

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.
Petitioner,

v.

OPTIMNET LLC,
Patent Owner,

U.S. Patent No. 8,780,887

Issue Date: July 15, 2014

Title: METHOD OF NETWORK-BASED COMMUNICATION IN VIRTUAL
NETWORK ENVIRONMENT

Inter Partes Review No.: IPR2026-00320

**DECLARATION OF DR. A.L. NARASIMHA REDDY IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF
U.S. PATENT 8,780,887**

Mail Stop "PATENT BOARD"

Patent Trial and Appeal Board
U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. MY BACKGROUND AND QUALIFICATIONS	3
III. RELEVANT LEGAL STANDARDS	6
IV. OVERVIEW	10
V. BACKGROUND	11
A. Summary of the '887 Patent.....	11
B. Claim Construction	14
C. Priority Date of Claims	15
D. Person of Ordinary Skill in the Art (POSITA)	15
E. State of the Art	16
VI. CHALLENGES	21
A. Identification of Challenges	21
B. Claims 1-10 are unpatentable over Pankratov.	22
Claim 1	23
Claim 2	32
Claim 3	34
Claim 4	37
Claim 5	38
Claim 6	40
Claim 7	44
Claim 8	46

Claim 9	48
Claim 10	51
CONCLUSION AND DECLARATION.....	i
APPENDIX A – MATERIALS RELIED UPON.....	A
APPENDIX B – <i>CURRICULUM VITAE</i>	B

I, Dr. A.L. Narasimha Reddy, declare as follows:

I. INTRODUCTION

1. I am over the age of eighteen (18) and otherwise competent to make this declaration.

2. I have been retained by Cisco Systems, Inc. (Petitioner) as an independent expert consultant in this proceeding before the United States Patent and Trademark Office. Although I am being compensated at my rate of \$650 per hour for the time I spend on this matter, no part of my compensation depends on the outcome of this proceeding, and I have no other interest in this proceeding. To the best of my knowledge, I have no financial interest in Client.

3. This Petition for *inter partes* review involves U.S. Patent No. 8,780,887 (“the ’887 Patent”, EX-1001), titled “METHOD OF NETWORK-BASED COMMUNICATION IN VIRTUAL NETWORK ENVIRONMENT” and lists Sun Cheul Kim, Seong Moon, Seungwoo Hong, Ho Sun Yoon, and Ho Yong Ryu, as the inventors.

4. For the purposes of this *inter partes* review as discussed later, I have been instructed to assume that the effective filing date of the Claims of the ’887 Patent challenged by the Petitioner in this *inter partes* review is no earlier than October 29, 2010, the filing date of Korean Patent Application No. 10-2010-0107107.

5. The '887 Patent is directed to network-based communications in a virtual network environment. One problem identified in the '887 Patent is that when virtual nodes are connected to the internet through a network address translation (NAT) device or firewall, there are problems with establishing direct communication between the nodes because the internal IP addresses are not known. The '887 Patent's proposed solution was to utilize a tunnel mapping controller to establish tunnels between the nodes to allow direct communication. I am familiar with the technology described in the '887 Patent as of the earliest possible priority date of October 29, 2010.

6. In preparing this Declaration, I have reviewed the '887 Patent (EX-1001), the file history of the '887 Patent (EX-1002), and each of the documents cited herein, and I have considered these documents in light of the general knowledge in the art as of October 29, 2010. In formulating my opinions, I have relied upon my experience in the relevant art. I have also considered the viewpoint of a person of ordinary skill in the art ("POSITA") in the relevant field, as of October 29, 2010, who I describe below (§V.D). For convenience, the materials I considered in forming my opinion are listed in Appendix A.

7. The opinions and comments formulated during this assessment are based on observations and information available at the time of the investigation. The findings presented herein are made to a reasonable degree of scientific

certainty. I have made every effort to accurately and completely investigate all areas of concern identified during our investigation.

8. I have been asked to provide my technical expertise, analysis, insights, and opinions regarding the '887 Patent and relevant references that form the basis of the arguments set forth in the accompanying Petition for *inter partes* review of the '887 Patent. As described in detail below, I offer the following opinion in this Declaration:

- A POSITA would have found Claims 1-10 of the '887 Patent obvious over Pankratov.

II. MY BACKGROUND AND QUALIFICATIONS

9. I am currently the Truchard Foundation Chair Professor of Electrical and Computer Engineering at Texas A&M University in College Station, Texas. I have over 25 years of experience in a wide variety of technologies and industries relating to data communications, storage systems, and distributed systems, including packet-switched network communications.

10. My academic credentials include a Bachelor of Technology degree in Electronics and Electrical Communications Engineering from the Indian Institute of Technology, Kharagpur, India, in August 1985. I then received a Master of Science and a Ph.D. in Computer Engineering from the University of Illinois at Urbana-Champaign in May 1987 and August 1990, respectively.

11. I have worked for over 25 years in the field of Electrical Engineering. My primary focus and research interest has been on computer networks, storage systems, multimedia systems, and computer architecture. I have authored and co-authored over a hundred technical papers and several book chapters related to several of these interests, including on such topics as multipath routing, route control, high-speed networks, network congestion, packet management, quality of service regulation, network security, network modeling, differentiated services, storage system enhancements, caching strategies, and multimedia system enhancements to name a few examples. I am listed as an inventor on six U.S. patents in the field of multi-node communication networks and computer systems.

12. My employment history following my graduation from the University of Illinois at Urbana-Champaign began at the IBM Almaden Research Center in San Jose, California in 1990. At IBM, I worked on projects related to disk arrays, multiprocessor communication, hierarchical storage systems and video servers.

13. In 1995, I joined the faculty of the department of Electrical Engineering at Texas A&M University initially as an Associate Professor and was later promoted to full, tenured Professor. At Texas A&M, during 2016-21, I was the Associate Agency Director for Strategic Initiatives and Centers for the Texas A&M Engineering Experiment Station (TEES), which engages in engineering and

technology-oriented research and educational collaborations. Further, I also served as Associate Dean for Research in the College of Engineering during 2016-21.

14. At Texas A&M, I have taught dozens of courses related to computer networking and communications, as well as computer architecture, multimedia systems and networks, topics in networking security, multimedia storage and delivery, as well as networking for multimedia applications. I have also served on various committees for the benefit of the scientific community and the Texas A&M University community.

15. I am a member of a number of professional societies, including the Institute of Electrical and Electronic Engineers (IEEE), where I have been elevated to an IEEE Fellow, and the Association for Computing Machinery (ACM). I have been responsible for chairing or co-chairing numerous conferences and programs, as well as presenting research at major IEEE and ACM conferences. For example, I served as program co-chair for the 2008 5th International Conference on Broadband Communications, Networks and Systems, panels co-chair for the 2008 3rd International Conference on Communication Systems Software & Middleware, and panel chair of the IEEE Conference of High Performance Computer Architecture.

16. My presentations include a Keynote speech at International Conference on Information Technology-New Generations in 2013, a Keynote

speech at IEEE International Symposium on Computers and Communications 2010, and several invited talks including at Georgia Tech (2013), COMSNETS Conference (2013), Int. Conf. on Networking and Communications (2012), Samsung (2011), Korea University (2011), Aijou University (2011), Catedra Series at University of Carlos III, Madrid (2009), Thomson Research, Paris (2009), Telefonica Research, Barcelona (2009), and a Distinguished Seminar at IBM Austin Research Lab (2008).

17. I have received multiple awards in the field of networks and computer architecture. I received the NSF Career Award from 1996-2000. I received an outstanding professor award by the IEEE student branch at Texas A&M during 1997-1998, an outstanding faculty award by the department of Electrical and Computer Engineering during 2003-2004, a Distinguished Achievement award for teaching from the former students' association of Texas A&M University, and a citation "for one of the most influential papers" from the 1st ACM Multimedia conference.

18. My curriculum vitae, which includes a more detailed summary of my background and experience, is attached as Appendix B.

III. RELEVANT LEGAL STANDARDS

19. I am not a lawyer and do not provide any legal opinions in this Declaration. Although I am not a lawyer, I have been informed and understand

that certain legal standards are to be applied by technical experts in forming opinions regarding the meaning and validity of patent claims. I have been asked to provide my opinions regarding whether the claims of the '887 Patent are anticipated or would have been obvious to a person having ordinary skill in the art at the time of the alleged invention, in light of the prior art.

20. I have been informed and understand that a patent claim is not patentable under 35 U.S.C. § 103 if the differences between the patent claim and the prior art are such that the claimed subject matter as a whole would have been obvious at the time the claimed invention was made to a person having ordinary skill in the relevant art. Obviousness, as I have been informed and understand, is based on the scope and content of the prior art, the differences between the prior art and the claim, the level of ordinary skill in the art, and, to the extent that they exist, certain objective indicia of non-obviousness.

21. I understand that objective indicia can be important evidence regarding whether a patent is obvious or nonobvious, if it has an appropriate nexus to the claimed invention, i.e., is a result of the merits of a claimed invention (rather than the result of design needs or market-pressure advertising or similar activities). Such indicia include: commercial success of products covered by the patent claims; a long-felt need for the invention; failed attempts by others to make the invention; copying of the invention by others in the field; unexpected results

achieved by the invention as compared to the closest prior art; praise of the invention by the infringer or others in the field; the taking of licenses under the patent by others; expressions of surprise by experts and those skilled in the art at the making of the invention; and the patentee proceeded contrary to the accepted wisdom of the prior art.

22. I have been informed that whether there are any relevant differences between the prior art and the claimed invention is to be analyzed from the view of a person of ordinary skill in the relevant art at the time of the invention. As such, my opinions below as to a person of ordinary skill in the art are as of the time of the invention, even if not expressly stated as such; for example, even if stated in the present tense.

23. In analyzing the relevance of the differences between the claimed invention and the prior art, I have been informed that I must consider the impact, if any, of such differences on the obviousness or non-obviousness of the invention as a whole, not merely some portion of it. The person of ordinary skill faced with a problem is able to apply his or her experience and ability to solve the problem and also look to any available prior art to help solve the problem.

24. I have been informed that a precise teaching in the prior art directed to the subject matter of the claimed invention is not needed. I have been informed that one may take into account the inferences and creative steps that a person of

ordinary skill in the art would have employed in reviewing the prior art at the time of the invention. For example, if the claimed invention combined elements known in the prior art and the combination yielded results that were predictable to a person of ordinary skill in the art at the time of the invention, then this evidence would make it more likely that the claim was obvious. On the other hand, if the combination of known elements yielded unexpected or unpredictable results, or if the prior art teaches away from combining the known elements, then this evidence would make it more likely that the claim that successfully combined those elements was not obvious.

25. I have been informed and understand that there are recognized, exemplary, rationales for combining or modifying references to show obviousness of claimed subject matter. Some of the rationales include the following: combining prior art elements according to known methods to yield predictable results; simple substitution of one known element for another to yield predictable results; use of a known technique to improve a similar device (method or product) in the same way; applying a known technique to a known device (method or product) ready for improvement to yield predictable results; choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success; known work in one field of endeavor may prompt variations of it for use in either the same field or a different one based on design incentives or other

market forces if the variations are predictable to one of ordinary skill in the art; and some teaching, suggestion, or motivation in the prior art that would have led one of ordinary skill to modify the prior art reference or to combine prior art teachings to arrive at the claimed invention.

IV. OVERVIEW

26. The '887 Patent is directed to network-based communications in a virtual network environment where communication nodes are behind a network address translation (NAT) device. Using a known hole-punching mechanism, the '887 Patent described using a tunnel mapping controller to establish a tunnel for direct communications between the nodes.

27. However, it was well-known to use hole-punching to establish direct communications using tunnels between communication nodes that were behind NAT devices. The Pankratov reference discloses using a hole-punching mechanism, and a meditation server as a tunnel mapping controller to establish a tunnel to establish direct communications between the nodes, and demonstrates that the purported invention recited in the Challenged Claims was well-known at the time of the '887 Patent and is unpatentable.

28. As detailed below, I explain how each limitation recited in the Challenged Claims are found in the prior art. The Challenged Claims are directed to prior art systems and combinations of conventional components that perform

conventional functions regarding network-based communications in a virtual network environment where communication nodes are behind a network address translation (NAT) device.

V. BACKGROUND

A. Summary of the '887 Patent

29. The '887 Patent is directed to network-based communications in a virtual network environment. One problem identified in the '887 Patent is that when virtual nodes are connected to the internet through a network address translation (NAT) device or firewall, there are problems with establishing direct communication between the nodes because the internal IP addresses are not known. EX-1001 3:19-24. The '887 Patent's proposed solution was to utilize a tunnel mapping controller to establish tunnels between the nodes to allow direct communication. EX-1001 6:8-14.

30. In one embodiment, virtual tunnel routers (TRs) operate as network components which form one or more virtual networks. EX-1001 5:39-40. A tunnel mapping controller serves to process and control tunneling between virtual TRs. The tunnel mapping controller establishes an IP tunnel in conjunction with each of the virtual TRs and connects each virtual IP prefix to the established tunnel, thus enabling direct communication between virtual networks. EX-1001 6:8-14.

31. The tunnel mapping controller includes a virtual network prefix information table. EX-1001 6:15-16. “Virtual network prefix information is addresses assigned to terminals or devices connected to respective virtual networks. For example, if a virtual network has a prefix and length of 1.1.1.0/24, each terminal connected to the virtual network is assigned one IP in the range from 1.1.1.1 to 1.1.1.255, which is virtual network prefix information.” EX-1001 6:48-53. The network prefix information is generated by the tunnel mapping controller or predetermined by an administrator, provided to the virtual TRs, and then stored in a database. EX-1001 6:53-61. When the virtual TRs request virtual network prefix information after registering themselves, the tunnel mapping controller allocates and forwards virtual network prefixes to the virtual TRs. EX-1001 6:16-20.

32. In the process of registering the virtual TRs in the tunnel mapping controller, a NAT traversal function or a Session Traversal Utilities for NAT (STUN) procedure is performed together to learn the incorporation of an NAT/firewall. The information learned through this procedure is used as information such as mapped IP/port across the NAT and network connection type or network operation mode to be used later in a virtual TR tunneling process. The NAT traversal function or STUN procedure is a well-known technology referred to

as “hole-punching” and the ’887 patent explains that a detailed description thereof was therefore omitted. EX-1001 50-60.

33. When the TRs are connected to the virtual network, the tunnel mapping controller creates information of a list of virtual TRs for tunneling and reports it to the TRs in order to notify the TRs that the currently registered first virtual TR and second virtual TR are targets to be tunneled for direct communication. The information of a list of virtual TRs for tunneling contains information such as the mapped IP/port, network connection type, and operation mode of a connected network to be used for the hole punching process. When each virtual TR receives the information of a list of virtual TRs for tunneling, a hole punching message for UDP/ TCP Hole Punching is exchanged in order to establish a tunnel with the virtual TR. Then, each virtual TR determines whether the transmission and reception of the hole punching message are successful. If the message transmission and reception are successful, a tunnel establishment message is exchanged. EX-1001 8:1-25.

34. The ’887 Patent issued with ten claims, with Claims 1 and 6 as independent. Claim 1 is representative:

1. A communication method that provides end-to-end communication in a virtual network environment, the method comprising:

a plurality of virtual tunnel routers comprising a first virtual tunnel router and a second virtual tunnel router receiving a message containing virtual network prefix information from a tunnel mapping controller;

determining whether the first virtual tunnel router has received a hole punching message from the second virtual tunnel router;

if the first virtual tunnel router has received a hole punching message, the first virtual tunnel router establishing a tunnel with the second virtual tunnel router based on the virtual network prefix information; and

the first virtual tunnel router performing direct communication with the second virtual tunnel router will be omitted in the exemplary embodiment of the present invention.

B. Claim Construction

35. I have been informed and understand that in order to properly evaluate the '887 Patent, the terms of the claims must first be interpreted.

36. I understand that Petitioner proposes that each claim term in the Challenged Claims be given its plain and ordinary meaning in this proceeding, and that no specific construction of any claim term is required because the prior art relied on in the Petition meets each of the claim terms under any reasonable construction. The plain meaning is the ordinary and customary meaning to those skilled in the art.

C. Priority Date of Claims

37. For the purpose of this *inter partes* review, I have been instructed to assume that the effective filing date of each of the Challenged Claims is no earlier than the filing date of Korean Patent Application No. 10-2010-0107107. Solely for the purpose of this IPR, I have been instructed to assume the priority date is October 29, 2010.

D. Person of Ordinary Skill in the Art (POSITA)

38. I am familiar with the knowledge and capabilities of one of ordinary skill in the art. Unless otherwise stated, my testimony below refers to the knowledge of one of ordinary skill in the art as of October 29, 2010, the earliest possible effective filing date of the '887 Patent.

39. I have been informed and understand that a Person of Ordinary Skill in the Art ("POSITA") is a hypothetical person who is presumed to be aware of all pertinent prior art, thinks along conventional wisdom in the art, and is a person of ordinary creativity.

40. A Person of Ordinary Skill in the Art ("POSITA") in October 2010 would have been someone knowledgeable and familiar with network based communications including virtual networks. A POSITA would have gained knowledge of these concepts through a mixture of training and work experience, such as by having at least a bachelor's degree in electrical engineering or computer

engineering or science, or related field, and at least two to three years of training or additional work experience in the area of network communications, or a related field. Additional hands-on and design experience could compensate for less formal education, and vice versa.

41. A POSITA may work as part of a multi-disciplinary team and draw upon not only his or her own skills, but also take advantage of certain specialized skills of others on the team to solve a given problem.

E. State of the Art

42. The following section describes the state of the art for network communications as of the priority date of the '887 Patent. The prior art references, and the discussions of what was known to a POSITA, provide the factual support for the general description of the state of the art at the time of the invention, provide additional motivation to modify or combine the references, and further demonstrates why doing so would involve a reasonable expectation of success.

43. Network Address Translation (NAT) causes well-known difficulties for peer-to-peer (P2P) communication, since the peers involved may not be reachable at any globally valid IP address. The Internet's original uniform address architecture, in which every node has a globally unique IP address and can communicate directly with every other node, has been replaced with a new de facto Internet address architecture, consisting of a global address realm and many private

address realms interconnected by Network Address Translators (NAT). EX-1005 p. 1.

44. In this address architecture, only nodes in the “main,” global address realm could be easily contacted from anywhere in the network, because only they have unique, globally routable IP addresses. Nodes on private networks could connect to other nodes on the same private network, and they could usually open TCP or UDP connections to “well-known” nodes in the global address realm. NATs on the path allocated temporary public endpoints for outgoing connections, and translated the addresses and port numbers in packets comprising those sessions, while generally blocking all incoming traffic unless otherwise specifically configured. EX-1005 p. 1.

45. The Internet's new de facto address architecture was suitable for client/server communication in the typical case when the client was on a private network and the server was in the global address realm. The architecture made it difficult for two nodes on different private networks to contact each other directly, however, which was often important to the “peer-to-peer” communication protocols used in applications such as teleconferencing and online gaming. EX-1005 pp. 1-2.

46. One prior art method of establishing peer-to-peer communication between hosts on different private networks was known as “hole punching.” This

technique was widely used in UDP-based and TCP-based applications. Hole punching enabled applications to function within the default security policy of most NATs, effectively signaling to NATs on the path that peer-to-peer communication sessions were “solicited” and thus should be accepted. EX-1005 p. 2.

47. UDP hole punching enabled two clients to set up a direct peer-to-peer UDP session with the help of a well-known rendezvous server, even if the clients were both behind NATs. Hole punching assumes that the two clients, A and B, already have active UDP sessions with a rendezvous server S. When a client registered with server S, the server recorded two endpoints for that client: the (Private IP address, UDP port) pair that the client believed itself to be using to talk with S, and the (Public IP address, UDP port) pair that the server observed the client to be using to talk with it. These endpoints were referred to as the client’s private endpoint and the client’s public endpoint, respectively. The server typically obtained the client’s private endpoint from the client itself in a field in the body of the client’s registration message, and obtained the client’s public endpoint from the source IP address and source UDP port fields in the IP and UDP headers of that registration message. EX-1005 p. 3.

48. Figure 5 illustrates hole punching where clients A and B had private IP addresses behind different NATs.




Figure 5: UDP Hole Punching, Peers Behind Different NATs

49. A and B each initiated UDP communication sessions from their local port 4321 to port 1234 on server S. In handling these outbound sessions, NAT A assigned port 62000 at its own public IP address, 155.99.25.11, for the use of A's session with S, and NAT B assigned port 31000 at its own public IP address, 138.76.29.7, to B's session with S. In A's registration message to S, A reported its private endpoint to S as 10.0.0.1:4321, where 10.0.0.1 was A's private IP address on its own private network. S recorded A's reported private endpoint, along with A's public endpoint as observed by S itself. A's public endpoint was 155.99.25.11:62000, the temporary endpoint assigned to the session by the NAT.

Similarly, when client B registered, S recorded B's private endpoint as 10.1.1.3:4321 and B's public endpoint as 138.76.29.7:31000. EX-1005 p. 5.

50. A and B each sent UDP datagrams directly to each of these endpoints. As A's outgoing message sent to B's public endpoint passed through A's NAT, this NAT noticed that this is the first UDP packet in a new outgoing session. The new session's source endpoint (10.0.0.1:4321) was the same as that of the existing session between A and S, but its destination endpoint was different. NAT A preserved the identity of A's private endpoint, and consistently translated all outbound sessions from private source endpoint 10.0.0.1:4321 to the corresponding public source endpoint 155.99.25.11:62000. A's first outgoing message to B's public endpoint thus, in effect, "punches a hole" in A's NAT for a new UDP session identified by the endpoints (10.0.0.1:4321, 138.76.29.7:31000) on A's private network (the network behind A's NAT), and by the endpoints (155.99.25.11:62000, 138.76.29.7:31000) on the main Internet. As this message traverses A's NAT, the NAT maps messages originating at 10.0.0.1:4321 to a destination of 138.76.29.7:31000, and vice-versa. EX-1005 p. 6.

51. If A's message to B's public endpoint reached B's NAT before B's first message to A had crossed B's own NAT, then B's NAT interpreted A's inbound message as unsolicited incoming traffic and dropped it—NATs do not automatically forward traffic sent to public endpoints to private endpoints. B's first

message to A's public address, however, similarly opened a hole in B's NAT, for a new UDP session identified by the endpoints (10.1.1.3:4321, 155.99.25.11:62000) on B's private network, and by the endpoints (138.76.29.7:31000, 155.99.25.11:62000) on the Internet. As this message traverses B's NAT, the NAT maps traffic originating from 155.99.25.11:62000 to a destination of 10.1.1.3:4321, and vice-versa. Once the first messages from A and B have crossed their respective NATs, holes are opened in each direction and UDP communication can proceed directly between A and B without assistance from S. This is because NATs A and B have created mappings between the appropriate public and private endpoints.

EX-1005 p. 6.

VI. CHALLENGES

A. Identification of Challenges

52. Based on my understanding of the law of obviousness discussed above, Claims 1-10 of the '887 Patent are unpatentable as being obvious over Pankratov alone.

Ground	Claims	Basis	Reference
1	1-10	§ 103	Pankratov

1. Prior Art

53. I have been informed that U.S. Patent Publication No. 2007/0157303 ("Pankratov") qualifies as prior art under Pre-AIA 35 USC §102(b).

B. Claims 1-10 are unpatentable over Pankratov.

1. The Pankratov Reference

54. Pankratov is generally directed to virtual private networking, network traffic tunneling, and network traffic relaying between internet-connected VPN members (“nodes” and “computers”). EX-1004 [0005-0007], [0023]. Specifically, the method taught by Pankratov “improves and extends UDP hole punching technology such that the chances of successfully connecting two computers, each using a NAT device, significantly increases.” EX-1004 [0022].

“In general, this advantage is provided by a secure mediated peer-to-peer system that comprises a mediation server and end-node peer clients. The mediation server tracks client locations and provides mediation services required for establishing direct peer-to-peer tunnels between client nodes.”

EX-1004 [0023].

55. Pankratov terms this method “assisted NAT and firewall traversal (“ANFT”), although this designation should not be construed to limit the invention, which can be practiced in association with clients that access the Internet through any type of traffic policing device.” EX-1004 [0047]. “In general, preparing a connection setup between two client computers using ANFT preferably includes discovery, prediction and confirmation. . . .” EX-1004 [0059].

Step (a) updates client address/port information for one or both clients. . . . [T]he mediation server performs this update by asking the

given client to send a single probe to the mediation server or other backend machine. If the server(s) do not receive this probe in a given time, the mediation server assumes the probe packet was lost

Step (b) occurs if the IP/port needs to be predicted. . . . [T]he mediation server uses the information [sic] has obtained during the earlier data collection processes to guess the next port the given client is going to use.

Step (c) prepares the tunnel In certain cases the mediation server will ask a given client to send a packet towards its peer.

Step (d), tunnel confirmation, is unconditional and is always executed. When a script reaches this step, the tunnel engine already knows what IP/port numbers both clients will be using towards each other; the mediation server uses this information to ask client A to send a probe to client B, and to ask client B to send a probe to client A. If, as described above, the mediation server then receives a probe reflection packet, the server has confirmed that the tunnel is operational in a respective direction. . . The mediation server then completes the setup sequence by notifying the clients.”

EX-1004 [0027]-[0030].

2. Detailed Application of Pankratov to the Challenged Claims

Claim 1

[1.0] A communication method that provides end-to-end communication in a virtual network environment, the method comprising:

56. To the extent that the preamble is limiting, Pankratov discloses a method that provides end-to-end communication in a virtual network environment. Specifically, Pankratov provides “a virtual private networking system that uses traditional VPNs, assisted NAT traversal (namely, UDP hole punching), proxy traversal (also known as HTTP and SSL tunneling), and/or relays.” EX-1004 [0038].

[1.1] a plurality of virtual tunnel routers comprising a first virtual tunnel router and a second virtual tunnel router receiving a message containing virtual network prefix information from a tunnel mapping controller;

57. Pankratov discloses that client computers [*virtual tunnel routers*] communicate with a server [*tunnel mapping controller*] to tunnel VPN traffic to/from other client computers through virtual networks. The client computers access a LAN or WAN and can operate as a server. EX-1004 [0040]. The client computers are responsible for emulating a LAN like environment on an end-user machine, communicating with the server and tunneling VPN traffic to/from other client computers through the virtual networks. EX-1004 [0107]. Pankratov further discloses that the client computer includes a virtual adapter. The virtual network

adapter is an operating system component that appears as a physical network card to the system, but does not correspond to any actual hardware. Instead, preferably all traffic sent by the system to this adapter is forwarded to an application process. This process then inspects the traffic and either drops it (discards it) or forwards it to the other VPN member(s) via the tunnel(s). EX-1004 [0110].

58. Pankratov further discloses that a mediation server [*tunnel mapping controller*] is responsible for configuring client computers, and that as part of configuring, it provides each computer with virtual IP address and a non-trivial subnet that covers the entire virtual IP address range as configured on the server [*receiving a message containing virtual network prefix information from a tunnel mapping controller*]. EX-1004 [0109], [0114]. The virtual adapter at the client computer is configured with the virtual IP address and the non-trivial subnet [*virtual network prefix information*]. The server uses the virtual IP address and the non-trivial subnet to group clients together. The mediation server decides which client computers require tunnels between themselves based on client grouping information. EX-1004 [0114]. Thus, a POSITA understood that the client computers would receive VPN traffic tunneled between VPN members for delivery to devices in their virtual sub-net and are therefore virtual tunnel routers.

59. Pankratov discloses that a mediation server (“server”) [*tunnel mapping controller*] provides mediation services to establish direct peer-to-peer

tunnels between client computers. EX-1004 [0023]. The mediation server tracks client locations and provides mediation services required for establishing direct peer-to-peer tunnels between client nodes. EX-1004 [0071]. In one example, Pankratov discloses that if client A (in a first peer) desires to have a secure tunnel established to client B (a second peer), both A and B execute commands from the mediation server's tunneling engine. The clients need to register for the service to obtain the tunnel engine. A connection process begins when client A establishes a connection to one of the mediation servers. Client A goes through a discovery process to enable the mediation server to understand exactly how the client's NAT device modifies the source port number when the packet passes through it from the LAN to the Internet. A similar discovery process may be carried for client B. Once discovery is complete, the mediation server knows each client's mapping/firewall type, and can generate a set of one or more scripts that will be used to carry out a tunnel setup sequence. EX-1004 [0072]. Thus, a POSITA understood that the mediation server is a *tunnel mapping controller*.

60. A POSITA understood a subnet to be a range of IP addresses associated with a network prefix to divide a larger network into smaller subnet. This understanding is consistent with Pankratov's disclosure that the virtual IP address and a non-trivial subnet covers a range of virtual IP addresses (EX-1004 [0114]), and is consistent with the disclosure of the '887 Patent that virtual

network prefix information includes the IP addresses and subnet that covers a range of virtual IP addresses:

Virtual network prefix information is addresses assigned to terminals or devices connected to respective virtual networks. For example, if a virtual network has a prefix and length of 1.1.1.0/24, each terminal connected to the virtual network is assigned one IP in the range from 1.1.1.1 to 1.1.1.255, which is virtual network prefix information.

Ex-1001 6:48-53.

61. Thus, Pankratov discloses this limitation.

[1.2] determining whether the first virtual tunnel router has received a hole punching message from the second virtual tunnel router;

62. Pankratov discloses this limitation. Pankratov discloses that messages for hole punching are sent for the purposes of creating a mapping in the NAT device that allows an internal computer to start receiving traffic from the peers outside the NAT. EX-1004 [0018]. Pankratov discloses that the ANFT process extends traditional UDP hole punching to support NAT types and performs this process by

63. sending messages in order to establish a tunnel. EX-1004 [0047].

64. Pankratov discloses that in one embodiment that the mediation server instructs the clients to go through a NAT type discovery process. The purpose of the discovery process is to enable the mediation server to understand exactly how

the client's NAT device modifies (or “mangles”) the source port number and IP source address when the packet passes through it from the LAN to the Internet. By performing discovery, the mediation server can determine the NAT mangling logic, which in turn enables the server to more accurately predict the port that will be used to set up the peer-to-peer tunnel. EX-1004 [0025]. As a result of the discovery process, the mediation server now knows each clients’ mapping/firewall type. EX-1004 [0025]. In other words, through the discovery process, the mediation server learns the mapping between the virtual IP address and port and the external IP and address of a port of a client device.

65. Specifically, the ANFT discovery process includes client computer A sending a ping packet [*hole punching message*] to computer B's address. Computer B then sends a ping reflection packet to the server containing the external source address and the port of client computer A as determined by computer B using the ping packet and the data from the ping packet. EX-1004 [0048]. The server uses this exchange to learn the state of the port mappings between the internal and external addresses on the NAT devices in front of client computer A and/or to verify that computer B is reachable from client computer A at a certain IP address and port [*determining whether the first virtual tunnel router*

has received a hole punching message from the second virtual tunnel router].¹

EX-1004 [0049].

66. Pankratov discloses that in a preferred embodiment the server again uses ping probing packets to instruct one or both client computers to communicate with the other, either sequentially or simultaneously, using the state of the port mappings the server learned or predicted based on received ping reflection packets. EX-1004 [0049], [0062]. It then observes the source information in the ping reflection packets to validate and possibly correct the predicted port mappings and confirm the tunnel between the clients. Thus, in a preferred embodiment, the server asks client A to ping client B; likewise, the server asks client B to ping client A. The server then observes the source information in the reflection packets to confirm the tunnel. EX-1004 [0067].

67. Thus, a POSITA understood that Pankratov's ANFT process is a hole punching method where both computer A and B sending ping packets [*hole punching messages*] and the server determines whether a hole punching message was received at one client from the other client.

¹ Pankratov discloses that the ping probing process can be performed for each client computer and thus either client A or client B can be the first virtual tunnel router, with the other being the second tunnel router.

[1.3] if the first virtual tunnel router has received a hole punching message, the first virtual tunnel router establishing a tunnel with the second virtual tunnel router based on the virtual network prefix information; and

68. Pankratov discloses this limitation. Pankratov discloses that its ANFT method and system extends traditional UDP hole punching to support NAT types that are, for example, not identity-preserving. EX-1004 [0047]. As discussed with respect to limitation [1.2], the discovery process is used to learn the mapping of the IP addresses in order to establish a tunnel between client A [*first virtual tunnel router*] and client B [*second virtual tunnel router*]. Pankratov discloses that once discovery is complete, a confirmation step is used to activate the tunnel between the client computers based on the discovery step including the mapping of the virtual IP addresses and subnets for client A and B [*virtual network prefix information*] to the external IP addresses of client A and B. [*if the first virtual tunnel router has received a hole punching message, the first virtual tunnel router establishing a tunnel with the second virtual tunnel router based on the virtual network prefix information*]. EX-1004 [0026], [0067]-[0069], [0110]. The confirmation step includes the server instructing one or both client computers to communicate with the other, either sequentially or simultaneously. It then observes the source information in the ping reflection packets to validate and possibly correct the prediction and confirm the tunnel between the clients. Thus, in a

preferred embodiment, the server asks client A to ping client B; likewise, the server asks client B to ping client A. The server then observes the source information in the reflection packets to confirm the tunnel. EX-1004 [0067].

69. Tunnel creation utilizing ANFT may be employed by a server to create tunnels between clients in fully-meshed and hub-n-spoke networks. The server groups clients and decides which clients require tunnels between themselves. EX-1004 [0107], [0115]-[0116]. When clients are configured in preparation for network creation, the server assigns clients a virtual IP address and a subnet. EX-1004 [0108]. Each client then configures a virtual network adapter with a non-trivial subnet that covers the entire virtual IP address range as configured by the server. EX-1004 [0110], [0114]. When a client on the network sends a packet through a tunnel, it uses its target's virtual IP address. When the target client receives the packet, the client's virtual network adapter passes the packet to the appropriate operating system. EX-1004 [0111]. The virtual network adapter may also forward the packet to other VPN member(s) via tunnel(s). EX-1004 [0110]. Thus, A POSITA would have understood that tunnels are established between clients based on the virtual network prefix information.

[1.4] the first virtual tunnel router performing direct communication with the second virtual tunnel router through the established tunnel.

70. Pankratov discloses this limitation. Pankratov discloses that if the tunnel setup protocol succeeds, the server provides each client with the IP address/port number of the other client so that the first client can start exchanging data with the second client directly through the tunnel and without passing through the server [*the first virtual tunnel router performing direct communication with the second virtual tunnel router through the established tunnel*]. EX-1004 [0031], [0070]. For example, Pankratov discloses that networks may be fully-meshed or hub-n-spoke. EX-1004 [0116]. In such networks, virtual network adapters provided to each client may receive packets and forward them to other VPN member(s) via tunnel(s). Clients may also utilize virtual network adapters to send packets to a target client using the target client's virtual IP address. EX-1004 [0110]-[0111]. Pankratov enables direct tunnels between members of a VPN. EX-1004 [0123].

Claim 2

[2.1] The method of claim 1, wherein the receiving of a message comprises: the first virtual tunnel router transmitting a registration request message to the tunnel mapping controller; and

71. Pankratov discloses this limitation. Pankratov discloses that , when the client computer is initially configured, it establishes a control session with the server EX-1004 [0109]. Pankratov further discloses that when the client computer is initially configured, it generates a cryptographic identity. It then establishes a

control session with the server, secures it using cryptography mechanisms, authenticates the server and submits its public identity (public key) for registration [*the first virtual tunnel router transmitting a registration request message to the tunnel mapping controller*]. In return, the server allocates a virtual IP address and subnet for the client computer and sends it to the client computer. The virtual IP address/cryptographic identity pair becomes part of the client computer's login information. EX-1004 [0109], [0114]. In the example discussed above, both client A and B would register with the mediation server and the server would allocate a virtual IP address for each client computer and send it to the client computers. Pankratov further discloses that each client via its virtual adapter is configured with a virtual IP address and a non-trivial subnet that covers the entire virtual IP address range as configured on the server. EX-1004 [0114].

[2.2] if the first virtual tunnel router is a virtual router not registered in the tunnel mapping controller, the first virtual tunnel router receiving a response message containing virtual network prefix information created from the first virtual tunnel router from the tunnel mapping controller.²

72. Pankratov discloses this limitation. As discussed with respect to limitation [2.1], a client registers with the mediation server. In return, the server

² The patent owner interprets this claim to read “. . . information created ~~from~~ [for] the first virtual tunnel router. . . .” EX-1006, p. 31.

allocates a virtual IP address and subnet for the client computer and sends it to the client computer. The virtual IP address/cryptographic identity pair becomes part of the client computer's login information. The client computer uses the virtual IP address and subnet to configure a virtual network adapter [*the first virtual tunnel router receiving a response message containing virtual network prefix information created from the first virtual tunnel router from the tunnel mapping controller*]. EX-1004 [0114]. A POSITA understands that the login information would be stored at the server, and if a client that has not previously registered with server, i.e., it does not have login information, the server would allocate a virtual IP address and subnet for the client computer and send it to the client computer [*if the first virtual tunnel router is a virtual router not registered in the tunnel mapping controller*]. EX-1004 [0109], [0114].

Claim 3

[3.1] The method of claim 2, comprising: if the first virtual tunnel router is a pre-stored virtual tunnel router and the first virtual tunnel router has moved, receiving an IP from a new virtual network to which the first virtual tunnel router has moved and registered;

73. Pankratov discloses this limitation. Pankratov discloses that a server in a virtual network registers a client computer in the virtual network by allocating a virtual IP address and cryptographic identity pair, which becomes part of the client computer's login information. EX-1004 [0109]. A POSITA understood that

the login information is prestored at the server so that the client computer can log into and off of the virtual network as desired. *[if the first virtual tunnel router is a pre-stored virtual tunnel router]*. EX-1004 [0115].

74. A POSITA would understand that if the client computer moves to a new virtual network, the client computer needs to register with a server in the new virtual network to receive a new virtual IP address and cryptographic identity pair, which becomes part of the client computer's login information for the new virtual network *[and the first virtual tunnel router has moved, receiving an IP from a new virtual network to which the first virtual tunnel router has moved and registered]*. EX-1004 [0109].

[3.2] the first virtual tunnel router connecting to the new virtual network using the received IP;

75. Pankratov discloses this limitation. Pankratov teaches that once the client has obtained the virtual IP address and subnet, it uses it to configure a virtual network adapter. EX-1004 [0114]. The virtual network adapter is an operating system component that appears as a physical network card to the system, but does not correspond to any actual hardware. All traffic sent by the system to this virtual network adapter is forwarded to an application process in the client computer, and data is sent to a target in the virtual network using the target's virtual IP address *[the first virtual tunnel router connecting to the new virtual network using the received IP]*. EX-1004 [0110].

[3.3] the first virtual tunnel router connected to the new virtual network, transmitting a registration and information update request message to the tunnel mapping controller; and

76. Pankratov discloses this limitation. Pankratov that after the client is connected to the server, it send updated client address/port information. The mediation server performs this update by asking the given client to send a single probe to the mediation server *[the first virtual tunnel router connected to the new virtual network, transmitting a registration and information update request message to the tunnel mapping controller]*. EX-1004 [0027], [0072].

Therefore, Pankratov discloses this limitation.

[3.4] performing a tunnel creation procedure with virtual tunnel routers adjacent to the first virtual tunnel router.

77. Pankratov discloses this limitation. Pankratov discloses that when two peer nodes like client A and client B desire to have a secure tunnel established between them, both client A and B can execute the commands of the server's tunneling engine. EX-1004 [0025]. Pankratov discloses that the tunnel creation process begins using ANFT. ANFT uses a packet exchange method referred to as ping probing. In one embodiment ping probing includes three transmissions. In the first transmission, the server transmits a ping request packet to client computer A. The ping request packet contains the address and the port of the computer B and some data, where computer B has a communication channel to the server S. The next transmission is client computer A sending a ping packet to the computer B's

address and with the payload (data) specified in the ping request packet. The third transmission is computer B sending a ping reflection packet to the server. The ping reflection packet contains the source address and the port of client computer A as determined by computer B using the ping packet and the data from the ping packet. EX-1004 [0048]. Likewise, the server may repeat ping probing originating with a ping request packet to computer B. The mediation server observes the source information in the reflection packets to confirm the tunnel and then activates the tunnel between the client computers using the predicted ports and to verify that the tunnel operates in both directions [*performing a tunnel creation procedure with virtual tunnel routers adjacent to the first virtual tunnel router*]. EX-1004 [0067].

78. Pankratov further discloses that the virtual network may be fully-meshed, which means that the tunnels are setup between all online network peers. EX-1004 0116]. Thus, in a fully meshed network, a POSITA understood that client A and client B would be *adjacent tunnel routers* connected by tunnels. Additionally, Pankratov discloses that the server decides which client computers require tunnels between themselves based on client computer grouping information. A POSITA understood that this grouping would include adjacent clients. EX-1004 [0115].

Claim 4

[4] The method of claim 3, wherein the performing of the tunnel creation procedure comprises the first virtual tunnel router

connected to the new virtual network receiving adjacent virtual tunnel router information from the tunnel mapping controller.

79. Pankratov discloses this limitation for the same reasons as described with respect to limitation [3.4]. Specifically, the server transmits a ping request packet that contains the address and the port of the computer B, which is an adjacent tunnel router in a fully-meshed network [*wherein the performing of the tunnel creation procedure comprises the first virtual tunnel router connected to the new virtual network receiving adjacent virtual tunnel router information from the tunnel mapping controller*]. EX-1004 [0048], [0116].

Claim 5

[5.1] The method of claim 1, comprising: if the first virtual tunnel router has not received a hole punching message from the second virtual tunnel router, the first virtual tunnel router performing a tunneling process with the tunnel mapping controller; and

80. Pankratov discloses this limitation. Pankratov discloses that a relay node controlled by the mediation server may be utilized if the ANFT tunnel setup protocol does not succeed [if the first virtual tunnel router has not received a hole punching message from the second virtual tunnel router, the first virtual tunnel router performing a tunneling process with the tunnel mapping controller]. EX-1004 [0031]-[0032]. Specifically, Pankratov discloses that a relay is used when the ANFT process fails, which a POSITA understood includes if the client

computer does not receive a hole punching message from the second computer.³

EX-1004 [0091]-[0095].

81. Pankratov describes that in one embodiment the relay node may be a separate computer controlled by the server, but also describes that the server may comprise the relay node:

“A representative mediation server machine is a network-based server running commodity (e.g. Pentium-class) hardware, an operating system (e.g., Linux, Windows, OS-X, or the like), an application runtime environment (e.g., Java, ASP), and a set of applications or processes (e.g., applets or servlets, linkable libraries, native code, or the like, depending on platform), that provide the functionality of a given system or subsystem. As described, the service may be implemented in a standalone server, or across a distributed set of machines.”

EX-1004 [0129].

³ Pankratov discloses that the AFNT can be implemented with both client computers A and B sending hole punching messages to the other computer and thus either computer A or B can be the recited *first virtual tunnel router* and *second virtual tunnel router*. EX-1004 [0018].

82. Thus, a POSITA understood that the relay may be included in the mediation server and thus if the hole punching is not successful the first virtual tunnel router performing a tunneling process with the tunnel mapping controller.

[5.3] when a tunnel is established between the first virtual tunnel router and the tunnel mapping controller, the first virtual tunnel router connecting the received virtual network prefix information to the tunnel.

83. As discussed with respect to limitation [1.1], Pankratov discloses that each client via its virtual adapter is configured with a virtual IP address and a non-trivial subnet that covers the entire virtual IP address range as configured on the server. EX-1004 [0114]. Thus, when a client is connected to tunnel, its adapter is linked to the tunnel so any traffic received by the client is delivered to the appropriate device associated with the virtual IP address and subnet *[when a tunnel is established between the first virtual tunnel router and the tunnel mapping controller, the first virtual tunnel router connecting the received virtual network prefix information to the tunnel]*. EX-1004 [0110]-[0111].

84. Therefore, Pankratov discloses this limitation.

Claim 6

[6.0] A communication method that enables direct communication between a plurality of virtual tunnel routers including a first virtual tunnel router and a second virtual tunnel router by working in conjunction with a tunnel mapping controller, the method comprising:

Pankratov discloses this limitation for the same reasons as discussed with respect to limitations [1.0] and [1.1].

[6.1] the first virtual tunnel router and the second virtual tunnel router requesting virtual network prefix information from the tunnel mapping controller;

85. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [1.1] and [2.1]. Specifically, Pankratov discloses that when the client computer is initially configured, it establishes a control session with the server a client computer and in return, the server allocates a virtual IP address and subnet for the client computer and sends it to the client computer [*the first virtual tunnel router and the second virtual tunnel router requesting virtual network prefix information from the tunnel mapping controller*]. EX-1004 [0109]-[0110].

[6.2] the tunnel mapping controller generating virtual network prefix information of the first virtual tunnel router and second virtual tunnel router, respectively, and transmitting the same along with information of a list of virtual tunnel routers for which previous connection information has been stored;

86. As discussed with respect to limitation [6.1], the server allocates virtual IP addresses and subnets to client computers. The client computer uses the virtual IP address and subnet to configure a virtual network adapter [*the tunnel mapping controller generating virtual network prefix information of the first*

virtual tunnel router and second virtual tunnel router, respectively, and transmitting the same]. EX-1004, [0114] . EX-1004 [0114].

87. Pankratov discloses that the server decides which client computers require tunnels in virtual network that may be fully-meshed, which means that the server sets up tunnels between all online peers. EX-1004 [0115]-[0116]. Thus, the server stores and sends to each computer the IP address and ports of other computers to establish tunnels with the computer as discussed with respect to Claim 1. For example, in a fully-meshed or hub-n-spoke network, a computer Client A would receive the IP address and port of computer client B from the mediation server, along with the IP address and port of any other client computer that computer client A will establish a tunnel with [*along with information of a list of virtual tunnel routers for which previous connection information has been stored*]. EX-1004, [0115]-[0116]. Therefore, Pankratov discloses this limitation.

[6.3] the first virtual tunnel router and the second virtual tunnel router creating a tunnel; and

88. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [1.3].

[6.4] performing direct communication between the first virtual tunnel router and the second virtual tunnel router by linking the virtual network prefix information received by the first virtual tunnel router and the second virtual tunnel router, respectively.

89. Pankratov discloses this limitation for the same reasons as discussed with respect to limitations [1.4] and [5.2]. A client's virtual adapter preferably is configured with a virtual IP address and a non-trivial subnet that covers the entire virtual IP address range as configured on the server. EX-1004 [0114]. Pankratov discloses that the server decides which client computers require a tunnel between them and groups client computers based on virtual configuration information, e.g., virtual IP address and subnet. EX-1004 [0115]. Further, the virtual network adaptor at the server receives network traffic and forwards it directly to the appropriate client computer via a tunnel associated with the subnet of the client computer. EX-1004 [0110]. Thus, a POSITA understood that the virtual IP address and subnet for sending and receiving client computers were linked to a specific tunnel associated with the subnet. *[performing direct communication between the first virtual tunnel router and the second virtual tunnel router by linking the virtual network prefix information received by the first virtual tunnel router and the second virtual tunnel router, respectively.*

90. When the mediation server provides computer A and computer B with the other's IP address and port *[received by the first virtual tunnel router and the second virtual tunnel router]*, the IP address and port of the two computers are linked to form a tunnel so that the first client can start exchanging data with the second client directly through the tunnel and without passing through the server

[performing direct communication between the first virtual tunnel router and the second virtual tunnel router]. EX-1004 Claim 1.

Claim 7

[7.1] The method of claim 6, comprising: before requesting for the virtual prefix information, the first virtual tunnel router and the second virtual tunnel router transmitting a registration request message to the tunnel mapping controller;

91. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [6.1]. The registration request message also occurs on *every* client computer to register with the server. Pankratov teaches that “A connection process [between client computers] begins when client A establishes a connection to one of the mediation servers.” EX-1004 [0072], [0109], [0110]. Furthermore, client computers establish a secure control session with the server before submitting a public identity for registration. EX-1004 [0109]. Thus, Pankratov teaches that both client computers must transmit a registration request message prior to receiving any virtual network prefix information, and discloses this limitation.

[7.2] the tunnel mapping controller determining whether the first and second virtual tunnel routers are pre-stored virtual tunnel routers based on the transmitted registration request message; and

92. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [3.1]. Pankratov discloses that a connection process

begins when client A establishes a connection to one of the mediation servers. If necessary, the mediation server instructs client A to go through a NAT type discovery process. EX-1004 [0025]. A POSITA understood that the discovery process is necessary if client A had not previously registered with the mediation server, since the mediation server uses this exchange to learn the state of the port mappings on the NAT devices in front of client computer A [*the tunnel mapping controller determining whether the first and second virtual tunnel routers are pre-stored virtual tunnel routers based on the transmitted registration request message*]. EX-1004 [0049]. This is further disclosed in Pankratov, which describes that a mediation server performs client presence tracking and location tracking which involves the client computers establishing sessions with the server to use the server's services. EX-1004 [0071], [0096]. Specifically, clients are assigned login information and may toggle their online presence individually. EX-1004 [0109], [0115]. Thus, a POSITA understood that Pankratov discloses that the server pre-stores the IP address and port number for each client computer in order to perform presence and location tracking, and enable online presence toggling.

[7.3] if the first and second virtual tunnel routers are not pre-stored virtual tunnel routers, the tunnel mapping controller storing information contained in the registration request message and registering the first virtual tunnel router and the second virtual tunnel router.

93. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [7.2]. Specifically, if the IP address and port of the first and second computers are not stored at the mediation server, the server registers the computer and executes the discovery process. EX-1004 [0025], [0049], [0109].

Claim 8

[8.1] The method of claim 6, wherein the creating of a tunnel comprises: the first virtual tunnel router and the second virtual tunnel router transmitting and receiving a hole punching message based on the information of a list of virtual tunnel routers;

94. Pankratov discloses this limitation for the same reasons as discussed with respect to limitations [1.2] and [1.3]. Specifically, Pankratov teaches the ANFT process, which extends traditional UDP hole punching to support various NAT types and utilizes hole punching messages in order to establish a tunnel between computers. EX-1004 [0047]. Additionally, Pankratov further discloses that in a fully meshed network each computer establishes a tunnel with every other computer. EX-1004 [0116]. In order to do so, a POSITA understood that a hole punching message must be exchanged between each pair of computers establishing a tunnel based on IP address and port information supplied by the mediation server *[based on the information of a list of virtual tunnel routers]*.

95. By way of further example, Pankratov teaches that in one embodiment, the server initiates tunneling based on client computer grouping information:

“The server decides which client computers require tunnels between themselves based on client computer grouping information. Client computers are grouped together using virtual network configuration objects. A client computer may belong to more than one network. The clients also have per-network online status, meaning that they can toggle their online presence for each network individually.”

EX-1004 [0115].

96. Pankratov also teach that “. . .the server completes a tunnel setup between two clients. . . .” EX-1004 [0117].

97. In other words, Pankratov teaches that hole punching messages between clients may be initiated by the server based on client grouping information, and discloses this limitation.

[8.2] determining whether the transmission and reception of the hole punching message between the first virtual tunnel router and the second virtual tunnel router are successful;

98. Pankratov discloses this limitation for the same reasons as described with respect to limitation [5.1]. Pankratov discloses that a relay is used when the ANFT process fails, which a POSITA understood includes determining if the

tunnel setup protocol does not succeed [*determining whether the transmission and reception of the hole punching message between the first virtual tunnel router and the second virtual tunnel router are successful*]. EX-1004 [0031]-[0032], [0091].

[8.3] if anyone of the virtual tunnel routers fails in the transmission and reception of the hole punching message, the first virtual tunnel router and the second virtual tunnel router each establishing a tunnel with the tunnel mapping controller; and

99. Pankratov discloses this limitation for the same reasons as described with respect to limitation [5.1]. Pankratov provides a relay node controlled by the server that connects via tunnels to clients and provides communications via relay when direct client-to-client tunneling is unavailable. EX-1004 [0020], [0091-0095]. Furthermore, as discussed with respect to claim [5.1], the server may comprise the relay node in some embodiments. EX-1004 [0129].

[8.4] the first virtual tunnel router and the second virtual tunnel router each linking the received virtual network prefix information to the established tunnel.

100. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [6.4].

Claim 9

[9.1] The method of claim 6, comprising: when the second virtual tunnel router moves after the performing of direct communication, registering the second virtual tunnel router to a virtual network to which the second virtual tunnel router has moved, and receiving an IP from the virtual network;

101. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [3.1].

[9.2] the second virtual tunnel router transmitting a registration/information update request message containing the received IP to the tunnel mapping controller;

102. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [3.3].

[9.3] the tunnel mapping controller updating the information of the second virtual tunnel router with the virtual tunnel router information contained in the registration/information update request message;

Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [3.3].

[9.4] the tunnel mapping controller searching for an adjacent virtual tunnel router having a tunnel established with the second virtual tunnel router; and

103. Pankratov discloses this limitation. Pankratov teaches that “[p]referably, the server also keeps track of all client computer pairs that have established tunnels with each other,” and that “[t]he server decides which client computers require tunnels between themselves based on client computer grouping information” [*the tunnel mapping controller searching for an adjacent virtual tunnel router having a tunnel established with the second virtual tunnel router*]. EX-1004 [0099], [0115].

104. As discussed above with respect to claim [3.4] and [6.2], Pankratov discloses a service with as few as two peer nodes, and a fully meshed networked where all nodes (*i.e.*, “adjacent” peers) are connected by tunnels, capable of connection, and a server which is aware of network topology and adjacency. *See, e.g.*, 1004 [0115], [0116], and [0024].

105. Therefore, Pankratov teaches a tunnel mapping controller with the capability to search for adjacent virtual tunnel routers and maintain in memory a listed of established tunnels and discloses this limitation.

[9.5] the tunnel mapping controller transmitting the updated virtual tunnel router information to the searched adjacent virtual tunnel router.

106. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [4.0].

107. Pankratov teaches that during the tunnel establishment process, the server will send the information of one client to an adjacent peer and request the client message that adjacent peer, and vice versa:

“Step (d), tunnel confirmation, is unconditional and is always executed. When a script reaches this step, the tunnel engine already knows what IP/port numbers both clients will be using towards each other; the mediation server uses this information to ask client A to send a probe to client B, and to ask client B to send a probe to client A.”

EX-1004 [0077].

Claim 10

[10] The method of claim 9, comprising, after the transmitting of the updated virtual tunnel router information, creating a tunnel between the second virtual tunnel router and the adjacent virtual tunnel router.

108. Pankratov discloses this limitation for the same reasons as discussed with respect to limitation [9.6]. Specifically, Step (c) prepares the tunnel, and Step (d) concludes with the server determining the tunnel is fully operational, and notifying the clients. EX-1004 [0076]-[0077]. When the tunnel is fully operational, Pankratov teaches that “client computers are provided with the other's confirmed external IP/port addresses and they may immediately start exchanging data through the tunnel.” EX-1004 [0069].

CONCLUSION AND DECLARATION

The findings and opinions set forth in this declaration are based on my research and analysis to date. I may continue my examinations in view of additional documents or other factual evidence over the course of this proceeding that may necessitate supplementing and/or refining my opinions. I reserve the right to add to, alter, or delete my opinions and my Declaration upon discovery of additional information. I reserve the right to make such changes as necessary.

In signing this declaration, I understand 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 acknowledge 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.

I declare that all statements made herein on my own knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

[SIGNATURE ON FOLLOWING PAGE]

Dated: Marh 27 2026

By: *A. L. Narasimha Reddy*

Dr. A.L. Narasimha Reddy

APPENDIX A – MATERIALS RELIED UPON

<i>Exhibit #</i>	<i>Description</i>
1001	U.S. Patent No. 8,780,887 (“the ’887 Patent”)
1002	Prosecution History of the ’887 Patent (“File History”)
1003	Declaration of Petitioner’s Expert
1004	U.S. Patent Publication No. 2007/0157303 to Pankratov (“Pankratov”)
1005	Brian Ford et al., <i>Peer-to-Peer Communication Across Network Address Translators</i> , ATEC '05: Proceedings of the USENIX Annual Technical Conference, Apr. 14, 2005.
1006	Ex. C of OptimNet’s Infringement Contentions Against Petitioner for the ’887 Patent in <i>OptimNet LLC v. Cisco Systems, INC</i> , No. 2:25-cv-00935-JRG (E.D. Tex.)

APPENDIX B – *CURRICULUM VITAE*

NARASIMHA REDDY
Dept. of Elec. & Comp. Engg.
Texas A & M University
333J Wisenbaker, College Station, TX 77843-3259
<https://cesg.tamu.edu/faculty/narasimha-reddy/>
<https://scholar.google.com/citations?user=en&user=NbHZDW4AAAAJ>
reddy@tamu.edu, 979-845-7598 (O), 979-845-2630 (FAX)

Educational Record:

University of Illinois (Urbana) : Ph.D. in Electrical and Computer Engineering. Aug. 1990.

University of Illinois (Urbana) : M.S. in Electrical and Computer Engineering. May 1987.

Indian Institute of Technology (Kharagpur, India) : Bachelor of Technology (Honors) in Electronics and Electrical Communication Engineering. May 1985.

Thesis Topic:

Ph.D Thesis : Issues in IO architectures for high-performance/ multiprocessor systems.

M.S. Thesis : Algorithm based Fault Detection Techniques for Specific Applications.

B.Tech Thesis (Project) : Studies on Residue Number System for Signal Processing

Honors and Awards:

Cited for one of the most influential papers of first ACM Multimedia Conference

Catedra de Excelencia, University of Carlos III, Madrid, Spain, Aug. 2009-Feb.2010.

Distinguished Achievement Award for teaching (College level), AFS, TAMU, 2005.

Ford Fellow, College of Engg., TAMU 2004, TEES Fellow, College of Engg., TAMU, 2001, Neely Faculty Fellow, College of Engg., TAMU, 1999.

Outstanding Professor Award, Dept. of ECE, TAMU, 2003-04, 2014-15.

Outstanding Professor Award, IEEE Student branch, TAMU, 1997-1998.

NSF Career Award (1996-2000).

5 Patents and an IBM Plateau Invention Achievement Award.

IBM Graduate Fellowship (1988-90).

Presidential Silver Medal, Indian Institute of Technology, Kharagpur, India (May 86).

Research Interests:

Storage Systems, Cybersecurity, Computer Networks and Computer Architecture.

Work Experience:

Associate Dean for Research, College of Engineering, Texas A&M Univ., Apr. 2016 -Present

Associate Agency Director for Strategic Initiatives and Centers, TEES, Aug. 2016 - Present.

Assistant Agency Director for National and Global Initiatives, TEES, June 2014 - Apr. 2016.

J.W. Runyon Professor, Dept. of ECE., Texas A&M University, Oct. 2008-Present.

Professor, Dept. of Elec. and Compu. Engg, Texas A & M University, Sep. 2004 -Sep. 2008.

Associate Professor, Dept. of Elec. Engg, Texas A & M University, Aug. 1995 - Aug. 2004.

Research Staff Member, IBM Almaden Research Center: Aug. 1990 - Aug. 1995.

Professional Activities:

- **Fellow** of the IEEE Computer Society, member of ACM SIGARCH.
- Associate Editor, IEEE/ACM Transactions on Networking, 2014-2016.
- Member, Advisory Board of Center for Hyper Scale Micro Data Center, South Korea.
- External Evaluator, Faculty Promotion Committee, University of Jordan, 2012.
- Program Committee Member, INFLOW Workshop, 2016.
- Program Committee Member, NVM Workshop, 2011, 2015, 2016.
- Program Committee Member, IEEE ANTS Conf., 2014.
- Program Committee Member, Parallel Data Storage Workshop, 2012-2014.
- Member, Information Technology Research Academy Strategy Formulation, India, Oct. 2011
- Review Chair, Los Alamos National Lab Research Information Institute program, Apr. 2011
- Program Committee Member, USENIX HotStorage Conf., 2010-2012.
- External Evaluator, Faculty Search Committee, ICS-FORTH, Crete, Greece, 2009
- Program Co-chair, IEEE Broadnets 2008
- Area Lead, IEEE Infocom 2008
- Program Committee Member, USENIX FAST Conf. 2009
- Panels Co-Chair, COMSWARE 2008
- Program Committee Member, ICDCS 2008, 2013
- Program Committee Member, IEEE ICNP 2008
- Program Committee Member, COMSNETS 2009-2013, 2015
- Program Committee Member, COMSWARE 2007-2008
- Program Committee Member, Int. Workshop Visualization for Computer Security, 2006
- Program Committee Member, CoNEXT Conference, 2006, 2007, 2009.
- Program Committee Member, Int. Workshop on QOS, 2006-2011.
- Program Committee Member, IEEE Infocom, 2005-2015.
- External Reviewer, Faculty promotion Committee, KAIST, Korea, 2003.
- Technical Consultant, EMC Corp., 2001-2003.
- Educational Content Expert on Computer Networking, Jones International University.

- Member, Technical Advisory Board, Netgeo Inc., Mountain View.
- Technical Adviser, TBD Networks.
- Proposal Reviewer, Netherlands NSF.
- Program Committee member, ACM Multimedia Conference, 2002.
- Program Committee member, MCW-2001, Multimedia on Clusters of Workstations, 2001
- Program Committee member, IEEE Int. Conf. on Multimedia Comp. and Systems, 1999
- Panel Chair, IEEE Conf. High Performance Computer Architecture, 1999.
- Program Committee member, High Performance Computing, 1999
- Program Committee member, Int. Conf. on Parallel Processing, 1998.
- Tutorial Chair, IEEE Conf. on Multimedia Computing Systems, 1998.
- Panel Chair, Int. Conf. on Parallel Processing, 1995.
- Program Committee member, Int. Conf. on Parallel Processing, 1995.
- Program track coordinator, Hawaii Int. Conf. on Systems and Sciences, 1994.

Patents

- N. Annapareddy, C-C Chou, C. Tirumalasetty, P. Gratz and A. Abouelwafa, "Faster Page Handling in Computer Systems", Provisional Patent Application, May 2020.
- N. Annapareddy, J. Brady, D. Finney, M. Hartung, M. Ko, J. Menon, and D. Nowlen, Conflict resolution in multi-node communication network, U.S. Patent number: 6,101,194, received Aug. 8, 2000.
- N. Annapareddy, J. Brady, D. Finney, R. Freitas, M. Ko, and M. Rayfield, Method and System for message status reporting in a multi-node communication network, U.S. Patent number: 5,717,862, received Feb. 10, 1998.
- N. Annapareddy, D. Finney, M. Jenkins, L. Kessler, A. Lang, S. Liang, D. Mora, D. Plomgren, P. Urbisci and A. Walls, Interconnection network for a multi-nodal data processing system which exhibits incremental scalability, US Patent number: 5,602,044, received Feb. 11, 1997.
- N. Annapareddy, J. Brady and D. Finney, Adaptive and Dynamic message routing system for multinode wormhole networks, US Patent number: 5,602,839, Received Feb. 11, 1997.
- N. Annapareddy, J. Brady, D. Finney, R. Freitas, M. Hartung, M. Ko, N. Mendelson, J. Menon, D. Nowlen and S-Y Tzou, System and method using chaines structured queues for ordering of message delivery between connected nodes wherein unsuccessful message portion is skipped and retried, US Patent number: 5,577,211, Received Nov. 19, 1996.

Archival Publications

- Mahesh, P., Tiwari, A., Jin, C., Kumar, P.R., Reddy, A. L. N., Bukkapatnam, S., Gupta, N., and Karri, R., "A Survey of Cybersecurity of Digital Manufacturing", Accepted for publication in Proceedings of IEEE, October 2020.

- Tiwari, A., Reddy, A. L. Narasimha, and Bukkaparnam, S., “Cybersecurity assurance in the emerging Manufacturing-as-a-Service (MaaS) paradigm: A lesson from the video streaming industry, ASTM Journal on Smart and Sustainable Manufacturing Systems, SSMS-2020-0066.R1, October 2020.
- Wen, F., Qin. M.¹, Gratz, P., and Reddy, A. L. Narasimha, “Hardware memory management for future mobile hybrid memory systems”, accepted for publication in IEEE Trans. on CAD, Nov. 2020
- Chou, C-C1, Jung, J., Reddy, A. L. Narasimha, Gratz, P., and Voight, D., “Virtualize and Share nonvolatile memories in user space”, CCF Transactions on High Performance Computing, Jan 2020
- Moon, S.¹, and Reddy, A. L. Narasimha, “Does RAID improve lifetime of SSD Arrays?”, accepted for publication at ACM Transactions on Storage., Apr. 2016.
- Fedorov, V.¹, Qiu, S.¹, Reddy, A. L. Narasimha and Gratz, P., “ARI: Adaptive LLC-memory traffic management”, ACM Transactions on Architecture and Code Optimization (TACO), Dec. 2013.
- Wu, X¹, Qiu, S.¹, and Reddy A. L. N., “SCMFS: A file system for Storage Class Memory and its extensions”, accepted for publication in ACM Transactions on Storage, Jan. 2013.
- Wu, X¹ and Reddy A. L. N., “A novel approach to manage hybrid storage system”, Journal of communications, vol.7., July 2012.
- Shin, S., Gu, G., Reddy, A. L. N., “A large-scale study of Conficker”, IEEE Trans. Information Forensics and Security, June 2012.
- Yadav, S¹, Reddy, A. K.¹, Reddy, A. L. Narasimha and Ranjan, S., “Detecting algorithmically generated malicious domain names”, ACM/IEEE Trans. on Networking, 2012.
- Lee, Y-O.¹ and Reddy, A. L. Narasimha, “Algorithms for disjoint path construction for failure recovery and multipath routing”, to appear in Journal on Computer Networks, 2012.
- Garrison, J.¹, and Reddy, A. L. Narasimha, “Umbrella File System: Storage management across heterogeneous devices”, ACM Transactions on Storage, vol.5, no.1, Mar. 2009.
- Kim, S.¹, and Reddy, A. L. Narasimha, “Statistical Techniques for detecting Traffic Anomalies through Packet header data” IEEE/ACM Transactions on Networking (ToN), July 2008.
- Liu, Y.¹, and Reddy, A. L. N., ”Multihoming route control among a group of multihomed stub networks”, Journal on Computer Communications, Nov. 2007.
- Bhandarkar, S.¹, Reddy, A. L. N., Allman, M., Blanton, E., ”Improving the robustness of TCP to non-congestion events”, IETF(Internet Engineering Task Force) RFC 4653, August 2006.
- Kang, S¹., Reddy, A. L. N., “An Approach to Virtual Allocation in Storage Systems”, in ACM Transactions on Storage, Nov. 2006, Vol.2, pages:371-399.
- Kim, S¹., Reddy, A. L. N., “Image based Anomaly detection: Design, Algorithms and Evaluation”, IEEE Journal on Selected Areas of Comm., vol. 24.10, October 2006, 1942-1954.

- Bhandarkar, S.¹, Reddy A. L. N., “LTCP: A layered framework for improving TCP’s performance in high-speed networks”, in ACM CCR, Jan. 2006.
- Kwon, D.W., Ko, M., Vannucci, M., Reddy, A. L. N., and Kim, S.¹, “Wavelet methods for the detection of anomalies and their application to network traffic analysis”, Quality and Reliability Engineering International (QREI) 2006, 22, pages: 1-17.
- Reddy, A. L. N., Wyllie, J., and Wijayarathne, R.¹, “Scheduling for multimedia servers”, invited submission for the premier issue of ACM Transactions on Multimedia, March 2005. Cited for being one of the most influential papers from ACM Multimedia conference 1993.
- Kim, S.¹, and Reddy, A. L. N., “Statistical techniques for detecting traffic anomalies through packet header data”, accepted for publication in ACM/IEEE transactions on Networking, January 2007.
- Garg, A.¹, and Reddy, A. L. N., “Mitigating Denial of service attacks through QOS regulation”, in the special issue of Journal of Microprocessors and Microsystems on secure platforms, May 2004.
- Bhandarkar, S.¹, Sadry, N.¹, Reddy A. L. N., Vaidya, N., “TCP-DCR: A novel protocol for tolerating wireless channel errors”, IEEE Transactions on Mobile Computing, March 2005.
- Zhao, Z.¹, Ametha. J., Swaroop. D. and Reddy A. L. N., “A method for estimating non-responsive traffic at a router”, ACM/IEEE Trans. on Networking, Aug. 2004.
- Yeom, I.¹. and Reddy A. L. N., “Adaptive Marking for Assured Forwarding Service” IEEE Communication Letters, Oct. 2002.
- Yeom, I.¹. and Reddy A. L. N., “Providing consistent delay differentiation” accepted by IEICE Transactions on Communications, Sept. 2002.
- Ma, X.¹. and Reddy A. L. N., “Multiview Storage Systems”, IEEE Trans. on Parallel and Distributed Systems, Apr. 2003.
- Yeom, I.¹. and Reddy A. L. N., “Marking for QOS Improvement”, Elsevier Journal of Computer Communications, pages 35-50, vol. 24, no.1, Jan 2001.
- Yeom, I.¹. and Reddy, A. L. N., “Modeling TCP behavior in a differentiated services network”, ACM/IEEE Transactions on Networking, vol. 9, no.1, pages 31-46, Feb. 2001.
- Ma, G.¹, Khaleel, A.¹. and Reddy, A. L. N., “Evaluation of storage systems based on network-attached disks”, IEEE Transactions on Parallel and Distributed Systems, September 2000, 956-968.
- Wijayarathne, R.¹. and Reddy, A. L. N., “Providing QOS guarantees for disk I/O”, Proc. of ACM/Springer Journal on Multimedia Systems, Jan 2000, vol.8, no.1, pages:57-68.
- Yeom, I.¹. and Reddy, A. L. N. , “ENDE: An End-to-end Network Delay Emulator” , Journal on Multimedia Tools and Applications, Kluwer Publishers, vol. 14, no.3, June 2001.
- Reddy, A. L. N. and Upfal, E., “Real-time communication scheduling in a multicomputer video server”, *Journal on Parallel and Distributed Computing*, Sep. 1999., vol.58, pages: 425-445.

- Reddy, A. L. N. and Wyllie, J., “I/O issues in a multimedia system”, *IEEE Computer*, Mar. 1994, vol.27, no.3, pages:69-74.
- Reddy, A. L. N., Chandy, J. and Banerjee, P., “Design and evaluation of gracefully degradable disk arrays”, *Journal of Parallel and Distributed Computing*, Jan. 1993, vol.17, pages: 28-40.
- Reddy, A. L. N. and Banerjee, P., “Algorithm Based Fault Detection Techniques for Signal Processing Applications”, *IEEE Transactions on Computers*, Oct. 1990, vol.39, no.10, pages: 1304-1308.
- Reddy, A. L. N. and Banerjee, P., “Design, Analysis and Simulation of IO Architectures for Hypercube Multiprocessors” *IEEE Transactions on Parallel and Distributed Systems*, April 1990, vol.1, no.2, pages:140-151.
- Reddy, A. L. N. and Banerjee, P., “An evaluation of multiple-disk IO systems” *IEEE Transactions on Computers*, Dec. 1989, vol. 38, no.12, pages:1680-1690.
- Reddy, A. L. N. and Banerjee, P., “A Study of Parallel Disk Organizations” *ACM Computer Architecture News*, Sept. 1989.
- Abraham, J., Banerjee, P., Chen, C.Y., Fuchs, W.K., Kuo, S. Y. and Reddy, A.L.N., “Fault Tolerance Techniques for Systolic Arrays”, *IEEE Computer*, July 1987, vol.20, no.7, pages:65-85.
- Chakraborti, N.B., Soundararajan, J.S. and Reddy, A. L. N., “An Implementation of Mixed-Radix Conversion for Residue Number Applications”, *IEEE Transactions on computers*, August 1986, vol. C-35, no.8, pages:762-764.

Publication of books or book chapters

- Bhandarkar, S.¹, and Reddy, A. L. N., “TCP in Wireless Networks”, invited book chapter in *Encyclopedia on wireless and mobile communications*, CRC Press, Dec. 2007.
- I. Yeom¹ and A. L. Narasimha Reddy, “Differentiated Services”, invited book chapter in *Wiley’s Encyclopedia on Communications*, 2002.
- A. L. Narasimha Reddy and J. Wyllie, “Disk Scheduling in a Multimedia I/O System”, reproduced as a book chapter in *Readings in Multimedia Computing and Networking*, Morgan Kaufmann Publishers, Editors: Kevin Jeffay and HongJiang Zhang, Summer 2001.
- Reddy, A. L. N. and Wijayarathne R., “Scheduling for integrated service in multimedia systems”, invited book chapter in *Multimedia Handbook*, CRC Press, Dec. 1997.
- Reddy, A. L. N. and Haskin R., “Video Servers”, invited book chapter in *Communications Handbook*, CRC Press, 1996.
- Reddy, A. L. N., “Improving the interactive responsiveness in a video server”, invited book chapter in *Multimedia Information Storage and Management*, Kluwer Academic Publishers, Dec. 1996.
- Reddy, A. L. N., “Scheduling in Multimedia Systems”, invited book chapter in *Design and Applications of Multimedia Systems*, Kluwer Academic Publishers, Aug. 1995.

Rigorously Refereed Conference Publications

- Qin, M.1, Reddy, A. L. Narasimha, Gratz, P., Pitchumani, R., Ki, Y-S., “KVRAID: High Performance, Write Efficient, Update Friendly Erasure Coding Scheme for KV-SSDs”, accepted at ACM Systor Conference., June 2021.
- Wen, F., Qin, M.1, Gratz. P., and Reddy, A. L. Narasimha, “OpenMem: Hardware/Software Cooperative Management for Mobile Memory System”, accepted at ACM/IEEE Design Automation Conference, Dec. 2021.
- Wen, F., Qin, M.1, Gratz. P., and Reddy, A. L. Narasimha, “FPGA-based Hybrid Memory Emulation System”, accepted at Field Programmable Logic and applications Conference, Aug. 2021.
- Qin, M.1, Lee. J-H., Pitchumani, R., Ki, Y-S., Reddy, A. L. Narasimha, Gratz, P., “A Generic FPGA Accelerator for Minimum Storage Regenerating Codes”, Proceedings of the 25th Asia and South Pacific Design Automation Conference, January, 2020.
- Chou, C-C.,1, Chen, Y., Milojicic, D., Reddy, A. L. Narasimha, Gratz, P., “Optimizing Post-Copy Live Migration with System-Level Checkpoint Using Fabric-Attached Memory, 2019 IEEE/ACM Workshop on Memory Centric High Performance Computing (MCHPC), Denver, CO, USA, 2019, pp. 16-24.
- Chou, Chih Chieh¹, Jung, J¹, Reddy, A.L. Narasimha, Gratz, Paul V., and Voight, Doug, “vNVML: An Efficient User Space Library for Virtualizing and Sharing Non-Volatile Memories”, IEEE 35th International Conference on Mass Storage Systems and Technology, 2019.
- Zheng, Z.¹, Webb, A.¹, Reddy, A. L. Narasimha, Bettati, Riccardo, “IoTAegis: A Scalable Framework to Secure the Internet of Things”, Invited Paper at IEEE ICCCN, July 2018.
- Zheng, Z.1, Reddy, A. L. Narasimha, “Safeguarding Building Automation Networks: THE-Driven Anomaly Detector Based on Traffic Analysis”, IEEE ICCCN (Invited Paper), July 2017.
- Fedorov, V¹., Kim, J., Qin, M¹., Reddy, A. L. Narasimha and Gratz, P., “Speculative Paging for future NVM and SSD”, Proc. of MEMSYS Conference, Oct. 2017.
- Zhang, Z¹, Jin, S., Bettati, R., Reddy, A. L. Narasimha, “Securing Cyber-Physical Systems with Adaptive Commensurate Response”, Proc. of IEEE CNS, Oct. 2017.
- Zhang, Z¹, and Reddy, A. L. Narasimha, “Safeguarding Building Automation Networks: THE-Driven Anomaly Detector Based on Traffic Analysis”, Invited Paper at ICCCN 2017, July 2017.
- Zhang, Z¹, and Reddy, A. L. Narasimha, “Towards Improving Data Validity of Cyber-Physical Systems through Path-Redundancy”, 3rd ACM Cyber-Physical Systems Security Workshop, April 2017.
- Kim, J., Pugsley, S., Gratz, P., Reddy, A. L. N., Wilkerson, C., and Chisti, Z., “Path Confidence based Lookahead Prefetching”, MICRO Conf., Oct. 2016. **Best Paper Nomination.**
- Webb, A.¹ and Reddy, A. L. Narasimha, “Finding proxy users at the service using anomaly detection”, IEEE Conf. on Communications and Network Security, Oct. 2016.

- Moon, S¹., and Reddy, A. L. Narasimha, “Adaptive policies for balancing performance and lifetime of mixed SSD arrays through workload sampling”, IEEE Mass Storage Symposium, April 2016
- Fedorov, V.¹, Reddy, A. L. Narasimha and Gratz, P., “Shared Last-Level Caches and The Case for Longer Timeslices”, MEMSYS Conference, Oct. 2015
- Kim, J., Fedorov, V.¹, Gratz, P., and Reddy, A. L. Narasimha, “Dynamic Memory Pressure Aware Ballooning”, MEMSYS Conference, Oct. 2015
- Fedorov, V.¹, Qiu, S.¹, Reddy, A. L. Narasimha and Gratz, P., “ARI: Adaptive LLC-memory traffic management”, HiPEAC Conf., Jan. 2014.
- Moon, S¹ and Reddy, A. L. Narasimha, “Don’t let RAID raid the lifetime of your SSD array”, Proc. of USENIX HotStorage Conf., June 2013.
- Qiu, S¹, and Reddy, A. L. Narasimha, “NVMFS: A hybrid file system for improving random write in NAND-flash SSD”, Proc. of Mass Storage Systems Symposium, May 2013.
- Lee, Y-O¹, and Reddy, A. L. Narasimha, “Multipath Routing for Reducing Network Energy”, IEEE Greencom Conf., Sep. 2012.
- Qui, S¹., Reddy, A. L. Narasimha, “Exploiting superpages in a nonvolatile memory file systems”, IEEE Mass Storage Symposium, Apr. 2012.
- Moon, S¹., Reddy, A. L. Narasimha, “Write amplification due to ECC on flash memory or leave those bit errors alone”, IEEE Mass Storage Symposium, Apr. 2012.
- Wu, X¹., Reddy, A. L. Narasimha, “SCMFS: A file system for storage class memory”, ACM/IEEE Supercomputing Conf. SC11, Nov. 2011.
- Yadav S.¹, Reddy, A. L. Narasimha, “Winning with DNS Failures: Strategies for faster botnet detection”, SecureComm Conf., Sep. 2011.
- Lee, Y-O¹, Rengarajan, B., and Reddy, A. L. Narasimha, “Reaching approximate Wardrop equilibrium at reduced cost of link state updates”, COMSNETS Conf., Jan 2011.
- Qian, B¹., and Reddy, A. L. Narasimha, “Measurement and performance study of PERT for on-demand video streaming”, Pfdnet workshop, Nov. 2010.
- Yadav, S¹, Reddy, A. K.¹, Reddy, A. L. Narasimha and Ranjan, S., “Detecting algorithmically generated malicious domain names”, Proc. of Internet Measurement Conference, Nov. 2010.
- Wu, X¹., Reddy, A. L. Narasimha, “Exploiting concurrency to improve latency and throughput in a hybrid storage system”, **Best Paper Nomination**, IEEE MASCOTS Conf., Aug. 2010.
- Devakota, P¹., Reddy, A. L. Narasimha, “Performance of Quantized Congestion Notification in TCP Incast Scenarios of Data centers”, IEEE MASCOTS Conf., Aug. 2010.
- Lee, Y-Oh¹., Reddy, A. L. Narasimha, “Disjoint Multi-Path Routing and Failure Recovery”, Proc. of ICC, May 2010.

- Yadav, S¹., Reddy, A. L. Narasimha, “DNS Anomaly detection with NS and AS information”, IASTED Conf. on Network Security, Nov. 2010.
- Wu, X¹., Reddy, A. L. Narasimha, “Managing storage space in a flash and disk hybrid storage system”, IEEE MASCOTS Conf., Oct. 2009.
- Podduturi, V., Kim, Y., Shakkottai, S., Reddy, A. L. Narasimha, “MultiTrack: A delay and cost aware P2P Overlay architecture”, ACM Sigmetrics, June 2009.
- Kotla, K¹., Reddy, A. L. N., “Making a delay-based protocol adaptive to heterogeneous environments”, Proc. of IwQOS, June 2008.
- Kang, S¹., Reddy, A. L. N., “User-centric data migration in networked storage systems”, Proc. of IPDPS, April 2008.
- Bhandarkar, S¹., Reddy, A. L. N., Zhang, Y., Loguinov, D., “Emulation of AQM at end hosts”, Proc. ACM Sigcomm, Aug. 2007.
- Bhandarkar, S¹., Reddy, A. L. N., “Robustness to packet reordering in high-speed networks”, Proc. of Pfdnet, Feb. 2007.
- Al-Najjar, C¹., Reddy, A. L. N., “A service provider’s approach to improving performance of aggregate voice-over-IP traffic”, Proc. of IWQOS, June 2006.
- Kim, S¹., Reddy, A. L. N., “Netviewer: A Network traffic visualization and analysis tool”, Proc. of USENIX LISA Conference, Dec. 2005, 185-196.
- Liu, Y¹, and Reddy A. L. N., “Routing among a group of multi-homed stub networks”, Proc. of Globecom, Dec. 2005.
- Bhandarkar, S¹., Reddy, A. L. N., “LTCP: A layered approach to TCP for high-speed networks”, PFLDNet, Feb. 2005.
- Kim, S¹., Reddy, A. L. N., “Modeling network traffic as images”, ICC, May 2005.
- Kim, S¹., Reddy, A. L. N., “A study of analyzing network traffic as images in real-time”, IEEE Infocom, March 2005.
- Kim, S¹., Reddy, A. L. N., “Real-time detection and containment of network attacks using QOS regulation”, Proc. of ICC, May 2005.
- Kang, S¹., Reddy, A. L. N., “Virtual Allocation: A scheme for flexible storage allocation” OASIS workshop at ASPLOS, Oct. 2004.
- Liu, Y¹, and Reddy A. L. N., “A fast rerouting scheme for OSPF/IS-IS networks”, Proc. of ICCCN, Oct. 2004.
- Achanta, P.¹, and Reddy A. L. N., “LRU-FQ: Design and Evaluation of a Partial State Router”, Proc. of ICC 2004, June 2004.
- Zhao, Z¹., Reddy, A. L. N., “Impact of bandwidth-delay product and non-responsive flows on queue management schemes”, Proc. of ICC 2004, June 2004.
- Bhandarkar, S¹., Reddy, A. L. N., “TCP-DCR: Making TCP robust to non-congestion events”, Proc. of Networking 2004, May 2004.

- Kim, S¹., Reddy, A. L. N., and Vannucci, M., “Detecting traffic anomalies at source through aggregate analysis of packet header data”, Proc. of Networking 2004, May 2004.
- Kim, S¹., Reddy, A. L. N., and Vannucci, M., “Detecting traffic anomalies using discrete wavelet transform”, Proc. of Int. Conf. on Information Networking, Feb. 2004.
- Ma, X¹., and Reddy, A. L. N., “SPIRAL: A novel third party transfer scheme for network attached devices”, IEEE Mass Storage Symposium, April 2003.
- Zhao, Z¹., Ametha. J., Swaroop. D. and Reddy A. L. N., “A method for estimating non-responsive traffic at a router”, ACM Conf. on Sigmetrics, June 2002.
- Garg, A¹., and Reddy, A. L. N., “Mitigating Denial of service attacks through QOS regulation”, IwQoS workshop, May 2002.
- Wijayarathne, R¹. and Reddy, A. L. N., “System support for providing integrated services from networked multimedia storage servers”, ACM Conf. on Multimedia, October 2001, 270-279.
- Smitha¹ and Reddy, A. L. N., “LRU-RED: An active queue management scheme to contain high bandwidth flows at a congested router”, in IEEE Globecom Conference, Nov. 2001.
- Smitha¹, Kim, I¹., and Reddy, A. L. N., “Identifying long term high rate flows at a router”, in Proc. of High Performance Computing, Dec. 2001.
- Yeom, I¹., and Reddy, A. L. N., “Adaptive marking for aggregated flows”, in Proc. of IEEE Globecom, Nov. 2001.
- Gopalakrishnan, S¹. and Reddy, A. L. N., “SACRIO: An active buffer management scheme for Differentiated Service Networks”, Proc. of NOSSDAV, June 2001.
- Ma, X¹. and Reddy, A. L. N., “MVSS: MultiView Storage Systems”, Proc. of Int. Conf. on Distributed Computer Systems (ICDCS), Apr. 2001.
- Khaleel, A¹. and Reddy, A. L. N., “Evaluation of Request and data distribution policies in a clustered server”, Proc. of High Performance Computing, Dec. 1999.
- Yeom, I¹. and Reddy, A. L. N., “Realizing throughput guarantees in a differentiated services network”, Proc. of IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Wijayarathne, R¹., Reddy, A. L. N., “Integrated QOS management for disk I/O”, Proc. of IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Yeom, I¹. and Reddy, A. L. N., “Impact of marking strategy on aggregated flows in a differentiated-services network”, Proc. of IEEE International Workshop on QOS (IwQOS 99), June 1999.
- Tong, D¹. and Reddy, A. L. N., “QOS Enhancement with Partial Sate”, Proc. of IEEE International Workshop on QOS (IwQOS 99), June 1999.
- Gupta, S¹. and Reddy, A. L. N., “IPRP: A Client-oriented IP Redirection mechanism”, Proc. of IEEE Infocom Conf., March 1999.
- Wijayarathne, R¹. and Reddy, A. L. N., “Techniques for improving the throughput of VBR streams”, Proc. of ACM/SPIE Multimedia Computing and Networking, Jan. 1999.

- Ma, G¹. and Reddy, A. L. N., “An evaluation of storage systems based on network-attached disks”, Proc. of ICPP Conf., Aug. 1998.
- Reddy, A. L. N., “Effectiveness of caching policies for a web server”, Proc. of High Performance Computing, Dec. 1997.
- Reddy, A.L.N., “Caching strategies for a multimedia server”, Proc. of IEEE Int. conf. on Multimedia Computing and Systems, June 1997.
- Reddy, A. L. N., “Improving Latency in an Interactive Video server”, Proc. of SPIE conf. on Multimedia Computing and Networking, Feb. 1997.
- Reddy, A.L.N., “Movie and communication scheduling on a multiprocessor video server system”, Proc. of *IEEE Int. Conf. on Multimedia Comput. and Systems*, May. 1995.
- Reddy, A.L.N., Wyllie, J., “Disk Scheduling in a Multimedia I/O System”, Proc. of *1st ACM Conf. on Multimedia*, Aug. 1993. Included as a chapter in “A guided tour of multimedia systems and applications”, IEEE Press, 1995. Translated into Japanese by *Nikkei Electronics*, a Japanese technical magazine. To be included as “one of the most influential papers” in ACM Transactions on Multimedia, Applications and Communications.
- Chandy, J., and Reddy, A.L.N., “Failure evaluation of disk array organizations”, Proc. of *Int. Conf. on Distributed Computing Systems*, May 1993.
- Reddy, A.L.N. and Freitas, R., “Fault-tolerance of adaptive routing algorithms in multicomputers”, Proc. of *4th IEEE Symp. on Parallel and Distributed Processing*, Dec. 1992.
- Reddy, A.L.N., “A Study of IO System Organizations”, Proc. of *ACM Int. Symp. on Computer Architecture*, May 1992.
- Reddy, A.L.N., “Reads and Writes: When IOs aren’t the same”, Proc. of Hawaii Int. Symp. on Systems and Sciences, Jan. 1992.
- Reddy, A.L.N., Banerjee, P. and Chen, D.K., “Compiler Support for Parallel I/O Operations”, Proc. of *Int. Conf. on Parallel Processing*, Aug. 1991.
- Reddy, A. L. N. and Banerjee, P., “Gracefully Degradable Disk Arrays”, Proc. of *Int. Symp. on Fault Tolerant Computer Systems*, June 1991.
- Reddy, A. L. N. and Banerjee, P., “A Study of IO Behavior of Scientific Applications on A Multiprocessor”, Proc. of *ACM Int. Symp. on Computer Architecture*, June 1990
- Reddy, A. L. N. and Banerjee, P., “Performance evaluation of multiple-disk IO systems”, Proc. of *Int. Conf. on Parallel Processing*, Aug. 1989.
- Reddy, A. L. N. and Banerjee, P., “IO issues in hypercubes”, Proc. of *ACM Int. Conf. on Supercomputing*, June 1989.
- Reddy, A. L. N., Banerjee, P. and Abrham, S.G., “IO embedding in Hypercubes”, Proc. of *Int. Conf. on Parallel Processing*, August 1988.
- Reddy, A. L. N. and Banerjee, P., “A Fault-Secure Dictionary Machine”, Proc. of *3rd International Conference on Data Engineering*, Feb. 1987.

Other Publications

- Kim, J., Gratz, P., and Reddy, A. L. Narasimha, “ Lookahead Prefetching with Signature Path”, ACM ISCA Workshop on Data Prefetch Competition, June 2015
- Huang, B¹., Wang. S¹., and Reddy, A. L. Narasima, “Who is Hiring Whom: A New Method in Measuring Graduate Programs”, ASEE conference, June 2015.
- Singh, A.¹, and Reddy, A. L. Narasimha, “Multi path PERT”, **Invited Paper** at Int. Conf. on Computer Communications and Networks, July 2013.
- Amaleswaram, A.¹, Reddy, A. L. N., Yadav, S.¹, Gu. G., Yang, C., “CATS: Characterizing Automation of Twitter Spammers”, **Invited Paper** at Int. COMSNETS conf., Jan. 2013.
- Wu, X.¹, and Reddy, A. L. N., “Data organization in a hybrid storage system”, **Invited Paper** at Int. Conf. on Computing, Networking and Communications, ICNC’12, Feb. 2012.
- Yin, Z.¹, Hasari, C., Al-Nuweiri, H., and Reddy, A. L.N., “Improving the Performance of Delay Based Protocol in Delivering Real Time Media via Early Retransmission”, IEEE ICT Conf., Apr. 2011.
- Yin, Z.¹, Al-Nuweiri, H., and Reddy, A. L.N., “Performance of Early retransmission scheme in streaming media”, IEEE ICT Conf., Apr. 2010.
- Kim, S¹., and Reddy, A. L. N., “An evaluation of the effectiveness of measurement-based anomaly detection techniques”, ADSN workshop, held in conjunction with ICDCS, June 2006.
- Vellanki, S¹., and Reddy, A. L. N., “Improving service availability during link failure transients through alternate routing”, Tech. Report, Texas A & M University, Feb. 2003.
- Ma, X¹., and Reddy, A. L. N., “Implementation and Evaluation of an active storage system prototype”, Workshop on novel uses of System Area Networks, Feb. 2002.
- Reddy, A.L.N., “ An evaluation of multiprocessor video server system”, *IBM Confidential report*, June 1994.
- Menon, J., Reddy, A.L.N., Treiber, K. and Tzou, S.Y., “Daisy: A Novel Heirachical Storage Management System, Functional Specification”, *IBM Confidential Report*, 77 pages, Jan 1994.

Recent Presentations:

- **Tutorial** on Network Security at Int. Conference on Information Security, Hyderabad, India, Dec. 2014.
- **Invited talk** at IEEE Communications Socieity, Hyderabad, India, August 2014.
- **Invited talk** at University of Texas at San Antonio, Nov. 2013.
- Presentation at Flash Summit, Aug. 2013
- **Invited talk** at Georgia Tech., May 2013.

- **Keynote talk** at Int. Conf. on Information Technology: New Generations, Las Vegas, April 2013.
- **Invited talk** at COMSNETS Conf., Bangalore, India, Jan. 2013.
- **Invited talk** at Int. Conf. on Networking and Communications, Maui, Feb. 2012.
- **Invited talk** at NVRAMOS 2011, Jeju island, South Korea, Nov. 2011.
- **Invited talks** at Korea University, Ajou University, Samsung and ADD, Korea, Nov. 2011.
- **Keynote talk** at IEEE Int. Symp. on Computers and Communications, Italy, June 2010.
- **Catedra Series talk** at University of Carlos III, Dec. 2009, Madrid, Spain.
- Invited talk at Thomson Research, Paris, France, October 2009.
- Invited talk at Telefonica Research, Barcelona, Spain, September 2009.
- IBM Austin Research Lab **Distinguished Seminar series**, May 1 2008.
- Invited talk at University of Illinois, November 2006.
- Invited talk at Cisco, January 2005.
- Invited talk at University of Birmigham, UK, June 2004.
- Invited talk at Cisco Systems, San Jose, June 2002.
- Invited talk at Juniper Networks, Mountain View, June 2002.
- Invited talk at Univ. of California at Riverside, June 2002.
- Invited talk at Nortel Networks, Santa Clara, June 2000.
- Invited talk at Juniper Networks, Mountain View, June 2000.
- Talk at Internet Engineering Task Force, Washington D.C., Oct. 1999.
- Invited talk at EMC Corp., Boston, Aug. 1999.
- Tutorial on “Scheduling for Multimedia Applications”, presented at IEEE Int. Conf. on Multimedia Computing and Systems, June 1999.
- Tutorial on High-Performance I/O at IEEE/ACM Int. Conf on Computer Architecture, Atlanta, May 1999.
- Invited talk at AT & T Research Center, Florham Park, March 1999.
- Invited talk at Bell Labs Research Center, Murray Hill, March 1999.
- Invited talk at EMC Corp., Boston, Nov. 1997.

Funding:

- “Cooprative Research Agreement”, Army Futures Command, Co-PI, 2019-24.

- “Multi-layer Cybersecurity and Situational Awareness to Enhance Resiliency in Qatars Power Grid, Qatar National Research Foundation, 2020-2025, co-PI.
- “Protecting Building Automation Networks from insider threats”, 2016, QNRF, lead PI.
- EMC, DELL, HPE and Samsung membership participation in NSF I/UCRC Center.
- “NSF I/UCRC Center on Intelligent Storage”, 2014, NSF, lead PI.
- “Equipment Donation”, Viking Technologies, 2013.
- “Emerging memory architectures for big memory applications”, National Science Foundation, 2013-2016, co-PI.
- “Exploiting social communication channels against cyber criminals”, National Science Foundation, 2012-2015.
- “Client interfaces to object storage systems”, Argonne National Lab, 2012-2015.
- “Detecting malicious activities by analyzing social communication contexts”, Qatar National Research Foundation, 2012-2015.
- “Network traffic analysis for botnet detection”, Narus Inc., Research gift, 2009-2010.
- “Design and Analysis of new generation of protocols for triple-play networks”, QNRF, Jan. 1. 2008 - Dec. 31 2010, Co-PI.
- “Dynamic allocation and data distribution in networked storage systems”, National Science Foundation, Oct. 2007 - Sept. 2010, sole PI.
- “Active Data Systems”, National Science Foundation, Sept. 2006 -Aug. 2009, sole PI.
- NSF REU grant, co-PI, 2004-2007.
- Network Elements based on partial state, Intel Corp., gift, Feb. 2003.
- Texas A & M University, “High Availability Infrastructure Laboratory”, TITF Award, Sept. 2002 - Dec. 2003, PI with R. Bettati, M. Vannucci and D. Swaroop.
- “Secure Networked Storage Systems”, National Science Foundation, Research Instrumentation grant, Sept. 2002- Sept. 2005.
- Texas Advanced Technology Program, “Network Architectures based on partial state”, Jan. 2001 - Dec. 2002.
- “Fellowships for Computer Engineering”, Department of Education, Aug. 2001- Aug. 2004, PI with R. Mercer and K. Watson.
- “QOS Enhancement with Partial state”, National Science Foundation, June 2001-June 2004, sole PI.
- “Multiview Storage Systems”, National Science Foundation, May 2001 - May 2004, sole PI.
- Equipment Donation, 3Ware Inc., Nov. 2000.

- “Scalable request-specific file service”, National Science Foundation, Aug. 1999 - Aug. 2002, sole PI.
- Research gift, EMC Corp., 1998.
- “Scalable Multimedia Storage Server”, Texas Higher Education Board, January 1996 - Dec. 1997, sole PI.
- “Scalable Storage Server for Multimedia and Scientific Applications”, National Science Foundation, 1996-2000, sole PI.
- “Development of Novel Systems Software for Multimedia and High-performance computing”, National Science Foundation, 1997-2000, PI along with R. Mercer, P. Cantrell, G. Choi and M. Lu.
- Equipment grant, Tandem Corporation, March 1997.
- Software donation, Microsoft Corporation, April 1997.

Teaching:

Redesigned the course ECEN 449 “Microprocessor Systems Design”, Fall 2007.

Developed a new course ELEN 627 “Multimedia Systems and Networks”.

Developed a freshman undergraduate course ENGR111 “Introduction to Electrical and Computer Engineering”.

Developing a new graduate course on “Network systems security”. Taught in Spring 2004, 2006.

Courses Taught: Undergraduate Introduction to Electrical and Computer Engineering, Undergraduate Computer Architecture, Graduate Network Systems security, Graduate Computer Architecture, Graduate Computer Networks, Graduate Multimedia Storage and Delivery, Graduate Parallel Computer Systems, Graduate I/O and Memory Systems.

Student Advising:

- Chris Baldwin, Ph.D. candidate, Aug 2021-, Storage systems.
- Susav Shresta, Ph.D. candidate, Aug 2021-, Memory systems.
- Chandrahas Tirumalasetty, Ph.D. candidate, Aug 2019-present, Storage Systems.
- Rishabh Singla, Ph.D. candidate, Aug 2017-Present, Network Security.
- Mian Qin, Ph.D. candidate, Aug 2016-Aug 2021, Storage Systems.
- Chih-Chieh Chou, Ph.D., Aug 2014-Jan. 2020, “Optimizing Emerging Memory Systems for Performance”, next at Amazon.
- Zhiyuan Zheng, Ph.D., Aug. 2013 -Dec. 2017, “Protecting Cyber-Physical Systems With Special Emphasis On Building Automation Networks”, next at Pinterest.
- Jaemin Jung, Postdoc, Oct. 2016-18, Storage systems, next at Samsung Research.

- Allen Webb, Ph.D. candidate, Jan. 2013-17, “Addressing Insider Threats From Smart Devices”, next at Google.
- Viacheslav Fedorov, Ph.D., 2010 - 2016, “Memory Management for emerging memory technologies”, later at NXP.
- Jongmin Gim, Postdoc, 2013-2015, Storage Systems Software, Later at Samsung Research Labs, California.
- Sangwhan Moon, Ph.D., 2010 - 2015, “Reliability of SSD Storage Systems”, later at Riverbed Technology.
- Sheng Qui, Ph.D. 2013 “Integration of nonvolatile memory into storage hierarchy”, later at Samsung Research Labs.
- Yong Oh Lee, Ph.D., May 2012, “Improving efficiency and effectiveness of multipath routing in computer networks”, later at Samsung.
- Jagadeesh Dyaberi, Ph.D., Aug. 2011, “Networking and Storage support for on-demand video delivery”, Purdue University, External thesis committee member.
- Sandeep Yadav, Ph.D., 2012, “Scalable techniques for anomaly detection”, Dec. 2012, later at Dumballa.
- Xiaojin Wu, Ph.D., Flash and Disk Hybrid storage systems, Aug. 2011, later at Marvell Semiconductor.
- John Garrison, Ph.D. , “Umbrella File System: Managing heterogeneous storage devices File systems”, Aug. 2009, later at DOD Labs, Tuscaloosa, Alabama.
- Sukwoo Kang, Ph.D., “Flexible space allocation and management in networked storage systems”, May 2007, later at IBM Almaden Research Center.
- Sumitha Bhandarkar, Ph.D., “Congestion control algorithms for TCP in emerging networks”, Aug. 2006, later at Motorola.
- Yong Liu, Ph.D., “Dynamic Routing protocols in multihomed networks”, Dec. 2006, later at EMC Corp.
- Seong Soo Kim, Ph.D., “Real-time analysis of aggregate network traffic”, May 2005, later at Samsung, Seoul, Korea.
- Zhili Zhao, Ph.D., May 2004, “Estimation and control of nonresponsive traffic”, later at Intel.
- Xiaonan Ma, Ph.D., Aug. 2003, “Augmenting functionality at storage systems”, later at IBM Almaden Research Center, San Jose.
- Ravi Wijayarathne, Ph.D., Dec. 2000, “PRISM: A file server architecture for providing integrated services”, later at Quantum, San Jose.
- Ikjun Yeom, Ph.D., Dec. 2000, “Bandwidth assurance in differentiated services networks”, later a faculty member at KAIST, South Korea.
- Daksh Vasistha, M.S., 2015-17, “ Detecting Anomalies in Controller Area Network for Automobiles”, next at Facebook

- Ayush Ruia, M.S., 2015, “Virtualization of nonvolatile memory”, later at Microsoft.
- Bolun Huang, M.S. 2014, “Sentiment based classification of tweeters and university programs”, later at Microsoft.
- Shyam Sambasivan, M.S. “Enhancements to SQLite library to improve performance on mobile platforms”, Aug. 2013, later at Cisco.
- Ankit Singh, M.S., “Multipath Probabilistic Early Response TCP”, Aug. 2012, later at Amazon.
- Cindy Cao, M.S., “SCMFS Performance Enhancement and Implementation on Mobile Platform”, Aug. 2012, later at Microsoft.
- Amit Anand, M.S., “Proactive spammer detection on Twitter”, Aug. 2012, later at Amazon.
- Bin Qian, M.S.. Aug. 2011, “Evaluation of Probabilistic Early Response TCP (PERT) for video streaming and extension with ACK coalescing”, later at Amazon Data Services.
- Ashwath Reddy, M.S., Aug. 2010, Detecting networks employing algorithmically generated names, later at Microsoft.
- Prajjwal Devakota, M.S. Aug. 2010, Performance of Quantized Congestion Notification in TCP Incast in Data Centers, later at Bloomberg.
- Chaitanya Tadi, M.S., May 2009, IP Prefix Hijack Prevention Service, later at Intel.
- Marcus Dunn, M.S., May 2009, A new I/O scheduler for Solid State Devices, later at Cisco.
- Kiran Kotla, M.S., “Adapting the delay-based protocol PERT to heterogeneous environments”, Aug. 2008, later at Cisco.
- Camelia Al-Najjar, M.S., August 2005, thesis on “An approach for improving performance of aggregate voice over IP traffic”, pursuing Ph.D. at Case Western University.
- Saurabh Jain, M.S., May 2005, thesis on “LTCP-RC: RTT compensation technique to scale highspeed protocols to high RTT links”, later at Cisco.
- Sridhar Vellanki, M.S., Feb. 2003, thesis on “Improving service availability through alternate routing”, later at Microsoft.
- Sadry, Nauzad, M.S., Dec. 2002, thesis on “Evaluation of Delayed Congestion Response protocols for tolerating channel errors in wireless networks”, later at Netezza.
- Achanta, Phani, M.S., Dec. 2002, thesis on “Design and Evaluation of Partial State Routers”, later at IBM.
- Smitha, M.S. student, May 2001, thesis on “An active queue management scheme to contain high bandwidth flows at a congested router”, later at Intel India.
- Sumitha Bhandarkar, M.S., August 2001, thesis on “Delayed Congestion Response Protocols”, Ph.D. at TAMU.
- Aman Garg, M.S., August 2001, thesis on “Mitigating Denial of Service using QOS regulation”, later at Tipping Point Technologies, Austin.

- Sai Gopalakrishnan, M.S., May 2000, later at Cisco, thesis on “TRIO: Active resource management in differentiated services networks”.
- Qian Ren, M.S., May 2000, later at Lucent, thesis on “MPEG-aware disk storage system”.
- Marcus Grande, M.S., Dec. 1999, later at IBM, thesis on “Multiported Storage Devices”.
- Adnan Khaleel, M.S., Aug. 1999, later at Motorola, thesis on “Request and data distribution on a clustered server”.
- Deying Tong, M.S., May 1999, later at Cisco, thesis on “Partial state based approaches to QOS enhancement”.
- Sumit Gupta, M.S., Aug. 1998, later at Sun Labs, thesis on “IPRP: A client-oriented IP redirection mechanism and its applications”.
- Ashwin Madhwaraj, M.S., May 1998, later at Cisco, thesis on “Adaptive video transport over networks”.
- Gang Ma, M.S., May 1998, later at EMC Corp., thesis on “Evaluation of storage systems based on network attached storage”.
- Mike Nichols, M.S., May 1998, later at UT Law School, thesis on “Personal informal retrieval agent”.
- Ikjun Yeom, M.S., May 1998, later a faculty member at KAIST, Korea, thesis on “ENDE: An End-to-end Network Delay Emulator”.

Service:

- Member, University IT Research Committee, 2018.
- Chair, Turbo Lab Director Search Committee, 2017.
- Chair, College Cybersecurity Search Committee, 2016-Present.
- Member, University Chief Communications Officer Hiring Committee, 2015-16.
- Member, University level Assoc. Former Students Awards Committee, 2016.
- Member, TAMUQ FASC Committee, 2013.
- Chair, Department Growth Committee, 2013.
- Chair, College Taskforce on National Security, 2013.
- Member, College Research Strategy Subcommittee, 2012.
- Member, Computer Science and Engineering Head Search Committee, 2010-11, 2013-14.
- Member College Awards Committee 2006-2008. 2010-2012.
- Actively participating in improving undergraduate research experience. Co-PI of NSF REU grant. Advised 7 undergraduate students in the last few years.

- Member, Computer Engineering Curriculum Committee, 2003-2012. Led the restructuring of Undergraduate curriculum in Computer Engineering.
- Member, Department Tenure & Promotion Committee, 2004-2006, 2010-2012.
- Member, Electrical and Computer Engineering Department Head Search Committee, 2004.
- Member, Department Advisory committee 2003-2004.
- Chair and Member, department Awards committee 2004-2006.
- Member, Department Awards committee, 1997-2003, 2008-2009.
- Member, Department Distinguished Speakers committee, 1996-1999, 2006-2008.
- Chair and Member, Department Distinguished Speakers committee, 2010-present.
- Member, Department undergraduate studies committee, 1997-2000.
- Member, Department web site design committee, 2000-2004.
- Group leader, Computer Engineering, 2005-2014.

Administrative Experience:

Group Leader, Computer Engineering, 2005-2014:

- Led evaluation, development and rearchitecting the Computer Engineering program, a joint program across two departments of Electrical and Computer Engineering and Computer Science Engineering.
- Led the hiring of about 10 Computer Engineering faculty.
- Led the development of ENGR 111 Freshman Introduction of Electrical and Computer Engineering, primarily to improve retention in Electrical and Computer Engineering programs. Increased retention rates from 45% to about 80%.
- Led the development of several other courses ECEN 449 Microprocessor Systems Design, ECEN 627 Multimedia Systems and Networks, ECEN689 Computer Networks Security.
- Rebuilt several existing courses including ECEN 350 Computer Architecture and ECEN 676 Advanced Computer Architecture, ECEN 602 Computer Networks.
- Contributed to the national ranking of both undergraduate and graduate Computer Engineering programs to top 20 programs in the country.

TEES Assistant Agency Director for National and Global Initiatives, 2014- April 2016:

- Developed and Established a TEES office in Yukatan, Mexico.
- Nurtured relationships in Brazil, France and India. The France relationship resulted in a collaborative center with ENSAM.
- Established realtionships with Kuwait Oil Company which resulted in a \$15 million research contract.

- Developed a strategy to engage with National Labs and Federal Agencies.

**Associate Dean for Research, College of Engineering and TEES Associate Agency
Director for Centers and Strategic Initiatives, April 2016-Present:**

- Established strategic focus on several emerging areas of research including Autonomy, Synthetic Biology, Soft Matter, Cyber Security and Quantum Computing.
- Developed a strategy for supporting large multi-PI proposals and increased large proposal funding from 5-6 per year to about 25 per year.
- Increased TEES Research funding from about \$110 million/year to \$210 million/year from 2016-20.
- Increased Federal funding percentage from 38% to 68% of total TEES funding.
- Nurtured relationships with Army Research Lab (ARL) and facilitated the establishment of ARL-south relationship with A & M. Direct impact: ARL Cooperative Agreements in Synthetic Biology, Electric Power, Materials and Manufacturing, Autonomy and Cybersecurity with about 15 ARL people stationed at Texas A& M.
- Supported establishment of \$65 million ARL/Army Futures Command (AFC) Cooperative agreement on Hypersonics, Autonomous vehicles and Networks.
- Contributed to \$26 million Distributed Autonomy related project funded by Army Research Lab.
- Nurtured and furthered collaborations with Air Force Research Lab, resulting in the establishment of AFRL fellowships at A & M. Direct Impact: AFRL \$4 million Minority Leaders Scholarships in Material Discovery and several projects with sensors, human performance and information directorates.
- Established relationships with Army Engineering and Research Development Center (ERDC). Resulted in \$3 million contract. Exploring further options for joint projects.
- Establishing close collaborative ties with DOE/DOD National Labs. Resulted in Master Research Agreements and joint projects with several labs including Idaho National Lab(INL), Lawrence Livermore National Lab(LLNL), Pacific Northwest National Lab (PNNL) and others.
- Contributed to successful \$40 million proposal for "Offshore Energy Safety Center 2.0" proposal to Department of Interior and Department of Energy.
- Started and organized "TEES Tech. Talk" series to facilitate hosting Labs and Agency researchers and program managers.
- Organized National Labs Day to facilitate interactions of graduate students with National Labs.
- Championed "Remote Ph.D. program for national labs".
- Established "Research Impact Awards" to facilitate a discussion and recognition of impactful research within the College.

- Established External Awards Support office to promote faculty for external professional awards.
- Supported young faculty applications to NSF, DOD/DOE/DARPA CAREER and Young Investigator awards. Increased NSF CAREER awards from 2/year to 8/year.
- Established NSA Masters Program.
- Expanded a weekly "Research Opportunities" bulletin to cover broader research opportunities.
- Interacted with the Research Deans of other colleges and supported inter-college collaborations.
- Established procedures and metrics for evaluating centers within college and reorganized centers after a thorough review of the existing centers. Managed about 30 Centers and Institutes.
- Established support to help graduate student applications for NSF and NDSEG Fellowships.

¹Reddy's Students.