UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

SERVICENOW, INC.,

Petitioner,

v.

HEWLETT-PACKARD COMPANY,

Patent Owner.

Case No. IPR2015-00631

U.S. Patent No. 7,392,300

DECLARATION OF DANIEL MENASCÉ, Ph.D. REGARDING CLAIMS 1, 7, 8, 10, 21, AND 22 OF U.S. PATENT NO. 7,392,300

TABLE OF CONTENTS

I. INTRODUCTION 1

II. MATERIALS CONSIDERED 12

III. UNDERSTANDING OF THE LAW 12

A. Obviousness 13

B. Claim Construction 14

IV. PERSON OF ORDINARY SKILL IN THE ART 15

V. INSTITUTED GROUNDS 16

VI. BRIEF DESCRIPTION OF THE TECHNOLOGY 17

A. Networks 17

B. Computer models 19

C. Event-driven systems 19

VII. THE INVENTION OF THE '300 PATENT 23

VIII. OVERVIEW OF THE CITED PRIOR ART REFERENCES 39

A. Matheny 39

B. XML in A Nutshell (Harold) 41

C. Hamner 41

D. Pitt 44

IX. THE CHALLENGED '300 PATENT CLAIMS ARE NON-OBVIOUS BECAUSE CLAIM LIMITATIONS ARE NOT TAUGHT BY THE PRIOR ART 45

A. Hamner does not disclose "a network event." 46

1. The term "network event" should be construed as "an action or occurrence within the network that is detected or received by the system." 47

2. The tasks described in Hamner are not "network events" under the proper construction of the term. 50

3. Hamner does not disclose "a network event" under the Board's construction of the term. 55

B. Neither Matheny nor Hamner alone, nor the combination of Hamner with Pitt, discloses "processing a network event using the network model . . . ." 56

1. The claimed "identifying" and "determining" sub-steps must use the network model and the network event must be detected or received before it is processed. 57

2. Matheny does not disclose "processing a network event using the network model" under the plain language of the claims. 60

3. Hamner does not disclose "identifying one or more network objects of the plurality of network objects" under the proper reading of the claims. 61

4. Hamner does not disclose "determining an order of operation on the one or more network objects" under the proper reading of the claims. 64

5. Pitt does not disclose "determining an order of operation on the one or more network objects" under the proper reading of the claims 65

C. Matheny does not disclose "a network model." 67

1. The claim term "network model" should be construed to mean "computer-based representation of a network comprising the objects in the network and the relationships between them." 68

2. Matheny does not disclose a "network model" under the proper construction of the term. 69

3. Matheny does not disclose a "network model" under the Board's construction of the term. 71

X. THE CHALLENGED CLAIMS ARE NOT OBVIOUS BECAUSE A PERSON OF ORDINARY SKILL IN THE ART WOULD NOT HAVE COMBINED MATHENY WITH HAMNER OR MATHENY AND HAMNER WITH PITT 72

A. A person of ordinary skill in the art would not have combined Matheny with Hamner. 72

B. A person of ordinary skill in the art would not have combined Matheny and Hamner with Pitt. 76

XI. CONCLUSION 80

XII. AVAILABILITY FOR CROSS-EXAMINATION 80

XIII. RIGHT TO SUPPLEMENT 81

I, Daniel A. Menascé, declare as follows:

I. INTRODUCTION

1. I am a University Professor of Computer Science at George Mason University ("Mason") in Fairfax, Virginia. University Professor is the highest rank conferred by Mason's President and Board of Visitors to "its faculty women and men of great national or international reputation. The rank of University Professor is reserved for such eminent individuals." *See* Section 2.2.5 of Mason's Faculty Handbook, available at http://www.gmu.edu/resources/facstaff/handbook/GMU_FACULTY_HANDBOOK-2014_Final.pdf. Only a very select group of Full Professors at Mason become University Professors. I have been a Professor of Computer Science at Mason since 1992.

2. I received a Ph.D. in Computer Science from the University of California at Los Angeles ("UCLA") in 1978. I obtained a Master of Science degree in Computer Science in 1975, as well as a Bachelor of Science degree in Electrical Engineering in 1974, both from the Pontifical Catholic University in Rio de Janeiro, Brazil ("PUC-Rio").

3. Prior to joining Mason, from 1978-1992, I was Professor of Computer Science and Chair of the Computer Science Department at PUC-Rio. During this time, I have also held visiting faculty positions at the University of Maryland

Institute for Advanced Computer Studies ("UMIACS"), University of Maryland, College Park, and at the University of Rome, Italy. From 1981 to 1991, I was the co-founder and CEO of Tecnosoft, a software company that specialized in the development of large computerized information systems for companies such as Brazilian oil company Petrobras and Brazilian telecommunications company Embratel. I designed and personally directed the development of these information systems for these and other customers. Tecnosoft also developed and commercialized two database management systems and a software system for capacity planning and Quality of Service ("QoS") prediction of computer systems.

4. I have devoted the past 40 years of my professional career to the area of computer science and in particular to the fields of analytical modeling and simulation of centralized and distributed computer systems, operating systems, computer systems architecture, communications networks, electronic commerce, Web-based systems, database design and management, service-oriented architectures, software performance engineering, secure computer systems, autonomic computing, and operating systems. My field of expertise includes the study and comparison of computer-based systems and software architectures for commercial applications, including information systems in a variety of settings, from PCs to secure networked and Web-based environments.

5. During my time at Mason, I was the lead designer of Mason's Executive Master of Secure Information Systems, the Founding Director of its Master of Science in E-commerce program, and the founding co-Director of Mason's E-Center for E-Business.

6. Also during my time at Mason, I co-founded the Center for the New Engineer ("CNE") in 1993, and was the Associate Director of CNE from 1993 to 1998. Under my direction, CNE created a library of Web-accessible tutorial modules that covered eight topics in computer science, one in general engineering, a refresher for high-school math, and a refresher for college statistics.

7. In 1998, CNE was renamed the HyperLearning Center ("HLC") (see http://denninginstitute.com/cne/cne_hist.html), and I became its director until 2001, when the Center ceased to exist. CNE and HLC received over $3.4 million in research funding from the United States Department of Defense Advanced Research Projects Agency ("DARPA"), the National Science Foundation ("NSF"), and the Association for Computing Machinery ("ACM").

8. From 2005 to 2012, I was the Senior Associate Dean of the Volgenau School of Engineering at Mason ("School of Engineering"). As Senior Associate Dean, I was in charge of research, graduate programs, graduate admissions, promotion and tenure of the faculty, and Web information systems for the entire

School of Engineering. As Senior Associate Dean of the School of Engineering, I was also the director of the school's Ph.D. degree program in Information Technology. In that role, I attended all doctoral dissertation defenses to make a final determination whether the doctorate should be awarded before appending my signature.

9. During my academic career, I have taught a variety of courses at the graduate and undergraduate level, including courses dealing with computer networks, distributed systems, e-commerce, Web services, performance modeling and analysis of computer systems, computer system architecture, database management system, operating systems, and autonomic computing. I have also been the dissertation advisor of 26 Ph.D. students and 52 M.S. students.

10. I am the author of more than 240 peer-reviewed technical papers that have appeared in journals and conference proceedings. My publications have received more than 9,280 citations and my h-index is 46. (The h-index is an index that attempts to measure both the productivity and impact of the published work of a scientist or scholar. The index is based on the set of a scientist's most cited papers and the number of citations that they have received in other publications.)

11. I am the chief author of several books listed below. These books contain extensive discussions on the modeling and analysis of a variety of

computer systems. The methods discussed in these books, as well as in my other published work, are applied to a variety of settings including computer networks, multi-tiered Web-based systems, help-desk systems, e-commerce systems, and software system architectures.

- *Performance by Design: Computer Capacity Planning by Example*, published by Prentice Hall in 2004 (cited 457 times).
- *Capacity Planning for Web Services: Metrics, Models, and Methods*, published by Prentice Hall in 2002 and translated into Russian and Portuguese (cited 805 times);
- *Scaling for E-business: Technologies, Models, Performance, and Capacity Planning*, published by Prentice Hall in 2000 and translated into Korean (cited 404 times);
- *Capacity Planning for Web Performance*, published by Prentice Hall in 1998; and
- *Capacity Planning and Performance Modeling: From Mainframes to Client-Server Systems*, published by Prentice Hall in 1994.

12. All my books come with accompanying software that I developed to solve the mathematical models for queuing theory discussed in the books. These

models are used to predict the performance of computer systems and to verify when their SLAs will be violated.

13. I have received several lifetime achievement awards and recognitions, including elevation to the rank of Fellow of the Institute of Electrical and Electronics Engineers ("IEEE") for "contributions to research and education in performance evaluation of computer systems"; induction as a Fellow of the Association of Computing Machinery ("ACM") for "fundamental contributions to education and practice of computer networks and performance evaluation, and material contributions to the establishment of a strong computing industry in Brazil"; a finalist (28 out of 115 nominations) in the 2014 statewide Outstanding Faculty Award competition among all faculty members of all disciplines in all public and private higher education institutions of Virginia; the 2001 A.A. Michelson Award, a lifetime achievement award given by the Computer Measurement Group, for my contributions to computer metrics; the 2009 Outstanding Research Faculty award by the Volgenau School of Engineering at Mason; the 2000 Teaching Excellence award from Mason; the 1999 Outstanding Teaching award from the School of Engineering at Mason; and several best paper awards.

14. The external funding for my research exceeds $7.4 million and has been provided by the United States Department of Defense Advanced Research Projects Agency ("DARPA"), the Air Force Office of Scientific Research ("AFOSR"), the United States National Aeronautic and Space Administration ("NASA"), the National Science Foundation ("NSF"), the National Geospatial-Intelligence Agency ("NGA"), the National Institute of Standards and Technology ("NIST"), Dominion Virginia Power, Virginia's Center for Innovative Technology ("CIT"), OPNET Technologies, TRW, Hughes Applied Information Systems, Embratel, the Brazilian Research Council ("CNPq"), the Brazilian Ministry of Science and Technology, and IBM Brazil.

15. I have consulted for many government organizations and private companies, including the U.S. Army, NASA, the U.S. Mint, the Defense Information Systems Agency ("DISA"), the Ballistic Missile Defense Organization, the National Institutes of Health, IBM, SABRE (travelocity.com), United Online (netzero.com), Lockheed Martin, Capital One, and the Inter-American Development Bank. The vast majority of these consulting engagements involved developing models for the network and computer systems of these organizations.

16. As an example, I designed and implemented for the National Institutes of Health (from 1994 to 1996) an object-oriented modeling tool, called CMWLan, to do capacity planning for their local area networks. The tool was deployed to all institutes, centers, and divisions of the NIH.

17. I have experience with the design of complex data-intensive distributed information systems in the commercial arena through Tecnosoft, the company I founded and managed from 1981 to 1991, and in the scientific domain where I helped NASA design the federated architecture of its Earth Orbiting System Data and Information System ("EOSDIS"). For the latter work, I received the outstanding paper award from the IEEE International Conference on Engineering of Complex Computer Systems, Southern Florida, USA, November 6-10, 1995, for the paper "A Performance-Oriented Design Methodology for Large-Scale Distributed Data Intensive Information Systems."

18. I have been invited to give keynote addresses at several conferences, universities, and companies around the world. Examples include:

- "Resource Optimization for Iaas and SaaS Providers," International Computer Measurement Group Conference, San Antonio, Texas, November 3, 2015.

- "On the Use of Performance Models in Autonomic Computing," Congress of the Brazilian Computer Society, Curitiba, Brazil, July 18, 2012;
- "Self-Architecting Software Systems," University at Buffalo, September 20, 2011;
- "Virtualization and the On-Demand Data Center," Green Computing Summit, Washington, DC, December 3, 2008;
- "Achieving QoS in Complex Distributed Systems through Autonomic Computing," Alcatel Technical Academy, Antwerp, Belgium, October 3, 2005;
- "Quality of Service Challenges for Web Based Systems and E-commerce," E-Quality Research Center, University of Twente, The Netherlands, September 30, 2005;
- "On the Use of Online Performance Models in Autonomic Computing," IBM Watson Research Center, Hawthorne, NY, July 15, 2004;
- "QoS Challenges and Directions for Large Distributed Systems," Workshop on Quality of Service for Geographically Distributed Systems, Rome, Italy, June 9, 2004;

- "Self-Managing E-commerce Sites," WWW/Internet 2003 IADIS International Conference, November 6, 2003, Algarve, Portugal;
- "Software, Performance, or Engineering?," Third International Workshop on Software and Performance (WOSP 2002), July 24-26, 2002, Rome, Italy;
- "QoS Issues in Web and E-commerce Services," Distinguished Lecturer Series, Computer Science and Engineering Division, University of Michigan, October 25, 2001;
- "Using Performance Models to Dynamically Control E-Commerce Performance," 2001 Aachen International Multiconference on Measurement, Modeling, and Evaluation of Computer-Communication Systems, Aachen, Germany, September 12, 2001; and
- "Understanding Workloads in E-Business," Microsoft Research, Seattle, WA, May 1, 2001.

19. I was the General Chair of ACM's 2007 Federated Computing Research Conference ("FCRC") held in June 2007 in San Diego. This is the largest and most prestigious research event in the computer science field and includes sixteen co-located conferences and many workshops with a total attendance of more than 2,000 researchers.

20. I am a member of the editorial board of ACM's Transactions on Internet Technologies (TOIT), ACM's Transactions on Autonomous and Adaptive Systems (TAAS), and of Elsevier's Performance Evaluation Journal. I was an Associate Editor of ACM's Transactions on the Web ("TWEB") journal, an Associate Editor of Elsevier's Electronic Commerce Research and Applications journal, a member of the Editorial Board of IEEE's Internet Computing, and an Associate Editor of Elsevier's Journal of the Brazilian Computer Society.

21. I am top secret qualified and currently cleared at the secret level by the U.S. Department of Defense.

22. In addition, I am the co-inventor of a U.S. patent entitled "Meta-Protocol" and of the pending U.S. patent application entitled "System and Method for Managing Insider Security Threats."

23. A copy of my latest curriculum vitae (CV) is attached as Exhibit A.

24. I am being compensated at my normal consulting rate of $675/hour for my time. My compensation is not dependent on and in no way affects the substance of my statements in this declaration.

25. I have no financial interest in HP. I similarly have no financial interest in the '300 patent, and have had no contact with the named inventors of the '300 patent.

II. MATERIALS CONSIDERED

26. I have reviewed the specification, claims, and file history of the '300 patent. I understand that this inter partes review is ongoing and other documents have been filed.

27. I have also reviewed the following documents:

- Petition for *Inter Partes* Review of U.S. Patent No. 7,392,300 ("Pet.") (Paper 1), including all exhibits cited therein;
- Decision: Institution of *Inter Partes* Review ("Decision on Institution") (Paper 12), including all exhibits cited therein;
- All documents cited in this declaration.

I have read and understood each of the above documents.

28. I have considered certain issues from the perspective of a person of ordinary skill in the art at the time the application for the '300 patent was filed.

III. UNDERSTANDING OF THE LAW

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

A. Obviousness

30. I have been informed that a patent claim is considered obvious if the subject matter as a whole would have been obvious to one of ordinary skill in the art at the time the invention was made. The obviousness analysis involves several factual inquiries: (i) the scope and content of the prior art; (ii) the differences between the prior art and the claimed invention; (iii) the level of ordinary skill in the art at the time of the invention; and (iv) the existence of objective indicia of non-obviousness ("secondary considerations"), such as unexpected results, long felt but unresolved need, failure of others, and industry skepticism followed by acceptance.

31. In connection with obviousness, I have been informed that there should be some reason that would have led one of ordinary skill in the art to combine or modify the relevant prior-art teachings to obtain the claimed invention. Furthermore, there must be a reasonable expectation of success that one of ordinary skill in the art would have obtained the claimed invention based on the teachings of the prior art. An invention is more likely to be deemed non-obvious to one of ordinary skill in the art if it yields unexpected results or if the prior art teaches away from the claimed invention.

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

33. In connection with objective indicia of non-obviousness, I have been informed that there must be a nexus between the claimed invention and the evidence of objective indicia.

B. Claim Construction

34. I understand that before a patentability analysis can be conducted, the claims of a patent must be interpreted. I have been informed and understand that a claim in *inter partes* review is given the broadest reasonable interpretation in light of the specification.

35. I understand that in the Decision on Institution the Patent and Trial and Appeal Board (the "Board") initially construed the following claim terms:

Claim Term	**Initial Construction**
"network representation"	"information about at least one object in the network or a relationship between objects in the network"
"network model"	"a computer-based representation of a plurality of objects in a network and the relationships between those objects"
"network event"	"an action or occurrence, including actions generated by users of devices on a network, that is received or detected by a network"

Decision on Institution at 5-9 (Paper 12).

36. I have applied the Board's construction of these terms in this declaration, unless noted otherwise.

37. With respect to the constructions of "network model" and "network event," in my opinion, as explained below in Section IX, the Board should reconsider its constructions of these terms. For each term, I have first applied what I believe to be the proper construction of the term, and then applied the Board's construction immediately thereafter when appropriate.

IV. PERSON OF ORDINARY SKILL IN THE ART

38. I have been advised that "a person of ordinary skill in the relevant field" is a hypothetical person to whom one could assign a routine task with reasonable confidence that the task would be successfully carried out.

39. The '300 patent application was filed in the USPTO on January 8, 2004, and does not claim priority to any earlier filing date, either in the U.S. or any foreign patent offices. Accordingly, I understand that the patent's claims must be construed from the perspective of a person of ordinary skill in the art as of January 8, 2004.

40. Based on my understanding of the '300 patent and my knowledge and experience, a person of ordinary skill in the art would have a Bachelor's degree in computer science, electrical engineering, or related discipline, and either (1)

approximately one year of relevant industry experience or (2) an advanced degree, such as a Master's degree or Ph.D., in computer science, electrical engineering, or related discipline.

As reflected in my qualifications set forth above, as of 2004, I would have met or exceeded that level of experience. I have also considered the definition of a person of ordinary skill in the art as set forth in the declaration of Dr. Lavian. Lavian Decl. (Ex. 1002) ¶ 19. Under that definition, my opinions are the same as set forth in this declaration.

V. INSTITUTED GROUNDS

41. I understand that *inter partes* review has been instituted by the Board on the following ground presented in the petition filed by ServiceNow, Inc. ("ServiceNow" or "Petitioner") (Paper 1):

- Ground 1: Claims 1, 7, 8, 10, 21, and 22 as obvious over the combination of U.S. Patent Publication Number 2002/0161883 to Matheny ("Matheny") (Ex. 1003) in view of excerpts from Harold et al., XML in a Nutshell (2001) ("Harold") (Ex. 1004), U.S. Patent No. 5,796,951 to Hamner ("Hamner") (Ex. 1005), and U.S. Patent No. 5,717,934 to Pitt ("Pitt") (Ex. 1007).

VI. BRIEF DESCRIPTION OF THE TECHNOLOGY

A. Networks

42. A network consists of a group of computer and other devices connected by physical links and other types of connections. Networks can include a variety of devices besides computers, such as printers, servers, storage devices, routers, switches, and gateways. '300 patent (Ex. 1001) at 2:36-40, 4:2-6. Common types of networks include local area networks ("LANs") and wide area networks ("WANs"). The devices on the network communicate using standard communication protocols that dictate the rules that each device must use to communicate.

43. A network allows applications running on computers in the network to communicate with each other, exchange data, and share resources such as printers and storage devices. Examples include a browser at a user's computer interacting with a web server that could be located in a different continent, or a printer that is shared by several computers on the network.

44. The many connections in a network and sharing of data and resources create numerous dependencies in the network. For example, a shared printer and a storage device may be connected only to a server, which manages access to the printer and the storage device. The printer and storage device therefore are

dependent on the server. Computers may also be connected to a server that provides access to other services, such as the Internet.

45. Some types of networks use additional types of connections called "virtual circuits," which are connections that are created on top of physical links and which may also receive data from other virtual circuits underneath them in a hierarchy. Although the details are not important for present purposes, these virtual circuits are dependent on the circuits (both virtual and physical) below them in the network.

46. Networks have to be designed to be reliable, i.e., robust in the face of failures including link and router failures and to exhibit a certain level of expected performance (e.g., measured in average packet delay or jitter). Many different design options may be available for a given network based on its reliability and performance requirements and on cost constraints.

47. Moreover, once a network is in operation, it needs to be managed (either automatically or semi-automatically) so that the operational status of its components is constantly assessed and corrective actions are taken when needed to restore the network to a state that meets its design requirements.

B. Computer models

48. A computer model is a computer program that is designed to provide a representation of a real-world system. A computer model can be used to simulate the behavior of the real-world system that has been modeled, which can help provide insight into the system and how it would react under various scenarios. Computer models and simulations are used in a wide range of disciplines, including physics, climatology, mathematics, economics, and engineering.

C. Event-driven systems

49. Event-driven systems are systems in which actions or occurrences such as sensor outputs, hardware signals, or messages are detected and then the part of the system programmed to process the event produces a reaction or response to the event. An event can occur at any time (i.e., it is unpredictable), and will be detected or received by a specific mechanism, such as a software listener, watcher, or by using a hardware interrupt. *See, e.g.*, Microsoft Computer Dictionary 198 (Fifth Ed. 2002) (Ex. 2004) (event: "An action or occurrence, often generated by the user, to which a program might respond - for example, key presses, button clicks, or mouse movements. See also event-driven programming."); Random House Webster's Computer and Internet Dictionary 199 (Third Ed. 1999) (Ex. 2005) (event: "An action or occurrence detected by a

program. Events can be user actions, such as clicking a mouse button or pressing a key, or system occurrences, such as running out of memory. Most modern applications, particularly those that run in Macintosh and Windows environments, are said to be event-driven, because they are designed to respond to events."); Sydow (Ex. 2010) at 69 ("Macintosh programs always have been event-based. An event is an action of some kind, such as a click of mouse or the press of a key. When it occurs, the program responds."), 73 (2002) (Ex. 2010) ("Sydow") ("One of the primary jobs of the Carbon Event Manager system software is to watch for events.").

50. The purpose of an event is to notify a system that something needs attention, and thus the event communicates information about what needs attention by creating an "event notification," which is typically then sent to an event handler. The part of the system or component that receives the event will be programmed to respond or react to the event. The responding part of the system or component does not know or control when the event will occur.

51. Examples of events include key presses, button clicks, and mouse movements. These inputs are events that occur at an input device (i.e., a mouse or keyboard) and are communicated locally to the processor of the computer, which interprets and reacts to them. The input generates a hardware interrupt (a signal),

which acts as an event notification. Wang et al., "Modeling and Integration of Peripheral Devices in Embedded Systems," Proceedings of the IEEE Design, Automation and Test in Europe Conference and Exhibition 3 (2003) ("Devices generate hardware events to indicate changes of hardware states. A typical hardware event is an interrupt."). The event is sent by an event handler to the operating system and appropriate application for processing. *Id.* ("The processing of events in the event driven state machines is specified in event handlers. The core functions provide the device services to the upper layers of the software – the OS and the application").

52. Events are also used in distributed middleware systems to allow computers to collaborate with each other. CORBA (Common Object Request Broker Architecture) is designed to facilitate communications between different systems. CORBA provides an event handling service "that *signals the occurrence* of an event" so that "[c]lients amenable to that event [can] *take appropriate action*." Tanenbaum, "Distributed Systems: Principles and Paradigms" 502 (2002) (Ex. 2012) ("Tanenbaum") (emphasis added). Because clients do not know and cannot control when an event will occur, clients can either passively wait for events to happen, or can poll an event supplier to check on whether an event has occurred. *Id.* at 503.

53. Many other examples of events used in hardware, software, and hardware-software systems abound. Some examples include:

- A divide by zero event detected by an Arithmetic and Logic Unit (ALU) in a CPU, which generates a notification and causes an exception handling routine to process the event (e.g., by replacing the result of the operation with the numerical representation of infinity and resuming the program).
- In a virtual memory system, a page fault event is detected by the CPU to indicate that the data and/or instruction requested by a program do not reside in main memory. The CPU generates an interrupt (the notification) and control is passed to the operating system, which brings the requested memory page from secondary storage (e.g., disk) to main memory (event processing).
- A Network Interface Card (NIC) receives a data frame from the network (the event) and generates an interrupt (the notification). The computer in which the NIC resides receives the interrupt and moves the data from the NIC's data buffer to the computer's memory (event processing).

54. Each of these examples of events includes the common characteristics of event detection, event notification, and event processing. An event in a computer system can happen at any time in a system and generates an event notification that provides information about what needs attention. The system must be programmed to detect the event (via the notification) and react to it (event processing).

VII. THE INVENTION OF THE '300 PATENT

55. U.S. Patent No. 7,392,300 is directed to a novel method and system for creating a flexible and easily-configurable model of a communications network using a structured language, such as eXtensible Markup Language ("XML"). '300 patent (Ex. 1001) at 1:6-8, 2:51-62, 3:55-62, 3:34-41.

56. An important consideration for network design is the evaluation of design and maintenance alternatives. As the inventors of the '300 patent recognized, maintaining, troubleshooting, and testing design and maintenance alternatives on an actual network is costly, time consuming, and potentially disruptive, at least in terms of testing the effect of alternatives in network design on the network itself. *Id.* at 1:28-31. One way to avoid these issues is to model the network and simulate its behavior. *Id.* at 1:31-34. Simulation is a technique that works on a model of the system being simulated. Using a computer model to

simulate the architecture and the behavior of the network can help identify potential problems before any changes are made to the network itself. *Id.* For example, in the event of a server failure, using the model to first identify devices connected to the server will help the network administrator determine which devices will need to be moved to a new server and create an orderly plan for implementing the move. Once a plan is in place, the actual replacement of the server should be more efficient and less disruptive.

57. The computer modeling system described in the '300 patent (also called the "system" or "network inventory adapter") models the network and then processes network events that occur in the network by using the model. '300 patent (Ex. 1001) at Abstract, 3:9-18, 6:17-26. The modeling system first creates and stores a model of the network called a "network model." The network model can then be used to simulate the effect of events that occur in the network. The network model can be rebuilt any time there are changes in the network. *Id.* at 3:42-45, 6:34-46. When a "network event" is detected, the modeling system uses the "network model"—which reflects the current configuration of the network—to determine how to respond to the network event. The invention of '300 patent therefore is an event-driven system in which the modeling system responds to and processes network events received from the network.

58. Independent claim 1 claims a method for modeling a network and using the network model to process a network event:

> 1. A method of modelling a communications network using a computer system, the method including:
>
> generating a network representation using computer-readable code, the computer-readable code representing structured information;
>
> parsing the network representation;
>
> generating a network model using the parsed network representation, the network model including a plurality of network objects and relationships between the plurality of network objects;
>
> storing the network model in memory; and
>
> processing a network event using the network model,
>
> > wherein the processing includes identifying one or more network objects of the plurality of network objects, and
> >
> > the processing further includes determining an order of operation on the one or more network objects.

'300 patent (Ex. 1001) at claim 1. Claims 10 and 21 have similar steps.

59. Figure 5 is a flow chart that shows the steps of modeling the network and processing a network event using the network model:

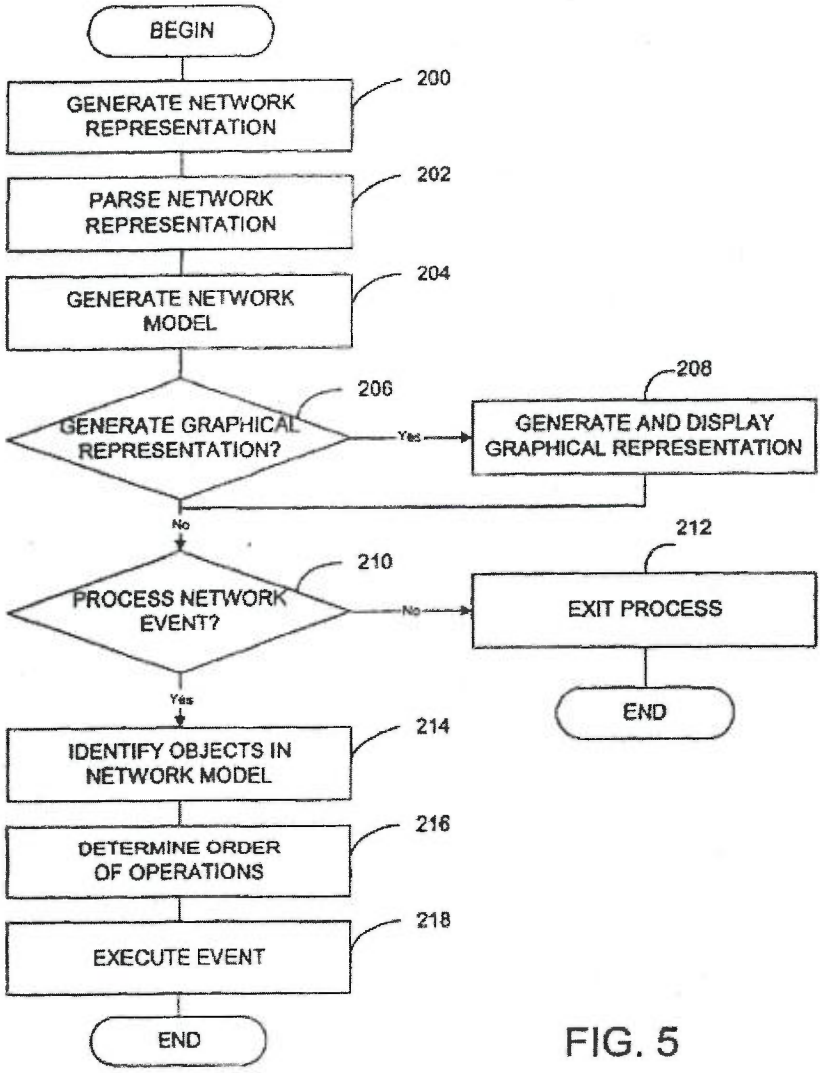

FIG. 5

'300 patent (Ex. 1001) at Fig. 5, 6:32-34. The system first generates a network representation in step 200 using structured code. One such structured language mentioned in the patent and some of the claims is eXtensible Markup Language

("XML"). *Id.* at 6:34-36. The "network representation" described by the XML code contains the inventory of devices on the network such as servers, printers, computers, routers, and other devices, and the relationships between them. *Id.* at 3:67-4:8 (listing types of devices), 6:42-44. The modeling system then "parses" the network representation—i.e., the XML code is read and divided into individual elements according to XML rules—in step 202 to create the network model in step 204. *Id.* at 6:44-46. The network model will then contain a representation of the devices on the network and the relationships between them.

60. Once the network model has been created, it can be used to process "network events" that occur in the network. The '300 patent specification primarily relies on the example of a local area network with dial-in access, and includes provisioning or deleting a network device as examples of network events. But the '300 patent invention can be used to model other types of networks and process other types of network events. *Id.* at 2:52-58 (listing examples relevant to a dial-up network context).

61. In step 210 of Figure 5, the modeling system determines whether a network event has been received for processing. *Id.* at 6:49-51 ("[T]hen the system determines whether an event is to be processed, step 210."). If a network event has been received, the modeling system will need to respond to the event, such as by

updating the network model in in response to the event. To determine what needs to be done and how, the modeling system first identifies the network objects (which represent, for example, the network devices) in the network model that are impacted by the network event and that may need to be updated or deleted or are otherwise affected by the network event. *Id.* at 6:52-53 ("If yes, then the system identifies the needed objects in the network model, step 214.") The modeling system then determines which objects need to be operated on (i.e., updated or deleted), the operation needed for each object, and the order in which those operations will occur. *Id.* at 6:53-55 ("In step 216, the system determines the order of operations needed to process the network event."), 5:9-11 ("The adapters 34 [system of the invention] may support specific operations and allow for various operations to be triggered by certain events and commands.").

62. After the objects are identified and the operations and their order determined, the operations will be executed in the determined order to finish processing the network event. *Id.* at 6:55-56. The network model will then be up to date because it will reflect the changes to the network caused by the network event. The network model therefore is used to process a network event. For example, when a new circuit, switch, or service is added to the network, the modeling system will receive a network event from the network. The network

model can then be used to determine which network objects will be affected by the network event and need to be updated, and the order in which the changes to the network objects must be performed.

63. A common and simple example illustrates how the network model in the '300 patent is used to process a network event and the importance of determining the order that the network objects will be operated on. Consider a network that has (among other components) a server with several devices, such as printers and computers, connected to it. The network model, which includes the server and connected devices, is generated and stored in a database. *Id.* at 8:48-49 ("In one embodiment, the model is generated in a network database."). The modeling system then receives a network event from the network notifying the system that the server has failed, which will then need to be replaced. Replacing the server in the network model requires deleting the server database record and adding the new server record to the database.

64. But the old server object cannot simply be deleted first because it has other devices connected to it, and those devices must be connected to the new server. If the old server is deleted first, the information in the database representing the connections will also be deleted. There would be no way to find the now-disconnected devices to connect them to the new server after it is added.

In other words, these devices, having lost their parent record in the database on which they depend, will be orphaned objects.

65. To avoid this problem, all the dependencies related to the failed server—such as the connected devices—are first located by following each connection from the server record until all devices connected to it (and anything else that depends on it, such as processes that run on it) are identified. In this example, the old server object in the network model is first located by server name, location, serial number or some other identifying information. Then each connection from the server to other records in the database is followed so that each network object connected to or communicating with the server is identified. The server is then replaced by performing the following steps in order: (1) the new server object is added to the database as a record; (2) the records for the devices connected to the old server are updated to instead connect to the new server object; and (3) the old server object (which no longer has any devices that depend on it) is deleted.

66. This orderly process of examining the connections in the network model to identify any objects dependent on the server and the nature of the dependencies before acting on any changes required by the network event will reduce or eliminate the possibility of creating orphaned devices. This is important

because although an orphaned device will be lost in the network model, it may still be physically connected to the actual network and will still need to be managed.

67. In addition, it is important to detect the network event before identifying network objects and determining the order of operations on the network objects because the most recent version of the network model must be used to accurately simulate how the network event will affect the physical network. The network model as it exists when the network event is received will be the most up-to-date version of the network model. And of course the network event must be known before the objects needed to process the network event can be identified.

68. The '300 patent also contains an example that illustrates these steps of receiving a network event, identifying the objects impacted in the network model, and determining an order of operations on the network objects using the network model.

69. Figure 3 of the '300 patent shows a graphical representation of a network model. This example shows the network objects of a network model that are needed to provide dial-in access to a local area network using an Internet Service Provider. *Id.* at 5:14-47, Figs. 2, 3. Although the details of this particular network are not important, understanding what is meant by "circuits" in the patent is helpful to understanding the example of processing a network event discussed in

column 7 of the patent. Figure 3 depicts virtual circuits (also known as virtual connections or circuit connections) created when two devices are connected using a specific communications protocol:

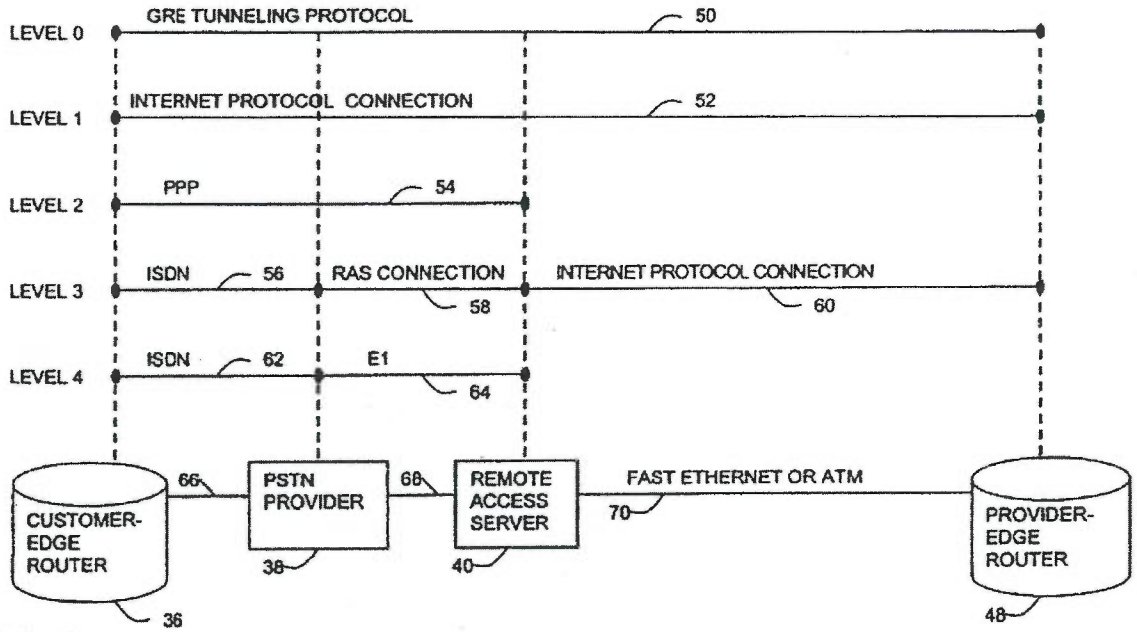

FIG. 3

Id. at 5:46-59, Fig. 3.

70. For example, the customer edge router 36 and the provider edge router 48 are connected by an IP tunnel circuit 50 created using the GRE (Generic Routing Encapsulation) tunneling protocol (*id.* at 5:48-49), and the customer edge router 36 is connected to the PSTN (Public Switched Telephone Network) provider 38 via the ISDN (Integrated Services for Digital Network) protocol (56, 62). *Id.* at

5:50-54. The customer edge router 36 is also connected to the PSTN provider via a physical link 66, such as a telephone line. *Id.* at 5:54-66.

71. When a connection is attempted (such as when a user dials in to the network) between the provider edge router 48 and the customer edge router 36—which is a network event called "provisioning"—each of the connection circuits are created (i.e., the connections are established) in the order shown from Level 4 to Level 0. The circuit level numbers in the network model reflect the order that the circuits must be provisioned. *Id.* at 5:40-45.

72. The '300 specification uses the example network shown in Figures 2 and 3 to explain how the modeling system processes a "network event." *See* '300 patent (Ex. 1001) at 6:57-8:33. In this example, the modeling system has received an event indicating that a provisioning operation—the process of establishing one of the circuit connections shown in Figure 3—has failed, thus requiring a "rollback" operation. A rollback operation will restore the network to its pre-provisioning state by deleting any circuits and related network objects in the network model that were added during the provisioning process. *Id.* at 6:58-60 ("A rollback is the restoring of the status of the network inventory whenever a network operation or provision operation fails."). To perform the rollback, the modeling system must first determine which network objects in the network model are

affected—i.e., which circuit connections and other network objects must be deleted or updated—and then determine the order that the objects are to be deleted or updated. In this manner, the modeling system is used to process network events.

73. The following XML code show in column 7 of the '300 patent shows part of the network model used in the example, which includes an Update And RollBack operation used to identify the network objects and determine the order of the operations on the network objects:

```
<?xml version="1.0" encoding="ISO-8859-1"?>
<!DOCTYPE UpdateAndRollBack (View Source for full doctype. . .)>
<UpdateAndRollBack>
  <Circuit index="3" Type="ISDN S0 Bearer" DeleteCircuit="YES"
    UnderlyingCircuitsIndex="NA" UnderlyingLinkIndex="NA"
    UseSame="NA" Delete="" />
  <Link Index="3" DeleteLink="Yes" ModifyPort="StartPortName"
    AssociatedNode="StartNodeName" DeleteDevice="NA" I>
  <ISDN>
    <Variant Type="GRE">
      <Circuit index="0" Type="GRE" DeleteCircuit="Yes"
        UnderLyingCircuitsIndex="1" UnderlyingLinkIndex="NA"
        UseSame="NA" Delete="Loopback" />
      <Circuit index="1" Type="IP Connectivity"
        DeleteCircuit="Yes" UnderlyingCircuitsIndex="2"
        UnderlyingLinkIndex="NA" UseSame="DestNodeName"
        Delete="NA" />
      <Circuit index="2" Type="PPP" DeleteCircuit="Yes"
        UnderlyingCircuitsIndex="3,4" UnderlyingLinkIndex="NA"
        UseSame="DestNodeName" Delete="NA" />
      <Circuit index="3" Type="RAS Connection"
        DeleteCircuit="NO" UnderlyingCircuitsIndex="NA"
        UnderlyingLinkIndex="NA" UseSame="SourceNodeName"
        Delete="NA" />
      <Circuit index="4" Type= "ISDN Connection"
        Deletecircuit="Yes" UnderlyingCircuitsIndex="PP"
        UnderlyingLinkIndex="NA" UseSame="DestNodeName"
        Delete="NA" />
    </Variant>
  </ISDN>
</UpdateAndRollBack>
```

'300 patent (Ex. 1001) at 7:6-39. The network model in the XML code above contains the network objects representing the circuit connections shown in Figure 3, from the Level 0 GRE circuit (shown in blue) to the Level 4 ISDN Connection

(pink). Id. at 7:39-44 ("'Circuit index' identifies the circuit level in terms of distance from the top level. In the model shown in FIG. 3, the GRE tunneling connection is the topmost level, therefore, having an index of zero (0). The Index field identifies the order in which the circuits are to be deleted.").

74. When the modeling system receives the network event from the network, such as via a middleware or messaging bus, the system will process the event using the network model. *Id.* at 3:52-54, 4:26-28, 6:49-55, claim 1. To process the network event, the modeling system first identifies the network objects—in this example, the circuits that have been provisioned—that are needed to process the rollback operation by locating each circuit under the UpdateAndRollback tag in the network model and identifying each circuit (using the "Type" tag) where the DeleteCircuit field is "YES." *Id.* at 7:44-45 ("'DeleteCircuit' identifies whether or not a particular circuit needs to be deleted."). If a circuit that is to be deleted also has an underlying link (i.e., a physical connection, such as telephone line), the network objects representing those links will also be identified because they will need to be deleted as well *Id.* at 7:53-55 ("'UnderlyingLinkindex' identifies whether or not the circuit has an underlying link. If an underlying link exists, the underlying link index has numerical value identifying the order in which it is to be deleted").

75. Once the network objects are identified, the modeling system will determine the order in which they will be deleted using the network model. The deletion order is indicated by the Index and UnderlyingLinkIndex fields (if any). *Id.* at 7:42-46 ("The Index field identifies the order in which the circuits are to be deleted. . . . "DeleteCircuit" identifies whether or not a particular circuit needs to be deleted."), 7:53-57 ("'UnderlyingLinkindex' identifies whether or not the circuit has an underlying link. If an underlying link exists, the underlying link index has numerical value identifying the order in which it is to be deleted."). In the example, the objects associated with the GRE circuit with index = 0 (blue) will be deleted first, and the objects associated with the ISDN Connection with index=4 (pink) will be deleted last. *Id.* at 8:25-33 ("If the DeleteCircuit or DeleteLink attribute is "YES," then the circuit object or link object corresponding to the index is retrieved. . . . If any objects associated with this 30 circuit or link are to be deleted, they are then deleted. After the associated objects are deleted, the circuit or link is deleted.").

76. The RAS Connection at Level 3 (orange), however, will not be deleted because it was not established during the provisioning process, and the modeling system is designed to preserve circuits that existed before the provisioning was attempted. *Id.* at 6:64-7:4. The value "NO" for "DeleteCircuit"

(purple) for the RAS Connection at Level 3 indicates that it will not be deleted. *Id.* at 7:44-48 ("'DeleteCircuit' identifies whether or not a particular circuit needs to be deleted. This field is used since certain circuits may not need to be deleted, such as circuits used in multiple models or circuits that may be part of a backbone.").

77. As previously discussed, identifying the correct objects is important to maintaining the integrity of the network. For example, in the network model above, the RAS Connection circuit is not deleted because it was not provisioned during the provisioning process that is being rolled back. The RAS Connection connects the PSTN (Public Switched Telephone Network) provider 38 to the Remote Access Server 40. '300 patent (Ex. 1001) at Fig. 3. The Remote Access Server authentication services each time a user attempts to dial in to the network, and therefore should not be deleted. *Id.* at 5:32-34, 6:64-65. The order in which the circuits are deleted is also important because each circuit depends on one or more of the circuits below it. For example, the GRE circuit receives data from—and is thus connected to—the Internet Protocol circuit. If the Internet Protocol circuit is deleted first, there would be no way to identify the GRE circuit in the network model because it would be an orphan – i.e., an object that was not connected to anything.

VIII. OVERVIEW OF THE CITED PRIOR ART REFERENCES

78. The Petition relies on four prior art reference, discussed in turn below.

A. Matheny

79. The Matheny reference is primarily used in the Petition for the first four limitations of the challenged claims – generating a network representation, parsing the network representation, generating a network model, and storing the network model. Petition at 17, 24-34. Matheny is the only reference that contains any mention of XML or any other type of structured code as required by the step of "generating a network representation."

80. Matheny describes collecting and coalescing data discovered on a network for use by a management server. Matheny (Ex. 1003) at Abstract, 3:30-35. Matheny uses discovery "agents" to poll targeted devices on the network and collect information about them. *Id.* ¶ 0011. The discovered information is saved in files.

81. Matheny describes two types of files. When data is collected about a device, the discovery agent places the collected data in a file created for that device. *Id.* ¶ 0021. The discovery agent may also create a "relationship file," which contains information about how devices relate to each other. *Id.* ¶ 0022.

82. The data in the discovery files relating to the devices is then aggregated and coalesced – i.e., copied – into a "discovery document." *Id.* ¶¶ 0021-0022, 0024. Duplicate data is removed during the coalescing process so that only one entry remains in the discovery document for each discovered device. *Id.* ¶ 0025, claim 1.

83. Matheny contains no description of the data in the discovery document. Importantly, however, the discovery document is created only from the files that contain information "for the discovered devices," not from data in the relationship file. Matheny explains that the relationship file "may" be created, but never mentions the data in any meaningful detail again or how this relationship file can later be put to use.

84. In addition, although Matheny explains that network management tools "may be used to determine the operational status of equipment and transmission facilities and to obtain notification of faults and threshold conditions, e.g., network traffic bottlenecks" (Matheny (Ex. 1003) ¶ 0001), Matheny does not explain how its "discovery document" is used after it has been created. Matheny is primarily concerned with the discovery process.

B. XML in A Nutshell (Harold)

85. Harold is a 2001 desktop reference manual that explains the rules of XML that apply to all XML documents. The excerpt of Harold used in the Petition contains an introduction to how XML works, discusses some of the fundamentals of XML (including tags and elements), and discloses how to write simple XML documents. Harold (Ex. 1004) at 9-37.

C. Hamner

86. The Hamner reference is used in the Petition for the steps of "processing a network event." Petition at 17, 34-43. Hamner describes a network management system that provides network management services that allows a user (such as a network administrator) to display devices on a network, display "tasks" that can be performed on those devices, and perform selected "tasks." Hamner (Ex. 1005) at 2:46-47, 3:57-59, 4:33-38. Some examples of tasks in Hamner include "viewing the screen of a particular PC; displaying packet counts; running a report; executing a remote virus scan; rebooting selected workstations; displaying print jobs; or, displaying non-functioning printers." *Id.* at 3:51-56.

87. As shown in Figure 3, the software described in Hamner consists of several modules, including a discovery manager 301, a database engine 302, a physical network model 303, a view generator 304, and a task manager 305. *Id.* at

4:59-63. The discovery module 301 collects data about the devices on the network and adds the information to the physical network model 303, which is stored in the physical model database. *Id.* at 5:27-33. The task manager 304 determines which tasks can be performed on the network, and associates each task with the devices and groups of devices on which the task can be performed. *Id.* at 5:33-37, Fig. 10. Each of these features, including the management console of Figure 2A, are implemented in a management server 12, which is connected to the managed network 10. *See id.* at Fig. 1, 3:25-40

88. Figure 2A of Hamner depicts the management console used by the network administrator. The management console has a device window 201 on the left to display devices and groups of devices on the network. On the right is a task window 202.

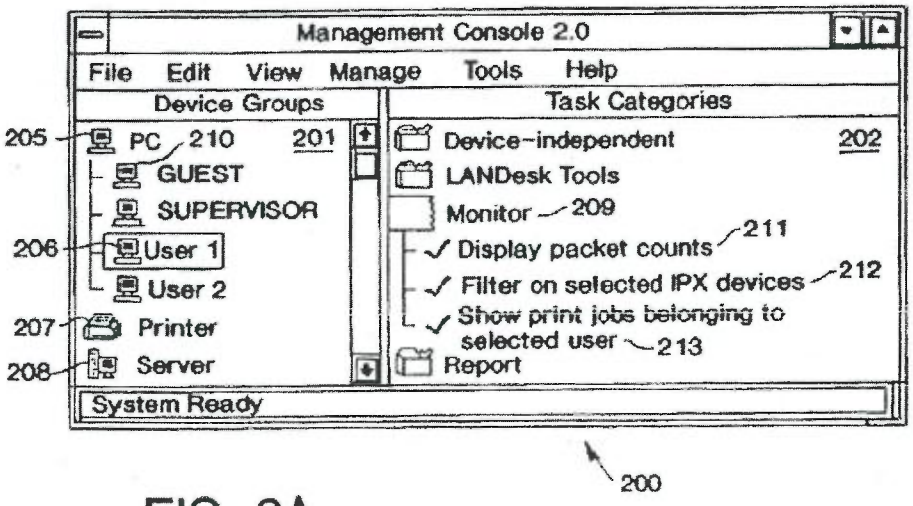

FIG. 2A

Id. at Fig. 2A, 4:2-4.

89. When the network administrator selects a specific device or device group from the device window 201 on the left, the task manager 305 of Hamner will calculate the tasks that are associated with the device or device group and then display them in the task window 202 on the right. *Id.* at 4:17-32, 11:22-25, Fig. 2A. The user can then select an available task (211, 212, 213) associated with the selected device to be performed on the device. *Id.* at 4:33-39, Fig. 2A. The user can select a task to be initiated. In other words, the user – the network administrator – first views the devices on the network, then selects a device, and then the devices are displayed. It is not until after a device has been selected that the tasks assigned to the device are displayed and available for selection and execution. The network administrator has complete control over when the tasks in Hamner are initiated.

90. Hamner also states that a "task consists essentially of an atomic script and any associated parameters." *Id.* at 11:3-4. Hamner explains that the selected device on which the task will be performed is one such parameter. *Id.* at 11:4-6. Hamner does not disclose how tasks are assigned to devices, written, or executed.

91. Notably, Hamner makes no mention of "events" or receiving events from the network and processing them. A "task" in Hamner does not originate in

the network and is not "detected or received" by the management server. Rather, "tasks" originate in the management server itself and the associated physical network model 303 saved within the management server.

D. Pitt

92. The Pitt reference is used in the Petition for a single limitation: "determining an order of operation on the one or more network objects." Petition at 18. Pitt describes a sequential computer network shutdown system and process in which a user can create a plan for shutting down network devices in the event of a power failure or scheduled or manual shutdown. Pitt (Ex. 1007) at 1:21-22.

93. Pitt consists of a two-part system: (1) user interface software for configuring a rules-based shutdown plan that will be installed on each network device; and (2) a process for executing the shutdown plan when a power failure occurs or a shutdown is scheduled or manually activated. *Id.* at 2:43-56, 3:34-37, Figs. 2, 3.

94. In the configuration process, the user interface software builds a list of computers on the network, identifies any interdependencies between devices, and recommends a shutdown schedule and rules for each device based on the interdependencies and the available run time provided by the UPS (uninterruptible power supply) connected to each device. *Id.* at 4:13-36. The user can make

changes to the interdependencies using the user interface. *Id.* at 4:33-40. Once the shutdown plan is created and approved, it is installed on each network device to complete the configuration process. *Id.* at 4:66-5:4.

95. The shutdown plan will then be executed on each device in response to a manual request, a scheduled shutdown, or a power failure. *Id.* at 5:12-18. By the time the plan is executed, however, the order in which the devices will be shut down has already been determined.

IX. THE CHALLENGED '300 PATENT CLAIMS ARE NON-OBVIOUS BECAUSE CLAIM LIMITATIONS ARE NOT TAUGHT BY THE PRIOR ART

96. I understand that, in *inter partes* review, a claim term is given its "broadest reasonable construction in light of the specification of the patent in which it appears." 37 C.F.R. § 42.100(b). Petitioner proposed constructions, and the Board construed, three terms from the '300 patent claims—"network representation," "network model," and "network event." The Board also noted that the phrase "using the network model" does not limit the "identifying" and "determining" sub-steps of "processing a network event using the network model." *Id.* at 17. For my opinions, I do not dispute the Board's construction of "network representation." In my opinion, however, the term "network event" should be construed in a different way from that given by the Board.

97. It is my understanding that, in order to prevail on obviousness grounds, all limitations of the challenged claims must be taught or suggested by the prior art. Under the proper claim constructions, the prior art relied upon by the Petitioner fails to disclose several limitations, including: (a) Hamner does not disclose a "network event," (b) neither Matheny nor Hamner alone, nor the combination of Hamner and Pitt, discloses "processing a network event using the network model, wherein the processing includes identifying one or more network objects of the plurality of network objects, and the processing further includes determining an order of operation on the one or more network objects," and (c) Matheny does not disclose a "network model."

A. Hamner does not disclose "a network event."

98. Petitioner argues that the "tasks" disclosed in Hamner are "network events." Petition at 36. Petitioner's argument is based on an overly-broad construction of "network event" as "one or more operations that can be performed on or by a network or network object" – i.e., essentially anything that can be done in a network. The Board did not adopt that construction. Although I disagree with parts of the Board's construction, Hamner does not disclose "network events" under either the correct construction or the Board's construction.

1. The term "network event" should be construed as "an action or occurrence within the network that is detected or received by the system."

99. The Board construed the term "network event" as "an action or occurrence, including actions generated by users of devices on a network, that is received or detected by a network." Decision on Institution (Paper 12) at 9. In my opinion, the correct construction of "network event" is "an action or occurrence within the network that is detected or received by the system." Only this construction is consistent with the ordinary meaning of the term "network event" and the '300 specification.

100. As discussed above in Section VI.C, the term "network event" incorporates the term "event," which is well-known in computer science, and is used in many different types of systems. An event is an action or occurrence that can occur at any time. For this reason, an event will generate an event notification that will be sent to an event handler, which in turn will send the event information to the part of the system or the component that will process and react to the event. The event is detected or received by a specific mechanism, such as a software listener, watcher, or by using a hardware interrupt. The event notification communicates that something needs attention, and the part of the system or component that is notified of the event will be programmed to respond or react to

the event. The receiving part of the system does not know or control when the event will occur. That is, the timing of when an event will occur is unpredictable.

101. I have described several systems that use events in Section VI.C above. Events in each of these examples have the same characteristics: an event that occurs at an unpredictable moment, an event notification, an event handler, and a system or component that processes and reacts to the event.

102. The '300 patent describes an event-driven system, and the "network event" in the patent is an "event" that occurs in the network and is detected and processed by the modeling system. The '300 patent describes an example of a "network event": the failure of a provisioning operation, which requires an automated rollback. '300 patent (Ex. 1001) at 6:58-61. The failure event is received by the modeling system via the middleware bus. *See id.* at 2:57-58 ("The network event may be received from the middleware bus."), 3:14-18 ("The applications may also be interfaced with the network using a middleware bus. Upon ***receiving events from the network***, the adapter reads and parses the network representation to determine which network objects are to be operated on and the order of operation."), 3:52-54 ("The network inventory adapter [system of the invention] may also receive events from the middleware bus to provision a specific scenario in the network inventory."). When the network event is received,

it triggers operations to be performed in response. *See, e.g., id.* at 5:9-11 ("The adapters 34 [system of the invention] may support specific operations and allow for various operations to be triggered by certain events and commands."). The "network event" has the same characteristics as an "event" in other event-driven systems—the "network event" occurs at any time and is detected and processed by the modeling system. *Id.* at 6:29-31. The modeling system does not know or control when the event will occur.

103. The Board's construction does not include the phrase "*within the network*" and requires that the "network event" be detected or received "by the network," which is where the event occurred. I disagree with this construction. The inclusion of the word "network" in "network event" indicates that the event must occur in the network. There would be no other reason to add "network" to the term "network event." As discussed above, the specification explains that the network event occurs in the network and is then received by the modeling system via the middleware bus. In other words, it is the modeling system —not the network—that processes the network event. '300 patent (Ex. 1001) at 3:14-18, 3:52-54, 6:50-55. This is a key feature and purpose of the invention, and is reflected in the claim language. *See id.* at claim 1 ("A method of modelling a

communications network . . . including . . . processing a network event using the network model . . . ").

104. For these reasons, it is my opinion that the broadest reasonable construction of "network event" in light of the plain language of the claims and the specification is "an action or occurrence within the network that is detected or received by the system."

105. The Board also added "including actions generated by users of devices on a network" to the construction of "network event." Decision on Institution (Paper 12) at 9. This addition is unnecessary because HP's proposed construction does not exclude the generation of an event by a user. To the extent that phrase is considered to be necessary, the following alternative construction of "network event" still reflects that the network event occurs in the network and is detected or received by the modeling system: "an action or occurrence within the network, including actions generated by users of devices on a network, that is detected or received by the system."

2. The tasks described in Hamner are not "network events" under the proper construction of the term.

106. Under the proper construction of "network event," the tasks in Hamner are not "network events" because they are not actions or occurrences

within the network that are ***received or detected by the modeling system*** from the network.

107. As discussed above, a "task" in Hamner is a network management "action[] that can be taken to troubleshoot, monitor, or report on devices in the network." Hamner (Ex. 1005) at 3:49-51; *see also id.* 3:4-8 ("The core services 300 provide a user, such as a LAN administrator, with customizable views of the layout of a network, the various network management tasks that can be performed on devices in the network, and the particular devices on which each task can be performed."). Hamner lists a few examples, such as displaying print jobs, displaying non-functioning printers, or viewing the screen of a particular computer in the network. *Id.* at 3:51-56. Tasks as disclosed in Hamner are very different from the network events described in the patent.

108. A task in Hamner behaves differently than a "network event" – it does not originate in the network to be later detected or received by the management server, as is the case with a "network event." Rather, a task originates in the management server itself and is not executed until after the user selects it. To initiate (i.e., execute) a task, the user (such as a network administrator) first displays the network devices on the left side of the user interface (the management console shown in Figure 2A), selects a device, causing the tasks available for the

device to be displayed on the right side of the management console, and then selects one of the tasks to be performed. Hamner (Ex. 1005) at 3:64-67 ("The ***user may initiate any of the displayed tasks*** by applying a user input via a 'control panel' screen (e.g., by double-clicking on a task icon with a mouse)."), 4:33-39 ("The ***user can cause any displayed task to be performed*** upon a device. A ***task is initiated by the user's selecting*** the bitmaps of the desired task and the device or group upon which the task is to be performed."), 11:25-29 ("A ***user may initiate a task*** by, for example, selecting a device or group and then double-clicking on the available (displayed) task.") (emphasis added throughout), Fig. 2A. There is no need for the management server to "detect[] or receive[]" the tasks in Hamner. A task in Hamner is not an action or occurrence "within the network" and the tasks are not "detected or received by the system," as required by the proper construction of "network event." For these reasons, a "task" in Hamner is not a "network event."

109. I note that Petitioner does not argue that the use of the mouse to select a task in Hamner is a "network event." It would not be a "network event" because it did not occur in the network.

110. Tasks in Hamner also are not "network events" because they are not unpredictable, as are network events (that is, a task is not an "action or occurrence"

as is a network event). The network administrator in Hamner has complete control over which "task" is selected and when the task is performed. Hamner (Ex. 1005) at 3:64-67, 4:33-39, Fig. 2A. The tasks are assigned to each device and stored when the device is discovered, and the user needs only to display the devices, select a device, and then select an available task for execution whenever desired. *See, e.g.,* Hamner (Ex. 1005) at 3:64-67. In the '300 patent, the modeling system does not know anything about a network event – including when it will occur – until after it occurs and is detected. The modeling system is programmed to detect or receive such a "network event" from the network (such as via a listening process) and to respond to it at any time. The user of the modeling system cannot select a network event and has no control over when it will occur or be received. Unlike a "task" in Hamner, a "network event" is not known until it is detected or received at an unpredictable time whenever it occurs. The modeling system does not know when a network event will occur and need to be processed, unlike a task of Hamner.

111. A "task" in Hamner is implemented by an "atomic script," which is a series of commands or instructions.[1] A "network event," in contrast, is simply an

[1] A script is "atomic" if it is designed to either execute completely or not execute at all. If it is interrupted, an atomic script will return the system to its state prior to

action or occurrence, such as the failure of a server on the network, that is communicated to the modeling system. A "network event" can trigger an operation, command, or program, but is not itself an operation, command, or program and is not executed or performed.

112. Hamner and the '300 patent are fundamentally different systems with different uses and purposes. Hamner is not an event-driven system because tasks do not originate in the network before being received or detected for processing by the management server. Hamner is a network management tool that allows a network administrator to select routine network management tasks for performance, whereas the '300 patent is directed to a modeling system used for *event-driven* processing of network events. Hamner simply is not an event-drive system, and in fact, never uses the term "event" or describes any of the characteristics of an "event," even though this was a well-known term in the computer science field well before Hamner's filing date. *Id.* at passim

---

the initiation of the script. *See* Microsoft Computer Dictionary 40 (Fifth Ed. 2002) (Ex. 2013) ("atomic operation": "An operation considered to be ***guaranteed to be indivisible*** . . . . Either the operation is uninterruptible or, if it is aborted, a mechanism is provided that ensures the return of the system to its ***state prior to initiation of the operation***.") (emphasis added).

3. Hamner does not disclose "a network event" under the Board's construction of the term.

113. In my opinion, Hamner does not disclose a "network event" even under the Board's construction of the term. The Board stated that "Hamner's tasks are actions that may be initiated by a user on various network devices, the execution or initiation of which may be detected by the network." Decision on Institution (Paper 12) at 17. I disagree. Hamner does not disclose "the execution or initiation" of "tasks" that are "received or detected by the network," as required by the Board's preliminary construction. The "tasks" in Hamner are functions, implemented by scripts, that already exist as part of Hamner's management server and are stored in the physical network model on the management server. Hamner (Ex. 1005) at 11:3-4. The user selects a task on the management server to execute it, as described above, so there is no need for the task to be received or detected. *Id.* at 4:18-20, 4:34-36, 8:46-67 (describing the physical model network as storing devices and tasks and stating that "[e]ach device 701 can have one or more tasks associated with it..."), 10:66-11:21.

114. In other words, network events in the '300 patent originate in the network, whereas tasks in Hamner do not. The Board did not explain how a "task" in Hamner is "detected or received" by the network. It is not.

115. Although the user in Hamner can click on a displayed task in order to select it, this user input is not a claimed "network event" because, as set forth above, Hamner has no details whatsoever about processing such a user input with a "network model." A user input can be detected and dealt with by the computer on which it was entered, and there is no reason to use a "network model" to process such an input. For these reasons, the user input is not a "network event."

B. Neither Matheny nor Hamner alone, nor the combination of Hamner with Pitt, discloses "processing a network event using the network model . . . ."

116. Claim 1 of the '300 patent requires "processing a network event using the network model, wherein the processing includes identifying one or more network objects of the plurality of network objects, and the processing further includes determining an order of operation on the one or more network objects." The other challenged claims of the '300 patent contain almost identical language. In my opinion, the correct reading of the "processing" step requires the "identifying" and "determining" sub-steps to use the network model and the network event to be detected or received before it is processed. In my opinion, however, the prior art references cited in the Petition do not disclose these steps under either this correct reading or the Board's reading of the claims.

1. The claimed "identifying" and "determining" sub-steps must use the network model and the network event must be detected or received before it is processed.

117. The Board suggested that the claims "do not recite that the identifying and determining steps require the use of the network model." Decision on Institution (Paper 12) at 17. I disagree.

118. The plain language of the challenged claims requires use of the "network model" for the "identifying" and "determining" sub-steps. I understand the "identifying" and the "determining" steps to be sub-steps of "processing a network event using the network model," and by extension, any limitation on "processing a network event using the network model" must apply to both of those sub-steps. That would mean that both sub-steps "us[e] the network model." *See, e.g.*, '300 patent (Ex. 1001) at 9:51-56 (claim 1).

119. The language of the "identifying" and "determining" sub-steps also shows this. Both sub-steps require use of the "network objects." The "network objects" are part of the "network model." *Id.* at 9: 47-50. The use of the word "the" before "plurality of network objects" in "identifying one or more network objects of the plurality of network objects" and before "one or more network objects" in "determining an order of operation on the one or more network objects"

is a reference to the "network objects" in the "network model." For this reason, both the "identifying" and "determining" sub-steps must use the "network model."

120. This reading of the "processing" step and the "identifying" and "determining" sub-steps is consistent with the '300 specification. The '300 specification describes steps in which the system (1) first determines "whether an event is to be processed;" (2) if so, "identifies the needed objects in the network model" and (3) "determines the order of operations [on the network objects] needed to process the network event." '300 patent (Ex. 1001) at 6:49-55, Fig. 5 (flow chart showing that step 214 is to "identify objects in network model"). The purpose of the '300 patent invention is to use the network model to simulate how a network event will affect the network. *See, e.g.*, *id.* at 1:28-36, claim 1 ("A method of modelling a communications network using a computer system"), 8:50-52 ("[T]he model can be used for any service function such as, for example, provisioning, assurance, usage, modify, delete, and rollback."), 8:64-67 ("[T]he network model provides a system such as, for example, a single window system for performing any of the various operations of the service provider."). There would be no reason to create a network model at all in the claimed invention if it were not to be used later in the claim.

121. I also disagree with the Board's statement that the "network event" need not be detected before it is processed. Decision on Institution (Paper 12) at 17. As I explain in Section IX.A.1, the '300 patent repeatedly describes receiving the network event before it is processed. *Id.* at 3:15-18 ("Upon ***receiving events*** from the network, the adapter [system of the invention] reads and parses the network representation to determine which network objects are to be operated on and the order of ***operation***."), 6:46-55, 7:64-67 ("When the adapter ***receives*** an event to rollback a line, the adapter gets a Service Instance ID (SIID) as input. For this particular SIID, the associated circuits are retrieved and the correct circuit is chosen.") (emphasis added throughout).

122. This reading makes sense – the "network event" must be known to the modeling system before it can be processed. The modeling system cannot identify the objects needed to process the network event until the nature of the network event is known to the system. Moreover, the modeling system must use the most current version of the network model – i.e., the one that exists at the time the network event occurs – to process the network event to ensure that the correct network objects needed to process the event are identified. This reading is also consistent with event-driven systems (such as the '300 patent invention), where an event first occurs and is then processed.

123. Under the Board's preliminary construction, a "network event" is "an action or occurrence" – i.e., something that has already happened – "received or detected by a network" – i.e., after the action or occurrence has happened. This construction recognizes that an action or occurrence could not be received or detected until it happens. *Id.*

2. Matheny does not disclose "processing a network event using the network model" under the plain language of the claims.

124. Matheny states in its Background section that information gathered by the discovery "agents" "may be used to evaluate network performance and possible faults, as well [as] provide information needed to reconfigure the network." Petition at 34 (citing Matheny (Ex. 1003) at ¶ 0002). According to Petitioner, a person of "ordinary skill in the art would have understood the acts of 'evaluat[ing] network performance and possible faults,' at a minimum, included identifying one or more devices in the network." Petition at 34. I disagree: a person of ordinary skill in the art would find nothing in Matheny to suggest "processing a network event using the network model."

125. The Petitioner argues that the "discovery document" in Matheny is the "network model" (another conclusion with which I disagree). But even if the "discovery document" is a "network model," Matheny does not explain any way in

which the discovery document is used, much less how it could be used to "process a network event." Matheny (Ex. 1003) at ¶ 0002. The statement from Matheny quoted by Petitioner (from Matheny at ¶2) about "evaluat[ing] network performance" does not even mention Matheny's "discovery document." Matheny only discloses that its discovery document is created through the "coalescing" process. Matheny (Ex. 1003) ¶¶ 0024, 0028.

3. Hamner does not disclose "identifying one or more network objects of the plurality of network objects" under the proper reading of the claims.

126. Petitioner makes two short arguments about the sub-step of "identifying one or more network objects of the plurality of network objects" of claims 1, 10, and 21. Petitioner first asserts that Hamner discloses this sub-step because a "task" in Hamner "could include network events that would identify one or more objects." Petition at 37-38. Petitioner next asserts that this sub-step is also disclosed in Hamner because Hamner "uses the physical model database" to "identify the device groups and the devices in them." *Id.* at 38. These arguments are premised on the argument that a "task" is a network event, which it is not. But even assuming that a "task" is a network event, I disagree with both arguments.

127. As an initial matter, Petitioner has not pointed to anything in Hamner that supports the argument that a "task" in Hamner "could include network events

that would identify one or more objects." For this reason, Petitioner cannot prove this argument.

128. In any event, as I explained in Section IX.B.1, a "network event" in the '300 patent must first be detected, and then the network objects that are needed to process the network event must be identified by locating those objects in the network model. Hamner does not contain any disclosures that meet these requirements of the "identifying" sub-step. At best, tasks in Hamner are handled in the opposite order that network events are processed in the '300 patent.

129. For example, selecting a "task" (Petitioner's first argument) does not result in "identifying one or more network objects" using the network model because the network objects associated with the task are already known when the task is selected. As previously discussed, Hamner describes a network management tool with a user interface where the user must first view and select a device on the network, which causes the tasks available for the device to display, only after which a task can be selected for execution. *Id.* at 11:23-26 ("the task manager 305 can provide a display, ***in response to a user's selection of one or more devices or groups, showing the effective tasks*** associated with the selected devices or groups"), 11:25-27 ("A user may initiate a task by, for example, ***selecting a device or group and then double-clicking on an available (displayed)***

task."), 11:4-5 ("A parameter associated with a task may be a device upon which a task will be performed."), Figs. 2A, 10. Hamner never mentions identifying any device after a task is selected, or any need to do so. Because the device is already known when the task is selected, there is no need to identify any devices (the "network objects") as part of processing the task after the task has been selected.

130. The act of expanding a device group in the user interface (Figure 2A) of Hamner so that the individual devices can be displayed (Petitioner's second argument) simply involves retrieving the devices within the group from the physical model database. No tasks have yet been displayed or selected at this time. *Id.* at 4:10-16, 9:61-67 ("The logical group view generator 341 fills an outline for display within the device window 201 with a list of devices grouped according to their group memberships. The group view generator 341 obtains the group definitions from the physical model database 333, while device information is obtained from the discovery manager 301."). Expanding a device group does not involve tasks at all, and is therefore also not part of "processing a network event using the network model."

131. Tasks in Hamner are handled in the reverse order from that required by the claimed steps of "processing a network event using the network model" including "identifying one or more network objects of the plurality of network

objects." The '300 patent discloses first detecting or receiving a network event, and then processing it by consulting the network model to identify the affected objects. It is important to identify the network objects ***after*** the network event is received or detected because only the most current network model will reflect the correct and current state of the network. Hamner, on the other hand, is not an event-driven system that *reacts* to events in the network whenever they occur. In Hamner, the user of Hamner's user interface controls which devices are selected and which tasks are performed and when, and thus knows which device or group of devices relates to the task to be selected. For this reason, unlike the unpredictable network events that occur in the network in the '300 patent, there is no need in Hamner to model how a task will affect the network.

4. Hamner does not disclose "determining an order of operation on the one or more network objects" under the proper reading of the claims.

132. Hamner does not disclose "determining an order of operation on the one or more network objects" using the network model. Petitioner asserts that the "determining" sub-step corresponds to (1) "identifying *or* executing a 'script' on a network device" because (2) a task is implemented by a script and (3) a script is "a computer program that defines a sequence of operations" for execution in order. Petition at 41, 42-43. I disagree.

133. Hamner contains a single sentence about scripts, stating that "[a] task consists essentially of an atomic script and any associated parameters" (Hamner (Ex. 1005) at 11:3-4), and Hamner does not explain how a script is created, assigned to a device, or executed. Hamner does not disclose or suggest that the scripts determine the order of operation on the network objects.

134. In any event, the order of operations on the network objects does not need to be determined when the task is selected or executed. At the time the script is selected or executed in Hamner, its instructions have already been written and any order of operations is pre-determined. There is no need to consult the network model to make this determination. In other words, unlike the '300 patent claims, Hamner does not teach detecting or receiving a task (the purported network event) and then determining an order of operation in response to the task.

5. Pitt does not disclose "determining an order of operation on the one or more network objects" under the proper reading of the claims.

135. Petitioner asserts that "Pitt discloses the ability to create a 'shutdown plan'" with ordering rules that execute in response to a "shutdown event," such that "all computers are . . . shut down in the preprogrammed sequence." Petition at 52-53. I disagree. Pitt does not teach "determining an order of operation on the

one or more network objects" as part of processing a networking event that has been detected.

136. Pitt describes a "sequential network computer shutdown system" ((Pitt (Ex. 1007) at 1:21-22) with two parts: (1) software implementing the "configuration process" for creating a shutdown schedule (*id.* at 4:21-46; Lavian Dep. (Ex. 2009) at 102:17-19, 103:16-21); and (2) software for executing the shutdown plan on each device on the network (Pitt (Ex. 1007) at 5:12-63).

137. The configuration process is not part of "processing a network event" because no "event" has occurred at the time the configuration process is used to create the shutdown. During the configuration process, the user interface software builds a list of computers on the network, identifies any interdependencies between devices, and recommends a shutdown schedule for each device. *Id.* at 4:21-36. There is no shutdown event (the purported network event) at this time.

138. When a shutdown event occurs in Pitt, the shutdown schedule is executed. Pitt (Ex. 1007) at 5:12-18. At this time, however, the order in which the computers will be shut down has already been determined in the schedule and is preprogrammed into the software. *See id.* at 5:52-56 ("Thus, regardless of the particular cause of shutdown, each computer continues to run from its own power supply in its UPS for the time set by the instructions preprogrammed in its own

software 30, and shuts down at the conclusion of the set time period."). For this reason, Pitt does not disclose determining the order of operation on the network objects as part of processing a network event using the network model. Once again, unlike the '300 patent claims, Pitt does not teach detecting or receiving a shutdown event (the purported network event) and then determining an order of operation in response to the shutdown event.

139. As already discussed, it is important in the '300 patent invention to use the network model to determine the order of operations and to do so only after the network event occurs, because the current state of the network model and network objects is needed to process the network event accurately. Pitt does not satisfy this requirement. In fact, a significant amount of time may pass between the creation of the shutdown schedule and the actual shutdown event in Pitt, which could render the shutdown rules obsolete.

140. Accordingly, Pitt does not disclose "determining an order of operation on the one or more network objects."

C. Matheny does not disclose "a network model."

141. The Board construed the term "network model" to mean "a computer-based representation of a plurality of objects in a network and the relationships between those objects" (Decision on Institution (Paper 12) at 7), and agreed with

Petitioner that Matheny discloses "generating a network model." *Id.* at 14. I do not agree.

1. The claim term "network model" should be construed to mean "computer-based representation of a network comprising the objects in the network and the relationships between them."

142. In my opinion, the broadest reasonable interpretation of "network model" from independent claims 1, 10, and 21 of the '300 patent is a "computer-based representation of a network comprising the objects in the network and the relationships between them." The Board's construction of the term "network model" as "a computer-based representation of a plurality of objects in a network and the relationships between those objects" (Decision on Institution (Paper 12) at 7) does not require the representation to be of a "network." In my opinion, this omission is incorrect.

143. A "model" is a representation of a real world system. *See* IEEE Standard Dictionary of Electrical and Electronics Terms 594 (Fourth Ed. 1988) (Ex. 2003) (model: "A representation of a real world process, device, or concept."). In the '300 patent, that system is the network. For this reason, the construction of the term "network model" should reflect that the model is a "representation of a *network*." A random set of objects and their relationships do not necessarily form a network model.

2. Matheny does not disclose a "network model" under the proper construction of the term.

144. Petitioner and the Board conclude that Matheny's "discovery document"—the purported "network model"—is a "network model" because Matheny describes "loop[ing] through the information in *all* of the discovery files . . . and cop[ying] the data into the discovery document." Decision on Institution (Paper 12) at 14. I disagree that this disclosure supports the conclusion that the discovery document is a network model.

145. Matheny states only that "all of the discovery files ***for the discovered devices***" are used to create the discovery document, which excludes the relationship file. Matheny (Ex. 1003) ¶ 0024 (emphasis added). In Matheny, a "relationship file" is a separate type of file from the files containing information about discovered devices. *Id.* at ¶¶ 0021, 0022 (stating the "discovery agent may ***also create a relationship file***."). In discussing how the "discovery document" is formed, Matheny never mentions use of the relationship file, and in fact, creation of the relationship file is optional to begin with. *See* Matheny (Ex. 1003) ¶ 0022 ("The discovery agent ***may*** create a relationship file . . . ."), ¶ 0024 (describing the "coalescing" operation to form the "discovery document," but mentioning only "all of the discovery files for the discovered devices," but not the "relationship file") (emphasis added); *see also id.* ¶¶ 0024-28 (failing to state that the relationship file

are used to create the "discovery document"). The entire purpose of the coalescing process – during which each discovery file is examined – is to eliminate duplicate data about a discovered device so that only one entry about a discovered device exists in the discovery document. *Id.* at claim 1 ("coalescing the discovery data in a discovery document, said discovery data including two or more duplicate data entries; and removing all but one of the duplicate data entries from the discovery document."). No relationship file is used or needed during this process.

146. In addition, the bare disclosure that data is copied into a document, without more, is not sufficient to disclose a "network model" as required by the '300 patent claims. The discovery document could be nothing more than a data file with information about discovered devices, without actually forming a network model using that information.

147. In sum, without more information about what exactly is in its "discovery document," there is no way to determine if Matheny's discovery document contains a "computer-based representation of a network comprising the objects in the network and the relationships between them." In my opinion, a person of ordinary skill in the art would conclude that Matheny's "discovery document" is not a "network model."

3. Matheny does not disclose a "network model" under the Board's construction of the term.

148. Under the Board's construction of "network model" as "a computer-based representation of a plurality of objects in a network and the relationships between those objects" (Decision at 7), Matheny does not disclose a "network model" for two reasons.

149. As explained in Section IX.C.2 above, Matheny does not disclose the use of its "relationship file" in forming the "discovery document." The fact that Matheny discloses looping through "all of the discovery files for the discovered devices" in forming the "discovery document" does not mean that the relationship file is used for this process. The relationship file is separate from the discovery files for the discovered devices, and there is no basis for concluding that the "discovery document" includes "the relationships between those objects" as required by the Board's construction. Matheny contains no such disclosure.

150. The "discovery document" also does not include "objects in a network" because the "discovery document" of Matheny is not a representation of a network. As previously discussed, it is possible to aggregate information about discovered devices in a network without actually forming a model of the network. Matheny does not disclose that the "discovery document" represents a network.

X. THE CHALLENGED CLAIMS ARE NOT OBVIOUS BECAUSE A PERSON OF ORDINARY SKILL IN THE ART WOULD NOT HAVE COMBINED MATHENY WITH HAMNER OR MATHENY AND HAMNER WITH PITT

151. It is my opinion that a person of ordinary skill in the art would not have combined Matheny with Hamner, or Matheny and Hamner with Pitt, to arrive at the claimed invention.

A. A person of ordinary skill in the art would not have combined Matheny with Hamner.

152. There is no reason to combine Matheny with Hamner. In fact, for a person of ordinary skill in the art, there are several reasons not to do so.

153. Matheny discloses discovery "agents" to poll targeted devices on the network and collect information about the devices. Matheny (Ex. 1003) ¶ 0011. The discovered information about devices is saved into discovery files that are then "aggregated" and "coalesced"—i.e., copied—into a "discovery document." *Id.* ¶¶ 0021-0022, 0024. Petitioner contends that this "discovery document" in Matheny is a "network model" and that "[o]ne of ordinary skill in the art would understand that a network management system like Hamner could take over where Matheny leaves off" and use "the 'network model' to process network events." Petition at 40. Petitioner also argues that "Hamner could have provided a user interface to use

the network model in Matheny to process network events." Petition at 41. I disagree.

154. Matheny focuses on the discovery agents, aggregating the discovered data, and removing duplicate data, but contains no description of the data that is actually in the "discovery document" (the purported "network model"). Matheny (Ex. 1003) ¶¶ 0002, 0024, 0025. Without such a description, a person of ordinary skill in the art would not look to Hamner to provide a user interface to use with the discovery document, because such a person would not have considered the discovery document to be sufficiently described in Matheny to be usable with any network management system, much less Hamner's system.

155. If anything, the information that *is* disclosed about the discovery document teaches away from combining it with Hamner. As discussed above in Section VIII.A, the relationship file is not added to the discovery document. Matheny (Ex. 2003) ¶ 0025. Without the relationship data, the discovery document cannot be a "network model" and cannot be used with a system that needs a network model. Moreover, Hamner already discloses a physical network model that includes the devices and the relationships between them. Hamner (Ex. 2005) at 7:43-53. Hamner's physical network model contains more information than what is included in the discovery document. Therefore, use of Matheny would

not cure any deficiency in Hamner. Thus, a skilled artisan would not be motivated to combine Matheny with Hamner.

156. Petitioner relies on Hamner's statement that "[t]he management services are designed from a flexible, device-oriented perspective in order to provide the user with information on the layout of the network, tasks that can be performed, and the devices on which each particular task can be performed" as an "express motivation to combine" Hamner with the "network model generator of Matheny," as Petitioner asserts. Petition at 40; Hamner (Ex. 1005) at 12:22-27. I disagree. The quoted statement describes the invention of *Hamner itself*, and does not indicate any need to combine Hamner with any other reference. *See id.* at 12:21-27 ("Thus, a method and apparatus for providing management services for a computer network has been described. The management services are designed."). And there is simply no need to combine Hamner with Matheny. Hamner already discloses a discovery manager that collects information about devices on the network and stores the information in the physical network model. Hamner (Ex. 1005) at 5:28-33. A person of ordinary skill in the art would not consider Hamner to have any deficiencies in this respect, much less any that could be cured by Matheny's discovery document.

157. Finally, Petitioner's combination of Hamner with Matheny would not work. Petitioner argues that the combination of Matheny and Hamner "would have predictably resulted in a system in which the network management system of Hamner [would] use the 'network model' in Matheny (the 'discovery document' in Matheny) to process the 'network events.'" Petition at 38-39; Lavian Decl. (Ex. 1002) at ¶ 98. I disagree. In Hamner, the physical network model 303, and more particularly, its physical model database 333, stores the tasks that can be performed on devices and is used to populate the available tasks for a device. *See* Hamner (Ex. 1005) at 11:8-21, 8:47-51 ("available tasks are stored in the physical model database"), Figs 4 and 6; Petition at 37 (explaining that the tasks of Hamner are stored in the physical model database). Matheny contains no disclosure whatsoever of tasks or of processing tasks (or network events, for that matter). If a person of ordinary skill in the art were to use Matheny's discovery document in place of Hamner's physical network model—as Petitioner suggests for its obviousness combination—the system of Hamner would not work correctly because the discovery document of Matheny does not have tasks associated with it that could be used to populate the user interface of Hamner. Even under Petitioner's view that the "discovery document" of Matheny would include information about devices in a network and the relationships between those devices

(a conclusion that I believe to be incorrect), the discovery document still does not contain information about tasks that can be performed within the network. Without tasks, or without the scripts for these tasks, Matheny's discovery document would render the Hamner system inoperable. There would be no tasks to populate the user interface of Figure 2A when a user selects a device. There would be no scripts to be performed even if, somehow, a task were to be listed in Hamner's user interface of Figure 2A (although there would be no tasks to list to begin with). There would be no purpose of providing a user interface for task management if there were no tasks. There is no reason a person of ordinary skill in the art would form such an inoperable system. Hamner's task management system would serve no purpose if the discovery document of Matheny were substituted into Hamner's system. Accordingly, Matheny would not be combinable with Hamner to render the asserted claims obvious, and because Petitioner's only combination includes both Hamner and Matheny, the challenged claims are not obvious in view of Petitioner's purported combination.

B. A person of ordinary skill in the art would not have combined Matheny and Hamner with Pitt.

158. The Board stated that "at least a combination of Hamner and Pitt" teaches the determining step and that "an ordinarily skilled artisan would have had

reason to combine the teachings of Pitt with those of Matheny and Hamner." Decision at 19. I disagree.

159. Petitioner contends that "[t]he task database in Hamner could have been adapted to incorporate order of operation information, such as the sequenced shutdown rules in Pitt." Petition at 54. I disagree: one of ordinary skill in the art would have found no reason provided in these references for combining them, nor has Petitioner shown how such a combination would have been made.

160. Nothing in Hamner suggests any need for the additional functionality of Pitt. The shutdown plan described in Pitt is purportedly needed to prevent data loss in the event of a power failure or scheduled shutdown. Pitt (Ex. 1007) at 5:11-17. Petitioner cites the task of "rebooting [of] selected workstations" as potentially requiring execution in a particular order. Petition at 54. But Hamner does not need Pitt to perform this function because the network administrator using the management server user interface in Hamner already has complete control over which workstations will be rebooted and when, and thus already has the ability to delay the reboot to avoid data loss. Hamner (Ex. 1005) at 3:64-67, 4:33-39, Fig. 2A.

161. Petitioner also asserts that "a 'reboot' as mentioned in Hamner would have involved shutting down all applications and devices running on the computer

(including network devices)," such that "[t]he shutdown plan in Pitt would therefore have provided a benefit to the devices managed by the system of Hamner." Petition at 54-55; Lavian Decl. (Ex. 1002) ¶ 131. This is incorrect. Pitt could not have affected the shutdown order of applications and devices in a computer during reboot because this order is programmed in the operating system and cannot be changed by a user. Thus, there is no need for an automated shutdown plan for the rebooting task.

162. To the extent the other tasks described in Hamner—viewing the screen of a particular PC, displaying packet counts, running a report, executing a remote virus scan, displaying print jobs, displaying non-functioning printer (Hamner (Ex. 1005) at 3:52-56), or otherwise "troubleshooting, monitoring, and reporting on devices" (*id.* at 11:6-7)—pose any risk of data loss or otherwise need to be performed in a particular order, this too can be completely controlled by the network administrator in Hamner and would not require a system for recommending the sequence as described in Pitt. For these reasons, a person of ordinary skill in the art would not add Pitt to Hamner to form the claimed combination.

163. Petitioner also does not explain how Hamner would be modified to incorporate the "order of operation teachings" of Pitt. Petition at 54. The only

modification proposed by Petitioner is the suggestion of adapting the task database in Hamner "to incorporate order of operation information, such as the sequenced shutdown rules in Pitt." Petition at 54. Petitioner does not explain how the sequenced shutdown rules of Pitt would be stored in the Hamner system, whether the rules would be associated with devices, tasks, or both, or how they would be associated. Petitioner makes no attempt to explain how Pitt would be incorporated into Hamner, and a person of ordinary skill in the art would not do so in order to form the claimed combination.

164. Implementing the Pitt shutdown schedule into Hamner's system would not have been simply a matter of "incorporate[ing] order of operation information," such as the sequenced shutdown rules of Pitt, into the physical model database of Hamner, as Petitioner suggests. Petition at 54. To the extent that the "order of operations on the network objects" is determined in Pitt, that process is done with the user interface software, as discussed in Section VIII.D. This process of determining these rules is more complex, as Pitt describes, and would be an entirely new add-on (and unneeded) functionality to Hamner. Adding these features of Pitt to Hamner would require a new user interface and user interface software program to create and validate the shutdown plan. Pitt (Ex. 1007) at 4:13-65. In addition, the shutdown rules in Pitt are installed on each device, not in

a database, and are executed automatically when a shutdown event occurs. In Hamner, however, the user must manually select each device to perform a shutdown task. Petitioner never mentions how the combination of these two very different systems would be accomplished. Nor would a person of ordinary skill in the art have formed such a combination.

165. For these reasons, a person of ordinary skill in the art would not have combined Pitt with Hamner to form the claimed combination.

XI. CONCLUSION

166. For the foregoing reasons, it is my opinion that there is no showing that any of the challenged claims are unpatentable as obvious.

XII. AVAILABILITY FOR CROSS-EXAMINATION

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

XIII. RIGHT TO SUPPLEMENT

168. I reserve the right to supplement my opinions in the future to respond to any arguments that Petitioner raises and to take into account new information as it becomes available to me.

169. I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1101 of Title 18 of the United States Code.

Date: November 9, 2015

Daniel Menascé, Ph.D.

Exhibit A

Daniel A. Menascé, PhD
University Professor of Computer Science
www.cs.gmu.edu/faculty/menasce.html

Daniel A. Menascé is a University Professor of Computer Science at George Mason University where he has been since 1992. He was the Senior Associate Dean at the Volgenau School of Engineering at Mason from 2005 to 2012. He received a Ph.D. degree in Computer Science from UCLA in 1978 and held visiting positions at the University of Rome, Italy, and at the University of Maryland at College Park. Prior to joining Mason, he was a faculty member at the Department of Computer Science at the Pontifical Catholic University of Rio de Janeiro, Brazil, for 14 years, where he was also chair of Computer Science.

Menascé was elevated to the rank of Fellow of the IEEE in 2014 and was elected a Fellow of the Association of Computing Machinery (ACM) in 1997. He was elected a member of IFIP's Working Group 7.3 (Performance Evaluation) in 1998. The Computer Measurement Group awarded him the 2001 lifetime A.A. Michelson Award for "outstanding contributions to computer metrics." Menascé was selected a 2014 Outstanding Faculty Award Finalist by the State Council of Higher Education for Virginia (SCHEV).

He received the 2009 Outstanding Research Faculty Award from the Volgenau School of Engineering, the 2000 Mason Teaching Excellence Award, and the 1999 Outstanding Teaching Award from the Volgenau School of Engineering. Menascé was inducted in 2009 as an honorary member of the Golden Key International Honour Society, "in recognition of outstanding scholastic achievement and excellence." He was the recipient of various best paper awards at various conferences including 2014 and 2013 Computer Measurement Group (CMG) Conferences, 2011 SECURWARE, 1997 Computer Measurement Group (CMG) Conference, and 1995 IEEE International Conference on the Engineering of Complex Computer Systems. He received numerous service awards from the ACM and from the IEEE for his work as General Chair of major conferences.

He published over 240 refereed papers that received more than 9,200 citations. Menascé's h-index[1] is 46. He was the chief author of five books dealing with web technologies, e-commerce, and capacity planning. These books were published by Prentice Hall and translated into Russian, Korean, and Portuguese and have been adopted as text books in many universities in at least the following countries: US, Australia, Brazil, Canada, Dubai, Finland, Germany, Iran, Italy, and Malaysia. Menascé is an inventor on one U.S. patent and on two pending patents. He graduated 26 PhD students.

Menascé has been invited to give keynote addresses at many conferences and research events in the US and abroad (e.g., Belgium, Brazil, Chile, Germany, the Netherlands, Italy, and Portugal) and in Distinguished Lecture Series in US universities and research centers (e.g., IBM T.J. Watson and Microsoft Research). He

[1] The h-index h is the largest number h such that h publications have at least h citations.

gave an invited advanced doctoral course at the Gran Sasso Science Institute, L'Aquila, Italy.

His research has been funded by DARPA, the AFOSR, NASA, NSF, NIST, National Geospatial-Intelligence Agency (NGA), Virginia's Center for Innovative Technology, Dominion Virginia Power, OPNET Technologies, TRW, and Hughes Applied Information Systems. Menascé obtained more than $7.4 million in research grants.

Menascé served as the first elected Vice-Chair of ACM's Special Interest Group in E-commerce from 2003 to 2005. He was the General Chair of ACM's 2007 Federated Computing Research Conference (FCRC), the most prestigious research event organized by the ACM. This event aggregates 18 conferences over 10 days with the participation of around 2,000 people and a budget of $700,000. He also served as Program Committee Chair and General Chair of the major conferences in his field including IFIP's Performance, ACM Sigmetrics, ACM E-Commerce, IEEE ICCAC, and ACM WOSP. Menascé is a member of the editorial board of ACM Transactions on Internet Technologies, and of Elsevier's Performance Evaluation journal. He was an Associate Editor of ACM's Transactions on the Web (TWEB), an associate editor of Elsevier's Electronic Commerce Research and Applications journal from 2001 to 2006, and a member of the Editorial Board of IEEE Internet Computing for several years until 2008. During that time he guest edited two special issues.

Menascé led the design and implementation of several web-based distance education and training systems at Mason in the mid 90's in a project sponsored by the Department of Defense for the Defense Acquisition University. While Senior Associate Dean he designed and supervised the team that implemented several web-based information systems for faculty performance tracking, student evaluation, and research portfolio tracking and reporting at the Volgenau School of Engineering.

He consulted for the US Army, the Defense Information Systems Agency (DISA), the Ballistic Missile Defense Organization, NASA, the US Mint, the National Institutes of Health, the Center for Excellence in Space Data Information Systems, SABRE (travelocity.com), IBM, Lockheed Martin, United Online (netzero.com), TIS Labs at Network Associates, Hughes Applied Information Systems, and the InterAmerican Development Bank. Menascé served as an expert witness in various patent infringement cases and in an alleged misappropriation of trade secrets and confidential information case. He was deposed and testified several times in federal and state court as an expert witness in these litigations.

His areas of expertise include autonomic computing, databases and data storage systems, distributed systems, service oriented architectures, performance modeling and analysis, computer security, software performance engineering, e-commerce, and web technologies.

Menascé is a resident of the state of Maryland, a citizen of the United States, and is top secret qualified but currently cleared at the secret level by the U.S. Department of Defense.

Dr. Daniel A. Menascé
University Professor of Computer Science
George Mason University, Fairfax, VA 22030-4444
www.cs.gmu.edu/faculty/menasce.html

1. Education 2
2. Prizes and Distinctions Obtained 2
2.1 Research/Lifetime Distinctions 2
2.2 Teaching Distinctions 3
2.3 Service Distinctions 3
3. Regular Activities 4
4. Consulting Activities 6
4.1 Expert Witness Work 6
4.2 Other Consulting Work 10
5. Main System Development and Coordination Activities 13
6. Patents 14
7. Research Awards/Grants 14
8. Publications 17
8.1 Books 17
8.2 Journal Publications 18
8.3 Book Chapters and Papers in Encyclopedia 23
8.4 Publications in Conference Proceedings 24
8.5 Technical Reports and Unrefereed Publications 40
8.6 Theses 43
9. Scholarships Obtained 43
10. Teaching Experience 43
10.1 Graduate Courses 43
10.2 Undergraduate Courses 45
11. Theses Supervised and Under Supervision 45
11.1 Ph.D. Students Supervised 45
11.2 MS Students Supervised 47
11.3 MS Projects Supervised 51
11.4 Current Ph.D. Students 51
12. Presentations in Colloquia and Seminars 51
13. Participation in Committees 58
14. Professional Membership 65
15. Editorial and Other Services 65

1 Education

- Ph.D. in Computer Science, University of California at Los Angeles (UCLA), 1978.
- MS in Computer Science, Pontifical Catholic University (PUC-Rio), Rio de Janeiro, Brazil, 1975.
- BS in Electrical Engineering, Pontifical Catholic University (PUC-Rio), Rio de Janeiro, Brazil, 1974.

2 Prizes and Distinctions Obtained

2.1 Research/Lifetime Distinctions

- *Best Paper Award*, 2014 Computer Measurement Group Conference, November 2014.
- **Fellow of the IEEE**, for "contributions to research and education in performance evaluation of computer systems," 2014.
- **2014 Outstanding Faculty Award Finalist**, State Council of Higher Education of Virginia (SCHEV), (28 finalists out of 115 nominations).
- *Best Paper Award*, 2013 Computer Measurement Group Conference, November 2013.
- *Faculty Study Leave Award*, George Mason University, Fall 2013-Spring 2014.
- **University Professor**, designation given by George Mason University's Board of Visitors for "men and women of unusually great stature and eminence from the world of national and international achievement." January 2012.
- *Best Paper Award*, 2011 SECURWARE Conference, August 2011.
- *Outstanding Research Award.* Department of Computer Science, Volgenau School of Engineering, George Mason University, May 2009.
- *Outstanding Research Faculty Award,* Volgenau School of Engineering, George Mason University, April 2009.
- **A.A. Michelson Award**, Computer Measurement Group (CMG), lifetime achievement award given "for outstanding contributions to computer metrics," 2001.
- *Faculty Study Leave Award*, George Mason University, Fall 1999.

- *Member Elected to the IFIP Working Group 7.3* for "contributions and accomplishments in performance evaluation," 1998.
- **Fellow of the ACM**, for "fundamental contributions to education and practice of computer networks and performance evaluation, and material contributions to the establishment of a strong computing industry in Brazil." Inducted in 1997.
- *Best Paper Award*, 1997 Computer Measurement Group Conference, December 1997.
- *Outstanding Paper Award in the Systems Engineering Track*, First IEEE International Conference on Engineering of Complex Computer Systems, November 8, 1995.
- *Winner of the V National Prize in Informatics (Hardware Category)*, awarded by the Presidential Science and Technology Secretariat, Roberto Marinho Foundation and Moddata Inc., 1991, Brazil.
- *Best Monograph in Computer Science*, Special Secretariat of Informatics Contest, Brazil, 1983.
- *Best Technical Paper*, XIV National Informatics Congress, SUCESU, Brazil, October 1981.

2.2 Teaching Distinctions

- *GMU Teaching Excellence Award*, George Mason University, 2000.
- *Outstanding Teaching Award*, School of Information Technology and Engineering, George Mason University, 1999.
- *Outstanding Person in Education*, Manchete Magazine, Brazil, October 1987 (previous award recipients included Brazilian Ministers of Education).

2.3 Service Distinctions

- *Recognition of Service Award*, IEEE, in "recognition and appreciation of contributions to ICCAC 2015 as General Chair." September 2015.
- *Distinguished Service Award*, Computer Science Department, Volgenau School of Engineering, in recognition of accomplishments as Senior Associate Dean, May 2012.
- *Honorary Member*, Golden Key International Honour Society,"in recognition of outstanding scholastic achievement and excellence," inducted on November 6, 2009.

- *Recognition of Service Award*, Association for Computing Machinery, "in appreciation for contributions to ACM as the General Chair of the 2007 Federated Computing Research conference, June 2007.
- *Recognition of Service Award*, Association for Computing Machinery, "in appreciation for contributions to ACM as the General Chair of ACM Fourth Conference on E-commerce," June 2003.
- *Outstanding Reviewer*, IEEE Internet Computing, 2001.
- *Recognition of Service Award*, Association for Computing Machinery, "in appreciation for contributions to ACM as Program co-chair of WOSP 2000: 2nd Workshop on Software and Performance," September, 2000.
- *Recognition of Service Award*, Association for Computing Machinery, "in appreciation for contributions to ACM as the General Chair of ACM Sigmetrics Conference," May 1999.
- *Leadership Award*, Walt Whitman High School Media Center, Bethesda, MD, for his role in connecting the school to the Internet and providing each student with an individual e-mail account, June 1996.

3 Regular Activities

- University Professor of Computer Science, George Mason University, January 2012 to present.
- Professor of Computer Science, Department of Computer Science, George Mason University, September 1993 to present.
- Senior Associate Dean, Volgenau School of Engineering, George Mason University. Direct supervisor of research, graduate programs, promotion and tenure, graduate admissions, graduate student services, and the school's web site and web-based information system. From August 2007 to May 2012.
- Associate Dean for Research and Graduate Studies, Volgenau School of Engineering, George Mason University. Direct supervisor of research, graduate programs, promotion and tenure, graduate admissions, graduate student services, and the school's web site and web-based information system. From August 2005 to July 2007.
- Co-founder, Center for Smart Power Grids, George Mason University, April 2009 to present.
- Founding Director of the MS in E-commerce program, George Mason University, September 2001 to August 2005.

- Founding co-Director, E-center for E-business, School of Information Technology and Engineering, George Mason University, May 2001 to August 2005.
- Director, Hyperlearning Center (former Center for the New Engineer), School of Information Technology and Engineering, George Mason University, March 1998 to May 2001.
- Associate Director, Center for the New Engineer, School of Information Technology and Engineering, George Mason University, August 1993 to March 1998.
- Member of the Center for Information Systems Integration and Evolution, George Mason University, May 1994 to May 2001.
- Core Faculty Member, Computational Sciences Institute, George Mason University, September 1993 to 2003.
- Special Member of the Graduate Faculty, University of Maryland Graduate School, Baltimore, December 1994 to November 1997.
- Visiting Professor of Computer Science, Department of Computer Science, George Mason University, August 1992 to August 1993.
- Visiting Faculty, UMIACS, University of Maryland at College Park, July 1991 to July 1992.
- Professor of Computer Science, Computer Science Department, Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Brazil, June 1992 to July 1992.
- Associate Professor of Computer Science, Computer Science Department, Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Brazil, August 1981 to May 1992.
- Principal Investigator of the Project *Research in Parallel and Distributed Systems* sponsored by IBM Brazil, 1989 to July 1991.
- Director and Founder of *Tecnosoft Tecnologia de Software*, Brazil, December 1982 to May 1991.
- Director and Founder of *Capacidade Sistemas de Computação*, Brazil, 1980 to July 1991.
- President (Elected) of the Brazilian Computer Society, March 20, 1987 to April 20, 1989.
- Chairman of the Computer Science Department, Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Brazil, October 1981 to October 1983.
- Visiting Professor, Istituto Matematico Guido Castelnuovo, Universita di Roma, Italy, January to February 1983.

- Principal Investigator of the Project *Research and Development in Computer Networks*, sponsored by the Brazilian Telecommunications Company (Embratel), October 1980 to September 1984.
- Assistant Professor of Computer Science, Computer Science Department, Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Brazil, November 1978 to July 1981.
- Postgraduate Research Engineer in the Secure Systems and Software Architecture Group sponsored by the Advanced Research Projects Agency (ARPA), Computer Science Department, UCLA, January 1977 to November 1978.
- Software Engineer, in charge of basic software development in *Projeto Guaranys* (the project of the first Brazilian minicomputer), October 1973 to August 1975.
- Research Assistant, Computer Science Department, Pontifical Catholic University of Rio de Janeiro (PUC-RIO), Brazil, July 1970 to October 1973.

4 Consulting Activities

4.1 Expert Witness Work

- *EMC Corporation, EMC International Company, and EMC Information Systems International v. Pure Storage, Inc.,* Civil Action No. 13-1985 (RGA), District Court of Delaware, expert witness for a patent case against Pure Storage, Inc. Prepared two declarations in support of claim construction brief, an infringement report, a response to an invalidity report, a reply report to a rebuttal to an non-infringement report, and a declaration in support of Motion for Sumary Judgement. Was deposed on infringement and validity on August 20, 2015. September 30, 2014 to present.
- *Hewlett-Packard Co.* Technical expert for the case Hewlett-Packard Company v. ServiceNow, Inc., Case No. 14-cv-00570-BLF, Northern District of California, assisted attorneys with technical aspects of the case, including claim construction, and prepared a declaration in support of an HP's brief. June 25, 2014 to present.
- *EMC Corp. and Decho Corp.* Expert witness for non-infringement in the case "Oasis Research, LLC. v. AT&T Corp. et. al.," US Patents 5,771,354, 5,901,228, and 7,080,051, Case No.: 4:10-CV-435-ALM, Eastern District of Texas Sherman Division, Law firm for defense: Orrick, Herrington, Sutcliffe, LLP. Prepared two non-infringement reports (one for Mozy and one for Atmos), prepared a response report to the plaintiff's induced infringement report, assisted attorneys in the preparation of motion for summary judgement, and prepared a response to a supplemental report. Was deposed on November 9, 2012 and on July 22, 2015. July 2012 to September 2015. Case settled.

- *Apple Inc.* Technical expert for Inter-Partes Review related to U.S. Patents 7,496,854 and 7,917,843. Law firm: Morrison Foerster. Assisted attorneys with review of invalidity charts and prepared three declarations in support of petitioner's petitions for inter-partes review. Was deposed on August 7, 2014. November 2013 to March 2015. Case settled.

- *Apple Inc.* Technical expert for Covered Business Method Review related to U.S. Patent 8,533,860. Law firm: Morrison Foerster. Prepared a draft of a declaration in support of petitioner's petition for covered business method review. November 20, 2014 to March 2015. Case settled.

- *Microsoft Corporation.*, Expert witness for invalidity and non-infringement in the case "Sentius International LLC v. Microsoft Corporation," U.S. Patents RE43,633, RE40,731, and 7,672,985. Case No. 5:13-cv-00825 PSG, District of Delaware, Law firm for defendant: Fish & Richardson. Prepared invalidity report and non-infringement report. March 2014 to January 2015. Case settled.

- *Apple Inc.* Expert witness for invalidity and non-infringement in the case "Robocast Inc. v. Apple Inc.," U.S. Patent No. 7,155,451 B1, Case No. 11-235 (RGA), District of Delaware, Law firm for defense: Simpson Thacher & Bartlett LLP. Assisted attorneys with questions on claim construction, prepared invalidity and non-infringement reports. Was deposed on invalidity on September 13, 2013 and on non-infringement on October 3, 2013. June 2012 to October 8. 2014. Case dismissed by the Court on October 8, 2014.

- *Oracle Corp.*, Expert witness for invalidity and non-infringement in the case IpLearn LLC v. Oracle Corp. et al., US Patents 6,126,448; 6,213,780; 6,398,556; 6,685,478; 5,779,486; 6,118,973; 6,688,888; RE38,432; and RE39,942. Law Firm: WilmerHale. Prepared invalidity and non-infringement reports and assisted attorneys on various technical issues. Was deposed on invalidity on May 15, 2014 and on non-infringement on May 16, 2014. May 2013 to June 2014. Case settled prior to trial.

- *Microsoft Corporation.* Expert witness for non-infringement in the case "Walker Digital, LLC v. Activision Blizzard, Inc. et al.," U.S. Patent No. 5,970,143, Case no. 11-CV-322-SLR, District of Delaware. Law firm for defense: Fish & Richardson. Prepared non-infringement report and was deposed on July 10, 2013. September 2012 to January 2014. Case dismissed on January 2014.

- *Unified Messaging Solutions, LLC.*, Consultant in the case "Unified Messaging Solutions, LLC v. Facebook et. al," Law Firm: Nelson Bumgardner Casto, P.C.. Assisted attorneys in the interpretation of technical terms and prepared a declaration for the Court regarding claim construction. March 2012 to present.

- *Dassault Systemes Enovia Corp. and Dassault Systemes SolidWorks, Corp.* Expert witness in the case "Sky Technologies v. Microsoft Corporation et al." US Patent 6,141,653. Case 1:11-cv-10833-WGY. Law firm for the defense: Quinn Emanuel Urquhart & Sullivan, LLP. Assisted attorneys in preparation of draft of invalidity report. February 2012 to April 2012. Judge W.G. Young entered order of stipulation and dismissal.

- *Manhattan Associates, Inc.* Expert witness in the case "Sky Technologies v. Microsoft Corporation et al." US Patent 6,141,653. Case 1:11-cv-10833-WGY. Law firm for the defense: Kilpatrick Townsend & Stockton. Assisted attorneys in preparation of draft of invalidity report. February 2012 to June 2012. Claims dismissed with prejudice.

- *JDA Software Group.* Expert witness in the case "Sky Technologies v. Microsoft Corporation et al." US Patent 6,141,653. Case 1:11-cv-10833-WGY. Law firm for the defense: DLA Piper. Assisted attorneys in preparation of draft of invalidity report. February 2012 to April 2012. Judge W.G. Young entered order of stipulation and dismissal.

- *Siemens USA Holdings, Siemens Corporation, and Siemens Product Lifecycle Management Software, Inc.* Expert witness in the case "Sky Technologies v. Microsoft Corporation et al." US Patent 6,141,653. Case 1:11-cv-10833-WGY. Law firm for the defense: McDermott Will & Emery LLP. Assisted attorneys in preparation of draft of invalidity report. February 2012 to June 2012. Claims dismissed with prejudice.

- *Microsoft Corporation.* Expert witness in the case "Sky Technologies v. Microsoft Corporation et al." US Patent 6,141,653. Case 1:11-cv-10833-WGY. Law firm for the defense: Orrick, Herrington & Sutcliffe LLP. Assisted attorneys in invalidity chart preparation. July 2011 to December 2011. Case settled.

- *SellerBid, Inc.* expert witness for infringement and invalidity in the case "SellerBid Inc., v. Groupon, Inc.; Hungry Machine , Inc.; D/B/A LivingSocial.Com; BuyWithMe.Com; and OpenTable, Inc." US Patents 7,674,024 and 7,983,616. Assisted attorneys in claim construction, assisted attorneys in preparation of infringement charts, and prepared partial infringement report. From August 2011 to December 2011. Case settled.

- *Virginia E-commerce Solutions*, expert witness for infringement and invalidity in the case "Virginia E-commerce Solutions v. eBay and Paypal," US Patent US RE40,753. Law firms for the plaintiff: The Simon Law Firm, P.C. and Harney, Dickey & Pierce, PLC. Prepared infringement report, rebuttal to invalidity report, supplemental to infringement report, assisted attorneys in claim construction and summary judgement motion. From February 2011 to June 2011.

- *Web Tracking Solutions, LLC and Daniel Wexler*, technical advisor for infringement and invalidity in the case "Web Tracking Solutions, LLC and Daniel Wexler v. Google, Inc.," US

Patent 5,960,409. Law firms for the plaintiff: Simmons Browder Gianaris Angelidis & Barnerd LLC, Eckert Seamans Cherin & Mellot, LLC, Hanly Conroy Bierstein Sheridan Fisher Hayes LLP, From December 2009 to September 2012. Case settled.

- *Performance Pricing*, non-testifying consultant for infringement and invalidity in the case "Performance Pricing v. Google et al.". US Patent 6,978,253 B2. Law firm involved: Dovel & Luner, LLP. From February 2009 to November 2009. Case settled.

- *SAP AG, SAP America, Inc. and Oracle Corporation*, expert witness in the patent infringement case "Sky Technologies LLC v. SAP AG, SAP America, Inc, and Oracle Corporation." Eastern District of Texas Marshall Division, US Patents 6,141,653 and 7,162,458. Oracle settled. Prepared invalidity report, assisted attorneys on various matters, prepared supplemental report, was deposed on May 14, 2008. Law firm involved: Day, Casebeer, Madrid & Batchelder, Cupertino, CA. From January 16, 2008 to August 2010. Case settled.

- *SAP America, Inc. and Business Objects Americas*, expert witness in the patent infringement case "Diagnostics Systems Corp. v. SAP and Business Object," Central District of California, US Patent 5,537,590. Case settled before trial. Prepare draft of non-infringement report. Law firm: Jones Day. From August 2008 till August 2009.

- *Ariba, Inc.*, expert witness in the patent infringement case "Emptoris Inc. v. Ariba, Inc.", Eastern District of Texas Lufkin Division, U.S. Patent no. 6,519.590. Assisted in the preparation of invalidity charts, reviewed prior art, assisted in the preparation of preliminary invalidity contentions, and assisted in claims construction. Law firm involved: Heller Ehrman, Menlo Park and San Francisco offices. In October, 2008, the case moved to Covington & Burlington, LLP. From April 2008 to November 2008. Emptoris dismissed all its claims against Ariba without prejudice.

- *Ariba, Inc.*, expert witness in patent infringement case "Sky Technologies v. Ariba, Inc.", Eastern District of Massachusetts, U.S. Patents nos.: 6,141,653 and 7,162,458. Assisted in claims construction, prepared invalidity and non-infringement reports, helped in preparing questions for opposing expert witness' deposition, helped with motion for summary judgement, prepared supplemental report to Court, and testified at trial on December 2007. Law firm involved: Heller Ehrman, Menlo Park and San Francisco offices. From March 1, 2007 to December 2007.

- *SAP America, Inc. and SAP AG*, expert witness in patent infringement case "ePlus, Inc. v. SAP America, Inc. and SAP AG", Eastern District Court of Virginia, Richmond, U.S. Patent nos: 6,023,683, 6,055,516, and 6,505,172. Prepared invalidity and non-infringement reports and testified at trial. Law firm involved: Day, Casebeer, Madrid & Batchelder, Cupertino, CA. From September 2005 to April 2006.

- *Ford Motor Co., Honda Corporation, and DaimlerChrysler*, expert witness in the patent infringement case "Orion IP v. For Motor Company et al., 2:04-CV-313," Eastern District of Texas, U.S. Patent No. 5,615,342. From April 2005 to February 2006. Case settled before trial. Law firms involved: Jones Day, Fish & Neave IP Group of Ropes & Gray, LLP.

- *Toyota Motor Company*, expert witness in the patent infringement case "Orion IP v. Staples et al., 2:04-CV-297," Eastern District of Texas, U.S. Patent No. 5,615,342. From April 2005 to May 2005. Case settled before trial. Law firm involved: McDermott, Will & Emery.

- *ApplyYourself*, Fairfax, VA, expert witness in the alleged patent infringement case CollegeNet v. ApplyYourself, Tried in Oregon's Federal Court. Prepared invalidity and non-infringement reports, testified at Markman Hearing and at trial. September 2002 to April 2003.

- *Actuate Corporation*, CA, expert witness in the alleged trade secret and confidential information misappropriation case Microstrategy vs. Actuate. Prepared report and testified at trial in VA's state court. September 2002 to March 2003. Law firm: Sonnenschein, Nath & Rosenthal, LLP.

4.2 Other Consulting Work

- *United States Army, Global Command and Control System-Army*, capacity planning study for database server for GCCS-A, Fort Belvoir, VA, January-March 2005.

- *United Online*, California, analysis of the scalability of Netzero's Highspeed ISP service, May/June 2003.

- *Jones International University*, Colorado, content expert in charge of developing the course "IT 430 - Innovation, Integration, and Technology in the Business," March 2002 to August 2002.

- *Keynote Systems*, San Mateo, CA, development of a white paper on benchmarking, load testing, and application performance management, February to April, 2002.

- *WebOS*, Columbia, MD, Senior Consultant for Performance and Scalability, February 2001 to May 2001.

- *Peakstone Corporation*, Sunnyvale, CA, technological assessment of Peakstone's eAssurance product, February 2001.

- *NASA Goddard Space Flight Center*, Greenbelt, MD, development of performance models and a tool for the analysis of FTP performance between SCFs and NASA DAACs, October 1999 to January 2001.

- *NASA Goddard Space Flight Center*, Greenbelt, MD, performance analysis of EOSDIS Core System, September 1998 to December 2000.
- *US Mint*, Washington, D.C., Capacity Planning for their e-commerce site, April 2000 to August 31, 2000.
- *United Arab Emirates University*, member of the team that designed a curriculum for their new School of Information Technology. I was responsible for designing the degree program in e-commerce, May 2000 to July 2000.
- *Lockheed Martin Corporation*, Fairfax, VA, design of a Knowledge Management Framework, November 1999 to July 31, 2000.
- *IBM*, consultant for the design of models to be used in a capacity planning tool for e-commerce, October 1999 to January 2000.
- *Trusted Information Systems Lab at Network Associates*, Tysons Corner, VA, formulation of a conceptual framework to model, measure, and analyze the survivability of security services, September 1998 to February 1999.
- *Defense Information Systems Agency*, DC, member of the grey-beard team responsible for capacity planning for the JOPES 2000 system, November 1998 to January 1999.
- *SABRE*, Dallas, TX, evaluated the capacity planning methodology used by www.travelocity.com—an e-commerce site powered by SABRE—October 1998 to December 1998.
- *Ballistic Missile Defense Organization (BMDO)*, through Futron Corporation, Washington, DC, development of a methodology for capacity planning for BMDO's LAN and intranet environment, October 1997 to June 1998.
- *US Army*, through Raven, Inc., VA, capacity planning for the migration to client server system of a US Army Personnel system, June 1996 to April 1998.
- *National Institutes of Health*, through Scientific Applications International Corporation (SAIC), VA, network capacity planning, June 1994 to February 1996.
- *Center for Excellence in Space Data and Information Sciences (CESDIS)*, NASA Goddard Space Flight Center, Greenbelt, MD, development of performance models of mass storage systems and consulting to the CESDIS director on performance related matters, April 1995 to August 1997.
- *Hughes Applied Information Systems (HAIS)*, Landover, MD, development of analytic performance models of NASA's EOSDIS Core System, August 1995 to November 1995.

- *University of Maryland Graduate School, Baltimore*, Maryland, external reviewer of the graduate program in Computer Science, February 1994.
- *Technology Transfer Institute*, Santa Monica, CA, Analysis of the Brazilian Market for Training in Computer Science, April 1993.
- *Interamerican Development Bank (IDB)*, Washington, D.C., USA, capacity planning study for IBM MVS/ESA DB2 environment, November 1992 to February 1993.
- *Ashton-Tate*, USA, Report on the Similarity between dBASE III Plus and other products. February 1990 to date.
- *Brahma*, Rio de Janeiro, Brazil, Capacity Planning Study for IBM VM/VSE Environments, May 1991 to July 1991.
- *National Social Security Data Processing Company (DATAPREV)*, Rio de Janeiro, Brazil, Capacity Planning Study for the Workers and Employers Database System, February 1991 to May 1991.
- *Lotus Desenvolvimento de Software*, Rio de Janeiro, Brazil, report on the similarity between two spreadsheet systems. December 1989 to January 1990.
- *DYTZ Informatica*, Brasilia, Brazil, report on the Automation Project of the Steel Plant COSIPA S.A., July 1986 to August 1986.
- *Medidata Informatica e Tecnologia S/A*, Rio de Janeiro, Brazil, design of the software of a X.25 front-end processor, July 1982 to January 1983.
- *Federal Data Processing Service (SERPRO)* , Brasilia, Brazil. Member of a committee that visited several educational and research institutions in Germany. May 1982.
- *Brazilian Telecommunications Agency (EMBRATEL)*, Rio de Janeiro, Brazil, design and implementation of a training program in computer networks (applied to over 250 engineers), January 1980 to 1983.
- *Medidata Informatica e Tecnologia S/A*, Rio de Janeiro, Brazil, design of a reliable B-Tree based access method for the MUMPS system, August 1980 to February 1981.
- *Planed - Planning and Engineering*, Rio de Janeiro, Brazil, Preparation of an Electrical Engineering curriculum with emphasis in Computer Science for FESP-SP and Estacio de Sa colleges, December 1980.
- *CITIBANK N.A.*, Rio de Janeiro, Brazil, Planning of the data processing activities of the bank for the 80's. January to March 1981.

5 Main System Development and Coordination Activities

- Chief designer of the eAdmit system, a web-based system for paperless processing of graduate applications, Volgenau School of Engineering, George Mason University, 2011-2012.
- Chief designer and coordinator of the implementation of the eResearch system, a web-based system for tracking and reporting on the research proposals and awards, Volgenau School of Engineering, George Mason University, 2009-10.
- Chief designer and coordinator of the implementation of the eGrad system, a web-based system for tracking and reporting on the progress of graduate students, Volgenau School of Engineering, George Mason University, 2008.
- Chief designer and coordinator of the implementation of the ePAR system, a web-based system that supports the Professional Activity Reports (PAR) of the faculty of the Volgenau School of Engineering, George Mason University, 2006 to 2007.
- Chief designer and leader of the development team that implemented the Hyperlearning Meter system, a Web-based tool for self-assessment and certification, George Mason University, 1995 to 1999.
- Design and implementation of CLISSPE—a system for software performance engineering of C/S systems. CLISSPE is a language used to describe C/S systems under development. The compiler for CLISSPE generates parameters for queuing network models and solves the corresponding model. The system is being used at the US Army, 1996.
- Design and implementation of CMWLan — a tool for capacity planning in LAN environments. The tool is an add-in to Excel 5.0 and was developed in Visual Basic. The tool is being deployed to all institutes, centers, and divisions of the National Institutes of Health, 1994-1995.
- Coordination of a distributed database project for student records on a client-server environment using Lotus Notes, George Mason University, 1993.
- Design and implementation of QSolver/1, a capacity planning tool based on analytic models that runs on IBM PC compatible computers, Tecnosoft Tecnologia de Software, 1989-1990.
- Coordination of the project of an electronic directory system for the Brazilian Telecommunications Company. The system runs on a dedicated VAX 780 and is connected to the public X.25 network, to the public telex network and to a public X.400 service, Tecnosoft Tecnologia de Software, 1987-1990.

- Design and participation in the implementation of ARCOIRIS, a database management system for IBM PC compatible computers, Tecnosoft Tecnologia de Software, 1987-1990.
- Coordination of the project of a retrieval information system about the artwork and life of the Brazilian painter Candido Portinari, Tecnosoft Tecnologia de Software, 1985-1987.
- Coordination of the project of a retrieval information system about the works of the most important Portuguese language writers, Tecnosoft Tecnologia de Software, 1984-1985.

6 Patents

- *Meta-Protocol*, by I.S. Abdullah and D.A. Menascé, US Patent 8,086,744 issued on December 27, 2011.
- *System and Method for Managing Insider Security Threats*, US Patent Application 13/180,151, US PTO, Ghassan Jabbour and Daniel A. Menascé, filed on July 11, 2011.

7 Research Awards/Grants

- Resilient Autonomic Software Systems (RASS), $1,016,041, Air Force Office of Scientific Research (AFOSR), PI, co-PI: Hassan Gomaa, October 15, 2015 to October 14, 2018.
- A Process Analytics Framework for Sustainable Manufacturing Modeling, Analysis, and Decision Optimization, $398,494, National Institute of Standards and Technology, co-PI, PI: Alex Brodsky, September 1, 2012 to February, 2016.
- Decision Guidance Approach to Power Optimization and Management, $25,000, Dominion Virginia Power, co-PI, PI: Alex Brodsky, co-PI: Bob Simon, May 25, 2013 to December 31, 2013.
- SASSY: Self-Architecting Software Systems, US $479,962, National Science Foundation grant CCF-0820060, PI: Daniel A. Menascé, co-PIs: Hassan Gomaa, Joao Sousa, and Sam Malek, June 15, 2008 to May 31, 2012.
- Knowledge Sifter: Ontology-based Search over Corporate and Open Sources using Agent-Based Knowledge Services, US $ 640,000, National Geospatial-Intelligence Agency (NGS) NURI program, PI: Larry Kerschberg, co-PI: Daniel Menascé, October 1, 2003 to September 30, 2007.
- Analysis of E-commerce Activities in the Home Building Market, Maintainum LLC, $6,000.00, 6/13/2005-10/31/2005.

- A Framework for the Dynamic Composition and Reconfiguration of QoS-Aware Next Generation Software System, NSF, $59,939.00, PI: Daniel Menascé, co-PI: H. Gomaa, 9/15/2002 to 1/31/04.
- A First Course Based on Complex Problem-Solving Within a Simulated Entrepreneurial Market, $180,000, National Science Foundation, PI (since 01/01/2003): Daniel Menascé, Peter Denning (PI till 12/31/02), other co-PIs: Peter Paris and Nada Dabbagh, December 1, 2000 to August 30, 2003.
- Designing Self-Tunable E-Business Servers Using Data Warehousing and Performance Modeling Techniques, $20,000, TRW, PI: Daniel Menascé, co-PI: Daniel Barbara, December 2000 to November 2001.
- The E-Center of Excellence for Research and Education in E-Business, $75,000, Virginia's Center for Innovative Technology (CIT), co-PI: Daniel Menascé with Larry Kerschberg, July 15, 2000 to July 14, 2001.
- A Study on Integrating Simulation Models with Analytic Models for Server Performance, $61,991, OPNET Technologies, PI: Daniel Menascé, May 2000 to October 31 2000.
- Designing and Evaluating Reusable Component Interconnection Patterns for Evolvable Distributed Software Architectures, $85,000, NSF, CCR-9804113, co-PI: Daniel Menascé with Hassan Gomaa, September 1998 to August 1999.
- MetaWorld: A Quality-of-Service Agent-Based Information Repository, $1,479,192, co-PI: Daniel Menascé, with Kerschberg, Gomaa, Motro, Bose, and Brodsky, DARPA, selectable for funding.
- Web-based Training, Self-Assessment, and Certification for Meeting Professionals, $8,800, CIMPA, PI: Daniel Menascé, July 1998 to February 1999.
- Scalability Analysis of ECS's Data Server, $105,600, PI: Daniel Menascé, NASA through the Center of Excellence in Space Data and Information Sciences (CESDIS) at NASA Goddard, PI, December 1996-August 1998.
- Innovative Ideas in Higher Education: HPC in the Curriculum, $3,228,656, ARPA, DoD, Grant no. DABT63-93-C-0026, co-PI: Daniel Menascé with Peter Denning, 1993-1997.
- A Prototype of a Web-based Course on Network and Data Security, PI: Daniel Menascé, January to April, 1997, US$20,510, ACM.
- Educating Engineers to Design Complex Systems, $156,000, National Science Foundation, NSF Grant EEC-9315476, co-PI: Daniel Menascé, 1993-1995.

- A Scientific and Technical Evaluation of EOSDIS Core System (ECS) for Earth Science Research, US$ 342,406, Hughes Applied Information Systems, 1994.
- Support for the Implementation of the Design Exhibition Course, $5,900, Zero-Based Curriculum Project, GMU, PI: Daniel Menascé, Spring'94.
- Capacity Management Education and Training Program, Defense Information Systems Agency (DISA), Senior Engineer in the grant # 5-25815, Donald Gantz (PI), Dec. 15, 1992 to January 15, 1993, $12,157 applied to academic year salary.
- Scheduling in Heterogeneous Parallel Architectures, $20,000, National Science Foundation, Supplement to Grant CCR-9002351 awarded to Satish Tripathi, 1991.
- Educational Award for Sabbatical Research, $18,000, PI: Daniel Menascé, IBM Brasil, 1991.
- Educational Award for Sabbatical Research, $24,470, PI: Daniel Menascé, CAPES, Ministry of Education, Brazil, 1991.
- Organization of the III Brazilian Symposium on Computer Architecture and Parallel Processing, PI: Daniel Menascé, US$100,000, Research Support Foundation of the State of Rio de Janeiro (FAPERJ), 1991.
- Research in Parallel and Distributed Systems, $25,000.00, PI: Daniel Menascé, IBM Brasil, 1989 to 1991.
- Travel Grant to Participate in the ACM-IEEE Supercomputing'90 Conference, PI: Daniel Menascé, New York, NY, $2,200, Research Support Foundation of the State of Rio de Janeiro (FAPERJ), 1990.
- Travel Grant to Participate in the RIAO'88 Conference, Boston, USA, $2,300, PI: Daniel Menascé, Brazilian Research Council, 1988.
- Travel Grant to Participate in the CD-ROM Conference, New York, NY, $2,200, PI: Daniel Menascé, Brazilian Research Council, 1987.
- Travel Grant to Participate in the Second International Symposium on the Performance of Computer Communication Networks, Zurich, Switzerland, $3,200, PI: Daniel Menascé, Brazilian Research Council, 1984.
- Travel Grant to Participate in the 1984 ACM SIGMETRICS Conference, Boston, August 22-24, $2,000, PI: Daniel Menascé, Brazilian Research Council (CNPq), 1984.
- Research and Development in Computer Networks, $250,000, Brazilian Telecommunications Company PI: Daniel Menascé, (Embratel), 1980-1984.

- Travel Grant to Participate in the 10th European Computer Manufacturers Association (ECOMA) Conference on Computer PI: Daniel Menascé, Measurement, Munich, Germany, $3,000, Brazilian Research Council, 1982.
- Travel Grant to Participate in the ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, Los Angeles, USA, PI: Daniel Menascé, $2,500, Brazilian Research Council, 1981.
- Travel Grant to Participate in the XII Computer Measurement Group Conference (CMG), New Orleans, USA, PI: Daniel Menascé, $2,200, Brazilian Research Council, 1981.
- Travel Grant to Participate in the IEEE Symposium on Reliability in Distributed Software and Database Systems, Pittsburgh, PA, PI: Daniel Menascé, $2,300, Brazilian Research Council, 1981.
- Travel Grant to Participate in the Sixth International Conference on Very Large Databases (VLDB), Montreal, Canada, $3,000, Brazilian Research Council, PI: Daniel Menascé, 1980.
- Travel Grant to Participate in the 1979 IEEE Computer Software and Applications Conferences (COMPSAC), Chicago, Illinois, PI: Daniel Menascé, $2,100, Brazilian Research Council, 1979.

8 Publications

8.1 Books

1. *Performance By Design: Computer Capacity Planning by Example*, D.A. Menascé, V.A.F. Almeida, and L. Dowdy, Prentice Hall, Upper Saddle River, NJ, 2004, 462 pages.
2. *Capacity Planning for Web Services: metrics, models, and methods*, D.A. Menascé and V.A.F. Almeida, 2002, 572 pages, Prentice Hall, Upper Saddle River, New Jersey. Translated into Russian and Portuguese.
3. *Scaling for E-business: technologies, models, performance, and capacity planning*, D.A. Menascé and V.A.F. Almeida, 2000, 449 pages, Prentice Hall, Upper Saddle River, NJ. Translated into Korean and published by Sung Woo Publishing Co.
4. *Capacity Planning for Web Performance: metrics, models, and methods*, D.A. Menascé and V.A.F. Almeida, 321 pages, 1998, Prentice Hall, Upper Saddle River, New Jersey.
5. *Capacity Planning and Performance Modeling: from mainframes to client-server systems*, D.A. Menascé, V.A.F. Almeida, and L. Dowdy, 412 pages, 1994, Prentice Hall, Upper Saddle River, NJ.

6. *Capacity Planning of Computer Systems: Operational Analysis as a Tool*, D.A. Menascé and V. A. F. Almeida, 83 pages, Editora Campus, Rio de Janeiro, Brazil, 1985 (in Portuguese).

7. *Computer Networks – Technical and Operational Aspects*, D.A. Menascé and D. Schwabe, 160 pages, Editora Campus, Rio de Janeiro, Brazil, 1984 (in Portuguese). Also translated into Spanish and published in Madrid, Spain, by Editora Paraninfo, 1988.

8.2 Journal Publications

1. *Autonomic Smart Manufacturing*, D.A. Menascé, M. Krishnamoorthy, and A. Brodsky, Journal of Decision Systems, Taylor & Francis, special issue on Integrated Decision Support Systems, eds. Isabelle Linden et al., Vol. 24(2), June 2015, pp. 206-224. DOI: 10.1080/12460125.2015.1046714.

2. *Predicting the Effect of Memory Contention in Multi-core Computers Using Analytic Performance Models*, S. Bardhan and D.A. Menascé, IEEE Transactions on Computers, in press, available online at http://dx.doi.org.mutex.gmu.edu/10.1109/TC.2014.2361511.

3. *Efficient Response Time Approximations for Multiclass Fork and Join Queues in Open and Closed Queuing Networks*, F. Alomari and D.A. Menascé, IEEE Transactions on Parallel and Distributed Systems, Vol. 25 (6), June 2014, pp. 1437–1446.

4. *The Meta-Protocol Framework*, I.S. Abdullah and D.A. Menascé, The Journal of Systems & Software, Elsevier ScienceDirect, 86(11), November 2013, pp. 2711-2724.

5. *A Scalability Analysis of an Architecture for Countering Network-Centric Insider Threats*, F. Sibai and D.A. Menascé, Intl. J. Advances in Security, IARIA, Vol. 5, Nos. 1 and 2, July 2012, pp. 16–27.

6. *SASSY: A Framework for Self-Architecting Service-Oriented Systems*, D.A. Menascé, Hassan Gomaa, Sam Malek, and João P. Sousa, IEEE Software, November/Dec. 2011, pp. 78–85.

7. *On Optimal Service Selection in Service Oriented Architectures*, D.A. Menascé, E. Casalicchio, and V. Dubey, Performance Evaluation Journal, Elsevier, August 2010, Vol. 67, Issue 8, pp. 659–675.

8. *QoS Management in Service-Oriented Architectures*, D.A. Menascé, H. Ruan, and H. Gomaa, Performance Evaluation Journal, North Holland, Elsevier, Vol. 64, Nos. 7-8, August 2007, pp. 646-663. Ranked #1 among the Top Hottest papers in Performance Evaluation July to September 2007.

9. *Improving the Performance of Online Auctions Through Server-side Activity-Based Caching*, D.A. Menascé and V. Akula, World Wide Web Journal, Springer Verlag, Vol. 10, No. 2, June 2007, pp. 181–204.

10. *Two-Level Workload Characterization of Online Auctions*, V. Akula and D.A. Menascé, Electronic Commerce Research and Applications Journal, Elsevier, Vol. 6, No. 2, pp. 192–208, Summer 2007.

11. *Student Perceptions of Engineering Entrepreneurship: An Exploratory Study*, Nada Dabbagh and D.A. Menascé, Journal of Engineering Education, America Society of Engineering Education, April 2006, Vol. 95, No. 2, pp. 153–163.

12. *Scaling the Web: Wrapping it Up*, D. Menascé, Internet Computing, Vol. 9., No. 4, July/August 2005, pp. 92–95.

13. *Scalable Access to Scientific Data*, D. Menascé, Internet Computing, Vol. 9., No. 3, May/June 2005, pp. 78–80.

14. *MOM vs. RPC: Communication Models for Distributed Applications*, D. Menascé, Internet Computing, Vol. 9., No. 2, March/April 2005, pp. 90–93.

15. *Allocating Applications in Distributed Computing*, D. Menascé, Internet Computing, Vol. 9., No. 1, January/February 2005, pp. 90–92.

16. *Composing Web Services: A QoS View*, D. Menascé, Internet Computing, Vol. 8., No. 6, November/December 2004, pp. 88–90.

17. *Mapping Service Level Agreements in Distributed Applications*, D. Menascé, Internet Computing, Vol.8, No. 5, September/October 2004, pp. 100-102.

18. *MARVIN: A Web-Based System for Representing, Retrieving, and Visualizing Analogies,*, H.J. Foxwell and D.A. Menascé, World Wide Web: Internet and Web Information Systems Journal, Kluwer Academic Publishers, Vol. 7, 2004, pp. 385–419.

19. *QoS in Grid Computing*, D. Menascé and E. Casalicchio, IEEE Internet Computing, Vol.8, No. 4, July/August 2004, pp. 85–87.

20. *Performance and Availability of Internet Data Centers*, IEEE Internet Computing, D. Menascé, Vol. 8, No. 3, May/June 2004, pp. 94–96.

21. *QoS-aware software components*, D. Menascé, IEEE Internet Computing, March/April 2004, Vol. 8, No. 2, pp. 91–93.

22. *Response Time Analysis of Composite Web Services,* D. Menascé, IEEE Internet Computing, January/February 2004, Vol. 8, No. 1.

23. *Web Server Software Architectures*, D. Menascé, IEEE Internet Computing, November/December 2003, Vol. 7, No. 6, pp. 78–81.

24. *Workload Characterization*, D. Menascé, IEEE Internet Computing, September/October 2003, Vol. 7, No. 5.

25. *A Hierarchical and Multiscale Approach to Analyze E-Business Workloads*, D. Menasce, V. Almeida, R. Riedi, F. Ribeiro, R. Fonseca, and W. Meira Jr., Performance Evaluation Review, Volume 54, Issue 1, September 2003, pp. 33–57.

26. *Scaling Web Sites Through Caching*, D. Menascé, IEEE Internet Computing, July/August 2003, Vol. 7, No. 4, pp. 86–89.

27. *Security Performance*, D. Menascé, IEEE Internet Computing, May/June 2003, Vol. 7, No. 3.

28. *Scalable P2P Search*, D. Menascé, IEEE Internet Computing, March/April 2003, Vol. 7, No. 2, pp. 83–87.

29. *Automatic QoS Control*, D. Menascé, IEEE Internet Computing, January/February 2003, Vol. 7, No. 1, pp. 92–95.

30. *QoS Issues in Web Services*, D. Menascé, IEEE Internet Computing, November/December 2002, Vol. 6, No. 6, pp. 72–74.

31. *Characterizing E-Business Workloads Using Fractal Methods*, D.A. Menascé, B. Abrahao, D. Barbara, V. Almeida, F. Ribeiro, Journal of Web Engineering, Rinton Press, Vol 1, No. 1, 2002, pp. 74–90.

32. *A Methodology for Analyzing the Performance of Authentication Protocols*, A. Harbitter and D.A. Menascé, ACM Transactions on Information Systems Security, vol 5, no. 4, November 2002, pp. 458-491.

33. *Trade-offs in Designing Web Clusters*, D.A. Menascé, IEEE Internet Computing, pp. 76–80, September/Oct. 2002.

34. *Load Testing of Web Sites*, D. A. Menascé, IEEE Internet Computing, July/August 2002, pp. 70–74.

35. *Simple Analytic Modeling of Software Contention*, D. Menascé, ACM Sigmetrics Performance Evaluation Review, Vol. 29, No. 4, pp. 24–30.

36. *Capacity Planning: an Essential Tool for Managing Web Services*, V. Almeida and D. Menascé, IEEE IT Professional, July/August 2002.

37. *TPC-W: a benchmark for E-commerce*, D.A. Menascé, IEEE Internet Computing, May/June 2002, pp. 73–77.

38. *Web Performance Modeling Issues*, D.A. Menascé, International Journal of High-Performance Computing Applications, Volume 14, Number 4, Winter 2000, pp. 292–303.

39. *Business-oriented Resource Management Policies for E-commerce Servers*, D.A. Menascé, V. Almeida, R. Fonseca, and M. A. Mendes, Performance Evaluation, Vol. 42, Nos. 3-4, Oct. 2000, pp. 223-239.

40. *A Reference Model for Designing a Curriculum for E-commerce*, D.A. Menascé, IEEE Concurrency, March 2000.

41. *Analytic Modeling of Distributed Hierarchical Mass Storage Systems with Network-Attached Storage Devices*, O. Pentakalos, D. A. Menascé, and Y. Yesha, to appear in the IEEE Transactions on Parallel and Distributed Systems.

42. *A Method for Design and Performance Modeling of Client/Server Systems*, D.A. Menascé and H. Gomaa, IEEE Transactions on Software Engineering, Vol. 26, No. 11, Nov. 2000, pp. 1066–1085.

43. *Analytical Performance Modeling of Hierarchical Mass Storage Systems*, O.I. Pentakalos, D.A. Menascé, M. Halem, and Y. Yesha, IEEE Transactions on Computers, Vol. 46, No. 10, October 1997, pp. 1103-1118.

44. *Pythia and Pythia/WK: Tools for the Performance Analysis of Mass Storage Systems*, O.I. Pentakalos, D.A. Menascé, and Y. Yesha, Software Practice and Experience, Vol. 27 (9), September 1997, pp. 1035-1054.

45. *A Software Architectural Design Method for Large-Scale Distributed Data Intensive Information Systems* , H. Gomaa, D.A. Menascé, and L. Kerschberg, Journal of Distributed Systems Engineering, Vol. 3, 1996, pp. 162-172.

46. *Static and Dynamic Processor Scheduling Disciplines in Heterogeneous Parallel Architectures*, D.A. Menascé, D. Saha, S.C. da Silva Porto, V.A.F. Almeida, and S.K. Tripathi, Journal of Parallel and Distributed Computing, Vol. 28 (1), July 1995, pp. 1-18.

47. *A Theoretical and Experimental Assessment of Data Partitioning Strategies on Networks of Heterogeneous Workstations*, D.A. Menascé and A. Bangalore, special issue on Parallel Computation of the Journal of the Brazilian Computer Society, No. 1, Vol. 2, July 1995, pp. 5-12.

48. *On a Unified Framework for the Evaluation of Distributed Quorum Attainment Protocols*, D.A. Menascé, Y. Yesha, and K. Kalpakis, IEEE Transactions on Software Engineering, Vol. 20, No. 11, November 1994.

49. *Static Heuristic Processor Assignment in Heterogeneous Multiprocessors*, D.A. Menascé, S.C.S. Porto, and S.K. Tripathi, International Journal of High Speed Computing, Vol. 6, No. 1 (March 1994).

50. *Scheduling on Heterogeneous Message Passing Architectures*, Daniel A. Menascé and Stella C.S. Porto, Journal of Computer and Software Engineering, Ablex Publishing Co., New Jersey, Volume 1, Number 3, 1993.

51. *A Methodology for Performance Evaluation of Parallel Applications in Shared Memory Multiprocessors*, Daniel A. Menascé and L.A. Barroso, Journal of Distributed and Parallel Computing, Vol. 14, No.1, January 1992.

52. *Heterogeneous Supercomputing: Why is it Cost-Effective?*, D.A. Menascé and V.A.F. Almeida, Supercomputing Review, August 1991, Vol. 4, No. 8.

53. *Performance Evaluation with Stochastic Petri Nets*, D.A. Menascé and N.L.S. Fonseca, Brazilian Computing Journal, SBC, September 1990.

54. *Analytic Models of Supercomputer Performance in Multiprogramming Environments*, Daniel A. Menascé and V.A. F. Almeida, The International Journal of Supercomputer Applications, The MIT Press Journals, Vol 3.2, 1989.

55. *Correctness and Performance Evaluation of a Two Phase Commit Based Protocol for DDBs*, T. Nakanishi and D.A. Menascé, Computer Performance, IPC Press, Vol. 5, No. 1, March 1984.

56. *Performance Evaluation of Broadcast Type Local Networks*, L.L.P. Leite and D.A. Menascé, Brazilian Computing Journal, SBC, Vol. 4, No. 1, 1984.

57. *Operational Analysis of Multiclass Systems with Variable Multiprogramming Level and Memory Queueing*, D.A. Menascé and V.A.F. Almeida, Computer Performance, IPC Press, Vol. 3, No. 3, September 1982.

58. *Optimistic Versus Pessimistic Concurrency Control Mechanisms in Database Management Systems*, D.A. Menascé, and T. Nakanishi. Information Systems, Pergamon Press, Vol. 7, No. 1, 1981.

59. *Specification of a Reliable Storage Component for a Distributed Database Management System*, O.E. Landes and D.A. Menascé, Brazilian Computing Journal, SBC, Vol. 1, No. 2, 1981.

60. *A Study on Crash Recovery Techniques in Databases*, O.E. Landes and D.A. Menascé, Brazilian Computing Journal, SBC, Vol. 1, No. 1, 1981.

61. *Locking Protocol for Resource Coordination in Distributed Databases*, D.A. Menascé, G.J. Popek and R.R. Muntz, ACM Transactions on Database Systems (TODS), Vol. 5, No. 2, June 1980. Also published in the Proceedings of the 1978 ACM/SIGMOD International Conference on the Management of Data, Austin, Texas, 1978.

62. *Locking and Deadlock Detection in Distributed Databases*, D.A. Menascé and R.R. Muntz, IEEE Transactions on Software Engineering, Vol. 5, 1979, also published in the Proceedings of the Third Berkeley Workshop on Distributed Data Management and Computer Networks, San Francisco, Ca., 1978. Also published in the book "Tutorial: Distributed Data Base Management", edited by P.A. Bernstein, J.B. Rothnie and D.W. Shipman, IEEE Computer Society, Long Beach, Ca., November 1978. Also published in the book "Tutorial: Centralized and Distributed Data Base Systems", edited by W.W. Chu and P.S. Chen, IEEE Computer Society, Long Beach, Ca., October 1979.

63. *Modelling and Evaluation of Computer Systems: A Case Study*, D.A. Menascé, Brazilian Technology Journal (CNPq), September 1973.

8.3 Book Chapters and Papers in Encyclopedia

1. *Performance Management of Composite Applications in Service Oriented Architectures*, V. Dubey and D.A. Menascé, in Performance and Dependability in service computing: Concepts, Techniques and Research Directions, eds. V. Cardellini, E. Casalicchio, K. Castello Branco, J.C. Estrella, and F.K. Monaco, IGI Global, July 2011, ISBN13: 9781609607944.

2. *Dynamic Server Allocation for Autonomic Service Centers in the Presence of Failures*, D.A. Menascé and M.N. Bennani, in Autonomic Computing: Concepts, Infrastructure, and Applications, eds. S. Hariri and M. Parashar, CRC Press, invited, December 2006, pp. 361-376.

3. *On the Use of Online Analytic Performance Models in Self-Managing and Self-Organizing Computer Systems*, D.A. Menascé, M.N.Bennani, and H. Ruan, in the book Self-Star Properties in Complex Information Systems, O. Babaoglu, M. Jelasity, A. Montresor, C. Fetzer, S. Leonardi, A. van Moorsel, and M. van Steen, eds., Lecture Notes in Computer Science, Vol. 3460, Springer Verlag, 2005, invited, pp. 128–142.

4. *Performance Engineering of Component-Based Distributed Software Systems*, H. Gomaa and D. A. Menascé, chapter in the book "Performance Engineering: state of the art and current trends," eds. R. Dumke, C. Rautenstrauch, A. Schmietendorf, and A. Scholz, LNCS State of the Art Series, Vol. 2047, Springer-Verlag, invited, May 2001.

5. *Performance of Client/Server Systems*, D.A. Menascé and V.A.F. Almeida, invited chapter in the book "Performance Evaluation: Origins and Directions", eds. Guenter Haring, Christoph Lindemann, Martin Reiser, LNCS State of the Art Series, Springer-Verlag, 2000.

6. *Computer Systems*, D.A. Menascé, 1998 Yearbook of Science and Technology, October 1997, McGraw-Hill, NY, pp. 79-81.

7. *Locking and Deadlock Detection in Distributed Databases*, D.A. Menascé and R.R. Muntz, published in the books:

   - "Tutorial: Centralized and Distributed Data Base Systems", edited by W.W. Chu and P.S. Chen, IEEE Computer Society, Long Beach, Ca., October 1979.
   - "Tutorial: Distributed Data Base Management", edited by P.A. Bernstein, J.B. Rothnie and D.W. Shipman, IEEE Computer Society, Long Beach, CA, November 1978.

8. *Centralized and Hierarchical Locking in Distributed Databases*, D.A. Menascé, G. J. Popek, and R. R. Muntz, published in the book "Tutorial: Distributed Data Base Management", edited by P.A. Bernstein, J.B.Rothnic, and D.W. Shipman, IEEE Computer Society, Long Beach, CA, November 1978.

8.4 Publications in Conference Proceedings

1. *Modular Modeling and Optimization of Temporal Manufacturing Processes with Inventories*, M. Krishnamoorthy, A. Brodsky, and D.A. Menascé, Proc. 49th Hawaii Intl. Conf. System Sciences (HICSS), Jan. 5-8, 2016, Kauai, Hawaii.

2. *Prediction-Based Admission Control for IaaS Clouds with Multiple Service Classes*, M. Carvalho, D.A. Menascé, and F. Brasileiro, 7^{th} IEEE Intl. Conf. Cloud Computing Technology and Science (CloudCom), Vancouver, Canada, Nov. 30-Dec. 3, 2015.

3. *Modeling the Tradeoffs Between System Performance and CPU Power Consumption*, D.A. Menascé, Proc. 2015 Computer Measurement Group Conf, November 2-5, 2015, San Antonio, Texas.

4. *Automatic Workload Characterization Using System Log Analysis*, M. Awad and D.A. Menascé, Proc. 2015 Computer Measurement Group Conf, November 2-5, 2015, San Antonio, Texas.

5. *Analysis and Optimization in Smart Manufacturing: Toward Standards on Reusable Knowledge Base of Process Performance*, A. Brodsky, G. Shao, M. Krishnamoorthy, A. Narayanan, D.A. Menascé, and R. Ak, Proc. 2015 IEEE Intl. Conf. Big Data (Big Data), Special session on "From Data to Insight: Big Data and Analytics for Smart Manufacturing Systems," October 29-Nov. 1, 2015, Santa Clara.

6. *Near-optimal Allocation of VMs from IaaS Providers by SaaS Providers*, A. Aldhalaan and D.A. Menascé, Proc. 2015 IEEE Intl. Conf. Cloud Autonomic Computing (ICCAC), Cambridge, MA, Sept. 21–25, 2015.

7. *Autonomic Metaheuristic Optimization with Application to Run-Time Software Adaptation*, J. Ewing and D.A. Menascé, The Eleventh International Conference on Autonomic and Autonomous Systems (ICAS 2015), Rome, Italy, May 24-29, 2015.

8. *Optimizing Stochastic Temporal Manufacturing Processes with Inventories: An Efficient Heuristic Algorithm based on Deterministic Approximations,* Mohan Krishnamoorthy, Alexander Brodsky, and Daniel Menascé, Proc. INFORMS Computing Society Conf., Richmond, VA, Jan. 11-3, 2015.

9. *Temporal Manufacturing Query Language (tMQL) for Domain Specific Composition, What-if Analysis, and Optimization of Manufacturing Processes With Inventories,* Mohan Krishnamoorthy, Alexander Brodsky, and Daniel Menascé, Proc. INFORMS Computing Society Conf., Richmond, VA, Jan. 11-3, 2015.

10. *Dynamic Derivation of Analytical Performance Models in Autonomic Computing Environments*, M. Awad and D.A. Menascé, Performance and Capacity Conference, Computer Measurement Group, Atlanta, Georgia, November 3–6, 2014. (**Best Paper Award**)

11. *Toward Smart Manufacturing using Decision Analytics*, Alexander Brodsky, Mohan Krishnamoorthy, Daniel Menascé, Guodong Shao, and Sudarsan Rachuri, IEEE NIST Data to Insight: Big Data and Analytics for Smart Manufacturing Systems Workshop, held jointly with the IEEE BigData 2014 Conf., October 27-30, 2014, Washington, DC, USA.

12. *A Contention Aware Hybrid Evaluator for Schedulers of Big Data Applications in Computer Clusters*, S. Bardhan and D.A. Menascé, 2nd Workshop on Scalable Cloud Data Management, co-located with the IEEE BigData Conference, Washington DC, October 27th 2014.

13. *On the Predictive Properties of Performance Models Derived Through Input-Output Relationships,*, M. Awad and D.A. Menascé, European Performance Engineering Workshop (EPEW 2014), Florence, Italy, September 11-12, 2014, Lecture Notes in Computer Science, LNCS 8721, Springer, pp. 89–103.

14. *Autonomic Allocation of Communicating Virtual Machines in Hierarchical Cloud Data Centers*, A. Aldhalaan and D.A. Menascé, 2014 IEEE Cloud and Autonomic Computing Conf. (CAC 2014), London, UK, September 8-12, 2014, pp. 161–171.

15. *A Meta-Controller Method for Improving Run-Time Self-Architecting in SOA Systems*, J. Ewing and D.A. Menascé, 5^{th} ACM/SPEC Intl. Conf. Performance Engineering (ICPE 2014), Dublin, Ireland, March 23-26, 2014.

16. *The Anatomy of MapReduce Jobs, Scheduling, and Performance Challenges*, S. Bardhan and D.A. Menascé, Computer Measurement Group Intl. Conf., La Jolla, CA, November 5-7, 2013.

17. *A Performance Model for Core System Monitoring Functions*, K. Mitchell and D.A. Menascé, Computer Measurement Group Intl. Conf., La Jolla, CA, November 5-7, 2013 (**Best Paper Award**).

18. *Analytic Performance Modeling and Optimization of Live VM Migration*, A. Aldhalaan and D.A. Menascé, 10^{th} European Workshop on Performance Engineering (EPEW 2013), Springer LNCS 8168, eds. M.S. Balsamo, W.J. Knottenbelt, Ca' Foscari University, Venice, Italy, September 16-17, 2013, pp. 28–42.

19. *Analytic Performance Models of Applications in Multi-core Computers*, S. Bardhan and D.A. Menascé, 21st Intl. Symp. Modeling, Analysis and Simulation of Computer Systems and Telecommunication Systems (MASCOTS 2013), San Francisco, CA, August 14-16, 2013.

20. *Autonomic Resource Provisioning in Cloud Systems with Availability Goals*, E. Casalicchio, D.A. Menascé, and A. Aldhalaan, The ACM Cloud and Autonomic Computing Conference (CAC 2013), Miami, FL, August 5-9, 2013.

21. *Self-Protecting and Self-Optimizing Database Systems: Implementation and Experimental Evaluation*, F. Alomari and D.A. Menascé, The ACM Cloud and Autonomic Computing Conference (CAC 2013), Miami, FL, August 5-9, 2013.

22. *Queuing Network Models to Predict the Completion Time of the Map Phase of MapReduce Jobs*, S. Bardhan and D.A. Menascé, 2012 Computer Measurement Group Intl. Conf., Las Vegas, NV, Dec. 3-7, 2012.

23. *Countering Network-Centric Insider Threats Through Self-Protective Autonomic Rule Generation*, F. Sibai and D.A. Menascé, IEEE Sixth International Conference on Software Security and Reliability (SERE 2012), June 20-22, 2012, Washington, D.C., USA.

24. *An Autonomic Framework for Integrating Security and Quality of Service Support in Databases,* F. Alomari and D.A. Menascé, IEEE Sixth International Conference on Software Security and Reliability (SERE 2012), June 20-22, 2012, Washington, D.C., USA.

25. *The Design of an Autonomic Controller for Self-Managed Emergency Departments*, S. Al-Momen and D.A. Menascé, HEALTHINF 2012 - International Conf. Health Informatics, February 1-4, 2012, Algarve, Portugal.

26. *A Methodology for Combinining GSPNs and QNs*, D.A. Menascé, 2011 Computer Measurement Group Conf, Dec. 5-9, 2011, Washington, D.C.

27. *A Scalable Architecture for Countering Network-Centric Insider Threats*, F. Sibai and D.A. Menascé, The Fifth International Conference on Emerging Security Information, Systems and

Technologies (SECURWARE), August 21-27, 2011, French Riviera, Nice/Saint Laurent du Var, France. Outstanding Paper Award.

28. *A Decision-Guided Energy Management for Sustainable Manufacturing*, G. Shao, A. Brodsky, J. Arinez, D.A. Menascé and P. Ammann, Proc. ASME 2011 International Design Engineering Technical Conference & Computers and Information in Engineering Conference (IDETC/CIE), American Society of Mechanical Engineers, Washington, DC, August 28–31, 2011.

29. *Decision-Guided Self-Architecting Framework for Integrated Distribution and Energy Management* , J. Arinez, S. Biller, A. Brodsky, D.A. Menascé, G. Shao, and J.P. Sousa, Second Conference on Innovative Smart Grid Technologies (ISGT 2011), IEEE Power Society, Anaheim, CA, January 17-19, 2011.

30. *Defeating the Insider Threat via Autonomic Network Capabilities*, F. Sibai and D.A. Menascé, The Third International Conference on COMmunication Systems and NETworkS (COMSNETS), Bangalore, India, January 4-8, 2011.

31. *An Autonomic Computing Framework for Self-managed Emergency Departments*, S. Al-Momem and D.A. Menascé, HEALTHIT 2011 - Intl. Conf. Health Informatics, Rome, Italy, January 26–29, 2011.

32. *Method and Model to Assess the Performance of Clustered Databases: The Oracle RAC Case*, K. Molloy and D.A. Menascé, 2010 Computer Measurement Group Conf., Orlando, FL, Dec. 6–8, 2010.

33. *Utility-based Optimal Service Selection for Business Processes in Service Oriented Architectures*, V. Dubey and Daniel A. Menascé, The Eight Intl. Conf. Web Services, July 5-10, 2010, Miami, Florida

34. *QoS Architectural Patterns for Self-Architecting Software Systems*, Daniel A. Menascé, João P. Sousa, Sam Malek, and Hassan Gomaa, The 7th IEEE International Conference on Autonomic Computing and Communications Washington, DC, USA June 7-11, 2010.

35. *Software Adaptation Patterns for Service Oriented Architectures,* Hassan Gomaa, Koji Hashimoto, Minseong Kim, Sam Malek, and Daniel A. Menascé, 25th ACM Symposium on Applied Computing, Dependable and Adaptive Distributed Systems, Sierre, Switzerland, March 22 - 26, 2010.

36. *A Framework for Utility-Based Service Oriented Design in SASSY*, Daniel A. Menascé, John M. Ewing, Hassan Gomaa, Sam Malek, and João P. Sousa, First Joint WOSP/SIPEW International Conference on Performance Engineering (WOSP/SIPEW 2010), San Jose, California, USA January 28-30, 2010.

37. *Understanding Cloud Computing: Experimentation and Capacity Planning*, D.A. Menascé and P. Ngo, 2009 Computer Measurement Group Conference, Dallas, TX, Dec. 6-11, 2009.

38. *Optimal Service Selection Heuristics in Service Oriented Architectures*, E. Casalicchio, D.A. Menascé, V. Dubey, and L. Silvestri, The 3rd International Workshop on Advanced Architectures and Algorithms for Internet DElivery and Applications (AAA-IDEA 2009), in conjunction with QShine 2009, Las Palmas de Gran Canaria, The Canary Islands, Spain November 26, 2009.

39. *A Modeling Language for Activity-Oriented Composition of Service-Oriented Software Systems*, N. Esfahani, S. Malek, J.P. Sousa, H. Gomaa, D.A. Menascé, ACM/IEEE 12th International Conference on Model Driven Engineering Languages and Systems, Denver, CO, October 2-4, 2009.

40. *Business-Oriented Autonomic Load Balancing for Multitiered Web Sites*, J. Ewing and D.A. Menascé, 17th ACM/IEEE Intl. Symp. on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2009), London, UK, September 21-23, 2009.

41. *The Insider Threat Security Architecture,* G. Jabbour and D.A. Menascé, 2009 IEEE Intl. Symposium on Secure Computing (SecureCom09), Vancouver, Canada, August 29–31, 2009.

42. *Stopping the Insider Threat: the case for implementing integrated autonomic defense mechanisms in computing systems,*, G. Jabbour and D.A. Menascé, Intl. Conf. on Security and Privacy (ISP'09), Orlando, FL, July 13-16, 2009.

43. *Self-Architecting Software SYstems (SASSY): From QoS-Annotated Models*, S. Malek, N. Esfahani, D.A. Menascé, J.P. Sousa, and H. Gomaa, Principles of Engineering Service Oriented Systems (PESOS 2009), Vancouver, Canada, May 18–19, 2009.

44. *Computing Missing Service Demand Parameters for Performance Models*, D.A. Menascé, 2008 Computer Measurement Group Conference, Dec. 7–12, 2008, Las Vegas, NV.

45. *Securing Security Policies in Autonomic Computing Systems*, G. Jabbour and D.A. Menascé, 2008 International Conference on Security and Management (SAM'08), July 14-17, Las Vegas, NV.

46. *A Heuristic Approach to Optimal Service Selection in Service Oriented Architectures*, D.A. Menascé, E. Casalicchio, and V. Dubey, Seventh ACM Intl. Workshop on Software and Performance (WOSP 2008), Princeton, NJ, USA, June 23-26, 2008.

47. *Policy-Based Enforcement of Database Security Configuration through Autonomic Capabilities*, G. Jabbour and D.A. Menasce, Proc. Intl. Conference on Autonomous and Autonomic Systems (ICAS'08), Gosier, Guadeloupe, March 16-21, 2008.

48. *An Empirical Evaluation of Communication and Processing Overhead in Service Oriented Architectures*, D.A. Menascé and V. Dubey, Proc. 2007 Computer Measurement Group Conference, San Diego, CA, Dec. 2-7, 2007.

49. *A Business oriented Load Dispatching Framework for Online Auction Sites*, D.A. Menascé and V. Akula, Proc. IEEE Intl. Conf. Quantitative Evaluation of Systems (QEST'07), Edinburgh, Scotland, Sept. 16-19, 2007.

50. *Utility-based QoS Brokering in Service Oriented Architectures*, D.A. Menascé and V. Dubey, IEEE 2007 Intl. Conf. Web Services (ICWS 2007), Application Services and Industry Track, Salt Lake City, Utah, July 9-13, 2007, pp. 422–430.

51. *Analytic Performance Models for Single Class and Multiple Class Multithreaded Software Servers*, D.A. Menascé and M.N. Bennani, Proc. 2006 Computer Measurement Group (CMG) Conf., Reno, NV, Dec. 3–8, 2006.

52. *Autonomic Virtualized Environments*, D.A. Menascé and M.N. Bennani, IEEE International Conference on Autonomic and Autonomous Systems (ICAS'06), July 19–21, 2006, Silicon Valley, USA.

53. *Virtualization: concepts, applications, and performance,* D.A. Menascé, Proc. 2005 Computer Measurement Group Conference, Dec. 4–9, 2005, Orlando, FL, pp. 407–413.

54. *Server-side caching strategies for online auction sites*, D.A. Menascé and V. Akula, 6th International Conference on Web Information Systems Engineering, New York, NY, November 20-22, 2005, Springer Verlag, LNCS 3806, pp. 231–244.

55. *Evaluating Caching Policies for Online Auctions*, D.A. Menascé and V. Akula, Proc. 2005 IFIP 7.3 Performance Conference, Juan-les-Pins, France, October 2005. Also published in a Special Issue of ACM Sigmetrics Performance Evaluation Reviews (PER).

56. *A Method for Evaluating the Impact of Software Configuration Parameters on E-Commerce Sites*, M. Sopitkamol and D.A. Menascé, Proc. 2005 ACM Workshop on Software and Performance, Palma de Mallorca, Spain, July 11–14, 2005, 53–64.

57. *Resource Allocation for Autonomic Data Centers using Analytic Performance Models,* M. Bennani and D.A. Menascé, Proc. Second IEEE International Conference on Autonomic Computing, Seattle, WA, June 13–16, 2005, pp. 229–240.

58. *Complex Problem-solving within a Simulated Entrepreneurial Market: An Exploratory Case Study,* N. Dabbagh and D.A. Menascé, Proc. Annual Meeting of the American Educational Research Association, Montreal, Canada, April 11-15, 2005.

59. *Quality of Service Aspects and Metrics in Grid Computing*, D. Menascé and E. Casalicchio, Proc. 2004 Computer Measurement Group Conference, Las Vegas, NV, December 5-10, 2004.

60. *A Framework for Resource Allocation in Grid Computing*, D. Menascé and E. Casalicchio, Proc. 12th Annual Meeting of the IEEE / ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Volendam, The Netherlands, October 5-7, 2004.

61. *Improving the Performance of Online Auction Sites through Closing Time Rescheduling*, D.A. Menascé and V. Akula, First International Conference on Quantitative Evaluation of Systems (QEST), Twente, Netherlands, September 27-30, 2004.

62. *An analysis of bidding activity in online auctions*, V. Akula and D. Menascé, 5th International Conference on Electronic Commerce and Web Technologies (EC-Web'04), August 30-September 3, 2004, Zaragoza, Spain.

63. *Assessing the Robustness of Self-Managing Computer Systems under Highly Variable Workloads*, M. Bennani and D.A. Menascé, Proc. International Conf. Autonomic Computing (ICAC-04), New York, NY, May 17-18, 2004.

64. *A Framework for QoS-Aware Software Components*, D. A. Menascé, H. Ruan, and H. Gomaa, Proc. 2004 ACM Workshop on Software and Performance, San Francisco, CA, January 14–16, 2004.

65. *On the Use of Performance Models to Design Self-Managing Computer Systems*, D. A. Menascé and M. Bennani, Proc. 2003 Computer Measurement Group Conference, Dallas, TX, December 7-12, 2003.

66. *Protocol Specification and Automatic Implementation Using XML and CBSE*, I. S. Abdullah and D. A. Menascé, IASTED conference on Communications, Internet and Information Technology (CIIT 2003), November 17-19, Scottsdale, AZ.

67. *Towards Workload Characterization of Auction Sites*, D.A. Menascé and V. Akula, Proc. 6th Workshop on Workload Characterization (WWC-6), Austin, Texas, October 27, 2003.

68. *A Unified Architecture for the Implementation of Security Protocols*, I. S. Abdullah and D. A. Menascé, International Conference on Computer Applications in Industry and Engineering (CAINE03), Las Vegas, Nevada, November 11-13, 2003.

69. *Designing a Complex E-commerce Simulation*, R. Dodge and D. A. Menascé, Proc. Seventh World Multi-conference on Systemics, Cybernetics and Informatics, Orlando, FL, July 27-30, 2003, pp. 243–248.

70. *Load Testing, Benchmarking, and Application Performance Management for the Web*, D. Menascé, Proc. 2002 Computer Measurement Group Conference, Reno, NV, Dec. 2002.

71. *Web-based Representation and Visualization of Analogies*, H. Foxwell and D. A. Menascé, Proc. IADIS International Conference WWW/Internet 2002, Nov. 13-15, 2002, Lisbon, Portugal.

72. *Two-level Iterative Queuing Modeling of Software Contention*, D. A. Menascé, Proc. Tenth IEEE/ACM International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS 2002), Oct. 12-16, 2002, Fort Worth, Texas.

73. *Probabilistic Scalable P2P Resource Location Services*, D. A. Menascé and Lavanya Kanchanapalli, 2002 Practical Aspects of Performance Analysis, held jointly with ACM Sigmetrics 2002, Marina del Rey, CA, June 15, 2002, also published in the ACM Sigmetrics Performance Evaluation Review, Vol. 30, No. 2, September 2002, pp. 48–58.

74. *Software, Performance, or Engineering?*, D. Menascé, Proc. WOSP 2002 Conference, Rome, Italy, July, 2002, keynote address paper.

75. *Implementation of CCSDS Lossless Data Compression in HDF*, Pen-Shu Yeh, W. Xia-Serafino, L. Miles, B. Kobler, and D.A. Menascé, NASA's Earth Science Technology Conference 2002, June 11-14 in Pasadena.

76. *Fractal Characterization of Web Workloads*, D.A. Menascé, B. Abrahão, D. Barbará, V.A.F. Almeida, and F. Peligrinelli, WWW 2002 Conference, Web-Engineering Track, Honolulu, Hawaii, May 7-11, 2002, (best paper finalist).

77. *Testing E-commerce Site Scalability with TPC-W*, R. Dodge, D.A. Menascé, and D. Barbará, Proc. 2001 Computer Measurement Group Conference, Anaheim, CA, December 2-7, 2001.

78. *The Performance of Public Key Enabled Kerberos Authentication in Mobile Computing Applications*, A. Harbitter and D.A. Menascé, Eighth ACM Conference on Computer and Communications Security (CCS-8), Philadelphia, PA, November 5-8, 2001.

79. *Preserving QoS of E-commerce Sites Through Self-Tuning: A Performance Model Approach*, D.A. Menascé, D. Barbará, and R. Dodge, 2001 ACM Conference on E-commerce, Tampa, FL, October 14-17, 2001.

80. *Using Performance Models to Dynamically Control E-Business Performance*, D.A. Menascé, Proc. 11^{th} GI/ITG Conf. on Measuring, Modeling and Evaluation of Computer and Communication Systems, Aachen, Germany, September 12-14, 2002.

81. *Analyzing Web Robots and their Impact on Caching*, V.A.F. Almeida, D.A. Menascé, R. Riedi, F. Pelegrinelli, R. Fonseca, and W. Meira, Jr., Proc. Sixth Workshop on Web Caching and Content Distribution, Boston, MA, June 20-22, 2001.

82. *Designing an IT College*, P. J. Denning, R. Athale, N. Dabbagh, D.A. Menascé, M. Pullen, S. Ruth, and R. Sandhu, IFIP World Congress on Computing Education (WCCE), Copenhagen, July 29, 2001.

83. *Characterizing and Modeling Robot Workload on E-Business Sites*, V.A.F. Almeida, D.A. Menascé, R. Riedi, R. Fonseca, W. Meira Jr., and F. Ribeiro, Proc. 2001 ACM Sigmetrics Conference, Cambridge, MA, June 16-20, 2001.

84. *Reusable Component Interconnection Patterns for Distributed Software Architectures*, H. Gomaa, D.A. Menascé, and M.E. Shin, Proc. Symposium on Software Reusability, Toronto, CA, May 18-20, 2001.

85. *Performance of Public Key-Enabled Kerberos Authentication in Large Networks*, A. Harbitter and D.A. Menascé, Proc. 2001 IEEE Symposium on Security and Privacy, Oakland, CA, May 13-16, 2001.

86. *In Search of Invariants for E-Business Workloads*, D.A. Menascé, V.A.F. Almeida, R. Riedi, F. Pelegrinelli, R. Fonseca, and W. Meira Jr., Proc. Second ACM Conference on Electronic Commerce, Minneapolis, MN, October 17-20, 2000.

87. *Challenges in Scaling E-Business Sites*, D.A. Menascé and V.A.F. Almeida, 2000 Computer Measurement Group Conference, Orlando, FL, December 10-15, 2000.

88. *Design and Performance Modeling of Component Interconnection Patterns for Distributed Software Architectures*, H. Gomaa and D.A. Menascé, Proc. Second Workshop on Software and Performance, Ottawa, Canada, Sept. 17-20, 2000.

89. *E-representative: a scalability scheme for e-commerce*, W. Meira Jr., D.A. Menascé, V.A.F. Almeida, and R. Fonseca, Second International Workshop on Advanced Issues of E-commerce and Web-Based Information Systems (WECWIS 2000), Milpitas, CA, June 8-9, 2000, pp. 168-175.

90. *Design and Implementation of a Tool for Measuring the Performance of Complex E-commerce Sites*, G.T. Paixao, W. Meira Jr., V.A.F. Almeida, D.A. Menascé, and A.M. Pereira, Tools 2000, March 27-31, 2000, Chicago, IL. Published in LNCS Vol. 1786/2000, pp. 309–323.

91. *A Methodology for Workload Characterization of E-commerce Sites*, D.A. Menascé, V. Almeida, R. Fonseca, and M.A. Mendes, 1999 ACM Conference on Electronic Commerce, Denver, CO, November, 1999.

92. *Planning the Capacity of a Web Server: An Experience Report*, D.A. Menascé, B. Peraino, N. Dinh, and Q. Dinh, Proc. CMG 99 Conference, Reno, Nevada, December 1999.

93. *An Analytical Model of a Bandwidth Allocation Scheme for a Network Service Provider (NSP)*, D.A. Menascé, A.E. Wang, B. Steele, Greg Weidman, Proc. CMG 99 Conference, Reno, Nevada, December 1999.

94. *Resource Management Policies for E-commerce Servers*, D.A. Menascé, V. Almeida, R. Fonseca, and M.A. Mendes, Second Workshop on Internet Server Performance, in conjunction with ACM SIGMETRICS 99/FCRC, Atlanta, GA, May 1st, 1999. Also published in the ACM Performance Evaluation Review, Vol. 27, Number 4, March 2000.

95. *Performance Issues in Large Distributed System Security*, A. Harbitter and D.A. Menascé, Proc. of the Computer Measurement Group Conference, Anaheim, CA, December 6-11, 1998.

96. *Prefetching Results of Web Searches*, H. Foxwell and D.A. Menascé, Proc. of the Computer Measurement Group Conference, Anaheim, CA, December 6-11, 1998.

97. *Performance and Availability Modeling of Replicated Database Servers*, D. A. Menascé and Nagesh Kakarla, Proc. of the Computer Measurement Group Conference, Anaheim, CA, December 6-11, 1998.

98. *Prefetching Inlines to Improve Web Server Latency*, R. Dodge and D. A. Menascé, Proc. of the Computer Measurement Group Conference, Anaheim, CA, December 6-11, 1998.

99. *Web Performance Modeling Issues*, D. Menascé, First NASA Information Systems Performance Modeling Workshop, NASA Ames Research Center, Moffett Field, CA, Sept. 28-29, 1998.

100. *Educational Challenges and Opportunities in the Web Era*, D. A. Menascé, Proc. of the Workshop on CS Education, 1998 Annual Congress of the Brazilian Computer Society, Belo Horizonte, Brazil, August 2-6, 1998 (invited paper).

101. *Web Performance: the Most Obvious Challenge to Electronic Commerce*, Proc. of the 1998 Annual Congress of the Brazilian Computer Society, Belo Horizonte, Brazil, August 2-6, 1998 (invited paper).

102. *On a Language Based Method for Software Performance Engineering of Client/Server Systems*, D. A. Menascé and H. Gomaa, First International Workshop on Software Performance Engineering, Santa Fe, NM, October 12-16, 1998.

103. *A Scalability Model for ECS's Data Server*, D. A. Menascé and M. Singhal, Joint Sixth NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies and Fifth IEEE Symposium on Mass Storage Systems, College Park, MD, March 23-26, 1998.

104. *A Framework for Software Performance Engineering of Client/Server Systems*, D.A. Menascé, 1997 Computer Measurement Group Conference, Orlando, December 1997.

105. *Performance Modeling of a Unix Communications Server*, D.A. Menascé, K. Nguyen, and D. Nguyen, 1997 Computer Measurement Group Conference, Orlando, December 1997 (**Best Paper Award**).

106. *Hyperlearning Meter: A Web-based tool for Certification and Self-Assessment*, D.A. Menascé, P.J. Denning, and H. Le, Third International Conference on Asynchronous Learning Networks, October 3-5, New York, NY 1997.

107. *Pythia: A Performance Analyzer of Hierarchical Mass Storage Systems*, O.I. Pentakalos, D.A. Menascé, and Y. Yesha, Performance Tools'97, 9^{th} International Conference on Modelling Techniques and Tools for Computer Performance Evaluation, St. Malo, France, June 2-6, 1997.

108. *Performance Prediction of Parallel Applications on Networks of Workstations*, D.A. Menascé and A. Rao, Computer Measurement Group Conference, San Diego, CA, December 1996.

109. *Toward Hyperlearning*, P.J. Denning and D.A. Menascé, Conference on Study of Postsecondary, Pedagogy, Mohonk Mountain House, Nov 20-22, 1996.

110. *Automated Clustering Based Workload Characterization for Mass Storage Systems*, O.I. Pentakalos and D.A. Menascé, Fifth NASA Goddard Space Flight Center Conference on Mass Storage Systems and Technologies, College Park, MD, September 17-19, 1996.

111. *Data and Information Architectures for Large-Scale Distributed Data Intensive Information Systems,* L. Kerschberg, H. Gomaa, D.A. Menascé, and J.P. Yoon, Proc. of the Eighth IEEE International Conference on Scientific and Statistical Database Management, Stockholm, Sweden, June 18-20, 1996.

112. *An Analytic Model of Hierarchical Mass Storage Systems with Network-Attached Storage Devices*, D.A. Menascé, O.I. Pentakalos, and Y. Yesha, Proc. of the 1996 ACM Sigmetrics Conference, Philadelphia, PA, May 1996.

113. *On the Design and Implementation of a Capacity Management Tool for LAN Environments*, D.A. Menascé and D. Dregits, 1995 Conference of the Computer Measurement Group, Nashville, Tennessee, December 1995.

114. *A Federation Oriented Capacity Management Methodology for LAN Environments*, D.A. Menascé, D. Dregits, R. Rossin, and D. Gantz, 1995 Conference of the Computer Measurement Group, Nashville, Tennessee, December 1995.

115. *A Performance-Oriented Design Methodology for Large-Scale Distributed Data Intensive Information Systems*, D.A. Menascé, H. Gomaa, and L. Kerschberg, First IEEE International Conference on Engineering of Complex Computer Systems, Southern Florida, USA, November 6–10, 1995 (**outstanding paper award**).

116. *Capacity Planning in Client/Server Environments*, D.A. Menascé, Proc. of the 8th International Conference on Modelling Techniques and Tools (PERFORMANCE TOOLS '95), Heidelberg, Germany, September 18-21, 1995.

117. *An Approximate Performance Model of a Unitree Mass Storage System*, O. I. Pentakalos, D.A. Menascé, Y. Yesha, and M. Halem, Fourteenth IEEE Symposium on Mass Storage Systems (Second International Symposium), Monterey, CA, September 11-14, 1995.

118. *Educating New Engineers*, P.J. Denning, D.A. Menascé, and J. Gerstner, Proceedings of the Annual Conference of the American Society for Engineering Education (ASEE), Anaheim, CA, June 25-28, 1995.

119. *Two-level Performance Models of Client-Server Systems*, D.A. Menascé and V.A.F. Almeida, 1994 Computer Measurement Group Conference, Orlando, Florida, December 1994.

120. *An Experimental Study of the Effects of Heterogeneity and Data Partitioning on the Performance of Parallel Applications*, A.K. Bangalore and D.A. Menascé, First International Workshop on Parallel Processing, Bangalore, India, December 27-30, 1994.

121. *A Methodology for the Performance Prediction of Massively Parallel Applications*, D.A. Menascé, S.H. Noh, and S.K. Tripathi, 5^{th} IEEE Symposium on Parallel and Distributed Processing, Dallas, TX, December 1-4, 1993.

122. *Processor Assignment in Heterogeneous Message Passing Parallel Architectures*, S.C.S. Porto and D.A. Menascé, HICSS-26 Hawaii International Conference on System Science, Kauai, Hawaii, January 5-8, 1993.

123. *Using Random Task Graphs to Investigate the Potential Benefits of Heterogeneity in Parallel Systems*, V.A.F. Almeida, I.M.M. Vasconcelos, J.N.C. Árabe, and D.A. Menascé, 1992 Supercomputing Conference, MN, November 16-20, 1992.

124. *Task Scheduling on Heterogeneous Multiprocessors*, S.C.S. Porto and D.A. Menascé, XIX SBC Integrated Seminar on Hardware and Software (SEMISH), Rio de Janeiro, Brazil, Sept. 20-Oct.2, 1992.

125. *On the Performance Prediction of Parallel Applications on Massively Parallel Machines*, D.A. Menascé, S.H. Noh, and S.K. Tripathi, Realistically Dependable Parallel Computing Workshop (READPAC), Stony Brook, NY, August 28-29, 1992.

126. *On Task Scheduling in Heterogeneous Multiprocessors*, S.C.S. Porto and D.A. Menascé, 10^{th} National Symposium on Probability and Statistics, August 3-7, 1992, Rio de Janeiro, Brazil, pp.215-220.

127. *On the Performance Evaluation of Reliable Distributed Quorum Attainment Protocols*, D.A. Menascé, Y. Yesha, and K. Kalpakis, 1992 Workshop on Fault-Tolerant Parallel and Distributed Systems, Amherst, Mass, USA, July 6-7, 1992.

128. *Scheduling Issues in Heterogeneous Multiprocessor Systems*, S.K. Tripathi, and D.A. Menascé, Transputer'92, Salines royales d'Arc et Senans, France, May 20-22, 1992.

129. *A Performance Comparison between Paradigms for Distributed Access to Databases in LANs*, C.A. M. Bastos, D.A. Menascé, and L.F.G. Soares, 10^{th} Brazilian Symposium on Computer Networks, Recife, Brazil, April 13-15, 1992.

130. *Processor Assignment in Heterogeneous Parallel Architectures*, D.A. Menascé, S.C. da Silva Porto, and S.K. Tripathi, Sixth International Parallel Processing Symposium (IPPS'92), Beverly Hills, CA, March 23-26, 1992.

131. *Heterogeneity in High Performance Computing*, D.A. Menascé and V.A.F. Almeida, Proceedings of the 2^{nd} Symposium on High Performance Computing, October 7-9, Montpellier, France, 1991, published by Elsevier Science Publishers, eds: M. Durand and F. El Dabaghi.

132. *A Methodology for the Construction of Heuristic Task Scheduling Algorithms in Heterogeneous Architecture Multiprocessors*, S.C. S. Porto and D.A. Menascé, Proceedings of the XVIII SBC Integrated Seminar on Hardware and Software (SEMISH), Santos, SP, Brazil, August 1991.

133. *Performance Modelling of Parallel Applications in the ACP II Multiprocessor*, Gledson Elias and Daniel A. Menascé, Proceedings of the XVIII SBC Integrated Seminar on Hardware and Software (SEMISH), Santos, SP, Brazil, August 1991.

134. *On Scheduling Disciplines for Multiprocessors with Heterogeneous Architectures*, D. A. Menascé, Proceedings of the Joint TIMS XXX - SOBRAPO XXIII International Meeting, Rio de Janeiro, Brazil, July 15-17, 1991.

135. *Scheduling Algorithms for Multiprocessors with Heterogeneous Architectures*, S. C. S. Porto, and D.A. Menascé, Proceedings of the III SBC Symposium on Computer Architectures and Parallel Processing, Rio de Janeiro, RJ, Brazil, November 1990.

136. *Cost-Performance Analysis of Heterogeneity in Supercomputer Architectures*, D.A. Menascé and V.A.F. Almeida, Proceedings of the ACM-IEEE Supercomputing'90 Conference, New York, NY, USA, November 12-16, 1990.

137. *Correctness Proof of Cache Coherence Control Mechanisms in Multiprocessors*, A.C.G. Nunes and D.A. Menascé, Proceedings of the III SBC Symposium on Computer Architectures and Parallel Processing, Rio de Janeiro, RJ, Brazil, November 1990.

138. *An Integrated Coherence and Concurrency Control Mechanism*, A.C.G. Nunes and D.A. Menascé, Proceedings of the III SBC Symposium on Computer Architectures and Parallel Processing, Rio de Janeiro, RJ, Brazil, November 1990.

139. *On the Investigation of Supercomputer Architectures in Multiprogramming Environments*, D.A. Menascé and V.A.F. Almeida, Proceedings of the II SBC Symposium on Computer Architectures and Parallel Processing, Aguas de Lindóia, SP, Brazil, September 1988.

140. *An Object-Oriented Approach to Database Access to Multimedia Databases in CD-ROM*, D.A. Menasce and R. Ierusalimschy, Proceedings of the RIAO'88, Boston, USA, March 1988.

141. *A Virtual Arrays Component*, N.S. Fonseca and D.A. Menascé, Proceedings of the X SBMAC Congress on Computational and Applied Mathematics, Gramado, RS, Brazil, September 1987.

142. *RAINBOW: A New Concept in Database Access*, D.A. Menascé, Proceedings of the XIV SBC Integrated Seminar on Hardware and Software (SEMISH), Salvador, Bahia, Brazil, July 1987.

143. *Efficient Scheduling Policies for Two-Arm Disks*, D.A. Menascé, D.P. Bovet, and J.R.S. Nunes, Proceedings of the XIV SBC Integrated Seminar on Hardware and Software (SEMISH), Salvador, Bahia, Brazil, July 1987.

144. *Process Allocation in a Distributed Architecture: Performance Evaluation*, L.C. Trevelin and D.A. Menascé, Proceedings of the XIV SBC Integrated Seminar on Hardware and Software (SEMISH), Salvador, Bahia, Brazil, July 1987.

145. *Temporal Logic Specification of a Token Bus Based Access Protocol for LANS*, L.F.G. Soares and D.A. Menascé, Proceedings of the IASTED International Conference Telcon 84 Telecommunication and Control, Greece, August 27-30, 1984.

146. *Performance Evaluation of Isolated and Interconnected Token Bus Local Area Networks*, D.A. Menascé and L.L.P. Leite, Proceedings of the 1984 ACM SIGMETRICS Conference on Measurement and Modelling of Computer Systems, Boston, August 22-24, 1984.

147. *Specification of a Network Interconnect Protocol*, D.A. Menascé, Proceedings of the XI SBC Integrated Seminar on Hardware and Software (SEMISH), Viçosa, MG, Brazil, July 1984.

148. *Evaluating the Performance of Interconnected Networks*, L.L.P. Leite and D.A. Menascé, Proceedings of the XI SBC Integrated Seminar on Hardware and Software (SEMISH), Viçosa, MG, Brazil, July 1984.

149. *Alternatives for the Interconnection of Local Networks and X.25 Public Networks*, O.P. Coelho, L.F.G. Soares, C.H.C. Correia, and D.A. Menascé, Proceedings of the XI SBC Integrated Seminar on Hardware and Software (SEMISH), Viçosa, MG, Brazil, July 1984.

150. *Temporal Logic Specification of the Medium Access Protocol of the Local Network REDPUC*, L.F.G. Soares and D.A. Menascé, Proceedings of the XI SBC Integrated Seminar on Hardware and Software (SEMISH), Viçosa, MG, Brazil, July 1984.

151. *Detection and Removal of Deadlocks in Store and Forward Communications Networks*, G. Gambosi, D.P. Bovet and D.A. Menascé, Proceedings of the Second International Symposium on the Performance of Computer Communication Networks, Zurich, Switzerland, March 1984.

152. *Approximate Modeling of Bus-Oriented Local Area Networks*, L.L.P. Leite and D.A. Menascé, Proceedings of the X SBC Integrated Seminar on Hardware and Software (SEMISH), Campinas, SP, Brazil, July 1983.

153. *Approximate Modeling of CPU Preemptive Resume Priority Scheduling Using Operational Analysis*, V.A.F. Almeida and D.A. Menascé, Proceedings of the 10th European Computer Manufactures Association (ECOMA) Conference on Computer Measurement, Munich, Germany, October 12-15, 1982.

154. *On the Evolution of the Architectures of Packet Switches*, D.A. Menascé and L.F.G. Soares, Proceedings of the XV National Congress on Informatics, Rio de Janeiro, Brazil, October 1982.

155. *A Conceptual Schema for Office Automation*, S.L.Santos and D.A. Menascé, Proceedings of the XV National Congress on Informatics, Rio de Janeiro, Brazil, October 1982.

156. *Hardware and Software Description of the Local Network REDPUC*, L.F.G. Soares, D.A. Menascé, C.H.C. Correa, and F.J. Oliveira, Proceedings of the XV National Congress on Informatics, Rio de Janeiro, Brazil, October 1982.

157. *A Transport Protocol for Virtual Circuit Based Public Networks*, D. Schwabe, D.A. Menascé, J.R.B. de Marca, and R. Lobel, Proceedings of the XV National Congress on Informatics, Rio de Janeiro, Brazil, October 1982.

158. *On the Use of Analytic Models in Capacity Planning*, D.A. Menascé and V.A.F. Almeida, Proceedings of the IX Latin American Conference on Informatics, Lima, Peru, August 16-20, 1982.

159. *A Protocol for Broadcast Type LANs*, L.F.G. Soares and D.A. Menascé, Proceedings of the IX Latin American Conference on Informatics, Lima, Peru, August 16-20, 1982.

160. *Synchronization and Crash Recovery Aspects in Distributed Databases*, D.A. Menascé, Proceedings of the Second Latin American Symposium on Computer Networks (SLARC), São Paulo, Brazil, June 14-17, 1982.

161. *The Architecture of PUC-RJ Packet Switch*, R. Roenick and D.A. Menascé, Proceedings of the XIV National Congress on Informatics, São Paulo, Brazil, October 1981.

162. *Specification of a Gateway Between a Local Area Network and a Packet Switched Network*, D.A. Menascé et al, Proceedings of the XIV National Congress on Informatics, São Paulo, Brazil, October 1981.

163. *A Capacity Planning Tool or How to Forecast your Future Hardware Needs*, V.A.F. Almeida and D.A. Menascé, Proceedings of the XIV National Congress on Informatics, São Paulo, Brazil, October 1981.

164. *Performance Evaluation of A Two-Phase Commit Based Protocol for DDBs*, D.A. Menascé and T. Nakanishi, Proceedings of the ACM SIGACT-SIGMOD Symposium on Principles of Database Systems, Los Angeles, USA, March 29-31, 1981.

165. *Computing Performance Measures of Computer Systems with Variable Degree of Multiprogramming*, D.A. Menascé and V.A.F. Almeida, Proceedings of the XII Computer Measurement Group Conference (CMG), New Orleans, USA, December 1-3, 1981.

166. *Dynamic Crash Recovery of Balanced Trees*, D.A. Menascé and O.E. Landes, Proceedings of the IEEE Symposium on Reability in Distributed Software and Database Systems, Pittsburgh, PA, July 21-22, 1981. (invited paper).

167. *On the Design of A Robust Storage Component for Distributed Database Management Systems*, D.A. Menascé and O.E. Landes, Proceedings of the Sixth International Conference on Very Large Databases (VLDB), Montreal, Canada, October 1-3, 1980.

168. *Selective Reloading of Very Large Databases*, D.A. Menascé, Proceedings 1979 IEEE Computer Software and Applications Conferences (COMPSAC), Chicago, Illinois, USA, November 5-8, 1979.

169. *A Formal Model of Crash Recovery in Computer Systems*, D.A. Menascé, R. R. Muntz and G. J. Popek, Twelfth Hawaii International Conference on Systems Science, Vol. I: Selected Papers in Software Engineering and Mini-Micro Systems, Western Periodicals Company, Honolulu, Hawaii, USA, January 1979.

170. *Crash Recovery in Computer Systems*, D.A. Menascé, Proceedings of the Computing School, São Paulo, Brazil, January 1979.

8.5 Technical Reports and Unrefereed Publications

1. *Near-optimal Allocation of VMs from IaaS Providers by SaaS Providers*, A. Aldhalaan and D.A. Menascé, Technical Report GMU-CS-2015-12, Dept. of Computer Science, George Mason University, July 2015.

2. *Temporal Manufacturing Query Language (tMQL) for Domain Specific Composition, What-if Analysis, and Optimization of Manufacturing Processes With Inventories,* M. Krishnamoorthy, A. Brodsky, and D.A. Menascé, Technical Report GMU-CS-2014-03, Dept. of Computer Science, George Mason University, March 2014.

3. *Trace-Driven Analytic Modeling for Evaluating Schedulers for Clusters of Computers*, S. Bardhan and D.A. Menascé, Technical Report GMU-CS-2014-02, Dept. of Computer Science, George Mason University, March 2014.

4. *Epochs: Trace-Driven Analytical Modeling of Job Execution Times*, D.A. Menascé and S. Bardhan, Technical Report GMU-CS-2014-01, Dept. of Computer Science, George Mason University, March 2014.

5. *Bringing Architecture Back to Computing*, interview to Peter Denning, ACM Ubiquity, April 2012.

6. *Guest Editor's Introduction*, D.A. Menascé and J. Kephart, Special Issue on Autonomic Computing, IEEE Internet Computing, January 2007.

7. *Guest Editors' Introduction*, C.M. Woodside and D.A. Menascé, Special Issue on Application-Level QoS, IEEE Internet Computing, Vol. 10, No. 3, pp. 13–15.

8. *Guest Editor's Foreword*, D.A. Menascé, Special Issue on E-commerce, ACM Sigmetrics Performance Evaluation Review, Vol. 32, No. 3, page 2.

9. *Scaling your E-Business*, D. A. Menascé and V. Almeida, Software Magazine, February/March 2001, cover story.

10. *Scaling for E-Business*, Northern Virginia Technology Council Magazine, Vol. 10, No. 10, Nov. 2000, pp. 8–9.

11. *Application Performance Management for E-business*, D. A. Menascé, white paper available at http://www.cptsoftware.com/.

12. *Evaluating Web Server Capacity: The Fundamentals of Web Performance*, D. A. Menascé and V. Almeida, WEBTechniques, Miller Freeman Publications, April, 1999.

13. *Performance and Software*, A. Ferscha, D. Menascé, J. Rolia, B. Sanders, and M Woodside, in System Performance Evaluation: Origins and Directions, eds. M. Reiser, C. Lindemann, and G. Haring, Schloss Dagstuhl Report no. 9738, September 1997.

14. *A Performance-Oriented Design Methodology for Large-Scale Distributed Data Intensive Information Systems*, D. A. Menascé, H. Gomaa, and L. Kerschberg, Dept. of Information Systems and Software Engineering, George Mason University, Technical Report ISSE-TR-95-107, July 1995.

15. *A Software Architecture Design Method of Large-Scale Distributed Data Intensive Information Systems* , H. Gomaa, L. Kerschberg, and D. A. Menascé, Dept. of Information Systems and Software Engineering, George Mason University, Technical Report ISSE-TR-95-108, July 1995.

16. *Data and Information Architectures for Large-Scale Distributed Data Intensive Information Systems,* L. Kerschberg, H. Gomaa, D. A. Menascé, and J. P. Yoon, Dept. of Information Systems and Software Engineering, George Mason University, Technical Report ISSE-TR-95-107, July 1995.

17. *Solutions Manual to the book "Capacity Planning and Performance Modeling: from mainframes to client-server systems"*, D. A. Menascé, V. A. F. Almeida, L. W. Dowdy, and J. Padhye, July 1995.

18. *What will the capacity of your system be?* (in Portuguese), LAN Times Brasil, Vol. 1, No. 1, April 1995.

19. *On a Unified Framework for the Evaluation of Distributed Quorum Attainment Protocols*, D. A. Menascé, Y. Yesha, and K. Kalpakis, University of Maryland at College Park, Technical Report UMIACS TR-92-13 and CS TR 2833, February 1992.

20. *HFT: A Highly Fault Tolerant and Efficient Protocol for Distributed Mutual Exclusion*, D. A. Menascé, Y. Yesha, and K. Kalpakis, University of Maryland at College Park, Technical Report UMIACS TR-92-12 and CS TR 2832, January 1992.

21. *Static and Dynamic Processor Scheduling Disciplines in Heterogeneous Parallel Architectures*, D. A. Menascé, D. Saha, S. C. da Silva Porto, V. A. F. Almeida, and S. K. Tripathi, University of Maryland at College Park, Technical Report UMIACS TR-91-162 and CS TR 2807, December 1991.

22. *A Performance Model of the Architecture of the ACP II*, G. E. Silveira, and D. A. Menascé, Technical Report No. CCR-138, Scientific Center Rio, IBM Brasil, December 1991.

23. *Static Heuristic Processor Assignment in Heterogeneous Parallel Multiprocessor Systems*, D. A. Menascé, S. C. da Silva Porto and S. K. Tripathi, Technical Report UMIACS TR 91-131 and CS TR 2765, University of Maryland at College Park, September 1991.

24. *Heterogeneous Parallel Architectures: Cost-Performance Tradeoff Analysis*, D.A. Menascé and V.A.F. Almeida, Technical Report No. CCR-102, Scientific Center Rio, IBM Brasil, January 1990.

25. *Performance of Processor-Memory Interconnection Networks Under Unbalanced Traffic*, L.A. Barroso and D.A. Menascé, Technical Report No. CCR-085, Scientific Center Rio, IBM Brasil, August 1989.

26. *A Methodology for Performance Evaluation of Parallel Applications on Multiprocessors*, D.A. Menascé and L.A. Barroso, Technical Report No. CCR-084, Scientific Center Rio, IBM Brasil, August 1989.

27. *Stochastic Petri Nets*, D.A. Menascé and N.L.S. Fonseca, Technical Report No. CCR-082, Scientific Center Rio, IBM Brasil, August 1989.

28. *CD-ROM: Easier Than it Seems*, Daniel A. Menascé, Revista Dados e Ideias, January 1988.

29. *Distributed Systems: A Current Trend in Computing*, D.A. Menascé, Revista Dados e Ideias, No.2, May 1980.

30. *On the Problem of Optimal Capacity Assignment in Computer Networks which Carry Messages with Different Priorities*, D.A. Menascé and M.A. Monteiro, Monograph Series in Computer Science, Departamento de Informática, PUC-RIO, May 1979.

31. *A Proposed Architecture for the UCLA Distributed Secure System Base*, D.A. Menascé, G.J. Rudisin, G.J. Popek, and C.S. Kline, Technical Report (79-10) (UCLA-ENG-7957), Computer Science Department, UCLA, November 1978.

32. *A Locking Protocol for Resource Coordination in Distributed Systems*, D.A. Menascé, G.J. Popek, and R. R. Muntz, Technical Report UCLA-ENG-7808, SDPS-77-001 (DSS MDA 903-C-0211), Computer Science Department, UCLA, October 1977.

33. *G/PL/I: Extending PL/I for Graph Processing*, A.O.F. da Silva, D.A. Menascé, and P. Blanco, Departamento de Informática, PUC-Rio, 1973.

34. *An Introduction to Simulation*, D.A. Menascé and O.A.F. Tourinho, Departamento de Informática, PUC-Rio, 1972.

8.6 Theses

1. *Coordination in Distributed Systems: Concurrency, Crash Recovery and Database Synchronization*, Ph.D. Thesis, Computer Science Department, University of California at Los Angeles (UCLA), October 1978. Also published as Technical Report 78-12 (UCLA-ENG-7955), Computer Science Department, UCLA, December 1978.

2. *An Efficient Backtracking Algorithm for Obtaining Minimum Spanning Trees*, MS Dissertation, Departamento de Informática, Pontificia Universidade Católica do Rio de Janeiro (PUC-Rio), April 1975.

9 Scholarships Obtained

- Doctoral Scholarship, Brazilian Research Council (CNPq), September 1975 to November 1978.
- MS Scholarship, Brazilian Research Council (CNPq), March 1974 to March 1975.
- Scientific Initiation Scholarship, Brazilian Research Council (CNPq), April 1972 to February 1974.

10 Teaching Experience

10.1 Graduate Courses

- *Autonomic Computing (CS 895/CS 788)*, George Mason University, Fall 2006, Fall 2010, Spring 2013, and Fall 2015.
- *Computer System Performance Evaluation (CS 672)*, George Mason University, Spring 1993, Spring 1994, Fall 1995, Spring 1997, Spring 1999, Spring 2000, Fall 2001, Spring 2003, Spring 2004, Spring 2005, Spring 2006, Spring 2007, Spring 2008, Fall 2009, Fall 2012, and Fall 2014.
- *E-commerce Software Services (EC 512)*, George Mason University, Spring 2003, Spring 2004, Spring 2005.
- *Basic Infrastructure for E-commerce (EC 511)*, George Mason University, Fall 2002, Fall 2003.
- *Quantitative Methods and Experimental Design in Computer Science (CS 700)*, George Mason University, Spring 2001, Fall 2002, Fall 2004, Fall 2007, Spring 2009, Fall 2011, Fall 2012, and Spring 2015.

- *Scaling Technologies for E-Business (INFT 818/IT 809)*, George Mason University, Fall 2000, Fall 2001, Fall 2002, and Fall 2003.
- *Secure Electronic Commerce (ISA 767)*, Spring 2005.
- *Topics in High Performance Computer Systems (INFT 818)*, George Mason University, Spring 1998.
- *Advanced Topics in Parallel Computation (INFT 915)*, George Mason University, Spring 1995.
- *Operating Systems Theory and Practice (INFS 601)*, George Mason University, Spring 96.
- *Operating Systems (CS 571)*, George Mason University, Fall 1994, Spring 96, Fall 1996, Spring 1997, Fall 1997, Fall 1998, Spring 2000, Fall 2000, Spring 2001.
- *Parallel Computing (INFT 815)*, George Mason University, Fall 1993.
- *Dissertation Topic Presentation (IT 990)*, George Mason University, Spring 1993 and Fall 2005, Fall 2006, Fall 2007, Fall 2008, Fall 2009, Fall 2011.
- *Performance Evaluation of Computer Systems I*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1979, 1980, 1981, 1982, 1984, 1985, 1986, 1988, 1989, 1990, and 1991.
- *Performance Evaluation of Computer Systems II*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1989 and 1990.
- *Logical Database Design*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1987, and 1980.
- *Design of Local Networks*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1984.
- *Local Networks*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1982, and 1983.
- *Computer Networks*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1982, 1981, and 1980.
- *Distributed System Architecture*, University of São Paulo, Brazil, 1979.
- *Distributed Systems*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1979.
- *Discrete Event Simulation*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1973.

10.2 Undergraduate Courses

- *Computer Systems Architecture (CS 465)*, George Mason University, Spring 2013, Fall 2014, and Fall 2015.
- *Operating Systems (CS 471)*, George Mason University, Fall 1992, Spring 1993, Fall 1993, Spring 1995, Fall 1995, Spring 98, Fall 1998, Spring 1999, and Fall 2004.
- *Design Exhibition (CS 490)*, George Mason University, Spring 1994 and Fall 1994.
- *Engineering Fundamentals (ENG 107)*, George Mason University, Spring 1994 (team taught).
- *Introduction to Software Engineering (CS 421)*, George Mason University, Fall 1992.
- *Computer Science I (CS 112)*, George Mason University, Summer 1993.
- *Communication Networks (ENEE 426)*, Electrical Engineering Department, University of Maryland at College Park, Spring 1992.
- *Teleprocessing*, Pontifical Catholic University, Rio de Janeiro, Brazil, 1981.

11 Theses Supervised and Under Supervision

11.1 Ph.D. Students Supervised

1. *Autonomic Performance Optimization with Application to Self-Architecting Software Systems*, John Ewing, George Mason University, April 2015.
2. *Autonomic, Optimal, and Near-Optimal Resource Allocation in Cloud Computing,* Arwa Aldhalaan, George Mason University, April 2015.
3. *Design and Modeling of Schedulers for Multi-Task Jobs on Computer Clusters,* Shouvik Bardhan, George Mason University, April 2015.
4. *An Autonomic Framework for Integrating Security and Quality of Service Support in Databases*, Firas Alomari, George Mason University, May 2013.
5. *A Self-managed Healthcare Emergency Department System*, Serene Al-Momen, George Mason University, April 2012.
6. *Defeating Insider Attacks via Autonomic Self-Protective Networks*, Faisal Sibai, George Mason University, April 2012.

7. *Quality of Service Management of Business Processes in Service Oriented Architecture*, Vinod Dubey, George Mason University, April 2010.

8. *The Insider Threat Security Architecture: An Integrated, Inseparable, and Uninterrupted Self-Protection Autonomic Framework*, Ghassan "Gus" Jabbour, George Mason University, April 2010.

9. *Group-Centric Secure Information Sharing Models*, Ram Narayan Krishnan, Chair and co-director, co-director: Ravi Sandhu, George Mason University, November 2009.

10. *Fully Countering Trusting Trust through Diverse Double-Compiling*, David A. Wheeler, Chair and co-director, co-director: Ravi Sandhu, George Mason University, November 2009.

11. *Architecture and Models for Administration of User-Role Assignment in Role Based Access Control*, Venkata Bhamidipati, Chair and co-director, co-director: Ravi Sandhu, George Mason University, November 2008.

12. *Scalable Role and Organizational Based Access Control and its Administration*, Zhixiong "Jim" Zhang, George Mason University, Chair and co-director, co-director: Ravi Sandhu, April 2008.

13. *Workload Characterization and Business-oriented Performance Improvement Techniques for Online Auction Sites*, Vasudeva Akula, George Mason University, February 2007.

14. *Autonomic Computing Through Analytic Performance Models*, Mohamed Bennani, George Mason University, May 2006.

15. *Experimental Study of Performance Sensitivity of Configurable Parameters of Web-based Systems*, Monchai Sopitkamol, George Mason University, November 2004.

16. *Meta-Protocol: On-the-Fly Protocol Agreement and Code Generation*, Ibrahim S. Abdullah, George Mason University, March 2004.

17. *A Web-based System for Representing, Retrieving, and Visualizing Analogies*, Harry J. Foxwell, George Mason University, April 2003.

18. *A Methodology for Analyzing the Performance of Authentication Protocols*, PhD in Computer Science, Alan Harbitter, George Mason University, November 2002.

19. *Improving E-commerce System Performance with Dynamic System Tuning*, Ronald Dodge, George Mason University, April 2001.(co-director: D. Barbará)

20. *Hyperlearning: a non-linear, variable path, variable time, fixed outcome model for on-line self-assessment and certification*, Hai L. Le, George Mason University, December 1997.

21. *Performance Evaluation of Hierarchical Mass Storage Systems*, Odysseas I. Pentakalos, University of Maryland Baltimore County, April 1996.

22. *A Multicache Architecture with Integrated Coherence and Concurrency Control*, Adélia Cecília Gonçalves Nunes, PUC-RIO, May 1991.

23. *A Methodology for the Performance Evaluation of Real Time Applications in Parallel Architectures*, Luis Carlos Trevelin, PUC-RIO, April 1991.

24. *Design and Performance Evaluation of Isolated and Interconnected Local Networks*, Leonardo Lellis Pereira Leite, PUC-Rio, October 1985.

25. *Design and Development of Protocols for Local Computer Networks*, Luiz Fernando Gomes Soares, PUC-Rio, December 1983.

26. *Performance Evaluation of Concurrency Control Mechanisms for Database Systems*, Tatuo Nakanishi, PUC-Rio, September 1981.

11.2 MS Students Supervised

1. *Multi-threading Strategies for Web Servers*, Jason Petrone, George Mason University, May 2004.

2. *Adding Admission Control and Flow Control to Web Servers*, Jiehuan Li, George Mason University, June 1998.

3. *A Methodology and Technique for Practical Network Performance Analysis*, Khue Nguyen, George Mason University, December 1997.

4. *A Workflow System for Academic Business Process Reengineering*, Edgar Quisumbing, George Mason University, March, 1996.

5. *Performance Analysis and Prediction of Parallel Applications on Networks of Workstations*, Amar Rao, George Mason University, March, 1996.

6. *A Tool for the Performance Evaluation of Parallel Applications*, Paulo H. Schindler, PUC-RIO, August 1992.

7. *Heuristic Scheduling Algorithms for Task Scheduling in Heterogeneous Multiprocessor Architectures*, Stella Cavalcanti da Silva Porto, PUC-RIO, July 1991. **Winner of the V Brazilian Computing Society (SBC) Contest of Thesis and Dissertations, 1992.**

8. *Performance Evaluation of High Energy Physics Applications on the ACP II*, Gledson Elias, PUC-RIO, July 1991.

9. *A Tool for Software Performance Prediction*, Regina Elvira Machado Maia, PUC-RIO, July 1991.

10. *Database Servers for LANs: Specification and Prototype Implementation*, Mauro José Fridman Ferreira Pinto, PUC-RIO, June 1991.

11. *On the Automatic Integration of Generalized Stochastic Petri Nets and Queueing Networks*, Eduardo Castello Branco Bion, PUC-RIO, June 1991.

12. *Data Collection for Capacity Planning Models in VAX/VMS Environments*, Luiz Fernando de Barros Falcão Vergara, PUC-RIO, April 1991.

13. *MVS/CC: A Tool for Workload Characterization in MVS/XA Environments*, Eduardo Robson Tardin Costa, PUC-Rio, March 1990.

14. *Multimedia Databases in CD-ROM: A Text Management Environment*, Floriano Saad Mazini, PUC-Rio, February 1990.

15. *Multimedia Databases in CD-ROM: Access Primitives to Multimedia Objects*, Andréa Reis Ribeiro, PUC-Rio, February 1990.

16. *A Methodology for the Performance Evaluation of Scientific Applications in Multiprocessors*, Luiz André Barroso, PUC-Rio, August 1989.

17. *Authoring Systems for CD-ROM: An Architecture Proposal*, Elvira J. Vervloet, PUC-Rio, April 1989.

18. *A Portable Tool for Capacity Planning in Microcomputers*, N.L.S. Fonseca, PUC-Rio, February 1988. **Winner of the I Brazilian Computing Society (SBC) Contest of Thesis and Dissertations, 1988.**

19. *Satellite Data Networks*, Walter Araújo Diniz, PUC-Rio, April 1987.

20. *VM/CC: A Tool for Workload Characterization in VM Environments*, Ricardo Viegas, PUC-Rio, March 1987.

21. *An Experience in Evaluation of Protocols Applied to Local Networks*, Solon Benayon da Silva, PUC-Rio, April 1986.

22. *Simulation of A Deadlock Detection and Removal Algorithm for Store and Forward Computer Networks*, Mariza Carpenter Fraga Lourenço, PUC-Rio, March 1986.

23. *Implementation of a Class A Transport Protocol for the COBRA 540 Computer*, Mário César M. Rodrigues, PUC-Rio, August 1985.

24. *Specification and Implementation of a Gateway Between the Local Network REDPUC and the Brazilian Public Network RENPAC*, Otávio Pecego Coelho, PUC-Rio, May 1985.

25. *An Implementation of a Multilink Protocol for the COBRA 540 Computer*, Mário Roberto F. Benevides, PUC-Rio, March 1985.

26. *Specification and Implementation of a Name Server for a Local Computer Network*, José Laédio Medeiros, PUC-Rio, February 1985.

27. *A Telex Server for Local Computer Networks*, Luiz Geraldo Rocha Carvalho, PUC-Rio, February 1985.

28. *Modeling of Complex I/O Systems*, Adélia Cecília Gonçalves Nunes, PUC-Rio, February 1985.

29. *Specification of a Layered Architecture of a DBMS and Implementation of the Concurrency Control Mechanism*, José Danilo Silvestre Fernandes, PUC-Rio, December de 1984.

30. *Specification of a Layered Architecture of a DBMS and Implementation of the Crash Recovery Mechanism*, Marisol Meneses Rojas, PUC-Rio, July 1984.

31. *Specification and Implementation of an Electronic Folder System*, Mario André Guimarães, PUC-Rio, May 1984.

32. *A Distributed Agenda System for a Local Network*, Ricardo Soares Bigio, PUC-Rio, October 1983.

33. *Communication Processor for a Broadcast Type Local Network*, Carlos Henrique Cavalcanti Correa, PUC-Rio, April 1983.

34. *Development of a General Purpose Workstation for a Local Network*, Fernando Jefferson de Oliveira, PUC-Rio, April 1983.

35. *A Heuristic Algorithm for Closed Queueing Networks*, Maria Elisa Barroso M. Costa, PUC-Rio, October 1982.

36. *Considerations on the Specification and Implementation of a Transport Protocol*, Waldeck P. Araujo Jr., PUC-Rio, October 1982.

37. *Performance Evaluation of an Electronic Funds Transfer System*, Fernando Ribeiro Mendes, PUC-Rio, October 1982.

38. *Specification of an Electronic Funds Transfer System*, Roberto Luis M. Pereira de Castro, PUC-Rio, October 1982.

39. *Analytic Models of Link Control Protocols*, José Quenji Shitara, PUC-Rio, June 1982.

40. *An Analytic Model of a Transport Level Protocol for Computer Networks*, Eduardo Esteban Mendez Ortiz, PUC-Rio, May 1982.

41. *Specification and Implementation of an Office Automation Subsystem*, Susana Lent Santos, PUC-Rio, May 1982.

42. *A Communications Processor for the Connection of a Computer to a Packet Switched Network*, Selda Tereza Tribuzi Lula, PUC-Rio, October 1981.

43. *Specification of Intranetwork Protocols*, Izidério de Almeida Mendes, PUC-Rio, July 1981.

44. *Specification of a Banking Automation System*, Rodney Ferreira de Carvalho, PUC-Rio, April 1981.

45. *Distributed Architecture Packet Switch: Considerations and Preliminary Specification*, Ronaldo D'Avila Roenick, PUC-Rio, March 1981.

46. *Specification of the Implementation of the X.25 CCITT Recommendation in an IBM/370 System*, Yussef Farran Leiva, PUC-Rio, February 1981.

47. *Specification of High Level Protocols for Computer Networks*, Antonio Louro, Instituto Militar de Engenharia (IME), February 1980.

48. *Analytic Models and Simulation of Queues With and Without Priorities*, José Alberto Chahon, PUC-Rio, October 1980.

49. *A Model for Capacity Planning of Computer Systems*, Virgilio Augusto Fernandes de Almeida, PUC-Rio, September 1980.

50. *Study and Application of Analytic Techniques for the Performance Evaluation of Computer Systems*, Alcio José da Cintra Lapa, PUC-Rio, August 1980.

51. *Crash Recovery in Distributed Database Systems: Specification of a Reliable Storage Component*, Oscar Ernesto Landes, PUC-Rio, January 1980.

52. *Automatic Message Transmission in Computer Networks*, Mario Antonio Monteiro, PUC-Rio, July 1979.

11.3 Master Projects Supervised

1. *Prefetching Inlines to Increase Web Performance*, Ronald Dodge, GMU, August 1998.
2. *QNGed - A Graphical Editor for Specifying Queuing Networks*, Mahua Sinha, GMU, March 1994.
3. *A Client-Server System for a Distributed Student Record Database based on Lotus Notes*, Jeffrey Winkler, GMU, December 1993.

11.4 Current Ph.D. Students

1. Yong Xue, George Mason University, advanced to candidacy, (Dissertation title: An Architecture for End-to-end and Autonomic Control and Management of Virtual Service Networks in Multi-AS Environments)
2. Mahmoud Awad, George Mason University, advanced to candidacy, (Dissertation title: Dynamic Derivation of Analytical Performance Models in Autonomic Systems)
3. Mohan Krishnamoorthy, George Mason University, with Prof. Alex Brodsky, (Research Area: Modeling of Smart Manufacturing Processes)

12 Presentations in Colloquia and Seminars

1. *Autonomic Computing: a new design principle for complex systems*, 2-day advanced doctoral course, Gran Sasso Science Institute, L'Aquila, Italy, May 26-27, 2015.
2. *Academic Publishing Workshop*, panel member, George Mason University Libraries and Elsevier, March 3, 2015.
3. *Autonomic Computing and its Applications,* 3-hour tutorial presented at the USENIX/ICAC 2014 Conference, Philadelphia, PA, June 17, 2014.
4. *Autonomic Computing and Performance Modeling: Applications to Smart Manufacturing*, NIST Symposium on Smart Manufacturing, Gaithersburg, MD, March 19, 2014.
5. *Performance Body of Knowledge*, panel member, 2013 Computer Measurement Group Conference, La Jolla, CA, Nov. 5-8, 2013.
6. *Cloud and Autonomic Computing Hot Topics: Challenges, Research Opportunities, and Future Directions*, panel member, ACM 2013 Cloud and Autonomic Computing Conference, Miami, FL, August 6, 2013.

7. *On the Use of Performance Models in Autonomic Computing*, keynote address at WPerformance, Congress of the Brazilian Computer Society, Curitiba, Brazil, July 18, 2012.

8. *Cloud: magic and myth*, panel member, Open Source Industry Day, sponsored by the Open Source Software Institute and by the NSA, May 30, 2012, Applied Physics Lab, Johns Hopkins University.

9. *Ramping Up the Data Center*, panel member, Virtualization, Cloud Computing, and Green IT Summit, 1105 Government Information Group, Washington, DC, October 7, 2009.

10. *Virtualization at the Data Center: consolidation and self-management*, Symantec Government Symposium, July 16, 2009.

11. *Virtualization and the On-Demand Data Center*, Green Computing Summit, Washington, DC, December 3, 2008.

12. *Panel: Virtualization*, panel member, Symantec Government Symposium, Washington, DC, July 31, 2008.

13. *Web Workloads: Characterization, Modeling, and Application*, 3-hour tutorial with V.A.F. Almeida, 2007 World Wide Web Conference, Banff, Canada, May 8, 2007.

14. *Panel: The Future of Performance*, panel member, 2005 Computer Measurement Group Conference, Orlando, FL, Dec. 8, 2005.

15. *Achieving QoS in Complex Distributed Systems through Autonomic Computing*, invited talk, Alcatel Technical Academy, Antwerp, Belgium, October 3, 2005.

16. *Quality of Service Challenges for Web Based Systems and E-commerce*, keynote address, kick-off meeting of the E-Quality Research Center, University of Twente, Enschede, The Netherlands, September 30, 2005.

17. *Performance Engineering Principles*, bn.com, New York, NY, February 11, 2005.

18. *Quality of Service Aspects and Metrics in Grid Computing*, invited talk, NY Computer Measurement Group meeting, New York, NY, February 11, 2005.

19. *Fundamental Concepts in Performance Engineering*, 3-hour tutorial presented at the 2004 CMG Conference, Las Vegas, December 5, 2004.

20. *Performance by Design*, course taught at Mitre Corporation, 27-30, 2004.

21. *On the Use of Online Performance Models in Autonomic Computing*, invited talk at IBM Wastson Research Center, Hawthorne, July 15, 2004.

22. *QoS Challenges and Directions for Large Distributed Systems,*, **keynote** address, Workshop on Quality of Service for Geographically Distributed Systems, Rome, Italy, June 9, 2004.

23. *Performance by Design*, 3-day course offered through Demand Technology, Chicago, Illinois, May 10–12, 2004.

24. *Fundamental Concepts in Peformance and Capacity Planning*, 3-hour tutorial at the 2003 Computer Measurement Group Conference, Dallas, TX, Dec. 7, 2003.

25. *Self-Managing E-commerce Sites*, **keynote** address at the WWW/Internet 2003 IADIS International Conference, Algarve, Portugal, Nov. 6, 2003.

26. *Fundamental Concepts in Peformance and Capacity Planning*, 3-hour tutorial at the 2002 Computer Measurement Group Conference, Reno, Nevada, Dec. 8-13, 2002.

27. *Using Performance Models to Dynamically Control E-commerce Sites*, **keynote** address at the First Seminar on Advanced Research in Electronic Business to be held in Rio de Janeiro, Brazil, November 7-8, 2002.

28. *Software, Performance, or Engineering?*, **keynote** address at the Third International Workshop on Software and Performance (WOSP 2002), Rome, Italy, July 24-26, 2002.

29. *QoS Measurement and Control of Web and E-commerce Services*, North Carolina State University, March 20, 2002.

30. *QoS Measurement and Control of Web and E-commerce Services*, NY Computer Measurement Group Winter Meeting, February 8, 2002.

31. *Theoretical Aspects of Web Site Scalability*, panel presentation at the 2001 Computer Measurement Group Conference, Anaheim, CA, December 5, 2001.

32. *Technologies for Scalable Web Services*, Sunday workshop (3 hrs), 2001 Computer Measurement Group Conference, Anaheim, CA, December 2, 2001.

33. *How to Deal with Scalability in E-business*, Distinguished Lecturer Series, Computer Science and Engineering Division, University of Michigan, October 25, 2001.

34. *Using Performance Models to Dynamically Control E-Commerce Performance*, **keynote** speaker at the 2001 Aachen International Multiconference on Measurement, Modelling, and Evaluation of Computer-Communication Systems, Aachen, Germany, September 12, 2001.

35. *Why Should I Model My Servers?*, invited speaker at the panel "Server Performance Modeling and Tuning" at OPNETWork 2001, August 29, 2001, Washington, D.C.

36. *Understanding Workloads in E-Business*, Microsoft Research, Seattle, WA, May 1, 2001.

37. *Tools for Measuring E-commerce Performance*, Panel Chair, 2000 CMG Conference, Orlando, FL, December 10, 2000.

38. *Scaling for E-business*, Sunday Workshop, 2000 CMG Conference, Orlando, FL, December 10, 2000.

39. *Scaling for E-Business*, Northeast CMG, New York, NY, September 15, 2000.

40. *Scaling for E-Business*, invited speaker, 8th International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2000), September 1, 2000.

41. *Scaling for E-Business*, **keynote** address at the OPENTEWork 2000 Conference, Washington, DC, August 29, 2000.

42. *Capacity Planning for Web and E-commerce Applications*, CMG Italy, Milano, one-day course, March 15, 2000.

43. *Capacity Planning for Web and E-commerce Sites*, Sunday Workshop, 1999 CMG Conference, Reno, NV, Dec. 5, 1999 (morning and afternoon).

44. *Challenges in Capacity Planning for E-commerce*, Sun Microsystems, Richmond, December 17, 1997.

45. *Capacity Planning for Web and E-Commerce*, course organized by Demand Technologies, Reston, VA, November 1-2, 1999.

46. *Workload Characterization for E-commerce Servers*, Carleton University, Ottawa, Canada, October 14, 1999.

47. *Challenges for Capacity Planning in E-commerce*, National Capital Area Computer Measurement Group Fall Meeting, September 15, 1999 (invited speaker).

48. *Web-based Education: opportunities and limitations*, Annual Congress of the Brazilian Computer Society, July 22, 1999 (**keynote** speaker, video conference).

49. *Challenges for Capacity Planning in E-commerce*, 1999 International System Architecture Symposium, Oracle Corporation, **keynote** speaker, Las Vegas, Nevada, July 13, 1999.

50. *Capacity Planning for Web Performance*, Internet Fall'98, New York, NY, October 7, 1998.

51. *Web Performance Modeling Issues*, First Nasa Workshop on Performance-Engineered Information Systems, NASA Ames Research Center, Moffett Field, CA, September 29, 1998, invited speaker.

52. *Capacity Planning for Web Performance*, ISPCon, San Jose, CA, September 29, 1998.

53. *Capacity Planning for Web Performance*, one day seminar presented through GMU's Office of Continuing and Professional Education, August 26, 1998.

54. *Web and Intranet Performance: A Quantitative Analysis*, tutorial presented at the USENIX Association Conference, New Orleans, June 16, 1998.

55. *Software Performance Engineering of Client/Server Systems*, talk presented at the Workshop "System Performance Evaluation: Origins and Directions", Schloss Dagstuhl, September 18, 1997.

56. *Software Performance Engineering of Client/Server Systems*, tutorial presented at the 1997 ACM Sigmetrics Conference, Seattle, WA, June 1997.

57. *Performance Modeling of Next Generation Mass Storage Systems*, presentation to the Science Council Meeting of the Center for Excellence in Space Data Information Sciences (CESDIS), NASA Goddard Space Flight Center, Maryland, September 9, 1996.

58. *Distributed Internet-based Tools for Hyperlearning*, presentation to members of the Office of Science and Technology Policy, George Mason University, July 12, 1996.

59. *Member of a panel on Distance Learning/Internet Technologies*, CAETI Training Partnership Conference, Alexandria, VA, February 15-16, 1996.

60. *Capacity Planning in Client Server Environments*, tutorial presented at the 1995 Computer Measurement Group Conference, Nashville, TN, December 1995.

61. *Performance Modeling of Mass Storage Systems*, presentation to the Science Council Meeting of the Center for Excellence in Space Data Information Sciences (CESDIS), NASA Goddard Space Flight Center, Maryland, September 29, 1995.

62. *Capacity Planning in Client Server Environments*, tutorial presented at the 8th International Conference on Modelling Techniques and Tools (PERFORMANCE TOOLS '95), Heidelberg, Germany, September 19, 1995.

63. *Capacity Planning in Client Server Environments*, tutorial presented at the 1995 ACM Sigmetrics and PERFORMANCE '95 Conference, Ottawa, Canada, May 15, 1995.

64. *Capacity Planning for Client/Server Systems*, one day seminar presented at the DB Forum'95, São Paulo, Brazil, April 26, 1995.

65. *The Evolution of the Client/Server Model*, member of a panel at the DB Forum'95, São Paulo, Brazil, April 27, 1995.

66. *Hierarchical Mass Storage Systems*, DB Forum'95, São Paulo, Brazil, April 27, 1995.

67. *Capacity Planning and Performance Modeling: a key aspect of system design*, Center of Excellence in Space Data and Information Sciences (CESDIS), NASA Goddard Space Flight Center, Maryland, February 1, 1995.

68. *Capacity Planning and Performance Modeling: from mainframes to client/server systems*, Computer Literacy Bookshop, Vienna, VA, January 10, 1995.

69. *Ongoing Computer Science Efforts in the USA, Some Focused Research Results, and Collaboration Possibilities between Brazil and the US*, XIV Conference of the Brazilian Computer Society, Caxambu, MG, Brazil, August 2, 1994 (invited speaker).

70. *Capacity Planning: from Mainframes to Client-Server Systems*, ACM Washington DC Chapter Professional Development Seminars, April 1, 1994.

71. *Research in Distributed and Parallel Computing*, Bellcore, NJ, USA, May 7, 1992.

72. *Heterogeneity: the Key to High Performance Computing in the 90's*, Computer Science and Information and Software Systems Engineering Departments, George Mason University, USA, April 29, 1992.

73. *Research Issues in Heterogeneous Parallel Computing*, Computer Science Department, Purdue University, USA, April 24, 1992.

74. *Capacity Planning Techniques*, ACM Washington DC Chapter Professional Development Seminars, April 2, 1992.

75. *A Highly Efficient and Fault-Tolerant Protocol for Distributed Mutual Exclusion* Institute for Advanced Computer Studies, University of Maryland at College Park, USA, March 19, 1992.

76. *Research Issues in Heterogeneous Parallel Computing*, Space Telescope Science Institute, Baltimore, Maryland, USA, January 13, 1992.

77. *Processor Scheduling in Heterogeneous Parallel Machines*, University of Maryland at College Park, Maryland, USA, November 4, 1991.

78. *Research Issues in Heterogeneous Parallel Processing*, University of Maryland Baltimore County, Baltimore, Maryland, October 16, 1991.

79. *Research Issues in Heterogeneous Supercomputing*, National Science Foundation, Washington, D.C., USA, September 5, 1991.

80. *Performance Evaluation in Parallel and Distributed Computing*, Department of Computer Science, Georgetown University, Washington, USA, February 13, 1991.

81. *Cost-Performance Analysis of Heterogeneity in Supercomputer Architectures*, Computer Science Department, University of Maryland, College Park, USA, November 19, 1990.

82. *Capacity Planning in VAX/VMS Environments*, 8th Digital Equipment Users Association (DECUS) Seminar, Rio de Janeiro, Brazil, October 15, 1990.

83. *Multimedia Databases in Optical Disks*, DATAPREV Seminars, Rio de Janeiro, Brazil, March 30, 1989.

84. *How to Control Office Automation*, First International Informatics Congress da Sucesu, Rio de Janeiro, Brazil, August 1988.

85. *Optimistic Versus Pessimistic Concurrency Control Mechanisms in Database Management Systems*, Universita di Pisa, Italy, January 1983.

86. *Optimistic Versus Pessimistic Concurrency Control Mechanisms in Database Management Systems*, IASI, Rome, Italy, January 1983.

87. *An Application of the Local Network REDPUC: Packet Switching*, IBM Brasil Scientific Center, September 1982.

88. *Distributed Databases*, XII National Congress on Data Processing (SUCESU), Rio de Janeiro, Brazil, October 22, 1981.

89. *Optimistic Versus Pessimistic Concurrency Control Mechanisms in Database Management Systems*, Computer Science Department, University of California at Los Angeles (UCLA), Los Angeles, USA, October 9, 1980.

90. *Optimistic Versus Pessimistic Concurrency Control Mechanisms in Database Management Systems*, Computer Science Department, University of Southern California (USC), Los Angeles, USA, October 8, 1980.

91. *Short Course on Distributed Databases*, INTELCOM 1980 International Conference, Rio de Janeiro, Brazil, May 22-23, 1980.

92. *A Formal Model of Crash Recovery in Computer Systems*, Department of Computer Science, University of Pittsburgh, Pennsylvania, November 1979.

93. *Optimal Channel Capacity Allocation in Computer Networks with Priority Messages*, Sixth Integrated Seminar on Software and Hardware, São Paulo, SP, Brazil, July 24, 1979.

94. *Main Aspects in Distributed Database Systems*, Universidad Tecnica Federico Santa Maria, Valparaiso, Chile, July 1979.

95. *Main Aspects in Distributed Database Systems*, Empresa Nacional de Computación e Informatica (ECOM), Santiago, Chile, July 1979.

96. *Crash Recovery in Computer Systems*, Computing School, São Paulo, SP, Brazil, January 15, 1979.

97. *Member of a Panel on Concurrency Control in Distributed Database Systems*, Third Berkeley Workshop on Distributed Data Management and Computer Networks, San Francisco, California, USA, August 31, 1978.

98. *Locking and Deadlock Detection in Distributed Databases*, 12th Colloquium in Computer Science, Lawrence Berkeley Laboratory, Berkeley, CA, July 12, 1978.

99. *A Locking Protocol for Resource Coordination in Distributed Systems*, UCLA Computer Science Department Seminar, Los Angeles, USA, October 22, 1977.

13 Participation in Committees

- Program Committee Member,Seventh ACM/SPEC International Conf. Performance Engineering (ICPE 2016), Delft, The Netherlands, March 12-18, 2016.
- General Chair, IEEE International Cloud and Autonomic Computing Conference (ICCAC 2015), September 21-25, 2015, Massachusetts, USA.
- Program Committee Member, 23^{rd} IEEE Intl. Symp. Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2015), Atlanta, GA, October 5–7, 2015.
- Program Committee Member, Sixth ACM/SPEC International Conf. Performance Engineering (ICPE 2015), Austin, Texas, USA, January 31 to February 4, 2015.
- Member, Committee for the Selection of 2014 George Mason Emerging Researcher/Scholar/Creator Award.

- Panel Member, Science Foundation Ireland, reviewer of pre-proposals for the SFI Research Centres Programme 2013, Dublin, January 13-14, 2014.
- Member (appointed by the Provost) of a Tenure Appeal Board for the case of a faculty member in the College of Science, 2013-2014.
- Program Committee Member, 22^{th} IEEE Intl. Symp. Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2014), Paris, France.
- Program Committee Member, Fifth ACM/SPEC International Conf. Performance Engineering (ICPE 2014), Dublin, Ireland, March 23-26, 2014.
- Senator, Faculty Senate, George Mason University, Fall 2013 to Spring 2016.
- Member, Committee for the Selection of 2013 George Mason Emerging Researcher/Scholar/Creator Award.
- Program Committee Member, Fourth ACM/SPEC International Conf. Performance Engineering (ICPE 2013), Prague, Czech Republic, April 21-24, 2013.
- Chair, ACM Sigmetrics 2013 Test of Time Paper Award Committee.
- Program Committee Member, 21^{th} IEEE Intl. Symp. Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2013), San Francisco, CA, August 14-16, 2013.
- Program Committee Member, 20^{th} IEEE Intl. Symp. Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2012), Washington, D.C., August 7-9, 2012.
- Chair, Graduate Studies Committee, Computer Science Department, George Mason University, August 2012 to May 2013.
- Member, Research Council, Volgenau School of Engineering, George Mason University, August 2012 to May 2013.
- Member, Curriculum Committee for Graduate Programs in Data Analytics, Volgenau School of Engineering, George Mason University, August 2012 to May 2013.
- Member, Graduate Studies Committee, Volgenau School of Engineering, George Mason University, August 2012 to May 2013.
- Member, Faculty Recruitment Committee, Computer Science Department, George Mason University, January to May 2013.

- Program Committee Member, 19^{th} IEEE Intl. Symp. Modelling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2011), Singapore, July 25-27, 2011.
- Program Committee Member, 8^{th} Intl. Conf. Quantitative Evaluation of Systems (QEST 2011), Aachen, Gemany, Sept. 5-8, 2011.
- Program Committee Member, Second ACM/SPEC International Conf. Performance Engineering (ICPE 2011), Joint WOSP/SIPEW Conference, Karlsruhe, Germany, March 14-16, 2011.
- Member, Search Committee, Associate Director, Government Relations for Research, George Mason University, January to May 2011.
- Program Committee Member, 22nd Intl. Symp. Comp. Architecture and High Performance Computing, Brazilian Computing Society, Petropolis, RJ, Brazil, October 27-30, 2010.
- Program Committee Member, 1st International Workshop on Run-time mOdels for Self-managing Systems and Applications (ROSSA 2009), Pisa, Italy, October 19, 2009.
- Program Committee Member, Tenth ACM Conference on Electronic Commerce (EC'09), July 6-10, 2009, Stanford, California, USA
- Program Committee Member, 17^{th} IEEE Intl. Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2009), Imperial College, London, UK, September 2009.
- Program Committee member, 6^{th} International Conference on the Quantitative Evaluation of SysTems (QEST) 2009, third week of September 2009, Budapest, Hungary.
- Program Committee member, IEEE HPCC 2009 Conference on "Web Services and Internet Computing," sponsored by the IEEE Technical Committee on Scalable Computing, Seoul, Korea, June 2009
- Program Committee member, 7^{th} ACM International Workshop on Software and Performance (WOSP 2008), Princeton, NJ, USA, June 23–26, 2008.
- Program Committee member, The Fourth International Conf. Autonomic and Autonomous Systems (ICAS'08), March 16-21, 2008, Gosier, Guadeloupe, Spain.
- Program Committee member, 5^{th} International Conference on Quantitative Evaluation of SysTems, QEST 2008, St.Malo, France, September 2008.
- Program Committee member, The Third International Conference on Autonomic and Autonomous Systems (ICAS 2007), Athens, Greece, June 19–25, 2007.

- Program Committee member, 6^{th} ACM Workshop on Software and Performance (WOSP 2007), Buenos Aires, Argentina, February 2007.
- Program Committee member, The Second International Workshop on Advanced Architectures and Algorithms for Internet Delivery and Applications Pisa, Italy - October 10, 2006.
- Program Committee member, Autonomic Computing track of the 26th International Conference on Distributed Computing Systems (ICDCS 2006), Lisbon, Portugal, July 4-7, 2006
- Program Committee member, 2006 IEEE International Conference on Autonomic Computing (ICAC'06), Dublin, Ireland, June 12-16, 2006.
- Program Committee member, Seventh ACM Conference on E-commerce, Ann Arbor, MI, June 11-15, 2006.
- Program Committee member, Quantitative Evaluation of Systems (QEST'05), Torino, Italy, September 19-22, 2005.
- Program Committee Co-chair, Performance 2005, France, October 2005.
- Program Committee Member, 2005 International Conference on Web Engineering, Sydney, Australia, July 2005.
- Tutorials Chair and Program Committee Member, 2005 Workshop on Software and Performance, Palma de Mallorca, Spain, July 11-15, 2005.
- Program Committee Member, 2005 International Conference on Autonomic Computing (ICAC 2005), Seattle, Washington, June 2005.
- Program Committee Member, 25th IEEE International Conference on Distributed Computing Systems (ICDCS), Columbus, Ohio, June 5-9, 2005.
- Program Commitee Member, IEEE ISSPIT - International Symposium on Signal Processing and Information Technology, Rome, Italy, December 18–21, 2004.
- Program Committee Member, 1st. International Conference on Quantitative Evaluation of Systems (QEST), Enschede, The Netherlands, September 27–30, 2004.
- External Peer Reviewer for research projects submitted by Italian universities, Italian Ministry for Education University and Research (MIUR), since 2003.
- Program Committee Member, 2004 ACM Workshop on Software and Performance, San Francisco, CA, January 14-16, 2004.

- General Chair, 2003 ACM Conference on E-commerce, June 2003, San Diego, CA.
- Program Committee Member, 2003 IEEE/ACM Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), Florida, October 2003.
- Program Committee Member, The First Latin American World Wide Web Conference, 10-12 November 2003, Santiago, Chile.
- Program Committee Member, DSN-2003—Performance and Dependability Symposium The International Conference on Dependable Systems and Networks San Francisco, CA, June 22nd - 25th, 2003.
- Program Committee Member, Second Workshop on the Performance of Computer and Communication Systems (WPerformance 2003), Brazilian Computing Society, August 2003, Campinas, Brazil.
- Program Committee Member, Practice and Experience track, 2003 International World Wide Web conference (WWW2003), May 20-24, 2003 in Budapest, Hungary.
- Program Committee Member, special track on E-Commerce Technologies, ACM Symposium on Applied Computing (SAC), March 9 - 12, 2003, Melbourne, Florida, USA.
- Judge, 2002 Digital Government Awards, organized by Accenture and MIT.
- Program Committee Member, First Seminar on Advanced Research in Electronic Business (EBR'2002), Rio de Janeiro, Brazil, November, November 7–8, 2002.
- Program Committee Member, 10th IEEE/ACM Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), Fort Worth, Texas, Oct 11-16, 2002.
- Program co-chair and co-general chair of the 2002 Practical Aspects of Performance Analysis (PAPA 2002), June 15, Marina del Rey, CA, ACM Sigmetrics and CMG.
- Program Committee Member, 2002 ACM Sigmetrics Conference, Marina del Rey, CA, June 15-19, 2002.
- Program Committee Member, Third Workshop on Software and Performance, June 2002, Rome, Italy.
- Program Committee Member, Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation (Tools 2002), Imperial College, London, April 15-17.

- Program Committee Member, 2001 ACM Conference in E-commerce, October 15, 2001, Tampa, FL.
- Program Committee Chair, CMG 2000 Research Track, December 2000, Orlando, Florida.
- Program Committee Co-Chair, The Second International Workshop on Software and Performance, September 2000, Ontario, Canada.
- Program Committee Member, Tools 2000 Conference, Chicago, March 27-31, 2000.
- Member of the Program Committee, 7th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, October 24-28, 1999, Maryland, USA.
- General Chair, 1999 ACM Sigmetrics Conference, June 1999, Atlanta, GA.
- Member of the Program Committee, The First International Workshop on Software and Performance, Santa Fe, New Mexico, October 5-9, 1998.
- Tutorials Chair, 1998 ACM Sigmetrics Conference, Wisconsin-Madison, June 23-26, 1998.
- Member of the Program Committe of the International Workshop on Petri Nets and Performance Models 1997 (PNPM97), Saint Malo, France, 1997.
- Member of the Program Committee of the Performance'96 Conference, Lausanne, Switzerland, October 7-11, 1996.
- Member of the Program Committee of the Sigmetrics'95 Conference, Ottawa, Canada, May 1995.
- Member of the Mission, Structure, and Operations of the School of Information Technology and Engineering Committee, George Mason University, 1995.
- Member of the Program Committee of the VI Brazilian Symposium on Computer Architectures and High Performance Processors, Caxambu, Minas Gerais, Brazil, August 3-4, 1994.
- Member of the Search Committe for Director of Information Management and Analysis, George Mason University, 1994.
- Member of the Program Committee of the Second Workshop on Heterogeneous Processing of the Seventh International Parallel Processing Symposium, Newport Beach, CA, USA, April 13-16, 1993.

- Member of the Program Subcommittee on Modeling and Performance Evaluation of the 13^{th} International Conference on Distributed Computing Systems, Pittsburgh, Pennsylvania, USA, May 25-28, 1993.
- Member of the Program Committee of the International Conference on Data and Knowledge Management, Baltimore, MD, November 9-12, 1992.
- Member of the Program Committee of the First International Conference on Computer Communications and Networks (IC^3N), San Diego, CA, June 8-10, 1992.
- Chairman of the III Brazilian Symposium on Computer Architecture and Parallel Processing, sponsored by the Brazilian Computer Society, November 1990.
- Member of the Examination Committee for Assistant Professor at the Computer Science Department, Federal University of Minas Gerais, November 1989.
- Member of the Steering Committee of the Brazilian Computer Society, April 1989 to April 1991.
- Member of the Group for Software Similarity Analysis, Special Secretariat of Informatics (SEI), Ministry of Science and Technology. August 1988 to February 1989.
- Member of the Program Committee of the Third International Conference on Data Communication Systems and their Performance, Rio de Janeiro. June 1987.
- Member of the Committee *Human Resources in Computer Science and Computers in Education*, Federal Education Council, Ministry of Education. February 1987 to June 1987.
- Member of the Steering Committee of the Brazilian Computer Society, May 1983 to March 1987.
- Member of the Steering Committee in Computer Science of the Brazilian Research Council (CNPq). January 1984 to December 1986.
- Member of the Management and Budget Committee of the Technological and Scientific Center at PUC-Rio, September 1982 to March 1984.
- Member of the International Organization Committee of the Seventh International Conference on Very Large Databases, Cannes, France. September 1981.
- Member of the Program Committee of COMPDEC - Computer Data Engineering Conference, Los Angeles, USA. April 1984.
- Member of the Program Committee of the ACM SIGMOD Conference, 1984.

- Member of the Graduate Committee, Computer Science Department, PUC-Rio, March 1981 to April 1984.
- Member of the Program Committee of the Eighth International Conference on Very Large Databases, Mexico. September 1982.
- Member of the Organizing Committee of the Second Computer Science School, 1982.
- Coordinator of the Graduate Program, Computer Science Department, PUC-Rio, March 1979 to August 1979.

14 Professional Membership

- Association for Computing Machinery (ACM), Fellow
- Institute of Electrical and Electronics Engineers (IEEE), Fellow.
- Special Interest Group on Measurement and Evaluation (ACM Sigmetrics), Member.
- Special Interest Group in E-commerce (ACM SIGEcom), Member.
- Computer Measurement Group (CMG), Member.

15 Editorial and Other Services

- Member of the Editorial Board, ACM Transactions on Internet Technologies, 2013 to present.
- Member of the Editorial Board, ICST Transactions on Real-World Web, Institute for Computer Sciences, Social-informatics and Telecommunications Engineering, May 2009 to present.
- Member of the Editorial Board, Performance Evaluation Journal, Elsevier, from July 2010 to present.
- Member of the Editorial Board, Journal of the Brazilian Computing Society, Springer, UK, October 2009 to present.
- Member of the Editorial Board, ACM Transactions on the Web, 2012.
- Guest editor with J. Kephart, theme issue on Autonomic Computing, IEEE Internet Computing, January/February 2007.
- Member of the Editorial Board, ACM Transactions on the Web (TWEB), September 2005 to August 2007.

- Guest editor with C.M. Woodside, theme issue on Application-level QoS, IEEE Internet Computing, June/July 2006.
- First elected Vice-Chair, ACM's Special Interest Group on E-commerce (SIGecom), June 2003 to June 2005.
- Member, Editorial Board, IEEE Internet Computing, March 2003 to 2008.
- Associate Editor, Electronic Commerce Research and Applications journal, Elsevier Science, 2001 to August 2006.
- Guest co-editor, special issue on Application-level QoS, IEEE Internet Computing, Vol. 10, No. 3, May/June 2006.
- Guest editor, spccial isssue on E-commerce and services, ACM Sigmetrics Performance Evaluation Review, December 2004.
- Guest editor, special issue of the ACM PAPA 2002 workshop, ACM Sigmetrics Performance Evaluation Review, September 2002.
- Information Director, ACM Sigmetrics, July 1996 to July 1997.
- Member of NSF Panel, January 1997, March 1999, and May 2003.
- External examiner of the Ph.D. dissertation of Koustuv Dasgupta, University of Maryland, Baltimore County (UBMC), April 30, 2003.
- External examiner of the Ph.D. Dissertation of Hesham El-Sayed, Carleton University, Ottawa, Canada, October 14, 1999.
- External examiner of the Ph.D. dissertation of Jim (Zhanwen) Li, Carleton University, Ottawa, Canada, April 18, 2011.
- Technical Editor of the Brazilian Computing Journal (SBC), April 1983 to June 1986.
- Member of the International Advising Board of the journal IEEE DataBase Engineering.
- Reviewer for several journals including: Journal of the ACM, Journal of Parallel and Distributed Processing, ACM Transactions on Database Systems, Computer Networks, IEEE Internet Computing, IEEE Transactions on Software Engineering, IEEE Transactions on Computers, IEEE Transactions on Knowledge and Data Engineering, IEEE Transactions on Services Computing, Performance Evaluation, International Journal of High Speed Computing, Journal of Computer and Software Engineering, and Brazilian Computing Journal.

- Book reviewer for McGraw Hill and John Wiley.
- External reviewer for promotion, hiring, and tenure cases at Carleton University, College of William and Mary, Georgia Institute of Technology, University of Minnesota, University of New Orleans, Boston University, University of Waterloo, University of Maryland Baltimore County, University of New Hampshire, New Jersey Institute of Technology, Georgetown University, and The City University of New York.
- External reviewer for Canada's Premier's Research Excellence Award (PREA), 1999.