

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

CISCO SYSTEMS, INC.
Petitioner,

v.

UMBRA TECHNOLOGIES, LTD.
Patent Owner.

U.S. Patent No. 11,108,595
Issue Date: August 31, 2021
Title: SYSTEMS AND METHODS FOR PROVIDING A GLOBAL VIRTUAL
NETWORK (GVN)

Inter Partes Review No.: IPR2024-00270

**DECLARATION OF DR. A.L. NARASIMHA REDDY IN SUPPORT OF
PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO.
11,108,595**

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	3
II. MY BACKGROUND AND QUALIFICATIONS	5
III. RELEVANT LEGAL STANDARDS.....	8
IV. BACKGROUND	12
A. Overview	12
B. Summary of the '595 Patent.....	16
C. Claim Construction.....	21
D. Priority Date of Claims	22
E. Person of Ordinary Skill in the Art (POSITA).....	22
V. STATE OF THE ART	23
1. Communication Tunnels.....	24
2. Selection of Communication Paths By Controller.....	25
VI. CHALLENGES.....	27
A. Identification of Challenge.....	27
B. Challenge 1: Kommula alone, Teaches Every Element of Claims 1-14	27
1. The Kommula Reference	27
2. Detailed Application of Kommula to the Challenged Claims	30
C. Challenge 2: Hankins alone or in view of Kommula Teaches Every Element of Claims 15-20	45
1. The Hankins Reference.....	45

2.	Detailed Application of Hankins, alone or in view of Kommula to the Challenged Claims	49
	CONCLUSION AND DECLARATION.....	1
	APPENDIX A – MATERIAL RELIED UPON	3
	APPENDIX B – CURRICULUM VITAE	4

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. 11,108,595 (“the 595 Patent”) (EX-1001), entitled “SYSTEMS AND METHODS FOR PROVIDING A GLOBAL VIRTUAL NETWORK (GVN)” and lists Jorn Allan Dose Knutsen, Carlos Eduardo Ore, Thibaud Augste Bernard Jean Saint-Martin, Fred Broussard (US), Joseph E. Rubenstein, 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 ’595 Patent challenged by the Petitioner in this *inter partes* review is no earlier than April 7, 2015, the filing date of U.S. Provisional Patent Application No. 62/144,293.

5. I understand that according to USPTO records, the '595 Patent is currently assigned to UMBRA Technologies Ltd. ("UMBRA" or "Patent Owner").

6. The '595 Patent is directed to managing a network connection between an endpoint device and an access point server. I am familiar with the technology described in the '595 Patent as of the earliest possible priority date of April 7, 2015.

7. In preparing this Declaration, I have reviewed the '595 Patent (EX-1001), the file history of the '595 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 April 7, 2015. 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 April 7, 2015, who I describe below (§IV.D). For convenience, the materials I considered in arriving to my opinion are listed in Appendix A.

8. 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.

9. I have been asked to provide my technical expertise, analysis, insights and opinions regarding the '595 Patent and relevant references that form the basis of the arguments set forth in the accompanying Petition for *inter partes* review of the '595 Patent. As described in detail below, I offer the following opinion in this Declaration:

- A POSITA would have found Claims 1-14 of the '595 Patent are unpatentable over *Kommula*.
- A POSITA would have found Claims 15-20 of the '595 Patent are unpatentable over *Hankins*, alone, or in view of *Kommula*.

II. MY BACKGROUND AND QUALIFICATIONS

10. 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.

11. 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.

12. 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.

13. 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.

14. 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.

15. 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.

16. 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.

17. 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).

18. 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.

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

III. RELEVANT LEGAL STANDARDS

20. 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 '632 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.

21. 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.

22. 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.

23. 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.

24. 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.

25. 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.

26. 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. BACKGROUND


A. Overview

27. The Challenged Claims are unpatentable as obvious over the prior art. The claims are directed to prior art systems and combinations of conventional components that perform conventional functions regarding managing virtual network connections, all of which was well-known in the art of virtual networks.

28. In one aspect, some of the claims of the '595 Patent are directed to a high level method performed by a computer processor including the steps of (1) “receiving” a request for a list of available servers from a network device; (2) “retrieving” the list of available servers; (3) “retrieving” one or more records associated with the network device; (4) “dynamically ordering” the list of available servers based on the one or more records; and (5) “transmitting” the ordered list of available servers to the network device, all as recited in Claim 1. *EX-1001* 32:49-61. Independent Claim 8 is directed to a generic system with memory and a processor configured to perform the same steps recited in Claim 1. *EX-1001* 33:14-31.

29. These high level steps are described by the prior art reference

Kommula as being performed by the architecture shown in FIG. 1:




EX-1004 FIG. 1.

30. Kommula describes that the one or more processors performs each of the recited steps including providing the same example of dynamically ordering the list of available servers as described in the '595 Patent. Specifically, Kommula describes that the list of available servers can be dynamically ordered based on the proximity of the server to the network device, just like described in the '595 Patent.

31. In another aspect, some of the claims of the '595 Patent are directed to a system in which a controller dynamically orders a list of available access point servers, which can be used to form tunnels between first and second endpoint

performing the functions recited in Claim 15, the only difference being Hankins uses different terminology.



EX-1005 FIG. 2 (annotated).

33. FIG. 2 of Hankins illustrates the network system recited in Claim 15 including a first endpoint device (green) communicating with a second endpoint device (blue) via a first tunnel (purple) connecting an access point server (red) to one or more endpoint device, and a controller (orange) connected to an access point server via a second tunnel (brown) for dynamically ordering a list of available servers for the first endpoint device.

34. Thus, the Challenged Claims of the '595 Patent are directed to systems and methods that were fully disclosed in the Kommula and Hankins references.

B. Summary of the '595 Patent

35. The '595 Patent is directed towards managing a network connection between an endpoint device and an access point server. The '595 Patent discloses as a problem that in the prior art there is insufficient control over the routes between an endpoint device and access point server as it relies on the policies of various middle players on the internet. *EX-1001* 2:18-36.

36. The '595 Patent describes that it is not uncommon to connect from within one local area network (LAN) to the resources of another remote LAN using a direct connection through a wide area network (WAN), such as the internet. However, the '595 Patent describes a problem that arises when a gateway to the Internet is located at a LAN, for example, located at the corporate headquarters, which handles traffic from a remote LAN of a subsidiary located in a different country. Traffic from the remote LAN is routed through the gateway at the headquarter LAN to the Internet. A slowdown occurs as traffic flows through the Internet back to servers in the same country as the subsidiary. Traffic must then go from the remote LAN through the WAN to the headquarters LAN and then through the Internet back to the server used by the subsidiary in the different

country, then back through the Internet to the gateway of the headquarters, and then back down the dedicated line to the client device within the remote LAN. This type of connection in essence doubles or triples (or worse) the global transit time of what should take a small fraction of global latency to access a server close to the remote LAN. Thus, the prior art connection mechanisms may not take into account the context of the connection needed for a particular source endpoint device and available access point servers. *EX-1001* 1:59-2:17.

37. The '595 Patent then describes a number of specific cloud based solutions for establishing dedicated connections between endpoint devices and an access point server. *EX-1001* 2:40-57.

38. But, the specificity of these solutions are found nowhere in the claims, which are much broader and directed to well-known Internet connections and/or well-known network tunneling technology.

39. As examples, in one embodiment, the network system may include an endpoint device, an access point server, and a control server. The endpoint device and the access point server may be connected with a first tunnel, and the access point server and the control server may be connected with a second tunnel. *EX-1001* Abstract; 3:4-11.

40. The '595 Patent describes that in one embodiment, in order to build a tunnel, certain information is needed. *EX-1001* 10:65-11:3.

41. For example, a server availability module examines current allocation of IP address and port combinations while taking account actual use, load, and other factors of the live production environment. It further analyzes the potential impact of assigned but standby allocations to predict the impact of those resource demands going live. When a device makes a request to the server availability module, what it is really asking for is a list of available servers, and specifically an IP and port which it could utilize to in the present or in the future to build a tunnel to the available server. Each tunnel has specific information which needs to be known to each peer pair in the tunnel building process as defined by tables under tunnel information. *EX-1001* 23:11-24.

42. The list of server availability is a list which can be utilized by all devices but for which a contextual list for one certain device is easily retrieved with contextual records relevant to it. *EX-1001* 23:33-36.

43. The server availability module considers many factors balancing demand for resources by clients such as endpoint devices wanting to build tunnels against available resources of server access points which they would connect with. *EX-1001* 23:40-43.

44. The '595 Patent describes the use of a governing mechanism that ranks the available sever access points in order contextually for that specific endpoint device. The ranked list takes into account a holistic view and predictively

spreading load in anticipation of demand pressures on the production environment.

EX-1001 23:55-59.

45. This governing mechanism ensures that automatic jumps of tunnels from many endpoint devices to one specific device do not occur. For example, if all endpoint devices share the same list of server availability, without an influence on how an endpoint device picks which servers to connect with, there are scenarios where too many endpoint devices connect with some access point servers and too few to others. And when the endpoint devices connected to the oversaturated access point server realize that they are receiving less than optimal connectivity, if they share the same list, by jumping to the next server on the list, they will be causing a wobbly shifting of the problem down the line. *EX-1001 23:44-54.*

46. In one embodiment, server availability prioritizes the list of server access port IP address and port combinations based on expected best performance of the tunnels to be built while also looking at current load of available access point servers, balancing assigned lists given to other end point devices as well as other available information. *EX-1001 11:21-26.*

47. The server availability list is generated for each request received from an endpoint device and the list is shared with the endpoint devices and available access point servers *EX-1001 12:50-59.*

48. The '595 Patent describes that the server availability module's main benefit is to catalog the resources usage, IP addresses and ports of access point servers and to look at which devices, such as end point devices that would like to connect to server access points to build tunnels with them. The server availability module calculates and compiles a custom-made list of IP addresses and ports which are contextually designed to provide best options for endpoint devices to utilize. The contextual information includes connectivity tests based on current and historical conditions to weigh and rank best connectivity options for tunnels to be built. *EX-1001* 27:20-31.

49. Although the '595 Patent describes very specific embodiments of generating available access point servers to build tunnels between endpoint devices, the Challenged Claims are not so limited. For example, independent Claims 1 and 8 are method and system, respectively, broadly directed to dynamically ordering a list of available servers based on a record associated with a network device. They do not recite that the server is an access point server, do not recite selecting a server to connect to the network device or provisioning a tunnel and connecting the network device to the tunnel, nor do they recite that the recited records include contextual information regarding the network device. Claim 1 is representative:

1. A method comprising:

receiving, by one or more computer processors, a request for a list of available servers from a network device;

retrieving, by the one or more computer processors, the list of available servers;

retrieving, by the one or more computer processors, one or more records associated with the network device;

dynamically ordering, by the one or more computer processors, the list of available servers based on the one or more records associated with the network device; and

transmitting, by the one or more computer processors, the ordered list of available servers to the network device.

C. Claim Construction

50. I have been informed and understand that in order to properly evaluate the '632 Patent, the terms of the claims must first be interpreted.

51. 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 this 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.

D. Priority Date of Claims

52. 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 U.S. Provisional Application Nos. 62/144,293. Solely for the purpose of this IPR, I have been instructed to assume the priority date is April 7, 2015.

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

53. 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 April 7, 2015, the earliest possible effective filing date of the '595 Patent.

54. 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.

55. With respect to the '595 Patent, a POSITA in the April 2015 timeframe would have been familiar with using cloud based resources offered on top of regular Internet connections to provide virtual networks and the formation of tunnels and the associated protocols as demonstrated by the prior art references and description in the State of the Art below.

56. A POSITA would have gained knowledge of these concepts through a mixture of training and work experience, such as by having a Bachelor's degree in computer science and two years of experience; or by obtaining a Master's degree in computer science, but having only one year of experience; or by having no formal education but experience in storage systems of at least four years.

57. 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.


V. STATE OF THE ART

58. The following section describes the state of the art for using cloud based resources on top of regular Internet connections to provide virtual networks and the formation of tunnels and the associated protocols build tunnels between an end point device and an access point server as of April 2015.

59. 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 reference. Accordingly, these references should be considered by the Board.

1. Communication Tunnels

60. The use of “tunnels” was well known in the art. For example, Iyer (2010) disclosed a plurality of communication paths (tunnels 1 and 2) between wireless clients, which were identified as “peer-to-peer tunnels” established between “any two clients in the mobility domain” to route traffic between the clients. Traffic between the clients pass through the tunnels and intermediate access point servers. EX-1006 [0015], [0030].



EX-1006 Fig. 4.

61. The two communication paths connecting wireless clients 500a and 500b are “form[ed]” by access nodes 400a and 400b, because access nodes 400a and 400b “establish” the tunnel 610, or tunnels 600a and 600b.

[M]obility manager 280 instructs access node 400a to establish peer-to-peer tunnel 610 between access node 400a and access 400b.

EX-1006 [0030].

62. In addition, the access nodes 400a,b can also be connected to controller 200 via tunnels:

[T]he access nodes establish communications with the controller using tunnels.

EX-1006 [0013].

2. Selection of Communication Paths By Controller

63. It was well known in the art to provide a controller that selected the communication path, including the selection of tunnels, between client devices. For example, Iyer discloses that the Controller 200 identifies the traffic destination and selects the tunnel through access nodes that the traffic uses. EX-1006 [0025 - 26].

64. Access nodes 400a, 400b, 400c, 400d provide wireless and possibly wired services to clients. As shown in FIG. 1, wireless client 500a is connected to

access node 400a. Access node 400b supports wireless client 500b and wired client 510b. Access node supports wireless client 500c. EX-1006 [0018].

65. Iyer describes a controller establishes peer-to-peer tunnels between access nodes to route that traffic between the clients. EX-1006 [0030].

This routing is performed by controller 200 using the IP addresses of clients 500a and 500b, as well as the MAC (media access controller) addresses of clients 500a, 500b and access nodes 400a and 400b.

EX-1006 [0026].

When client 500a wishes to send data to client 500b, it in essence forms an IP packet with client 500b's IP address as the destination, and with client 500a's IP address and MAC address as the source. This information is encapsulated and sent to controller 200.

EX-1006 [0026];

Controller 200 identifies the traffic destination as client 500b, and sends the traffic though tunnel 600b to access node 400b and client 500b.

Ex-1006 [0026].

66. .

VI. CHALLENGES

A. Identification of Challenge

67. Based on my understanding of the law of obviousness discussed above, Claims 1-14 of the '595 Patent are invalid as being obvious over Kommula and Claim 15-20 are invalid as being obvious over Hankins alone, or in view of Kommula.


B. Challenge 1: Kommula alone, Teaches Every Element of Claims 1-14

1. The Kommula Reference

68. Kommula discloses the use of a global server for load balancing cloud resources and communicates with numerous site switches that are coupled to host servers serving specific applications. *EX-1005* Abstract.

69. Kommula identifies the same contextual problem in the prior art identified in the '595 Patent—that the prior art global server does not take into consideration the host server's locations relative to the client. *EX-1004* [0007]. Like the '595 Patent, Kommula recognizes that the geographical distance between the server and a client is a factor affecting the response time for the client's access to the host server. *EX-1004* [0007]. Kommula provides the example that a client from Japan could receive better response time from a host server in Japan than from a host server in New York. Further, the conventional domain name server (DNS) algorithm allows invalid IP addresses (*e.g.*, that corresponding to a downed

server) to persist in a local DNS server until the invalid IP address expires. *EX-1004* [0007].



71. In one embodiment, the global server load balancing switch provides IP addresses of available host servers (red) to a client program (blue) which requested access to an application hosted on one of the host servers. The available host services are based on a selected set of performance metrics. Each site switch (yellow) provides the GSLB switch with current site-specific information

regarding the host servers (red) associated with the site switch. In one embodiment, when a domain name server (DNS) (purple) resolves a host name in a query and returns one or more IP addresses, the GSLB switch filters the IP addresses using performance metrics compiled from the site-specific information collected from the site switches and/or other metrics. The GSLB switch then returns a ranked or weighted list of IP addresses to the requesting computer program (blue) inquirer. In one embodiment, the IP address that is estimated to provide the best expected performance for the client is placed at the top of the list. *EX-1004* [0008].

72. Kommula explains that examples of suitable performance metrics include availability metrics (*e.g.*, a server's or an application's health), load metrics (*e.g.*, a site switch's session capacity or a corresponding preset threshold), and proximity metrics (*e.g.*, a round-trip time between the site switch and a requesting DNS server, the geographic location of the host server, the topological distance between the host server and the client program). A topological distance is the number of hops between the server and the client. Another proximity metric is the site switch's "flashback" speed (*i.e.*, how quickly a switch receives a health check result). Yet another metric is a connection-load metric that is based on a measure of new connections-per-second at a site. The ordered list can also be governed by other policies, such as the least selected host server. *EX-1004* [0009].

2. Detailed Application of Kommula to the Challenged Claims

a. Claim 1

[1.0] A method comprising:

73. Kommula discloses a method for using a global server load balancing switch to resolve a DNS query by a client by providing an ordered list of servers that is expected to serve the client with a high performance in a given application. *EX-1004* [0003]. As annotated below, Kommula discloses the architecture used to perform the method recited in Claim 1.




Fig. 1

EX-1004 FIG. 1 (annotated)¹.

¹ This figure will be referred to in the claims analysis below as Annotated FIG. 1:

[1.1] receiving, by one or more computer processors, a request for a list of available servers from a network device;

74. Kommula discloses receiving by the GSLB switch 12 and DNS 16 (*one or more computer processors*) (orange dashed line) a query for resources from a host server (red) which is a request for a list of IP addresses of host servers (*a request for a list of available servers*) from a client program on a client machine 28 (*network device*) (dashed blue line). *EX-1004* [0003]; [0008]; [0026];

Annotated FIG. 1.

75. Specifically, Kommula discloses that client program 28 is connected to Internet 14 and communicates with local DNS server 30 (purple). When a browser on client program 28 requests a web page, for example, using a Universal Resource Locator (URL), a query is sent to local DNS server 30 to resolve the symbolic host name (*e.g.*, *www.foundrynet.com*) to an IP address of a host server. This query is either sent to local DNS server 30's cache to retrieve the list of IP addresses, or if the time to live (TTL) of the responsive IP addresses in the cache has expired, the query is sent to GSLB switch 12 (green) to retrieve the list of IP address from DNS 16. *EX-1004* [0019].

76. Kommula's disclosure is consistent with the disclosure of the '595 Patent that when a device makes a request to a server, what it is really asking for is a list of available servers. *EX-1001* 23:11-24.

[1.2] retrieving, by the one or more computer processors, the list of available servers;

77. Kommula discloses that the GSLB switch in combination with DNS 16 (*one or more computer processors*) retrieves a list of IP addresses corresponding to the resolved host name. *EX-1004* [0019], [0021] and [0027].

[1.3] retrieving, by the one or more computer processors, one or more records associated with the network device;

78. Kommula discloses several examples of records of performance metrics associated with the network device that can be retrieved by the GSLB. For example, Kommula describes that a site switch (brown) measures the round trip time between the client machine and a host server and records it in an entry of the cache database (*one or more records associated with the network device*). Periodically, each site switch reports the RTT database to the GSLB switch along with load conditions (*retrieving, by the one or more computer processors, one or more records associated with the network device*). The GSLB switch aggregates the RTTs reported into a proximity table indexed by network neighborhood. *EX-1004* [0135].

79. More specifically, Kommula discloses that the RTT is computed for the interval between the time when a client machine requests a TCP connection to a proxy server configured on a site switch, sending the proxy server a TCP

synchronization (TCP SYN) packet, and the time a site switch receives from the client program a TCP acknowledgment (TCP ACK) packet. (In response to the TCP SYN packet, a host server sends a TCP SYN ACK packet, to indicate acceptance of a TCP connection; the client machine returns a TCP ACK packet to complete the setting up of the TCP connection.) The GSLB switch 12 maintains a database of RTT, which it creates and updates from data received periodically from the site switches (e.g., site switches 18A, 18B, 22A and 22B). Each site collects and stores RTT data for each TCP connection established with a client machine. *EX-1004* [0095].

80. Kommula further describes that the GSLB switch can look up the RTT for a client machine (*network device*) to any specific host server, based on the client's network neighborhood specified in the client's IP address. *EX-1004* [0135]. Thus, the GSLB (*one or more computer processors*) retrieves from the proximity table the RTT for a client machine to any specific host (*one or more records associated with the network device*).

81. Further, Kommula discloses other examples of performance metrics associated with a network device that can be retrieved by the GSLB switch from a database or memory that is remote from the GSLB. *EX-1004* [0024]. Examples of performance metrics associated with a network device include the topological

distance between the host server and the client. A topological distance is the number of hops between the server and the client. *EX-1004* [0009].

[1.4] dynamically ordering, by the one or more computer processors, the list of available servers based on the one or more records associated with the network device; and

82. Kommula discloses this limitation. Kommula discloses that when a client requests a resolution of a host server name, the GSLB switch (*one or more computer processors*) returns one or more ordered IP addresses for the host name (*dynamically ordering the list of available servers*). The IP addresses are ordered using metrics that include the information collected from the site switches or based on other metric information (*based on the one or more records associated with a network device*). The GSLB switch places the address that is deemed “best” at the top of the list (*dynamically ordering*). *EX-1004* Abstract.

83. As discussed in limitation [1.3], one of the performance metrics associated with the network device is RTT.

84. Kommula discloses that the IP addresses can be reordered (*dynamically ordering*) by the GSLB switch (*one or more computer processors*) based upon a round-trip time (RTT) between the site switch for the IP address and the client (*one or more records associated with the network device*).

85. Specifically, the RTT is computed for the interval between the time when a client machine requests a TCP connection to a proxy server configured on

a site switch, sending the proxy server a TCP SYN packet, and the time a site switch receives from the client program a TCP ACK packet. (In response to the TCP SYN packet, a host server sends a TCP SYN ACK packet, to indicate acceptance of a TCP connection; the client machine returns a TCP ACK packet to complete the setting up of the TCP connection.) The GSLB switch 12 maintains a database of RTT, which it creates and updates from data received periodically from the site switches (e.g., site switches 18A, 18B, 22A and 22B). Each site collects and stores RTT data for each TCP connection established with a client machine. In one embodiment, the GSLB switch favors one host server over another only if the difference in their RTTs with a client machine is greater than a specified percentage, the default specified percentage value being 10%, for example (*dynamically ordering*). *EX-1004* [0095].

[1.5] transmitting, by the one or more computer processors, the ordered list of available servers to the network device.

86. Kommula discloses that the list of ordered IP addresses is transmitted by the GSLB switch to the client. *EX-1004* [0095].

87. More specifically, Kommula discloses that the list of IP addresses of host servers is ordered (*ordered list of available servers*) by GSLB switch 12 (*one or more computer processors*) based on metrics such as RTT and sent to the client 28 (*network device*). *EX-1004* [0019]; [0095].

b. Claim 2

(a) *The method of claim 1 wherein the network device is an endpoint device.*

88. Kommula discloses that a client program 28 running on a client machine is connected to Internet 14, and communicates with local DNS server 30 in order to establish a TCP connection to a host server running an application. *EX-1004* [0005]; [0019]; [0095].

89. A POSITA understood that client machine 28 is an endpoint device.

90. To the extent that the Patent Owner argues that the client machine is not an endpoint device, a POSITA understood that the client machine 28 would communicate through an endpoint device in order to access the Internet 14. *EX-1004* FIG. 1; *EX-1006* [0018].

91. In this alternative, the request from the client machine would go through (and thus be received from) the endpoint device (as *network device*) and be returned to the endpoint through that device (as the *network device*).

c. Claim 3

(a) *The method of claim 2, wherein at least one server among the list of available servers is an access point server.*

92. Kommula discloses that host server may host application server programs, and POSITA understood that the host server may provide access for a

wired or wireless client to access the network, and thus act as an access point server. *EX-1004* [0005, 0017] FIG. 1.

d. Claim 4

- (a) *The method of claim 1, further comprising anticipating a level of demand for at least one server among the list of available servers, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.***

93. Kommula discloses that anticipating a level of demand at a server is one of the performance metrics that can be used by the GSLB switch to order the IP address of the host servers.

In accordance with embodiments of the invention, a weighted site or a weighted IP metric can be used. The weighted site metric can be based on a weight that determines how much traffic a site will receive with respect to other sites in comparison, while the weighted IP can be based on a weight that determines how much traffic an IP address will receive with respect to other IP addresses in comparison. An embodiment provides an active bindings metric that is based on a measure of the number of active real servers bound to a virtual IP address residing on a GSLB site. The active bindings metric bases its selection of the best IP address on this number and prefers a VIP with the highest number of active bindings.

EX-1004 [0010].

94. Specifically, Kommula discloses that one performance metric is a measure of new connections-per-second at a site which is used to calculate an average load for a site. EX-1004 [0098-99].

95. A POSITA understood that this calculation of average load is used as an estimate of the anticipated load at the site, e.g., it will continue at the average rate, so that the list of IP addresses of the servers can be order based on the site's connection load. Thus, Kommula disclosing calculating an average load using new connections-per-second at the site (*anticipating a level of demand of at least one server*) and reordering the list of IP addresses of the servers based on the average load (*dynamically reordering the list of available servers based on the anticipated level of demand*).

e. **Claim 5**

(a) ***The method of claim 1, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second network devices.***

96. This limitation is disclosed for the same reasons as disclosed with respect to limitations [1.3], [1.4] and [1.5] but for a request from a second network device. Kommula discloses that the GSLB switch will develop a second ordered list of the IP addresses (*dynamically ordering*) based upon a round-trip time (RTT)

between a site switch for the IP addresses and the second client (*one second ordered lists of available servers*) and the GSLB switch will send the list to the second client (*transmitted to one or more second network devices*). *EX-1004* [0095], [0097].

97. Further, because Kommula retrieves a performance metric for each site, such as a site switch's session capacity threshold, the first client's use of an IP address of a selected host server, may impact that site switch's session capacity which could impact the dynamic reordering of the second ordered list for the second client, and vice versa. *EX-1004* [0025].

98. Thus, the dynamic ordering of a list for a first network device would be based on the list of ordered available servers provided to a second network device.

99. A POSITA understood that the same weighted list of servers cannot be provided to all the network devices for load balancing purposes. If a server 1 is returned to one network device, returning the address of a second server 2 to a second network device would improve load balancing. Hence a POSITA understood that the ordered list returned to a network device would depend on responses to other network devices.

f. Claim 6

(a) *The method of claim 1, wherein dynamically ordering the list of available servers comprises*

generating ranking information for each server in the list of available servers.

100. Kommula discloses that in dynamically ordering as described in limitation [1.4] Kommula can generate a ranked or weighted list of IP addresses (*generating ranking information for each server in the list of available servers*):

Under one aspect of the present invention, when an authoritative DNS server resolves a host name in a query and returns one or more IP addresses, the GSLB switch filters the IP addresses using the performance metrics compiled from the site-specific information collected from the site switches and/ or other metrics. The GSLB switch then returns **a ranked or weighted list of IP addresses** to the inquirer. In one embodiment, the IP address that is estimated to provide the best expected performance for the client is placed at the top of the list.

EX-1004 [0008].

101. Kommula also discloses ordering “availability metrics” and RTTs which are dynamic in nature. *EX-1004* [0009, 0025].

g. Claim 7

(a) The method of claim 1, wherein the one or more records are retrieved based on a unique device identity associated with the network device.

102. As discussed with respect to limitation [1.3], the GSLB switch periodically receives RTT between a client machine and any specific host server

aggregates the RTTs reported into a proximity table indexed by network neighborhood. The GSLB switch can thus look up the RTT for a client machine to any specific host server, based on the client's network neighborhood specified in the client's IP address. *EX-1004* [0135].

103. A POSITA understood that in order to maintain a RTT for each client machine to each host server, a *unique device identity* would be assigned to each client machine in order to distinguish the RTT associated with specific client devices. To the extent that Patent Owner argues that Kommula does not expressly state that the client machine has a unique device identity, a POSITA would be motivated to associate a unique device identity with the client machine to facilitate the lookup of the RTT in the proximity table associated with the client machine.

h. Claim 8

[8.0] A system composing:

104. As discussed with respect to Claim 1, Annotated FIG. 1 of Kommula shows a system for using a global server load balancing switch to resolve a DNS query by a client by providing an ordered list of servers that is expected to serve the client with a high performance in a given application. *EX-1004* [0003]; FIG. 1.

[8.1] a memory storing instructions; and one or more processors coupled to the memory, the one or more processors being configured to execute the instructions, the instructions when executed

causing the one or more processors to perform operations comprising:

105. Kommula describes that FIG. 3 illustrates a block diagram showing the functional modules of the GSLB switch for performing the operations described in FIGs. 2A-2D. Kommula discloses that the instructions to be executed by the GSLB switch to control its operations are embodied in software or other machine-readable instruction stored on one or more machine-readable storage media (*memory storing instructions*). *EX-1004* [0026].

106. For example, such software to perform portions of the algorithm may be present at the GSLB switch 12 in one embodiment and executed by the switch controller 401 (*the instructions when executed causing the one or more processors to perform operations*). *EX-1004* [0026], FIGs. 2A-2D.

[8.2] receiving a request for a list of available servers from a network device;

107. Kommula discloses this limitation for the same reasons as discussed for limitation [1.1].

[8.3] retrieving the list of available servers;

108. Kommula discloses this limitation for the same reasons as discussed for limitation [1.2].

[8.4] retrieving one or more records associated with the network device;

109. Kommula discloses this limitation for the same reasons as discussed for limitation [1.3].

[8.5] dynamically ordering the list of available servers based on the one or more records associated with the network device; and

110. Kommula discloses this limitation for the same reasons as discussed for limitation [1.4].

[8.6] transmitting the ordered list of available servers to the network device.

111. Kommula discloses this limitation for the same reasons as discussed for limitation [1.5].

i. Claim 9

(a) The system of claim 8, wherein the network device is an endpoint device.

112. Kommula discloses this limitation for the same reasons as discussed for limitation [2.0].

j. Claim 10

(a) The system of claim 9, wherein at least one server among the list of available servers is an access point server.

113. Kommula discloses this limitation for the same reasons as discussed for limitation [3.0].

k. Claim 11

- (a) *The system of claim 8, wherein the operations further comprise anticipating a level of demand for at least one server among the list of available servers, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.***

114. Kommula discloses this limitation for the same reasons as discussed for limitation [4.0].

l. Claim 12

- (a) *The system of claim 8, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second network devices.***

115. Kommula discloses this limitation for the same reasons as discussed for limitation [5.0].

m. Claim 13

- (a) *The system of claim 8, wherein dynamically ordering the list of available servers comprises generating ranking information for each server in the list of available servers.***

116. Kommula discloses this limitation for the same reasons as discussed for limitation [6.0].

n. Claim 14

- (a) *The system of claim 8, where in the one or more records are retrieved based on a unique device identity associated with the network device.***

117. Kommula discloses this limitation for the same reasons as discussed for limitation [7.0].

C. Challenge 2: Hankins alone or in view of Kommula Teaches Every Element of Claims 15-20

1. The Hankins Reference


118. Hankins describes managing networks, including connecting a private network to the Internet and communicating signal packets from a first device in a first logical broadcast domain to another network device in another logical broadcast domains using a tunnel as if the logical broadcast domains are not separate. *EX-1005* Abstract.

119. Hankins describes that broadcast domain refers to a set of devices occupying a network address space, in which any device is able to communicate with any other device in the broadcast domain without rerouting a transmission. *EX-1005* 6:33-40.

120. A logical broadcast domain means that portions of the broadcast domain may not be entirely related to a particular set of physical devices and instead can be associated with a cloud implemented functionality. *EX-1005* 1:6-7; 4:12-16; 7:45-50.

121. Hankins discloses that signal packets are transmitted between the network devices of separate logical broadcast domains via a tunnel server. *EX-1005* Abstract.

122. As shown below Hankins connects the network devices via one or more tunneling servers (red). A tunneling server facilitates connections between a gateway device (green) in the first virtual domain with another gateway device (blue) in a second virtual domain via a plurality of tunnels (purple):



EX-1005, Fig. 2.

123. Hankins identifies the same contextual problem identified in the '595 Patent—that prior art tunnel controllers did not take into consideration contextual information regarding the network, such as the location of a tunnel server's locations relative to the gateways. Like the '595 Patent, Hankins proposes that the

communication channel may be enhanced by using this contextual information, such as by provisioning a tunnel server located more physically proximate to endpoints. *EX-1005* 25:7-25.

124. Like the '595 Patent, Hankins discloses that the controller may generate a list of available tunnel servers based at least in part on a host of considerations, such as, for example, cost, capability, efficiency, compatibility, resources, proximity, latency, bandwidth (e.g., throughput), utilization, others or combinations thereof. *EX-1005* 26:55-65.

125. For example, the gateway 311 communicates via tunnel server one 360 or tunnel server two 370 with gateway 321 over a plurality of tunnels (purple). The tunnels are selected by the controller 350 (orange) using information obtained from the network devices (gateway, tunnel servers, etc.) *EX-1005* 24:39-57, 25:15-55. Gateway device 311 may transmit signal packets via tunnel server one 360 since tunnel server one 360 is more proximate to both gateway device 311 and gateway device 321 in this example. Hankins discloses that an additional server may be selected as a secondary route such as gateway device 311 may transmit signal packets via tunnel server two 370, in the event that tunnel server one 360 is unavailable for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue,

scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. *EX-1005* 25:56-67.

126. Fig. 4 of Hankins generally illustrates the process and tasks of the Gateways (Green), Controller (Yellow) and Tunnel Servers (Red) in establishing the communication paths and transmitting packets.




FIG. 4

EX-1005 Fig. 4.

2. Detailed Application of Hankins, alone or in view of Kommula to the Challenged Claims

a. Claim 15

[15.0] A network system for managing a global virtual network, comprising:


127. To the extent that the preamble is limiting, annotated FIG. 2 illustrates a global virtual network of network devices, including gateways, distributed in various logical broadcast domains. A logical broadcast domain means that portions of the broadcast domain may be implemented in the cloud and may not be entirely related to a particular set of devices. *EX-1005* 1:6-7; 4:12-16; 7:45-50; 13:8-26; 19:51-55.

128. A POSITA understood that each of the various logical broadcast domains include gateways that extend the network edge into the cloud at various geographic locations and would be considered part of a *global virtual network*. This is consistent with the '595 Patent's disclosure that one of the features of the global virtual network is that the endpoint device, such as a gateway, acts as a bridge between the broader network and a local area network and brings elements of the cloud as a local node extension into the edge of the LAN. *EX-1001* 2:53-57; 10:37-43; FIG. 9.

[15.1] an access point server configured to connect to one or more endpoint devices via one or more first tunnels; and

129. Hankins discloses that the tunnel server (red) (access point server) is configured to connect one or more gateways (green) (endpoint devices) via one or more tunnels (purple). Hankins discloses that a controller may provision one or more tunnel servers to provide access to different broadcast domains and facilitate communications between network devices communicating as part of different broadcast domains. EX-1004 8:63-66; 26:65-27:8.

130. This is consistent with the disclosure of the '595 Patent where gateways are the endpoints of a tunnel connecting the gateways, and the gateways are the edge of their respective networks. EX-1001 FIGs. 9, 15.



EX-1005 FIG. 2.

[15.2] a control server connected to the access point server via a second tunnel, the control server storing one or more records associated with each of the one or more endpoint devices,

131. Hankins discloses that a controller (yellow) (*a control server*) is connected to a tunnel server (red) (*access point server*) via a *second tunnel* (brown). Although Hankins does not expressly state that the communication path shown in FIG. 2 between controller 350 and tunnel sever 360 or tunnel server 370 is a tunnel, a POSITA understood this path could be a tunnel. For example, Hankins discloses that a gateway can operate as a controller, and further discloses

that the communication path between a gateway and a tunnel server is via a tunnel.


EX-1005 14:41-49; 12:1-25.

132. Further, as discussed in the State of the Art, it was known to use a tunnel between an access point and the controller. *EX-1006* [0013]. Thus, a POSITA understood that controller could be implemented in a gateway which is connected to a tunnel server via a second tunnel.

133. Hankins further discloses that a controller may provision one or more tunnel servers to facilitate communications between network devices communicating as part of different broadcast domains. *EX-1004* 8:63-66.

134. In order to do so, the gateway devices transmit respective generated broadcast configuration (*one or more records associated with each of the one or more endpoint device*) to controller. *EX-1004* 20:24-25.

135. Hankins explains that broadcast domain configuration information includes one or more network protocols, available addresses, used addresses, topologies, devices used, such as switches or hubs, historical settings, such as for security, for a network protocol, etc., modifications of the foregoing, user accounts, including status, authentication, etc., security settings of a broadcast domain, workgroup or domain names, device names including status, available device features, etc., services available and/or used, status of the network devices, as well as other features. *EX-1004* 11:64-12:8.



EX-1005 FIG. 2.

136. In another example, Hankins discloses that a discovery device may comprise a network device capable of communicating as part of a broadcast domain to discover features of the particular broadcast domain without human intervention to generate a broadcast domain configuration of the particular broadcast domain. A discovery device may engage in a variety of activities to accomplish this including, but not limited to any of the following: passively monitoring signal packets, actively probing other network devices, querying other network devices, querying servers, such as a domain controller or other server responsible for a directory of broadcast domain features, any combinations thereof, or other approaches. EX-1004 9:10-22.

137. Thus, configuration information for a particular broadcast domain associated with a gateway may be generated using a discovery device.

138. As shown in FIG 1, a gateway device transmits a generated broadcast domain configuration associated with the gateway to the controller 206, and the controller stores the generated broadcast domain configuration associated with the gateway device in database 207 (*the control server storing one or more records associated with each of the one or more endpoint devices*). *EX-1004 21:7-19.*

[15.3] wherein the control server is configured to receive a request for a list of available servers from a first endpoint device among the one or more endpoint devices, the list of available servers including the access point server;

139. Hankins discloses that the controller receives a request from a gateway 311 (*endpoint device*) for a list of available servers including tunnel servers (*access point server*). For example, Hankins discloses that gateway 311 in one logical broadcast domain, may request services available from gateway 321 in another logical broadcast domain. *EX-1005 12:9-16.*

140. The controller identifies potential tunnel servers that are available to provide a tunnel between gateway 311 to gateway 321. For example, as shown in Annotated FIG. 2, tunnel servers 360 and 370 can be used to establish communications between gateway 311 and 321. Thus, a POSITA understood that a request from gateway 311 for services available from gateway 321 is a request

for available tunnel servers that can provide a tunnel between gateway 311 and 321. This is consistent with the disclosure of the '595 Patent that when a device makes a request to a server, what it is really asking for is a list of available servers, and specifically an IP and port which it could utilize to in the build a tunnel. *EX-1001* 23:11-24.

141. Furthermore, via a controller, such as 350, a private network may be reconfigured without human intervention, for example, to improve performance and/or reduce cost. For example, if available bandwidth decreases, if latency increases, and/or if network utilization increases, controller 350 may provision additional tunnel servers or may provision tunnel servers in a manner to improve performance and/or reduce cost without human intervention. Examples of possible monitoring and/or measurement of network performance and/or cost-related parameter were previously discussed. *EX-1005* 25:45-55.

142. Hankins discloses a controller 350 that receives domain configuration information and identifies tunnel servers associated with one of the paths or channels for communications between the gateway devices. *EX-1005* 24:49-55.

[15.4] retrieve the list of available servers; retrieve the one or more records associated with the first endpoint device

143. Hankins discloses that controller *retrieves a list of available tunnel servers* based on the stored configuration of the broadcast domain of a gateway

(retrieve the one or more records associated with the first endpoint device). For example, Annotated FIG. 2 shows an embodiment with four logical broadcast domains in communication via more than one tunnel server. Gateway devices 311, 321, 331, and 341 may generate broadcast configurations for their respective logical broadcast domains which may be used by controller to identify a list of available tunnel servers. For example, controller 350 may provision tunnel server one 360 and initiate instructions to gateway devices 311, 321, 331, and 341 to communicate with tunnel server one 360. *EX-1004* 24:47-55.

144. In addition, Hankins discloses retrieving more than one tunnel server to create multiple tunnels between gateways. For example, in addition to identifying tunnel server one 360, controller may identify tunnel server two 370 which is more physically proximate to local broadcast domain c 330 and logical broadcast domain D 340. *EX-1004* 25:14-19.

145. Additional tunnel servers can be identified based on a variety of factors or considerations, including, signal packets communicated via a tunnel server, geography, latency, utilization, bandwidth (e.g., throughput), availability, efficiency, cost, etc. *EX-1004* 34-45.

[15.5] dynamically order the list of available servers for the first endpoint device based on the one or more records associated with the network device; and

146. Hankins discloses several examples of this limitation. Hankins discloses dynamically ordering a list of available tunnel servers for the gateway based on the one or more records associated with the gateway that provide contextual information. For example, gateway 311 is provided with a list of tunnel servers such as tunnel server one 360 and tunnel server two 370, based on information obtained from the network devices (gateway, tunnel servers, etc.) such as proximity of gateway device to tunnel server 360. *EX-1005* 24:39-57, 25:15-55.

147. In addition, Hankins discloses that additional server, tunnel server two 370, may be selected as a secondary route such that gateway device 311 transmits signal packets via tunnel server two 370, in the event that tunnel server one 360 is unavailable for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue, scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. *EX-1005* 56-67.

148. Thus, Hankins discloses identifying an additional tunnel server to provide a secondary route in the event that a primary tunnel server is unavailable, which a POSITA understood discloses dynamically ordering the availability of servers.

149. By way of another example, Hankins discloses an embodiment where broadcast domains communicate through multiple tunnel servers. Each *n*th packet goes to the *n*th tunnel server. Hankins describes that this approach can improve security, fault tolerance, redundancy and/or load balancing of tunnel servers. *EX-1005* 26:21-48.

150. A POSITA understood that providing load balancing and fault tolerance means that a list of tunnel services would be provided in preference of order.

151. Hankins further describes that different tunnel servers could provide different services. For example, one tunnel server may be better suited for use in file serving and the other tunnel server could be better suited for message serving. Hankins describes that the provisioning of tunnel servers may be generated based on the capability of the tunnel servers and thus a POSITA understood that Hankins orders the availability of the tunnel servers on the needs of the gateways in the broadcast domains. For example, if the needs include file service more than messaging, the tunnel server handling file service would be ranked above the tunnel server for message servicing.

152. Thus, Hankins discloses dynamically ordering the list of available servers.

153. To the extent that Patent Owner assert that Hankins does not expressly state that the available tunnel servers are dynamically ordered, Kommula discloses dynamically ordering a list of available servers to provide a communication path between network devices based on contextual information. A POSITA would be motivated to modify the identification of available tunnel servers as taught by Hankins to include dynamically ordering as taught by Kommula in order to make a more efficient use of network resources.

154. Hankins and Kommula are analogous art. Both Hankins and Kommula are directed to identifying available servers to provide a communication path through a network. *EX-1005* 12:9-16; 24:39-57; 25:15-55; *EX-1004* [0008], [0019], [0026], [0027].

155. Both Hankins and Kommula use contextual information to identify available servers such as the proximity of the server to the endpoint device or gateway. *EX-1005* 25:7-25; 26:55-65; *EX-1004* [0009].

156. Kommula expressly describes dynamically ordering the list of available servers based on the contextual information such as proximity using RTT. *EX-1004* [0009], [0095], [0135].

157. A POSITA would be motivated to modify the teaching of Hankins to identify available tunnel servers based on contextual information by including dynamically ordering the available servers as taught by Kommula because it results

in a more efficient use of network resources. For example, Hankins describes that the available servers in FIG. 2 include tunnel server one 360 and tunnel server two 370. Dynamically ordering server one 360 ranked ahead of tunnel server two 370 because tunnel server 360 is more proximate to gateway 311 results in a shorter RTT, as taught by Kommula, for communications using tunnel server one as compared to tunnel server two, which is one of the express goals of Hankins.

158. A POSITA would be confident of the success in modifying the teachings of Hankins with Kommula for several reasons.

159. First, Hankins and Kommula both teach using software implementations to identify available servers, which could be identified by accessing stored contextual data. *EX-1005* FIGs. 4, 5; *EX-1004* [0013], [0026], FIGs 2A-2D.

160. Both Hankins and Kommula teach using a software implementation to access stored data to find the “best” available server, e.g. closest proximity or shortest RTT. *EX-1005* FIGs. 4, 5; *EX-1004* [0013], [0026], *FIGs 2A-2D*.

161. Thus, it is only a simple software change to Haskins teachings to dynamically rank the results of the contextual information retrieved from the database.

162. Further, Hankins discloses that the tunnel servers can implement network policies and do other services. *EX-1005* 8:12-37.

163. Hankins further describes that the provisioning of tunnel servers are reconfigurable in real-time. *EX-1005* 18:46-63.

164. Moreover, it was known that ranking of tunnel servers can be established as a network policy, as taught by Kommula and Hankins contemplates that tunnel servers are readily capable of executing such policies.

[15.7] transmit the ordered list of available servers to the first endpoint device.

165. Hankins in combination with Kommula discloses this limitation. Hankins discloses that the available tunnel servers are provided to gateway 311. For example, as discussed with limitation [15.4], Hankins discloses that controller can identify tunnel servers 360 and 370 for gateway 311 to communicate with gateway 321. Gateway 311 may transmit signal packets to gateway device 321 via tunnel server one 360 since tunnel server one 360 is more proximate to both gateway device 311 and gateway device 321 in this example.

166. Additionally, in response to dynamically ordering the available servers, gateway device 311 may transmit signal packets via tunnel server two 370 if tunnel server one may be unavailable, for a variety of reasons, such as, but not limited to one or more of the following situations: a software issue, a configuration issue, a power issue, scheduled maintenance, insufficient permissions, a security issue, other issues, or combinations thereof. Thus, Hankins discloses transmitting an ordered list, i.e., tunnel server 360 if available, and then tunnel server 370.

b. Claim 16

- (a) The network system of claim 15, wherein the control server is further configured to predict a level of demand for the access point server, and wherein dynamically ordering the list of available servers is further based on the anticipated level of demand.*

167. Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 4.

c. Claim 17

- (a) The network system of claim 15, wherein dynamically ordering the list of available servers is further based on one or more second ordered lists of available servers transmitted to one or more second endpoint devices among the plurality of endpoint devices.*

168. Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 5.

d. Claim 18

- (a) The network system of claim 15, where in dynamically ordering the list of available servers comprises generating ranking information for each server in the list of available servers.*

169. Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 6.

e. **Claim 19**

- (a) *The network system of claim 15, the one or more records are retrieved based on a unique device identity associated with the first endpoint device.*

170. Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 7.

f. **Claim 20**

- (a) *The network system of claim 15, wherein the control server receives tunnel information associated with the one or more first tunnels via the global virtual network, and wherein dynamically ordering the list of available servers is further based on the tunnel information.*

171. Hankins as modified by Kommula discloses this limitation for the same reason as discussed with respect to Claim 4.

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: 26th Feb. 2024

By *A.L. Narasimha Reddy*

Dr. A.L. Narasimha Reddy

APPENDIX A – MATERIAL RELIED UPON

<i>Exhibit #</i>	<i>Description</i>
1001	U.S. Patent No. 11,108,595 (the '595 Patent)
1002	Prosecution history of the '595 (“’595 File History”)
1003	Declaration of Petitioner’s Expert Dr. A.L. Narasimha Reddy(“Reddy”)
1004	U.S. Pat. Appl. Publ. No. 2010/0153558 to Kommula (“Kommula”)
1005	U.S. Pat. No. 10,177,957 (“Hankins”)
1006	U.S. Pat. App. Publ. No. 2010/0275017 to Iyer (“Iyer”)
1007	Related Patents

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.