” Case No. 1:13-cv-00678-LPS (D. Del.)
27. (2014-2015) Technical consultant and expert witness (*not* deposed) for Kyocera Corp. and Kyocera Communications Inc. in the litigation case “Adaptix, Inc. v. Kyocera Communications Inc.” Case No(s) 5:14-cv-02894-PSG (and -02895-PSG), Northern District of California, and related IPR case(s).
28. (2015-2016) Technical consultant and expert witness (*deposed*) for IBM Corporation in the litigation case “Twin Peaks Software Inc. v. IBM Corporation” Case No: 14-CV-03933-JST.
29. (2016-2019) Technical consultant and expert witness (*deposed*) for Samsung in the litigation case “Huawei Technologies, Co. Ltd. Et al v. Samsung Electronics Co. Ltd. Et al.” Case No: 3:16-CV-02787 N.D. Cal.
30. (2016-2017) Technical consultant and expert witness (*not* deposed) for Bunjie, Inc. in the litigation case(s) “Bunjie, Inc. v. Acceleration Bay, LLC.” Case No: IPR 2017-01600 (U.S. Patent No. 6,910,069) and Case No: IPR 2017-01601 (U.S. Patent No. 6,829,634).

31. (2016-2018) Technical consultant and expert witness (*not deposed*) for BlackBerry in the litigation case “BlackBerry Limited v. Nokia Corporation, Nokia Solutions and Networks Oy, and Nokia of America Corporation,” Case No. 1:17-CV-00155 (RGA), United States District Court for the District of Delaware. (Also provided consultation services for related IPRs, but did not participate as a disclosed expert witness).
32. (2017) Technical consultant and expert witness (*not deposed*) for Amazon.com and Amazon Web Services, Inc. in IPR proceedings related to the litigation case “Broadcom Corporation, et al. v. Amazon.com, Inc. et al.” Case No: 8:16-cv-01774-JVS-JCG (C.D. Cal.)
33. (2017-2019) Technical consultant and expert witness (*not deposed*) for Wiko SAS in the patent litigation case “Koninklijke Philips N.V. vs. Wiko SAS,” Case No. 200.219.487/01, Court of Appeal of The Hague in the Netherlands; and technical consultant and expert witness (*not deposed*) for AsusTek Computer Inc., Asus Europe B.V. and Asus Holland B.V. in the related patent litigation case “Koninklijke Philips N.V. vs. AsusTek Computer Inc., Asus Europe B.V. and Asus Holland B.V.,” Case No. 200.221.250/01, Court of Appeal of The Hague in the Netherlands.
34. (2018-2019) Technical consultant and expert witness (*not deposed*) for Juniper Networks, Inc. in the patent litigation case “Parity Networks, LLC v. Juniper Networks, Inc.,” Case No. 6:17-CV-00495-RWS-KNM, United States District Court for the Eastern District of Texas.
35. (2019) Technical consultant and expert witness (*not deposed*) for Metaswitch Networks Ltd. and Metaswitch Networks Corporation in the patent litigation case “Sonus Networks, Inc. (d/b/a Ribbon Communications Operating Company) v. Metaswitch Networks Ltd. and Metaswitch Networks Corporation (counterclaim case),” Civil Action Nos. 2:18-CV-57-RWS, Eastern District of Texas.
36. (2020 -2021) Technical consultant and expert witness (*not deposed*) for Dell Technologies Inc., Dell Inc., and EMC Corp. in the patent litigation case “WSOU Investments, LLC v. Dell Technologies Inc. et al., No 6:20-cv-477 (W.D. Tex.).
37. (2022 -) Technical consultant and expert witness (*not deposed*) for Ericsson Inc., AT&T Corp., AT&T Communications LLC., AT&T Mobility, AT&T Mobility II LLC, AT&T Services Inc. in the patent litigation case “K. Mizra LLC v. AT&T Corp. et al., 2:21-cv-00241 (E.D. Tex.) (Lead Case).
38. (2022 -) Technical consultant and expert witness (*not deposed*) for Dell in the patent litigation case Corrigent Corp. v. Dell Technologies Inc. and Dell Inc., Case No. 22-cv-496 (D. Del.); and also for Cisco in the related patent litigation case Corrigent Corp. v. Cisco Systems, Inc. Case No. 22-cv-396 (W.D. Tex.)

B) OTHER CONSULTING, ETC. ENGAGEMENTS OF NICHOLAS BAMBOS

Additionally, I have had the following startup-founding, technical advising, consulting, and other engagements (with compensation, unless otherwise specified)) in the past.

1. (1994-1996) Research visitor to the AT&T Bell Labs Mathematics Center for two weeks in the summer of 1994 and two weeks in the summer of 1996.
2. (1996-present) Research proposal review panelist at the National Science Foundation.

3. (1998) Research consultant for Sun Microsystems for two weeks in 1998.
4. (1999-2001) Co-founder of “AON Networks” (optical networking technology startup).
5. (2001-2003) Member of the Technical Advisory Board of ChainCast (peer-to-peer audio streaming technology startup; now StreamAudio).
6. (2001-2003) Member of the Technical Advisory Board and consultant for Greenfield Networks (packet switching technology startup; eventually acquired by Cisco).
7. (2001-2004) Advisor at ITU Ventures, Los Angeles, CA (technology investment firm).
8. (2004-2004) Member (uncompensated) of the Advisory Board of Detecon USA (consulting branch of Deutsche Telekom).
9. (2008-2009) Co-founder of Dhaani (initially Verde) Systems (technology startup focusing on power management of desktop computers).
10. (2008) Presentation for Cyber Ark International, Inc. on “science and technology innovation in the Silicon Valley” to Chinese delegation in Oct. 2008.
11. (2010 -) Member (uncompensated) of the Scientific Committee, Laboratory for Information, Networking and Communication Sciences (LINCS), Paris, France.
12. (2012) Consulting for Quinn-Emanuel on patent analysis in Oct. & Nov. of 2012.
13. (2013) Consulting for Kilpatrick-Townsend-Stockton on prior art for defendant in Sept. & Oct. 2013, regarding multimedia messaging services (MMS) technology.
14. (2015) Consulting for Quinn-Emanuel on patent analysis in Jan. to May of 2015.
15. (2015) Consulting for Kirkland-Ellis on patent analysis in April 2015, regarding network resource management technology.
16. (2015) Consulting for Winston-Strawn on prior art for defendant in the period May-September 2015.
17. (2016) Consulting for Quinn-Emanuel on patent litigation in January 2016 (engaged, but performed no work and received no compensation).
18. (2018 - 2020) Member (uncompensated) of the Advisory Board, German-Turkish Advanced Research Center (GT-ARC) for Information and Communication Technologies, Berlin, Germany.
19. (2019 - 2021) Consulting for Desmarais LLP on a patent litigation matter beginning September 2019, regarding wireless network communications.
20. (2021 - 2022) Consulting for an international law firm on a litigation matter beginning February 2021, regarding networking technology.

Appendix B

APPENDIX B: CHALLENGED CLAIM LISTING

No.	Limitation
1[pre]	A method for communication, comprising:
1[a]	configuring a network node having a plurality of ports, and at least first and second line cards with respective first and second ports, to operate as a distributed media access control (MAC) bridge in a Layer 2 data network;
1[b]	configuring a link aggregation (LAG) group of parallel physical links between two endpoints in said Layer 2 data network joined together into a single logical link, said LAG group having a plurality of LAG ports and a plurality of conjoined member line cards;
1[c]	providing for each of said member line cards a respective forwarding database (FDB) to hold records associating MAC addresses with ports of said plurality of ports of said network node;
1[d]	receiving a data packet on an ingress port of said network node from a MAC source address, said data packet specifying a MAC destination address on said Layer 2 data network;
1[e]	conveying, by transmitting said data packet to said MAC destination address via said first port, said received data packet in said network node to at least said first line card for transmission to said MAC destination address;
1[f]	if said MAC destination address does not appear in said FDB, flooding said data packet via one and only one LAG port of said plurality of LAG ports;
1[g]	checking said MAC source address of the data packet against records in said FDB of said first line card; and
1[h]	if said FDB of said first line card does not contain a record of an association of said MAC source address with said ingress port, creating a new record of said association, adding said new record to the FDB of said first line card, and sending a message of the association to each member line card of said plurality of member line cards.

No.	Limitation
2	The method according to claim 1, wherein sending the message comprises sending messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.
3	The method according to claim 1, and comprising receiving the message at the second line card, and responsively to the message, adding the record of the association to the FDB of the second line card if the record does not already exist in the FDB of the second line card.
4	The method according to claim 3, and comprising marking the records in the respective FDB of each line card to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.
5[pre]	The method according to claim 4, and comprising
5[a]	associating a respective aging time with each of the records;
5[b]	refreshing the records in the FDB responsively to further packets transmitted by the line cards; and
5[c]	removing the records from the respective FDB if the records are not refreshed within the respective aging time.
6	The method according to claim 1, wherein sending the message comprises transmitting a synchronization packet from the first line card via a switching core of the network node to at least the second line card
7	The method according to claim 6, wherein sending the synchronization packet comprises, if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, changing the record in the FDB of the first line card and sending a synchronization update packet to at least the second line card to indicate that the record has been changed.

No.	Limitation
8	The method according to claim 1, wherein the network node is configured to operate as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances
9	The method according to claim 8, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein sending the message comprises identifying the VPLS instance in the message so as to inform all the line cards that serve the VPLS instance.
10[pre]	The method according to claim 1, and comprising:
10[a]	conveying a further data packet, received from a further MAC source address, to the second line card for transmission over the network;
10[b]	checking the further MAC source address against the records in the FDB of the second line card; and
10[c]	responsively to the further data packet, adding a further record with respect to the MAC source address to the FDB of the second line card and sending a further message to inform at least the first line card of the further record.
11[pre]	A node for network communication, comprising:
11[a]	a switching core;
11[b]	a plurality of ports;
11[c]	a plurality of member line cards conjoined in a link aggregation (LAG) group of parallel physical links between two endpoints in a Layer 2 data network joined together into a single logical link, having a plurality of LAG ports to forward packets through said switching core so that the node operates as a virtual media access control (MAC) bridge in said Layer 2 data network, said plurality of member line cards including at least first and second line cards, each line card having respective ports and having a respective forwarding database (FDB) to hold records associating MAC addresses with said respective ports of said line cards,

No.	Limitation
11[d]	wherein said line cards are arranged so that upon receiving a data packet on an ingress line card from a MAC source address, said data packet specifying a MAC destination address, said ingress line card conveys said data packet via said switching core to at least said first line card for transmission to said MAC destination address, whereupon said first line card checks said MAC source address of said data packet against records in said FDB of said first line card, and if said FDB database of said first line card does not contain a record of an association of said MAC source address with said ingress port, adds said record to the FDB of said first line card and sends a message to at least said second line card informing said second line card of said association, and arranged, when said MAC destination address does not appear in said FDB, to flood said data packet via one and only one of said LAG ports.
12	The node according to claim 11, wherein at least the first line card is adapted to send messages periodically at predefined times to inform at least the second line card of new associations between the MAC addresses and the respective ports.
13	The node according to claim 11, wherein responsively to the message, the second line card adds the record of the association to the MAC database of the second line card if the record does not already exist in the FDB of the second line card.
14	The node according to claim 13, wherein the records in the respective FDB of each line card are marked to distinguish a first type of the records, which are added in response to data packets transmitted via a port of the line card, from a second type of the records, which are added in response to messages received from another of the line cards.
15	The node according to claim 14, wherein a respective aging time is associated with each of the records, and wherein the line cards are operative to refresh the records in the FDB responsively to further packets transmitted by the line cards, and to remove the records from the respective FDB if the records are not refreshed within the respective aging time.
16	The node according to claim 11, wherein the message comprises a synchronization packet, which is transmitted from the first line card via the switching core to at least the second line card.

No.	Limitation
17	The node according to claim 16, wherein the line cards are operative so that if the record in the FDB associates the MAC source address with a port different from the one of the ports on which the data packet was received, the record in the FDB of the first line card is changed, and the synchronization packet comprises a synchronization update packet, which indicates to at least the second line card to indicate that the record has been changed.
18	The node according to claim 11, wherein at least some of the line cards are configured so that the node operates as multiple virtual MAC bridges in a Layer 2 virtual private network (VPN), wherein each virtual MAC bridge is configured to serve a respective VPN instance, and wherein the records associating the MAC addresses with the respective ports are maintained independently for each of the VPN instances.
19	The node according to claim 18, wherein the VPN instance is a VPLS instance among multiple VPLS instances served by the network node, and wherein the VPLS instance is identified in the message so as to inform all the line cards that serve the VPLS instance of the association.
20	The node according to claim 11, wherein the line cards are adapted to forward a further data packet, received from a further MAC source address, to the second line card for transmission over the network, whereupon the second line card checks the further MAC source address against the records in the FDB of the second line card, and responsively to the further data packet, adds a further record with respect to the MAC source address to the FDB of the second line card and sends a further message to inform at least the first line card of the further record.