

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

AYLO FREESITES LTD,
Petitioner,

v.

DISH TECHNOLOGIES L.L.C.,
Patent Owner.

IPR2024-00043, -00044, -00045, -00046, -00047, -00048
Attorney Docket No.: 081841.0167

**DECLARATION OF DR. KEVIN JEFFAY IN SUPPORT OF
PATENT OWNER'S PRELIMINARY RESPONSES
UNDER 35 U.S.C. § 313 AND 37 C.F.R. § 42.107**

I. Introduction

1. I, Dr. Kevin Jeffay, hereby declare as follows:

A. Engagement

2. I am an independent consultant with over 40 years of experience in technologies related to media streaming. I have been retained on behalf of DISH Technologies L.L.C. (“Patent Owner” or “DISH”), as an expert consultant on telecommunications technology related to streaming media content over networks including the Internet.

3. As part of this work, I have been requested by counsel for DISH to review the Petitions for *Inter Partes* Review and accompanying Declarations of Dr. Henry Houh, which were submitted by Petitioner Aylo Freesites LTD, in support of their Patent Owner Preliminary Responses (“POPR”), and to assess Dr. Houh’s analysis of US Patent Nos. 10,951,680 (the “’680 Patent”), 11,677,798 (“the ’798 Patent”), 9,407,564 (“the ’564 Patent”), 10,469,554 (“the ’554 Patent”), 8,868,772 (“the ’772 Patent”), 11,470,138 (“the ’138 Patent”) (the “Asserted Patents” or “ABR [Adaptive Bit-Rate] Patents”) in view of certain prior art references.

4. I have prepared this report to summarize certain aspects of my opinions regarding this subject matter and its relevance to the ABR Patents as concerning validity. At minimum, I have reviewed, relied on, and/or had access to the materials (*i.e.*, prior art references and documents) described herein.

5. I am being compensated at \$800 per hour for my services in this matter. My compensation is not dependent upon the outcome of this matter. If called upon to do so, I am prepared to testify as an expert witness in this regard.

B. Background and Qualifications

6. I am qualified by education and experience to testify as an expert in the field of telecommunications and video streaming. Attached, as Exhibit A, is a copy of my resume detailing

my experience and education. Additionally, I provide the following overview of my background as it pertains to my qualifications for providing expert testimony in this matter.

7. I am a full professor in the Department of Computer Science at the University of North Carolina at Chapel Hill where I currently hold the position of Gillian T. Cell Distinguished Professor of Computer Science. I have been a faculty member at UNC since 1989. From 2014-2022 I also served as the Chairman of the Department.

8. I received a Ph.D. in computer science from the University of Washington in 1989. Before that I received a M.Sc. degree in computer science from the University of Toronto in 1984, and a B.S. degree with Highest Distinction in mathematics from the University of Illinois at Urbana-Champaign in 1982.

9. I have been involved in the research and development of computing systems for over 40 years. I have been a faculty member at the University of North Carolina since 1989 where I perform research and I teach in the areas of multimedia networking, computer networks, distributed systems, real-time systems, and operating systems, among others. A major theme of my research has been the development of technology to improve the performance of data transfers on the Internet generally and for real-time applications specifically. My research has examined problems ranging from network and operating system support for real-time multimedia applications such as audio and video streaming, voice-over-Internet protocol (VoIP) and Internet videoconferencing, to the design of congestion control mechanisms in network routers, to measurements and analysis of network traffic to passively assess the performance of servers on the Internet.

10. For example, starting in the late 1980s, the focus of my research was the development of network and operating system technology to enable the real-time transfer of streams of audio and video data across the Internet. This involved, among other things, the development of media

encoders, media players, and network communication protocols for adaptive transmission of live audio and video data on the Internet. This work culminated in my research group developing some of the first videoconferencing systems for the Internet. Several of the papers authored by myself and members of my research group on this project won awards for their technical contributions.

11. The videoconferencing and adaptive streaming research attracted the attention of industry groups such as IBM, Intel, Digital Equipment Corporation, Cabletron, and AT&T Bell Laboratories. For example, starting in 1991 IBM supported aspects of my research at UNC and these efforts resulted in U.S. Patent No. 5,892,754 on adaptive media streaming being issued to IBM and UNC.

12. During this time, I was also collaborating with researchers at the Intel Architecture Labs in Hillsboro, OR, to modify an Intel product, the ProShare videoconferencing system, to use technology developed in my research lab at UNC for adaptive media transmission. By the late 1990s, the ProShare product included the ability to also transfer desktop information of a PC to a remote PC.

13. My research group also developed other “data conferencing” systems, also known as “shared window systems,” that were capable of transferring the desktop information on a PC to a remote PC in real-time. One system that we built, called “XTV,” was operational by 1991. The source code for XTV was made freely available and, by 1993, had been downloaded by over 600 users and institutions. The system was functionally and visually equivalent to today’s Cisco’s WebEx and Zoom’s screen sharing products and services. It provided a sophisticated means of transferring the desktop information of a PC to a remote site. In the XTV system, individual windows displayed on the desktop, or the entire desktop itself, could be selected for sharing with remote users. The XTV system also allowed remote users to remotely control and manipulate the

desktop whose information was being distributed, as well as remotely control the applications that were generating the desktop information that was being shared.

14. In much of my research, I regularly build and use clusters of computers interconnected by network switches, bridges, and routers to form and evaluate experimental and production networks. For example, in the late 1990s and early 2000s, my research evolved to consider router-based mechanisms for controlling the performance of network traffic. In much of this research, my students and I built and instrumented network routers and performed large scale experiments with this equipment. The instrumentation included, for example, the development of network monitors that would receive copies of packets flowing on a network link and analyze and store packets or the results of analyses on these packets, all in real-time. Based on these experiments, in 2003, my group at UNC won the most prestigious research award for original research in computer networking. UNC applied for, and received, U.S. Patent No. 7,447,209 for aspects of this research.

15. This project, and others, took place in a networking lab my students and I constructed at UNC over a number of years. The lab consists of several hundred computers and networking devices. Managing this lab required establishing and configuring VLANs, firewalls, and other security appliances to isolate the lab from the campus network (and vice versa).

16. More recently, my research group has considered the design and operation of next generation aerial networks. UNC applied for, and received, U.S. Patent No. 9,832,705 for aspects of this research.

17. I have authored or co-authored over 100 articles in peer-reviewed journals, conference proceedings, texts, and monographs in the aforementioned areas of computer science and others. In addition, I have edited and co-edited numerous published proceedings of technical conferences and have edited a book of readings in multimedia computing and networking (with Hong-Jiang

Zhang) published by Morgan Kaufman. I am a co-author (with Long Le and F. Donelson Smith) of a monograph related to computer network protocols, and a co-author (with Jay Aikat and F. Donelson Smith) of a second monograph related to experimental computer networking.

18. My research extends to leadership positions in several journals, conferences, and committees. I have previously served as Editor-in-Chief for the journal Multimedia Systems and Associate Editor for the journal Real-Time Systems. In addition, I have been an active participant in the Association for Computing Machinery (“ACM”) and the Institute of Electrical and Electronics Engineers (“IEEE”). Specifically, I have been a member of the steering committees for the ACM Special Interest Group on Multimedia, and the ACM Special Interest Group on Data Communications (“SIGCOMM”) and its subgroup on Internet Measurement. I have also served as a member of the IEEE Technical Committee on Real-Time Systems. As a result of this involvement, I have served as a program chair or member of the technical program committee for over 100 professional, international, and technical conferences, workshops, and symposia. I was previously on the steering committee for the ACM technical committee on Internet measurement, where we worked on community standards for measurement-based research. I have also served on numerous proposal review panels for the National Science Foundation and other international funding agencies in the aforementioned areas of computer science.

19. In addition to the U.S. patents referenced above, I am a named inventor on a fourth patent. Each of these patents is generally related to computer networking and the delivery of services over networks including audio and video transmission.

20. I have developed and taught undergraduate and graduate level courses on operating systems, computer networking, computer security, and multimedia computing. These courses have been developed for, and delivered to, students at UNC as well as for industry professionals.

21. I have served as an expert witness and technical consultant in litigation and *inter partes* review matters concerning videoconferencing, data conferencing and screen sharing, cellular and wireline telephony, voice over IP (VoIP) telephony, multimedia networking, distributed systems, operating systems, computer networks, datacenter networking, embedded systems and embedded software, and real-time systems, among others. I have testified in several trials, arbitrations, and claim construction hearings as an expert witness.

II. Relevant Legal Principles

22. I have applied the following legal principles provided to me by counsel in arriving at the opinions set forth in this report.

A. Prior Art

23. I understand that a person shall not be entitled to a patent if the claimed invention was patented, described in a printed publication, or in public use, on sale, or otherwise available to the public before the effective filing date of the claimed invention. I understand that disclosures made one year or less before the effective filing date of the claimed invention shall not be prior art if the disclosures were made by an inventor, joint inventor, or by another who obtained the subject matter directly or indirectly from the inventor or a joint inventor; or the disclosures were made by another after an initial disclosure by an inventor, joint inventor, or by another who obtained the subject matter directly or indirectly from the inventor.

24. I understand that prior art also includes issued patents or published applications in which the patent or application names another inventor and was effectively filed before the effective filing date of the claimed invention. I understand that such disclosures shall not be prior art if (1) the subject matter disclosed was obtained directly or indirectly from the inventor or a joint inventor; (2) the subject matter disclosed had been publicly disclosed by the inventor before the

effective filing date; or (3) the disclosed subject matter and the claimed invention were owned by the same entity no later than the effective filing date.

B. Presumption of Validity

25. It is my understanding that the claims of an issued U.S. Patent are presumed to be valid and enforceable. It is my further understanding that the basis for that presumption of validity is the fact that the allowed claims went through a rigorous examination process at the United States Patent and Trademark Office (“USPTO”). Moreover, here, at least certain claims of the ’554 and ’564 Patents were also found valid over at least nine prior art grounds asserted by third-party Respondents in a prior litigation at the International Trade Commission. I understand that this is further evidence of the validity of at least those claims, and other similar claims of the ABR Patents asserted here.

26. It is my further understanding that the MG Defendants bear the burden of proving invalidity through facts supported by clear and convincing evidence. I understand that the “clear and convincing evidence” standard means evidence which produces an abiding conviction that the truth of a factual contention is highly probable.

27. It is also my understanding that the burden of proving invalidity by clear and convincing evidence based on a reference or a combination of references is especially difficult when any of those references was before the USPTO Examiner during the prosecution of the relevant patent application, because it is presumed that the Examiner considered those references prior to allowing the claims.

C. Anticipation

28. It is my understanding that invalidation by “anticipation” only exists if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference. I understand that the references must be viewed from the perspective of a person of

ordinary skill in the art (a “POSITA”). I understand that two or more items of prior art cannot be combined to find anticipation. I also understand that the single prior art reference must include all elements arranged as required by the claim. I further understand that to be considered anticipatory, the prior art reference must be enabling and must describe the patentee’s claimed invention sufficiently to have placed it in the possession of a person of ordinary skill in the field of invention. I understand that the disclosure in the particular item of prior art does not have to be in the same words as the claim, but all of the requirements of the claim must be there, either expressly or inherently, so that a person of ordinary skill in the art looking at that particular item of prior art would be able to make and use at least one embodiment of the claimed invention.

29. I understand that a reference that does not expressly anticipate a claim or claim element can be found to inherently anticipate, but only if the reference necessarily includes the unstated limitation or necessarily functions in accordance with the claim or claim element. I understand that inherency cannot be established by probabilities or possibilities. And I further understand that a reference must enable someone to practice the invention in order to anticipate under § 102(b).

D. Obviousness

30. I understand that a patent claim is invalid if the claimed invention would have been obvious to a person of ordinary skill in the art at the time the invention was made. I understand a claimed invention is obvious when, even if all the requirements of the claim cannot be found in a single item of prior art that would otherwise anticipate the claim, it would have been obvious to a person of ordinary skill in the art, having knowledge of all of the analogous prior art, to have come up with the claimed invention.

31. I understand that a patent claim composed of several elements is not proved obvious merely by demonstrating that each of its elements was independently known in the prior art. Most, if not all, inventions rely upon building blocks of prior art, and claimed discoveries almost of necessity

will be combinations of what, in some sense, is already known. Accordingly, to determine whether an invention is obvious, I understand it can be important to determine whether there is a reason that would have prompted a person of ordinary skill in the art to combine the elements in the way the claimed new invention does. I understand that such teachings, suggestions, and motivations may be found in the prior art, the knowledge of a person of ordinary skill in the art, or the nature of the problem solved by the claimed invention.

32. In evaluating whether the Asserted Claims¹ would have been obvious, I have also considered the following factors:

- Whether the Petitioner has identified a reason that would have prompted a person of ordinary skill in the art to combine elements or concepts from the prior art in the same way as in the claimed invention;
- Whether the claimed invention applies a known technique that had been used to improve a similar device or method in a similar way; and
- Whether the claimed invention would have been obvious to try, meaning that the claimed innovation was one of a relatively small number of possible approaches to the problem with a reasonable expectation of success by those skilled in the art.

33. I understand that it is important to be careful not to determine obviousness using hindsight, because many true inventions can seem obvious after the fact. Obviousness is determined from the position of a person of ordinary skill in the art at the time the claimed invention was made, and I

¹ I understand that DISH submitted claim charts with its Complaint in this case that map each limitation of the following claims of each ABR Patent to certain accused websites operated by the MG Defendants: '554 Patent Claims 16, 17, and 20; '138 Patent Claim 14; '798 Patent Claim 1; '772 Patent Claim 1; '564 Patent Claims 1, 3, 4, 5, and 8. I refer herein collectively to these claims as the "Asserted Claims."

understand it is improper to consider what is known today or what is learned from the teaching of the patent. For example, I understand it is improper to use the Asserted Patents as a road map for selecting and combining prior art.

34. I understand that the ultimate conclusion of whether a claim is obvious should be based on a determination of the following factual issues:

- The level of ordinary skill in the art that someone would have had at the time the claimed invention was made;
- The scope and content of the prior art;
- The differences, if any, that existed between the claimed invention and the prior art; and
- Secondary considerations or objective evidence of non-obviousness.

35. I understand that objective evidence or secondary considerations of non-obviousness include the following:

- Commercial success of a product due to the merits of the claimed invention;
- A long-felt, but unsolved, need for the solution provided by the claimed invention;
- Unsuccessful attempts by others to find the solution provided by the claimed invention;
- Copying of the claimed invention by others;
- Unexpected and superior results from the claimed invention;
- Acceptance by others of the claimed invention as shown by praise from others in the field of the invention; and

- Disclosures in the prior art that criticize, discredit, or otherwise discourage the claimed invention and would therefore tend to show that the invention was not obvious.

36. I have been informed that a patentee is entitled to a rebuttable presumption of nexus between the asserted evidence of objective indicia and a patent claim if the patentee shows that the asserted objective evidence is tied to a specific product and that the product is the invention disclosed and claimed. In other words, presuming nexus is appropriate when the patentee shows that the asserted objective evidence is tied to a specific product and that product embodies the claimed features, and is coextensive with them. If on the other hand, the patented invention is only a component of a commercially successful machine or process, the patentee is not entitled to a presumption of nexus.

III. Level of Ordinary Skill in the Art

37. Dr. Houh contends that a person of skill in the art (“POSITA”) at the time of invention of each of the ABR Patents “would have had at least a B.S. degree in computer science or electrical engineering (or comparable degree) and two years of experience in networking or streaming media, or a M.S. in computer science or electrical engineering (or comparable degree).” ’772 Patent Houh Dec. ¶51.²

38. I disagree. Instead, it is my opinion that a POSITA at the time of invention of each of the Asserted Patents would have been someone with at least a bachelor’s degree in electrical engineering or computer science, or an equivalent degree or experience in a related field, with at

² References may be made throughout to papers or exhibits submitted with a particular IPR. However, the same information may be found in each of the IPR petitions and accompanying exhibits at different locations. Accordingly, my analysis applies equally for each petition.

least two years of practical experience or coursework in the design or development of distributed multimedia delivery systems.

39. Dr. Houh's proposed definition of a POSITA is improper because it does not require any experience with "multimedia delivery systems" or "streaming media," and instead may be satisfied by someone with "networking" experience alone, or someone with simply a Master's degree in a relevant field and no networking experience. Accordingly, the definition of a POSITA I propose above is more suitable. However, for purposes of my analysis presented below, Dr. Houh's proposed definition may be accepted as sufficient.

IV. Claim Construction

40. In general, I understand that claim terms are to be given their plain and ordinary meaning according to a POSITA at the time of the invention. Here I understand that the International Trade Commission ("Commission" or "ITC") in Investigation No. 337-TA-1265 (the "1265 Investigation") previously construed the term "streamlet(s)" as recited in two of the ABR Patents as "any sized portion of a content file each of which is stored separately." However, this construction is not required for purposes of rendering the opinions I express here.

41. I have not identified any other terms of the asserted claims of the ABR Patents that require construction for purposes of rendering the opinions I express here.

42. Dr. Houh proposes a construction of "virtual timeline" as "a playlist of entire content files." Houh Dec. ¶43. While I disagree³ with that construction, the opinions I render below do not depend on the claimed "virtual timeline." Accordingly, I do not address Dr. Houh's construction of this term here.

³ I submitted a report in the ITC Investigation, which advocated for application of the plain and ordinary meaning (*e.g.*, a file that defines a playlist for a user to view) to the term "virtual timeline." I understand the claims reciting the term "virtual timeline" were not brought to trial, and therefore the ITC did not issue a claim construction on this term.

V. Summary of Opinions

43. Based on the materials that I have reviewed, the analyses I have performed, and the positions summarized within this report, it is my opinion that Petitioner and their expert, Dr. Houh, fail to raise any substantial question of validity regarding the Asserted Patents.

VI. Dr. Houh Fails to Show a Substantial Question of Invalidity

44. I have reviewed Dr. Houh's declarations, and conclude that in each declaration Dr. Houh fails to raise a substantial question of invalidity.

45. I disagree with Dr. Houh's conclusion that each Asserted Claim is allegedly invalid over Ogdon and Shteyn. *Id.* At the onset here, I note that as part of my service as an expert witness in patent litigation matters related to video streaming, I have had occasion to research the state of the art in video streaming in the late 1990s and early 2000s. As part of this research, and prior to my engagement in these matters, I had previously encountered patents in the family of the ABR Patents, specifically, U.S. Patent No. 7,818,444 to Brueck and Hurst around the 2016 – 2017 time period. Having done original research in adaptive streaming in the 1990's, I recall being impressed by the '444 patent, in particular, its architecture for what it termed "multi-bitrate video and/or audio content streaming." The architecture included using a web server hosting sets of content files where each set represented segments ("streamlets") of a larger media presentation encoded in the same format but at a particular streaming bit-rate. End-user stations were able to adaptively stream the media presentation by requesting content files in the same format, but at different bitrates, so as to receive media at a particular bit-rate that the end-user station had estimated was sustainable. While this paradigm of client adaptation and requesting of content is common for streaming media today (in particular, requesting via HTTP), the '444 patent impressed me as being the earliest patent I had seen with such an architecture.

46. In addition, while HTTP is today the dominant transport means for streaming media, I was impressed by the use of HTTP in the ABR Patents. As discussed above, I had worked on a variety of streaming media projects throughout the 1990s and recall that during this time I was skeptical that effective streaming systems could be developed using HTTP (because of its overhead, its reliance on TCP for underlying transport, and because of the lack of support for real-time content delivery in commodity web servers). Servers supporting streaming (or even video-on-demand) were commonly thought to require special support for media adaptation (*e.g.*, transcoding) and/or real-time storage and retrieval systems (*e.g.*, real-time file systems).

47. Thus, when I was approached by counsel for DISH Network to work on a related matter, I agreed because I had a good faith basis for believing that the ABR Patents were valid (and useful). And nothing in my objective review of the evidence provided by Dr. Houh has changed this opinion.

48. Dr. Houh's claim charts and the Aylo IPR's fail to raise any substantial question of invalidity for numerous reasons, which I address in the sections that follow. My analysis here is preliminary in nature—it should not be read or understood to be exhaustive of all possible deficiencies, and I reserve the right to raise additional deficiencies if called upon to testify in this matter or in any related IPR or other proceedings.

49. In rendering my opinions below, I address, as examples, elements of each Asserted Claim that are not disclosed, suggested, or rendered obvious by any prior art reference or combination raised by Dr. Houh. In my opinion, the claims containing such elements are not invalid. Moreover, to the extent that I have provided a reason that a specific independent claim is not rendered invalid, my opinion is that none of the claims depending from such independent claim are invalid for at least the reasons given for that dependent claim's respective independent claim.

50. In addition, I note that the Aylo IPR's address certain claims of the Asserted Patents that I understand are not currently asserted in this case. The fact that I do not specifically address any claim does not constitute a waiver or concession of invalidity as to any such claims.

A. Dr. Houh's Obviousness Opinions Are Foreclosed by Strong Objective Evidence of Nonobviousness Identified by the ITC

51. The Aylo IPR's exclusively rely on obviousness theories, and Dr. Houh separately asserts invalidity based on obviousness in his Claim charts corresponding to the Ogdon reference. However, Dr. Houh fails to address the International Trade Commission's finding that strong evidence of secondary considerations supported a finding of non-obviousness.

52. I previously testified at length about the strength of this evidence in the 1265 Investigation. The Final Initial Determination in that Investigation (as adopted in relevant part by the ITC in its written opinion) credited my testimony in finding the following:

... Dr. Jeffay demonstrated through corroborating evidence that Move's success is tied to the "unique characteristics of the claimed invention." As Dr. Jeffay testified, the Move system was described as early as 2008 as offering "a smooth end user experience as the Move Media Player up-shifts and down-shifts in response to network and client CPU availability." Moreover, the Move Media Player's encoding technique allowed streaming via HTTP, and the use of "[s]imple HTTP protocol transfer of media files from standard Web servers rather than deployment of expensive media servers' reduces cost and complexity." DISH further adduced evidence demonstrating that the Move Media Player was recognized as superior in streaming performance relative to its competitors, rebutting respondents' argument that "advancements in compression and encoding technologies," etc. account for the Move Media Player's success.

In view of this record evidence, I find DISH is entitled to a presumption that there is a nexus between the objective evidence of non-obvious DISH proffers and the claimed invention. I find respondents have not rebutted that presumption.

... [analyzing evidence of objective indicia of nonobviousness]

Having considered the claimed invention, the prior art, the objective indicia of non-obviousness, and the record as a whole, ***I find that the objective indicia support a determination of non-obviousness. Thus, even if a hindsight combination of the prior art***

references identified by respondents would comprise every element of the claimed invention, I find that the objective evidence indicates a person of skill in the art at the time of the invention would not have found the invention to be obvious in view of those references.

Final ID at 213-17 (internal citations omitted, emphasis added); *see* Comm’n Op. at 2 (adopting same).

53. Setting aside the merits of the Aylo IPR’s, I simply do not see how any “substantial question” of invalidity could be raised by a series of IPR’s that fail to address or even acknowledge the ITC’s prior determination concerning secondary considerations, which I understand potentially foreclose the statutory basis for each challenge. It is therefore my opinion that the evidence and arguments attached to Dr. Houh’s declaration do not show a substantial question as to the invalidity of the Asserted Claims for the threshold reason that they fail to acknowledge or address the ITC determination.

B. Ogdon

54. I have reviewed U.S. Patent No. 6,161,137 to Ogdon *et al.* (“Ogdon”), as well as Dr. Houh’s Declaration, and the claim charts and IPR filings cited therein. I understand that Petitioner in the Aylo IPR’s relies exclusively on Ogdon as the primary reference in each and every alleged ground of invalidity. But Ogdon itself has numerous deficiencies that Dr. Hough fails to address in any purported obviousness combination. In fact, Ogdon directly teaches away from video streaming as claimed by the ABR Patents. Moreover, as discussed above, I understand Dr. Houh’s obviousness theories are barred by evidence of secondary considerations of non-obviousness.

1. Ogdon is Not Analogous Art; It Does not Disclose or Suggest, and Teaches Away from, Video Streaming

55. Ogdon does not disclose or suggest video streaming as claimed by the ABR Patents and, in fact, expressly *teaches away* from it. Specifically, Ogdon recognizes that “telepresentations”

such as videoconferences were problematic due to their “substantial expense” and the requirement for “specialized video conferencing systems with high data transmission lines” and specialized facilities with their incumbent high establishment and maintenance costs. Ogdon at 1:19-48. A POSITA would recognize that the dominant factor driving Ogdon’s conclusions was the requirement for supporting a video presentation within the videoconferencing or videoconferencing-like system where video was a media treated and managed separately from any corresponding audio (*i.e.*, “video” in Ogdon was pure video and did not include audio, *see, e.g.*, Ogdon at 3:64 – 4:10). Having personal familiarity with the systems and environments Ogdon discloses for telepresentations throughout the 1990s, I recognize and appreciate the issues Ogdon raises.⁴

56. Ogdon appears to claim priority to a 1997 provisional application, Ogdon at 1:4-6. As of 1997, Ogdon correctly argues that the dominant approach to live, interactive “telepresentations” such as videoconferences of requiring specialized physical locations (rooms) with specialized equipment and connectivity to remote sites (rooms) had “numerous drawbacks.” *Id.* at 1:19-42.

⁵ Instead, Ogdon discloses an alternative that uses a combination of the telephone network or the Internet for delivery of audio and uses the Internet for its conception of video. However, audio and video are provided through “different communication channels.” *Id.* at 3:64–4:1.

⁴ Throughout the 1990s, UNC generally, and the Department of Computer Science specifically, maintained a set of facilities for “distance learning” and remote conferences that supported live, 2-way audio and video streaming using dedicated telecommunications circuits to partner sites at Duke University and North Carolina State University (among others).

⁵ Ogdon does not disclose the word “stream.” However, Ogdon does use the term “video conference” in two places (Ogdon at 1:27-29, 9:9-15) in distinguishing from its disclosed approach of merely “viewing the presentation.” *Id.* at 9:9-15.

57. Specifically, Ogdon recognizes that most of the overhead⁶ required to receive and process video at a client could be obviated by distributing only the “presentation materials” (*i.e.*, a series of multimedia “segments” or “slides” in a slide deck) instead of distributing the live conference video in a stream. While distribution of presentation materials alone is less desirable than live video, Ogdon teaches that its proposed solution is beneficial because it “allow[s] individuals to access and/or participate in a presentation using standard telephony and Internet network connections found in most offices and many homes.” Ogdon at 1:43-48. Conference participants therefore would not need to drive to dedicated “telepresentation” or “video conference” centers to attend a remotely hosted meeting—they could simply view basic “presentation materials” over a low bandwidth connection and listen in to a teleconference bridge in lieu of streaming video of the meeting.

58. By contrast, the ABR Patents are directed to **video streaming** (*i.e.*, distributing a continuous live video feed (or large pre-recorded video content) for consumption over a network where video in this context is understood to be a single integrated stream of moving images and the corresponding audio. ’564 Patent at [Title], [Abstract]. While the ABR Patents also recognize the limitations of network infrastructure and client hardware, unlike Ogdon, the ABR Patents do not

⁶ Ogdon acknowledges that “telepresentations” and “video conference[es]” were growing in popularity in the 1997 time frame of Ogdon’s earliest claimed priority date. Ogdon at 1:19-23. These technologies were early ancestors of today’s “Zoom” meetings, where a live video feed, typically of a room, would be captured at one conference site and distributed to one or more remote sites via a network typically comprised of dedicated telecommunications circuits. Audio was typically disseminated via “conventional telephone” because it was very difficult to beat the quality of telephony audio, however it may also be distributed by a data network technology. *Id.* at 2: 4-6. However, these telepresentation video conferences required “expensive” and “specialized” video conferencing systems. *Id.* at 1:22-37. Indeed, a POSITA would have understood that video conference telepresentations would have been extremely limited and expensive in 1997, particularly in typical home or even office settings where network bandwidth and access to videoconferencing hardware was extremely limited. Ogdon confirms that the availability and use of such conferencing systems was limited. *Id.* at 1:30-42.

disclose a shortcut or alternative to streaming video. Instead, the ABR Patents disclose a solution that allows a large number of remote clients to each *stream video content*, at the same time even though each client may have different network resources available. Such different network resources may include different connection speeds and equipment. In other words, the ABR Patents disclose and claim a solution that enables video streaming for use in different applications, including videoconferencing through the transmission of a single integrated data stream.

59. The ABR Patents achieve this in part by creating sets of “streamlets,” or “multiple files,” from the original large content file or video feed, where a plurality of streamlets in each set are aligned by starting time and duration (typically a few seconds) but have different bitrates. By segmenting and encoding the streamlets, the claimed inventions allow for contiguous playback of the streamlets or files that collectively yields playback of the full content. The common alignment of the streamlets or files in each set allows a playback device to select one quality of streamlet or file from a particular set and, as needed, to adjust for changing bandwidth resources, to select a different quality of streamlet or file from the subsequent set. When the bandwidth of the user’s network is constrained, the client can select a lower bitrate to maintain playback continuity instead of stalling and re-buffering. By using streamlets or multiple files, the claimed inventions eliminate the need for users to download the full content file before beginning playback, thereby offering instantaneous viewing along with the ability to fast forward, rewind, direct seek, and browse multiple streams. Additionally, segmenting the media into streamlets or multiple sets of files enables users to retrieve and enjoy content at the most appropriate bitrate possible as the media is streamed, thereby reducing latency, providing the highest quality video playout that is sustainable, and thereby improving the reliability and efficiency of computer devices used to stream content data over a network. It is also well suited for live stream playback.

60. Ogdon’s “presentation” is not a video stream as that term is used in the ABR Patents, and Ogdon’s presentation “segments” are not “streamlets” as disclosed and claimed in the ABR Patents. Instead, Ogdon’s “segments” correspond to slides in a powerpoint slide show—they are discrete, stand-alone pieces of content that accompany a presenter-led performance of a presentation.⁷ To be sure, while the individual slides (or “segments”) may include multimedia content, including “video” (*i.e.*, a “single image [] complex combination of images, audio, animation and video ... [and] interactive questions that audience members answer *e.g.*, by clicking on their display screens” (*id.* 3:4-10)), Ogdon’s “segments” are distinguishable from “streamlets” because they are not portions of any continuous larger video or live video feed. Indeed, Ogdon does not disclose streaming any larger videos with integrated audio like “movies,” “sports events,” “music,” “concert,” or any of the other streaming content disclosed by the ABR Patents. *See* ’564 Patent at 6:35-39. In fact, the phrases “video stream” and “video streaming” does not appear in Ogdon; nor does “compression” or even “live video.” Ogdon’s only disclosure of distributing “live” content involves distributing “live voice audio” through a “phone bridge,” *i.e.*, a telephone network. Ogdon at 4:48-55. Similarly, Ogdon’s only disclosure of “stream” is in relation to an “audio file or stream.” *Id.* at 13:12-19, 27:17-19, 28:31. Furthermore, Ogdon describes this audio as providing “substantial continuity in presentation regardless of whether a portion of the video for the presentation is displayed or not.” *Id.* at 24:28-39. A POSITA would recognize Ogdon’s disclosures as simply affirming the well-known principle that in videoconferencing applications, audio carries the vast majority of the information content and as such care is always taken to

⁷ Ogdon refers to these discrete pieces of content as “segments.”

optimize audio delivery. This explains why Ogdon explicitly separates and separately manages audio and video.⁸

61. Ogdon also discloses that a presentation may also have “alternative segments” that can be retrieved “depending on, for example, performance characteristics of network transmissions.” Ogdon at 3:1-2, 36-38; *see also id.* at 3:38-44. However, Ogdon does not disclose these “alternative segments” as a solution to enable streaming video as that term is used in the ABR Patents. Instead, Ogdon discloses alternative segments as static images or slides that enable a participant to participate in a presentation or meeting by viewing at least some limited set of presentation content. For example, Ogdon discloses that a presentation may include “first and second ordered segments” in “full animation,” while “alternative segments” are provided in “slide show format.” Ogdon at 3:44-57. However, Ogdon’s alternative segments are “slides,” not “streamlets” or “multiple files” of a video. They are not copies of a larger integrated audio/video encoded at a different bitrate and do not contain video data. Instead, they are static slide presentation content. While they may be presented “synchronously” in place of the default “full animation” segments, the alternative slides are not aligned in time and duration with the video animation segments because they are not videos—they are static slides.⁹

⁸ An award-winning paper I authored in 1994 showed the importance and utility of separating and separately managing audio and video data transmission in videoconferencing over computer networks (*see*, “Transport and Display Mechanisms For Multimedia Conferencing Across Packet-Switched Networks,” K. Jeffay, D.L. Stone, F.D. Smith, Computer Networks and ISDN Systems, Volume 26, Number 10, July 1994, pages 1281-1304. Republished in “A Guided Tour of Multimedia Systems and Applications,” B. Furht, M. Milenkovic, editors, IEEE Computer Society Press, 1995).

⁹ Ogdon’s claim 9 discusses conveying either a “first video portion” or a “second video portion” based on data transmission rates; however, that claim states that the second video portion is “different” from the first video portion and only “corresponds in content” with the first video portion, rather than conveying the identical content at a different encoded bitrate. In the context of the Ogdon specification, a POSITA would understand this claim does not disclose time-

62. To further illustrate Ogdon's lack of relevance to video streaming, consider Ogdon's disclosure that "collections of presentation elements can also be resident at the client node 56 since some portions of a presentation can also be distributed on CD-ROMs. Accordingly, step 524 of FIG. 2C (as well as other steps in the flowchart of FIG. 2) also may access a CD-ROM drive or other transportable storage media for various portions of a presentation." Ogdon at 23:23-35. A POSITA would not understand Ogdon's "presentations" to be a video stream at all, let alone a stream of video as that term is used in the ABR Patents. Instead, Ogdon's presentations comprise a series of various images, slides, and multimedia content accessible from various locations (including locally via CD) and pre-assembled in "scripts" by a presentation leader for synchronizing relevant material with presenter audio.

63. In sum, Ogdon is not analogous art to the field of video streaming that the ABR Patents are directed to because Ogdon separates and separately manages audio and video, recognizes problems with video only delivery, and then teaches replacement solutions that circumvent the need to perform *any* video delivery. Ogdon also does not propose a solution to the problem faced by the inventors of the ABR Patents because Ogdon identifies a related problem (*i.e.*, bandwidth limitations on video-only delivery) and then proposes other *non-streaming* solutions for facilitating a meeting between remote presentation participants. Merely recognizing the futility of trying to solve a problem (*i.e.*, video delivery via limited network infrastructure) and then posing work-arounds that do not actually resolve that problem (*i.e.*, distributing presentation materials in lieu of video) does not present information germane to solving the underlying problem.

aligned copies of the identical underlying content encoded at different bitrates as the claims require, but rather involves conveying alternative content to the full-bandwidth video. In addition, as noted, Ogdon's conception of video is not that claimed in the ABR Patents.

64. Furthermore, the teachings of Ogdon described above, particularly the identification of problems and alternative solutions, provide further objective indicia of non-obviousness. The description of such alternative solutions separate from the one presented in the ABR Patents demonstrate both that there was a long-felt need for a good solution for video streaming which went unsatisfied and that experts at the time were skeptical that it was possible to have a pure video streaming solution where one had to consider bandwidth limitations over a network.

2. Ogdon Fails to Disclose Encoding a Video into a Plurality of Streams

65. Each Asserted Claim recites encoding a video in “different bitrates” to create “a plurality of streams” including low, medium, and high quality streams. (the “Plurality of Streams” limitations).

Asserted Claim	Limitation
'554 Patent Claim 1	the live event video encoded at a plurality of different bitrates creating a plurality of streams including a low quality stream, a medium quality stream, and a high quality stream,
'554 Patent Claim 16	wherein the live event video is encoded at a plurality of different bitrates to create a plurality of streams including at least a low quality stream, a medium quality stream, and a high quality stream
'138 Patent Claim 1	wherein the video is digitally encoded at a plurality of different bitrates creating a plurality of streams including a low quality stream, a medium quality stream, and a high quality stream,
'138 Patent Claim 14	wherein the video is encoded at a plurality of different bitrates to create a plurality of streams including at least a low quality stream, a medium quality stream, and a high quality stream
'798 Patent Claim 1	at least one storage device storing digital content, the digital content encoded at a plurality of different bit rates creating a plurality of streams including a first bit rate stream, a second bit rate stream, and a third bit rate stream
'772 Patent Claim 1	streaming by a media player operating on an end user station a video from a set of one or more servers, wherein each of a plurality of different copies of the video encoded at different bit rates is stored as multiple files on the set of servers,
'564 Patent Claim 1	wherein multiple different copies of the video encoded at different bit rates are stored on the video server as multiple sets of files,
'564 Patent Claim 8	wherein multiple different copies of the video encoded at different bit rates are stored as multiple sets of files on the server

'680 Patent Claim 1	wherein the video encoded at a plurality of different bitrates creating a plurality of streams including a low quality stream, a medium quality stream, and a high quality stream,
'680 Patent Claim 22	storing, by the one or more servers, one or more virtual timelines corresponding to a plurality of streams including a low quality stream, a medium quality stream, and a high quality stream, wherein the low quality stream, the medium quality stream, and the high quality stream each comprise a group of streamlets encoded at a respective one of a plurality of different bitrates

66. Dr. Houh relies primarily on Ogdon's disclosure of "one or more presentation versions" and/or "variations" as allegedly disclosing this limitation across all Asserted Claims. *See, e.g., '772 Patent Houh Dec. at 31-34.* However, neither concept discloses encoding a video (of any kind) at a plurality of different bitrates to create different copies of a video as claimed and there is no disclosure in Ogdon of video consisting of integrated audio and video data as required for video streaming in the ABR Patents.

67. Specifically, Ogdon explains that "each [presentation] **version** may be for a different group of audience members such as a group for *Japanese speaking audience members, or audience members affiliated with a particular organization.*" *Ogdon* 2:30-35. Here, Ogdon contemplates that different presentation variations, or versions, refer to different content, customized based on language preferences of the audience or organization affiliation. "For example, a first variation might be directed to the customers of a corporation, another variation directed to the sales representatives of the corporation, and yet another directed to the investors of the corporation." *Id* at 25:38-42. Thus, Ogdon teaches creating separate presentation materials as "versions" or "variations" of content that are customized to a given audience. Customized versions of presentation materials do not disclose or suggest the claimed copies of a video with integrated audio encoded at different bitrates at least because the underlying content in each presentation is different. Indeed, the customized versions are not **copies** of a source video at all. Moreover, Ogdon

is silent regarding encoded customized versions at different bitrates. Ogdon never uses the word “bitrate” and has no discussion or even mention of video encoding.

68. Dr. Houh also relies on the following portion of Ogdon as allegedly disclosing this limitation:

Ogdon discloses that each copy of the video presentation is made up of multiple files (“segments”) and each segment has alternative bitrate versions for the client nodes to choose from. Ogdon uses the term “segment” to mean a portion of the video presentation and explains that “one or more” segments may be grouped into an “element,” “collection,” or “subcollection.” *E.g.*, EX1004, 2:36-49, 4:29-36, 18:40-45, 22:10-29. Ogdon explains that this is “an indivisible portion” of the video presentation (EX1004, 2:36- 42). and that a segment can be a file. EX1004, 28:17-18 (“presentation data segment file”). Thus, Ogdon discloses that the video presentation is made up of a set of segments as files.

’772 IPR Houh dec. at ¶72. Here, Ogdon discloses creating multiple presentation versions that each include “segments (subcollections)” of content. Ogdon gives two examples of segments. First, Ogdon discloses a configuration where the segments of one version contain “real time or animated video” and segments for another version contain “graphic slides.” Of course, “graphic slides” are not a copy of a video encoded at a different bitrate. Instead, “slides” are an alternative presentation format that would appear to include different content. Ogdon’s first example here does not disclose or suggest a plurality of streams as claimed. Indeed, Ogdon does not contain any disclosure of the required video streams encoded at different bitrates.

69. Second, Ogdon discloses a first presentation version including segments for an “audio presentation via a telephone” whereas an alternative segment collection is a “multimedia presentation element that is a combination of one or more of the following types of HTML multimedia data” Notably, none of the “versions” or “segments (subcollections)” disclosed here include a video stream. To be sure, Ogdon discloses that part of an alternative segment

subcollection can be HTML (Hyper-Text Markup Language), which is a content formatting language that dictates the display of text and graphical elements in a web browser, with the potential to display some “animation and video” content. However, nowhere does Ogdon disclose or suggest a video stream as that term is used in the ABR patents let alone a plurality of video streams encoded at different bitrates. At best, part of one “segment” may include “video,” while the other version includes only “an audio presentation.” Thus, Ogdon’s different “versions” refer to entirely different content, not copies of the same video encoded at a plurality of different bitrates.

70. Dr. Houh also relies on Ogdon’s vague disclosure of retrieving “presentation elements of *greater size* (and corresponding *enhanced quality of presentation*)” ’564 IPR Houh Dec. at ¶65 (quoting Ogdon at 2:30-66)), and that “presentation quality may be *enhanced* from *limited or no animation to full animation*.” (*id.* (quoting Ogdon at 22:44-47)). Based on this limited disclosure, Dr. Houh concludes “a POSITA would have understood the smaller size and lower transmission rate for ‘limited or no animation’ video to be encoded at a lower bitrate than the larger size, higher quality, and higher transmission rate requirement ‘full animation’ video, which would be encoded at a higher bitrate.”’564 IPR Houh Dec. at ¶65.

71. I disagree with Dr. Houh. When Ogdon refers to enhancing the “quality of presentation” in the above quoted portions, it is referring to different presentation elements with different content. For example, Ogdon discloses presenting one element as a graphic slide with still images (*i.e.*, “no animation”) as a low-bandwidth option and animated video (*i.e.*, “full animation”) as a high-bandwidth option. A POSITA would not understand a slide with images (“no animation”) to be a video much less a video stream with any associated bitrate. In the context of streaming video, “bitrate” is a measure of the number of bits that are processed per unit time when the video is displayed. Since the number of bits required to display a static image is a constant and once the

image is displayed it will remain on the screen until it is erased and no additional bits are required to be received or processed. As such, a POSITA would understand the concept of “bitrate” does not apply to images. Instead, the image would simply have a file size (*i.e.*, the number of bits required to represent the image on disk).

72. Moreover, Ogdon’s reference to elements with “limited ... animation” would have no generally accepted meaning to a POSITA and I note that Dr. Houh also fails to render an opinion about what exactly that term would mean to a POSITA. Ogdon provides no explanation. For example, Ogdon does not disclose any process of, for example, removing certain frames of a full animation to produce a more “limited animation” version of the presentation element. After combing the specification, “limited animation” could possibly refer to “GIF’s” (Ogdon at 5:32-34) which is a standard media type for bitmap images that can be used to display short, looping animations. But Ogdon is also silent on encoding any larger high quality source video into a GIF. Moreover, the looping animation of GIF’s is entirely unsuitable for video streaming (imagine if each 3-second clip of a video stream ran in a loop until the next 3-second clip was available).

73. As a broader point, Ogdon references the terms “animation” and “video” throughout its disclosure, sometimes combining the two terms into a single concept and sometimes distinguishing animation from video. *See e.g.*, Ogdon at 2:64 (“animated video”), 3:5 (“animation or video”). When taken in context with Ogdon’s other teachings¹⁰ and the state of the art as of the presumed 1997 priority date of Ogdon,¹¹ Ogdon’s use of the term “video” strongly suggests to a POSITA that “video”-based presentation segments actually refer to an animated series of fixed images

¹⁰ Recall that Ogdon begins by outlining challenges in conducting “telepresentations” because they require special “*video* conference” equipment for streaming video and higher bandwidth than was readily available for streaming videos at the time. *See supra* §VI.B.1.

¹¹ For example, RealNetworks, a recognized early pioneer in the development of multimedia players, only introduced RealVideo in 1997.

shown sequentially in a loop without audio, like the disclosed “GIF animations.” Ogdon at 5:33. For example, a “GIF” might include a series of 3 images that are displayed sequentially in a loop. I recall from my personal experience, that the use of GIFs to show animated “videos” on the (relatively new) world wide web in 1997 was popular.

74. This is quite different from what the inventors of the ABR Patents considered to be a “video” for adaptive bitrate streaming. “Videos” in the ABR Patents integrated audio and video, were much larger in duration (*e.g.*, movies, sporting events, television programs, etc.), suitable for segmenting into smaller streamlets for delivery to a client, and capable of being encoded at different bitrates. Nothing in Ogdon discloses streaming such a “video” or creating a plurality of copies of a source video encoded at different bitrates.

3. Ogdon Fails to Disclose Streamlets Corresponding to the Same Portion of a Video Across Streams

75. Each Asserted Claim recites encoding each video stream into a plurality of streamlets or files corresponding to portions of the video (the “Corresponding Streamlets” limitations).

Asserted Claim	Limitation
'554 Patent Claim 1	wherein the first streamlet of each of the groups of streamlets has the same first duration and encodes the same first portion of the live event video in each of the low quality stream, the medium quality stream, and the high quality stream, and wherein the first streamlet of the low quality stream encodes the same first portion of the live event video at a different bitrate than the first streamlet of the high quality stream and the first streamlet of the medium quality stream.
'554 Patent Claim 16	16[c][4] each of the low quality stream, the medium quality stream, and the high quality stream comprising a group of streamlets encoded at the same respective one of the different bitrates, each group comprising at least first and second streamlets, each of the streamlets corresponding to a portion of the live event video
'138 Patent Claim 1	wherein the first streamlet of each of the groups of streamlets has the same first duration and encodes the same first portion of the video in each of the low quality stream, the medium quality stream, and the high quality stream, and wherein the first streamlet of the low quality stream encodes the same first portion of the video at a lower bitrate

	than the first streamlet of the high quality stream and the first streamlet of the medium quality stream.
'138 Patent Claim 14	each of the low quality stream, the medium quality stream, and the high quality stream comprising a group of streamlets encoded at the same respective one of the different bitrates, each group comprising at least first and second streamlets, each of the streamlets corresponding to a portion of the video
'798 Patent Claim 1	wherein the first bit rate stream, the second bit rate stream, and the third bit rate stream each comprise a group of streamlets encoded at a respective one of the plurality of different bit rates, each group of streamlets comprising at least first and second streamlets, each of the streamlets corresponding to a portion of the digital content
'772 Patent Claim 1	wherein each of the multiple files yields a different portion of the video on playback, wherein the multiple files across the different copies yield the same portions of the video on playback
'564 Patent Claim 1	1[A][4]-[5] wherein each of the files yields a different portion of the video on playback, wherein the files across the different copies yield the same portions of the video on playback
'564 Patent Claim 8	wherein each of the files yields a different portion of the video on playback, wherein the files across the different copies yield the same portions of the video on playback,
'680 Patent Claim 1	wherein the first streamlet of each of the groups of streamlets has the same first duration and encodes the same first portion of the video in each of the low quality stream, the medium quality stream, and the high quality stream, and wherein the first streamlet of the low quality stream encodes the same first portion of the video at a different bitrate than the first streamlet of the high quality stream and the first streamlet of the medium quality stream.
'680 Patent Claim 22	wherein the low quality stream, the medium quality stream, and the high quality stream each comprise a group of streamlets encoded at a respective one of a plurality of different bitrates, each group comprising at least first and second streamlets, each of the streamlets corresponding to a portion of the video

76. Dr. Houh relies on Ogdon's disclosure of a "subcollection of segments (i.e., presentation elements) *ordered according to their presentation sequence*" as allegedly disclosing these limitations.'564 IPR Houh Dec. at ¶67 (quoting Ogdon at 2:44-45, 3:46-47) (emphasis added). Specifically, Ogdon's "presentation sequence" is a list of presentation elements (*i.e.*, slides) that is assembled by the presenter in a "script," and can be changed on the fly (*i.e.*, "in the middle of a presentation"). Ogdon at 3:16-18. Ogdon's "presentation sequence" does not disclose or suggest

encoding a source video or live video feed into a plurality of streams, each having streamlets or files that yield different portions of the source video. This distinction from what is claimed in the ABR patents is critical because the patents are meant to provide seamless streaming of a video, in part by adapting video quality according to present and changing network conditions. Without the segments corresponding to the same portion of the video, the content would jump in accordance with file differences.

77. Dr. Houh also relies on Ogdon’s disclosure that “corresponding segments (or subcollections thereof) in different versions” are the same such that “any of the corresponding alternative segments (or subcollections thereof) from different versions can be presented as a replacement for another such corresponding segment (or subcollection) during the presentation.” ’772 IPR Houh Dec. at ¶76 (quoting Ogdon at 2:49-55) (emphasis omitted);’564 IPR Houh Dec. at ¶68 (citing same). However, as discussed above, Ogdon’s “presentation versions” are not a plurality of video streams encoded at different bitrates, nor are they copies of the same video, nor are they videos as that time is used in the ABR patents. Instead, Ogdon’s presentation versions are entirely different presentations in different formats, such as a collection of graphic slides for one version, presenter audio for another version, and animation for yet another version of the presentation.

78. While Ogdon does not explain what the “presentation duration” of each segment refers to, a POSITA would understand that it could define how long a segment is displayed to a user so that presenter audio synchronizes with the displayed content. *See, e.g.*, Ogdon at 4:36-40 (“the host node coordinates any accompanying audio portion for this segment so that the timing for the presentation of these audio and video portions of the segment(s) are synchronized.”); Ogdon at 3:1-10 (“an alternative segment (subcollection) may be a multimedia presentation element that is

a combination of one or more of the following types of HTML multimedia data: audio, images, animation or video, wherein such a multimedia element *plays over a set period of time and can be as simple as a single image or as complex as a combination of images*, audio, animation and video.” (emphasis added)). A POSITA would therefore understand the “duration” of the segment to include a presenter-defined term that refers to the amount of time the content is relevant to the corresponding presenter audio. “Presentation duration” does not disclose or suggest that segments across versions are streamlets of video that each have the same playback duration and correspond to the same portion of a source video.

79. Dr. Hough’s reliance on Ogdon’s disclosure of using alternative segments as “replacement[s] for another such corresponding segment” as meeting these limitations is also flawed. ’564 IPR Houh Dec. at ¶68. Specifically, “replacing” an animated segment with an alternative “graphic slide” segment as disclosed by Ogdon does not disclose or suggest video streamlets or files that yield the same portion of a larger source video as recited by these limitations because a “graphic slide” is simply not a video or a portion of a larger video.

80. Moreover, even if Ogdon were to disclose providing multiple alternative segments as *video segments* (it does not, and Dr. Houh fails to identify any such teaching or suggestion in Ogdon), such a disclosure would not disclose or suggest the claimed files or streamlets because Ogdon’s presentation segments are discrete media segments that are arranged by a presenter in advance to correspond to live presenter audio. These “segments” are slides used by a human presenter who is leading the change of slides and overall performance of the presentation. They are not sequential parts of a larger overall video stream such as a movie, sporting event, or television show. Thus, these “segments” are not “multiple files.” Ogdon’s “presentations” are incompatible with streaming a larger video, which the Asserted Claims of the ABR Patents expressly require.

4. Ogdon Fails to Disclose Client Selection of Streams

81. The Asserted Claims recite client selection of one of the video streams.

Asserted Claim	Limitation
'554 Patent Claim 4	receive at least one streamlet request over one or more internet connections from the one or more end user stations to retrieve the first streamlet storing a portion of the video, wherein the at least one streamlet request from the one or more end user stations includes a request for a currently selected first streamlet from one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the end user station to select a higher or lower bitrate version of the streams;
'554 Patent Claim 16	select a specific one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the end user station to select a higher or lower bitrate version of the streams;
'138 Patent Claim 6	receive at least one streamlet request over one or more internet connections from a respective one of the one or more end user stations to retrieve the first streamlet, wherein the at least one streamlet request from the respective one of the one or more end user stations includes a request for a currently selected first streamlet from one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the respective one of the one or more end user stations to select a higher or lower bitrate version of a respective one of the plurality of streams;
'138 Patent Claim 14	select a specific one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the end user station to select a higher or lower bitrate version of the streams;
'772 Patent Claim 1	requesting by the media player a plurality of sequential ones of the files of one of the copies from the set of servers over a plurality of Transmission Control Protocol (TCP) connections based on the time indexes; automatically requesting by the media player from the set of servers over the plurality of TCP connections subsequent portions of the video by requesting for each such portion one of the files from one of the copies dependent upon successive determinations by the media player to shift the playback quality to a higher or lower quality one of the different copies,
'564 Patent Claim 1	1[C][1] automatically requesting from the video server subsequent portions of the video by requesting for each such portion one of the files from one of the copies dependent upon successive determinations

	by the media player to shift the playback quality to a higher or lower quality one of the different copies
'564 Patent Claim 8	automatically requesting by the media player from the server subsequent portions of the video by requesting for each such portion one of the files from one of the copies dependent upon successive determinations by the media player to shift the playback quality to a higher or lower quality one of the different copies
'680 Patent Claim 7	receive at least one streamlet request over one or more internet connections from the one or more end user stations to retrieve the first streamlet storing a portion of the video, wherein the at least one streamlet request from the one or more end user stations includes a request for a currently selected first streamlet from one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the end user station to select a higher or lower bitrate version of the streams;
'680 Patent Claim 22	wherein the at least one streamlet request from the one or more end user stations includes a request for a currently selected first streamlet from one of the low quality stream, the medium quality stream, and the high quality stream based upon a determination by the end user station to select a higher or lower bitrate version of the video;

82. Ogdon fails to disclose or suggest client selection of a stream and/or client requests for individual files or streamlets for at least the reasons discussed above in §§VI.B.2 and VI.B.3 regarding Ogdon's failure to disclose a plurality of streams and files or streamlets of those streams. Since Ogdon fails to disclose a plurality of streams with corresponding streamlets at a server, Ogdon also fails to disclose or suggest requesting subsequent portions of a larger video by requesting, for each portion, files or streamlets from that stream.

83. Moreover, while Ogdon generally discloses benefits of "controlled client requests," the client requests are ultimately controlled by a "presentation leader" via a server or "host" that "transmit[s] presentation controlling commands." Ogdon at 3:30-38. For example, Ogdon discloses that a ***presentation leader*** controls which segment versions are available for each portion of the presentation, and that the presentation can manually change scripts in the middle of a presentation to require clients to view a different variation of the presentation. Ogdon at 3:16-23 ("Thus, multiple sequentially-ordered scripts can be created so that the leader can choose to change

scripts in the middle of a presentation based on user feedback. Accordingly, a presentation leader has the ability to *stop presentation of a particular script and its current subcollection of segments and change to a different subcollection of segments to be delivered to the audience*. Subsequently, the leader can then resume the initial script at any time.” (emphasis added)). Thus, Ogdon does not disclose “automatically requesting ... subsequent portions of the video” because retrieving each segment is dependent upon the server’s/leader’s instruction, which may change at any time to eliminate or replace the version being presented.

5. Ogdon Fails to Disclose or Suggest Streaming “Live” Video

84. Asserted Claim 16 of the ’554 Patent is directed to streaming a “live event video.” As discussed above, Ogdon does not disclose live video as that term is used in the ABR patents. Further, as discussed above in §VI.B.1, Ogdon teaches away from video streaming, and particularly streaming a “live event video” because Ogdon presupposes that clients available in “most offices and many homes” *cannot* receive a live video broadcast of a “telepresentation.”

85. Dr. Houh merely string-cites to a series of portions of Ogdon of dubious relevance to the live video limitation.’554 IPR; Houh Dec. at ¶59. At best, Dr. Houh relies on Ogdon’s disclosure of delivering presentation segments “while the presentation is being performed.” *Id.* (citing Ogdon at 6:2-5). But Ogdon’s presentation “segments” are pre-existing content—not live content. They are arranged ahead of time by a presentation leader and saved in a “script.” Thus, while segments are distributed to clients “while the presentation is being performed,” this just simply indicates that the *delivery* of segments occurs in real-time. There is nothing “live” about any videos contained in presentation segments.

6. Dr. Houh’s Obviousness Theories Fail

86. Dr. Houh relies on the combination of Ogdon and Shteyn as allegedly rendering obvious the Plurality of Streams and Corresponding Streamlets limitations of the '772 and '564 Patents.¹²

See, e.g., '564 IPR Houh Dec. ¶¶113-119. Specifically, Dr. Houh alleges that

A POSITA would have understood the above disclosures [of Shteyn] and from Ogdon's disclosures [] including disclosure of different versions of the same video presentation each having "different network requirements," to teach ***encoding the same raw uncompressed video content into two or more versions having different bitrates***. It was well-known to a POSITA at the time of the alleged Priority Date that streaming different versions of content at different quality levels depending on client network requirements meant encoding the content into different quality levels with different bitrates.

...

A POSITA would have been motivated to modify the apparatus of Ogdon ***such that each segment is stored as a file and thus the multiple different copies of the video encoded at different bit rates are stored as multiple sets of files***, to simplify streaming requests by requesting individual successive files rather than requiring the requests to include byte- or other ranges of the video.

Id. at ¶116 (emphasis added). Thus, Dr. Houh relies on Shteyn for the limited purposes of encoding Ogdon's "presentation segments" in two or more "versions" with different bitrates and storing the different versions as "files." *Id.*

87. As discussed above, Dr. Houh's obviousness theories fail for the threshold reason that Dr. Houh fails to rebut the strong evidence of non-obviousness found by the ITC. *See supra* § VI.A. Dr. Houh's combination theories can therefore be dismissed on this basis alone.

88. Moreover, I disagree with Dr. Houh's conclusion that this theory renders obvious the Plurality of Streams and Corresponding Streamlets limitations of the ABR Patents. Dr. Houh's

¹² Dr. Houh also asserts invalidity of non-asserted Claim 16 of the '798 Patent based on a ground involving Shteyn.

analysis fails because it overlooks Ogdon’s central flaw—namely that Ogdon’s presentation segments are not portions of a larger video file or live video feed.

89. Dr. Houh’s theory unquestionably relies on modifying the system of Ogdon. *Id.* (“A POSITA would have been motivated to ***modify the apparatus of Ogdon....***” (emphasis added)). However, nothing in Dr. Houh’s analysis demonstrates how or why a POSITA would modify Ogdon’s presentation distribution system to stream pre-recorded or live videos. Even assuming that the combination of Ogdon and Shteyn were properly motivated,¹³ Dr. Houh’s analysis would result in a modified version of Ogdon’s system that encodes each presentation element that happens to include video into multiple segment variations encoded in different bitrates. However, no aspect of the combination purports to change Ogdon’s presentation delivery system into a system for streaming video files with multiple sequential segments. More fundamentally, Dr. Hough ignores the fact that Ogdon and Shteyn have difference conceptions of video and replacing Ogdon’s video with Shetyn’s video would remove Ogdon’s ability to separately manage audio and video and thereby result in lower quality conferences/presentations.

90. For the ’798, ’138, ’680, and ’554 Patents, Dr. Houh contends that Ogdon in combination with Allen renders obvious the Plurality of Streams and Corresponding Streamlets limitations. However, Dr. Houh’s only relies on a single direct¹⁴ citation to Allen’s disclosure of encoding a stream in multiple bitrates.’554 IPR Houh Dec. at ¶63 (citing Allen 12:1-5). However, Dr. Houh’s reliance on Allen, Chou and SMIL suffer the same deficiencies as the combination with Shteyn

¹³ As discussed at length above, Ogdon “teaches away” from streaming such that a POSITA would not be motivated to modify Ogdon to achieve the inventions claimed by the ABR Patents.

¹⁴ Dr. Houh also cites to Allen at 25:43-44, but this does not appear to be a proper citation in Allen, and provides a “see also” cite to Allen at pages 3 and 5 for the proposition that video encoders were available for real-time encoding in a plurality of bit rates. ’554 IPR Houh Dec. at ¶¶63, 67.

discussed above. Namely, Dr. Houh does not rely on any theory or combination that would change Ogdon's primary purpose from a presentation distribution system to a streaming system. The combinations with Allen, Chou, and SMIL are therefore not properly motivated because Ogdon teaches away from the combinations and fails to disclose or suggest the Plurality of Streams and Corresponding Streamlets limitations.

91. Dr. Houh's alternative combinations with SMIL 2.0 for the '798, '138, '680, and '554 Patents suffer the same fate because Dr. Hough does not propose a modification that would somehow convert Ogdon's presentation distribution system to a streaming system. *See, e.g.,* '554 IPR Houh Dec. at ¶¶144-145. Moreover, the example from SMIL 2.0 relied upon by Dr. Houh also fails to disclose or suggest encoded video because it shows a sequence of ".jpg" elements and, alternatively, system color representations. *Id.* at ¶145 (citing SMIL 2.0 at 273). Neither of these formats are video data. In addition, no larger video stream is disclosed or suggested by the seemingly random sequence of purple, then yellow, and then red images.

92. In sum, Dr. Houh failed to address secondary considerations of non-obviousness; Ogdon expressly teaches away from video streaming; and the proposed combinations with Shteyn, Allen, and SMIL 2.0 fail to address Ogdon's central flaw or otherwise disclose or suggest the Plurality of Streams and Corresponding Streamlets limitations they are relied upon to teach.

C. Shteyn

93. I understand that Dr. Houh alleges that U.S. Patent No. 7,529,806 to Shteyn ("Shteyn") anticipates each Asserted Claim. '772 IPR Houh Dec. at ¶¶56, 72, 81, 88, 101. I highlight the fact that Dr. Houh does not contend that Shteyn renders any Asserted Claims obvious. Instead, his contentions are limited to invalidity by "anticipation." *Id.* However, Dr. Houh fails to explain what the teachings of Shteyn mean to a POSITA or how any single embodiment in Shteyn discloses each and every limitation of the Asserted Claims. I

94. However, as a general matter, the Shteyn reference is quite limited in its disclosures. It includes just a *single page* of disclosure of “preferred embodiments”¹⁵ and does not provide enabling support for many of the general concepts disclosed therein. Although Dr. Houh does not provide sufficient substantive analysis, Dr. Houh appears to leverage the vagueness of the general, limited disclosures of Shteyn as allegedly anticipating the concepts recited by the Asserted Claims (and disclosed in detail in the ABR Patents). This is improper. In my opinion Shteyn is not an enabling disclosure because it does not explain to a POSITA how to make and use the concepts it discloses without undue experimentation. Moreover, Dr. Houh fails to offer any interpretation, opinion, or analysis of Shteyn. It is therefore my opinion that Shteyn does not and cannot anticipate the Asserted Claims because it does not enable a POSITA to practice the claimed invention.

1. Background of Shteyn

95. Shteyn was filed on November 4, 1999 and is directed to “emulating streaming” by downloading sequential segments of a partitioned “electronic file, e.g., an MP3 file” at a client. *Shteyn*, (54) Title, (57) Abstract. Shteyn aims to address two shortcomings of prior art content delivery solutions. First, Shteyn recognizes that “current major technologies for delivering digital content” involve “proprietary” formats that utilize non-public encoding/decoding methods. *Id.*, at 1:18-32. “This approach mostly excludes third parties from developing custom server software (e.g., advertising services) and/or client applications (UI, special effects, etc.).” *Id.*, at 1:38-41.

96. Second, Shteyn addresses “[a]nother known method” for delivering digital content which it refers to as the “complete download approach.” *Id.*, at 1:42-51. Under this approach, a complete content file is downloaded from a remote computer and then subsequently played out on the client.

¹⁵ Other portions of the Shteyn reference provide “background” and “summary” information (Shteyn at 1:1-2:38), as well as information about other publication that are incorporated by reference but are not expressly tied to in recognizable way to the present disclosure of Shteyn (Shteyn at 4:45-5:42).

Id., at 1:42-44. While this method of delivering digital content allows “client and server software/hardware manufacturers [] to develop their own solutions/products”—a feature Shteyn refers to as an “open architecture”—the “complete download approach” is not without drawbacks. *Id.*, at 1:47, 51-58. For example, “[a] major problem with the complete download approach is the inherent latency: there is a delay between the beginning of the download and the start of the play-out.” *Id.*, at 1:51-53. If it takes 10 minutes or more to download a complete digital file for a song (remember, Shteyn was filed in 1999, when dialup Internet access was pervasive and bandwidth for most consumer-grade Internet connected devices was incredibly sparse), the user could not start listening to the song until the entire file was downloaded. That inherent delay is “latency,” and Shteyn teaches that it “is particularly undesirable in consumer electronics systems, where perceived delay is detrimental to market acceptance of an open architecture” solution. *Id.*, at 1:55-58.

97. Shteyn therefore sets out to address these problems by “provid[ing] an open architecture solution for content delivery in a download approach that allows for a low or negligible play-out latency.” *Id.* 1:62-64. Specifically, Shteyn is almost exclusively directed to the goal of reducing play-out *latency* (*i.e.*, the time it takes for content to begin playing at a remote client). This is different from the objects of the invention claimed by the ABR Patents, which achieve *sustained* streaming and playback at a client (without a singular focus on initial playout latency) through adaptive bitrate switching.

2. Shteyn Fails to Disclose Client Generation of a “Factor”

98. The Asserted Claims of the ’772 and ’564 Patents each recite limitations directed to “repeatedly generating a set of one or more factors indicative of the current ability to sustain the streaming of the video” at a client. Shteyn fails to disclose the generation of any such factor, let alone generating such a factor on the client.

99. Shteyn discloses testing network conditions in the context of determining how to split a content file into parts to enable low-latency client playback. Shteyn at 1:65-67. Specifically, Shteyn discloses that the particular “size of the individual part can be determined by the communications bandwidth, e.g., through pinging for a latency-check.” *Id.* at 2:1-3. This network “ping” test is the sole disclosure of determining “bandwidth” in Shteyn.¹⁶ As an initial matter, I note that a POSITA would understand that a “ping” is simply a small message sent to a destination that will simply return the message back to the sender. The sender records the difference between the time the message is received from the destination and the time it was sent to the destination and reports this as the latency of the of the network path to the destination. A POSITA would understand that a “ping” does not measure bandwidth.

100. Critically, though, Shteyn discloses that the “ping” test is performed at *servers* that store content, not at the *client*. Specifically, the latency-check is performed as the server is determining how to “*split*” a content file to enable low-latency playback at clients. Shteyn’s clients are not involved in any content “splitting.” Thus, a POSITA would understand Shteyn’s “ping” test to only be performed at a server. Accordingly, Shteyn does not disclose generating a “factor” at a client as claimed.

101. Moreover, the “ping” test disclosed by Shteyn does not test the ability to “sustain” any streaming activity and therefore does not disclose the claimed “factor” recited in the Asserted Claims of the ’772 and ’564 Patents. Instead, a POSITA would understand that a “ping” simply

¹⁶ Shteyn also generally references that “the client apparatus could select a next segment in a different format for the same content to adapt to changing circumstances, e.g., lower bandwidth due to network congestion.” Shteyn at 4:20-23. But Shteyn does not disclose how or where the system would assess “lower bandwidth” conditions. Moreover, Shteyn does not disclose that any such assessment would consider the ability to “sustain” streaming as claimed. And as discussed above, Dr. Houh does not provide any supporting analysis that would indicate what a POSITA would understand Shteyn to teach here.

tests the time it takes for a signal to reach another node on the network. Indeed, Shteyn refers to the “ping” as a “latency-check,” and “latency” in computing refers to network delay.¹⁷ As discussed in §VI.C.1 above, Shteyn is primarily concerned with reducing the “latency” or delay before playback starts—Shteyn does not disclose assessing playback time, the rate of data transfer, or any network performance indicators that could be used to determine the “ability to *sustain* streaming” after it begins as claimed. Thus, Shteyn also fails to disclose generating a “factor” for this reason.

3. Shteyn Fails to Disclose Encoding Streams at Different Bitrates

102. Each Asserted Claim recites a limitation that requires encoding a plurality of streams at different bitrates (*see supra* §VI.B.2 defining the “Plurality of Streams” limitations). As discussed in my opinion above regarding Ogdon, a POSITA would understand “bit rate” (or “bitrate”) in the context of the ABR Patents to refer to the *volume* or number of bits that are played *per second* when the video is displayed/played-out. Shteyn fails to use the term “bitrate” in relation to non-adaptive bit-rate streaming and instead only mentions it during a discussion of variable bit-rate encoding. Accordingly, Shteyn fails to disclose the Plurality of Streams limitations.

103. Dr. Houh asserts that:

Shteyn discloses “[a]n electronic file . . . **partitioned into a sequence of segments** at the **server** side.” EX1016, Abstract. Each segment (or “file”) is stored in the server as different encoded bitrate (bandwidth) and size (bytes), such as “part1” and “part1 alt.” *Id.*, 3:44-56 (“segment labeled ‘**part1**’ is in a preferred format . . . An alternative first part is labeled ‘**part1 alt**’ having a **different length**” **in bytes** and a “different bandwidth requirement.”), 1:65-2:7 (“the **content file is split into multiple parts** . . . location of the whole file as well as of its parts **at the server**”)

¹⁷ Shteyn appears to equate a “ping” to a “bandwidth” test. Shteyn at 2:1-3. However, a “ping” is not a measure of “bandwidth” because it does not measure a rate of data transfer per unit of time. Instead, a “ping” merely measures the amount of time it takes for a packet to reach a destination and be returned to the sender without considering the amount of data being transmitted.

'772 IPR Houh Dec. at ¶123. But this discussion in Shteyn about determining a file-size for video segments so as to realize a particular playout latency. It does not disclose encoding at a particular bitrate trying to match a bitrate to available bandwidth. While the concept of “bitrate” could be implicated in the time it takes to download a streamlet (download time is a function of the duration of a streamlet and its average size), Shteyn doesn’t disclose calculating download times to measure “bitrate” or encoding at any particular “bitrates” to achieve any particular download times.

104. Shteyn also discloses that “parts” can have an associated “length,” “format,” and “minimum bandwidth requirement” and there can be multiple alternative versions of each part. Shteyn at 3:44-56. But Shteyn discloses that the “bandwidth requirement” information is merely used to determine *how many parts to download before starting the playout*:

To this end the content file is split into multiple parts. Each part or segment requires a relatively short download time. Therefore, the play-out latency is determined by the download time of the first part. The size of the individual part can be determined by the communications bandwidth, e.g., through pinging for a latency-check.

...

Preferably, the information contains references to the file location as well as references to the locations of the parts. *The intended bandwidth information is associated with the parts. The client may make its own decisions regarding how many parts to download before the start of the play out (execution).*

Shteyn at 1:65-2:24. Thus, Shteyn’s “bandwidth information” does not disclose a “bitrate” of the content itself. Instead, it merely discloses an “intended bandwidth” that allows the client to determine “how many parts to download” before starting play out.

4. Shteyn Fails to Disclose Streaming “Live” Video

105. Claim 16 of the '554 Patent recites encoding a live event video at a plurality of bitrates to create a plurality of streams. Shteyn does not disclose streaming live video. Dr. Houh cites to various portions of Shteyn that generally disclose Shteyn’s disclosed process of using

control information to identify alternative segments and determining which file to retrieve based on system restraints. However, none of the cited portions of Shteyn disclose encoding or streaming a live event video as claimed.

VII. Conclusion


106. I reserve the right to add to or modify this expert report in the event that inaccuracies or omissions are discovered or if new or additional information is provided to me. I also reserve the right to make any such other changes as may be deemed necessary or appropriate.

107. I have not yet developed any demonstratives or illustrations to illustrate my opinions and the bases for those opinions; however, I reserve the right to present such demonstratives or illustrations at trial or a hearing.

108. I reserve the right to supplement and/or amend my analyses and opinions expressed herein in response to any new positions taken by MG Defendants' experts or witnesses in reports, declarations, affidavits, or testimonies whether in response to my report herein or in amendment to their previous reports and testimonies to the extent such positions are relevant to the subject matter I have addressed in my report herein.

109. I declare under I declare under penalty of perjury that the foregoing is true and correct.

Signature:
Kevin Jeffay, Ph.D.

A handwritten signature in blue ink, appearing to read 'K. Jeffay', written over a horizontal line.

Executed On

January 18, 2024 at Chapel Hill, NC

EXHIBIT A

CURRICULUM VITAE

Kevin Jeffay

University of North Carolina at Chapel Hill
Department of Computer Science
Chapel Hill, NC 27599-3175
(919) 590-6238
jeffay@cs.unc.edu
http://www.cs.unc.edu/~jeffay

Education

Ph.D. Computer Science, University of Washington.

Honors: IBM Graduate Fellowship.

M.Sc. Computer Science, University of Toronto.

Honors: University of Toronto Open Fellowship.

B.S. Mathematics with Highest Distinction, University of Illinois at Urbana-Champaign.

Honors: Phi Beta Kappa, James Scholar.

Academic Experience

Chairman, Department of Computer Science, University of North Carolina at Chapel Hill, Chapel Hill, NC, 2014–2022.

Gillian T. Cell Distinguished Professor in Computer Science, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, 2008–present.

S. S. Jones Distinguished Professor in Computer Science, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, 2001–2008.

S. S. Jones Distinguished Term Associate Professor of Computer Science, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, 2000.

Associate Professor, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, 1996–2000.

Visiting Professor, Carnegie Mellon University, School of Computer Science, Pittsburgh, PA, 1994.

Assistant Professor, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, 1989–1995.

Honors and Awards

Alumni Achievement Award, University of Washington, Department of Computer Science and Engineering, June 2015.

Excellence in Teaching Award, Computer Science Student Association, University of North Carolina at Chapel Hill, Department of Computer Science, May 2011.

Favorite Faculty Award, an award given by the Computer Science majors of the 2010 graduating class, May 2010.

Carolina Women's Center Women's Advocacy Award, University of North Carolina at Chapel Hill, Chapel Hill, NC, March 2010.

Gillian T. Cell Distinguished Professor in Computer Science, University of North Carolina at Chapel Hill, College of Arts and Sciences, Chapel Hill, NC, June 2008.

Outstanding Teaching Award, an award given by the Computer Science majors of the 2008 graduating class, May 2008.

Favorite Faculty Award, an award given by the Computer Science majors of the 2004 graduating class, May 2004.

Edward Kidder Graham Outstanding Faculty Award, an award given by the Class of 2004 and the General Alumni Association, April 2004.

Edward Kidder Graham Advisor of the Year Award, an award given by the Class of 2004 and the General Alumni Association, April 2004.

IEEE Service Award, Service to the IEEE Technical Committee on Real-Time Systems, December 2001.

Outstanding Teaching Award, Computer Science Student Association, University of North Carolina at Chapel Hill, Department of Computer Science, May 2000.

S. Shepard Jones Distinguished Term Professorship, University of North Carolina at Chapel Hill, College of Arts and Sciences, Chapel Hill, NC, March 1999.

Award Papers:

ACM SIGCOMM 2000, 2003.

IEEE International Symposium on High Assurance Systems Engineering 1999, 2000.

IEEE Real-Time Technology and Applications Symposium, 1995.

ACM International Conference on Multimedia, 1994.

Computer Networks and ISDN Systems, 1994.

International Workshop on Network and Operating System Support for Digital Audio and Video, 1991, 1992, 1993, 1997.

**Testifying
Experience
(Within Last
Four Years)**

WSOU Investment, LLC d/b/a Brazos Licensing and Development v. Salesforce.com, Inc., United States District Court for the Western District of Texas, Waco Division, Case No. 6:20-cv-01167-ADA.
Provided deposition testimony on behalf of defendant Salesforce.com, Inc.

Smart Path Connections, LLC, v. Nokia Corporation, et al., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:22-cv-00296-JRG-RSP.
Provided deposition testimony on behalf of defendant Nokia Corporation.

3G Licensing, S.A., Koninklijke KPN N.V., and Orange, S.A., v. HTC Corporation, United States District Court for the District of Delaware, Case Nos. 17-cv-83-GBW.
Provided trial testimony on behalf of defendant HTC Corporation.

In re: Inter Partes review of U.S. Patent Nos. 9,742,824 and 9,762,636.
Provided deposition testimony on behalf of Petitioners Amazon.com, Inc., Amazon Web Services, Inc., and Amazon.com Services LLC.

In re: Inter Partes review of U.S. Patent No. 9,878,243.
Provided deposition testimony on behalf of Petitioner DraftKings Inc.

Attila Csupo, et al., v. Google LLC, Superior Court of the State of California for the County of Santa Clara, Case No. 1:19CV352557.
Provided deposition testimony on behalf of defendant Google LLC.

WSOU Investment, LLC d/b/a Brazos Licensing and Development v. Dell Technologies Inc., Dell Inc., EMC Corporation, and VMware Inc., United States District Court for the Western District of Texas, Waco Division, Case Nos. 6:20-cv-00481 & -00486-ADA-DTG.
Provided deposition testimony on behalf of defendant VMware, Inc.

Voxer, Inc. and Voxer IP LLC v. Facebook, Inc. and Instagram LLC, United States District Court for the Western District of Texas, Austin Division, Case No. 1:20-cv-00655 [ADA].
Provided deposition and trial testimony on behalf of defendants Facebook, Inc., and Instagram LLC.

Koninklijke KPN N.V., v. Telefonaktiebolaget LM Ericsson and Ericsson, Inc., United States District

Court for the Eastern District of Texas, Marshall Division, Case No. 2:21-cv-113.

Provided trial testimony on behalf of defendants Telefonaktiebolaget LM Ericsson and Ericsson, Inc.

In re: Inter Partes review of U.S. Patent No. 7,783,788.

Provided deposition testimony on behalf of Petitioner Hewlett Packard Enterprise Company.

Egenera, Inc., v. Cisco Systems, Inc., United States District Court for the District of Massachusetts, Case No. 1:16-cv-11613-RGS.

Provided deposition and trial testimony on behalf of defendant Cisco Systems, Inc.

In re: Inter Partes review of U.S. Patent No. US 8,243,593.

Provided deposition testimony on behalf of Petitioner Cloudflare, Inc. and Sonicwall, Inc.

Flexiworld Technologies, Inc., v. Roku, Inc., United States District Court for the Western District of Texas, Waco Division, Case No. 6:20-cv-00819-ADA.

Provided deposition testimony on behalf of defendant Roku, Inc.

Longhorn HD, LLC v. Netscout Systems, Inc., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:20-cv-00349.

Provided deposition and trial testimony on behalf of defendant Netscout Systems, Inc.

ClearDoc, Inc. v. Riverside, Inc., United States District Court for the District of Delaware, Case No. 21-1422-RGA.

Provided deposition testimony on behalf of defendant RiversideFM, Inc.

AlterWAN, Inc. v. Amazon.com, Inc. and Amazon Web Services, Inc., United States District Court for the District of Delaware, Case No. 1:19-cv-01544-MN.

Provided deposition testimony on behalf of defendants Amazon.com, Inc. and Amazon Web Services, Inc.

In the Matter of: Certain Fitness Devices, Streaming Components Thereof, and Systems Containing Same, United States International Trade Commission, Investigation No. 337-TA-1265.

Provided deposition and trial testimony on behalf of complainants DISH DBS Corporation, DISH Technologies, LLC, and Sling TV LLC.

Huawei Technologies Co., Ltd., Huawei Device Co., Ltd., and Huawei Digital Technologies (Chengdu) Co., Ltd. v. Verizon Communications Inc., Cellco Partnership d/b/a Verizon Wireless, and Verizon Business Network Services, Inc., United States District Court for the Western District of Texas, Waco Division, Case No. 6:20-cv-00900.

Provided deposition testimony on behalf of defendants Verizon Communications Inc., *et al.*

In re: Inter Partes review of U.S. Patent No. US 7,450,507.

Provided deposition testimony on behalf of Petitioner NetApp, Inc.

In the Matter of: Certain Digital Video-Capable Devices and Components Thereof, United States International Trade Commission, Investigation No. 337-TA-1224.

Provided deposition and trial testimony on behalf of respondent Intel Corp.

Gigamon Inc., v. Apcon, Inc., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:19-cv-300.

Provided deposition and trial testimony on behalf of plaintiff Gigamon Inc.

In re: Inter Partes review of U.S. Patent Nos. 10,022,590, 10,322,315, & 10,486,026.

Provided deposition testimony on behalf of Petitioner Echelon Fitness Multimedia, Inc.

Citrix Systems, Inc., v. Workspot Inc., United States District Court for the District of Delaware, Case No. 18-588-GMS.

Provided deposition testimony on behalf of defendant Workspot Inc.

In the Matter of: Certain Audio Players and Controllers, Components Thereof, and Products Containing the Same, United States International Trade Commission, Investigation No. 337-TA-1191.

Provided deposition and trial testimony on behalf of respondent Alphabet Inc. and Google LLC.

Monarch Networking Solutions LLC, v. Cisco Systems, Inc., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:20-cv-00015-JRG.

Provided deposition testimony on behalf of defendant Cisco Systems, Inc.

Sprint Communications Company L.P., v. Cequel Communications, LLC d/b/a Suddenlink Communications, and CSC Holdings, LLC d/b/a/ Optimun-Cablevision, United States District Court for the District of Delaware, Case No. 18-1752-RGA.

Provided deposition testimony on behalf of defendants Suddenlink Communications and Optimun-Cablevision.

In re: Inter Partes review of U.S. Patent Nos. 7,710,978, 7,773,588 & 8,560,828.

Provided deposition testimony on behalf of Patent Owner directPacket Research, Inc.

Centripetal Networks, Inc., v. Cisco Systems, Inc., United States District Court for the Eastern District of Virginia, Norfolk Division, Case No. 2:18-cv-00094-HCM-LRL.

Provided deposition and trial testimony on behalf of defendant Cisco Systems, Inc.

In re: Inter Partes review of U.S. Patent Nos. 8,064,935 & 9,602,655.

Provided deposition testimony on behalf of Petitioner Resideo Technologies, Inc.

In the Matter of: Certain Wireless Communication Devices and Related Components Thereof, United States International Trade Commission, Investigation No. 337-TA-1180.

Provided deposition testimony on behalf of respondent HTC Corporation and HTC America, Inc.

Gummarus LLC, v. Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc., United States District Court for the Eastern District of Texas, Sherman Division, Case No. 4:19-cv-00251.

Provided deposition testimony on behalf of defendant Samsung Electronics Co., Ltd. and Samsung Electronics America, Inc.

Industry Experience

Consultant, Mälardalen University, Västerås, Sweden, 2001.

Member, Technical Advisory Board, Ganymede Software Inc., Research Triangle Park, NC, 1996–1999.

Consultant and Author, INTEROP Graduate Institute, SOFTBANK Inc., Foster City, CA, 1996–1997.

Consultant, National Science Foundation, Arlington, VA, 1995–2014.

Consultant, Monterey Technologies Inc., Cary, NC, 1994–1996.

Consultant, Hewlett-Packard Inc., Ink Jet Components Division, Corvallis, OR, 1993–1994.

Visiting Researcher, IBM T.J. Watson Research Center, Computer Systems Principles Group, Yorktown Heights, NY, 1986.

Consultant, Boeing Aerospace, Kent, WA, 1985–1986.

Software Engineer, R.R. Donnelley & Sons Company Inc., Technical Center, Chicago, IL, 1984.

Computer Programmer, US Army Corps of Engineers Construction Engineering Research Lab, Champaign, IL, 1981–1982.

Community Service

Consultant (unpaid), Children's Museum About the World (now called *Exploris*), Raleigh, NC, 1995–1996.

Professional Activities

Editor Associate Editor, *Real-Time Systems*, Kluwer Academic Publishers, The Netherlands, 2003-2012.
Editor in Chief, *Multimedia Systems*, ACM/Springer-Verlag, Heidelberg, Germany, 2000–2001.

Editorial Board *Journal of Multimedia Tools and Applications*, Kluwer Academic Publishers, 1994–1999.

Guest Editor *Computer Communications*, special issue on system support for multimedia computing, Volume 18, Number 10, October 1995.

Executive Committees ACM/SIGCOMM Internet Measurement Conference Steering Committee, 2005-2009.
College of Reviewers, Canada Research Chairs Program, 2004-2008.
Statistical and Applied Mathematical Sciences Institute (SAMSI) program on Network Modeling for the Internet, 2002-2004.
IEEE Technical Committee on Real-Time Systems, 2000-2007.
ACM Special Interest Group on Multimedia, 2000-2002.

Other Committees *IEEE Distinguished Visitors Program*, 2004–2007.

Conference Program Chair 11th International Workshop on Quality-of-Service (Co-Chair), Monterey, CA, June 2003.
Sixth IEEE Symposium on Computers and Communications, Hammamet, Tunisia, July 2001.
21st IEEE Real-Time Systems Symposium, Orlando, FL, November 2000.
Tenth International Workshop on Network and Operating System Support for Digital Audio and Video, Chapel Hill, NC, June 2000.
Seventh ACM International Conference on Multimedia (Co-Chair), Orlando, FL, November 1999.
SPIE/ACM Multimedia Computing and Networking 1999 (Co-Chair), San Jose, CA, January 1999.
Sixth ACM International Conference on Multimedia (Associate Chair), Bristol, UK, September 1998.
IEEE International Conference on Multimedia Computing Systems (Associate Chair), Austin, TX, June 1998.
SPIE/ACM Multimedia Computing and Networking 1998 (Co-Chair), San Jose, CA, January 1998.
SPIE/ACM Multimedia Computing and Networking 1997 (Associate Chair), San Jose, CA, February 1997.
IEEE Workshop on Resource Allocation Problems in Multimedia Systems, Washington D.C., December 1996.
Second IEEE Real-Time Technology and Applications Symposium, Boston, MA, June 1996.

Conference General Chair ACM SIGCOMM Internet Measurement Conference, San Diego, CA, October 2007.
22nd IEEE Real-Time Systems Symposium, London, UK, December 2001.
Tenth International Workshop on Network and Operating System Support for Digital Audio and Video, Chapel Hill, NC, June 2000.
Third IEEE Real-Time Technology and Applications Symposium, Montreal, Canada, June 1997.
IEEE Workshop on Resource Allocation Problems in Multimedia Systems, Washington D.C., December 1996.

Member of 21st IEEE International Conference on Network Protocols, Goettingen, Germany, October 2013.
Conference 18th IEEE International Conference on Network Protocols, Kyoto, Japan, October 2010.
Program ACM SIGCOMM 2010, New Delhi, India, August 2010.
Committee 17th IEEE International Conference on Network Protocols, Princeton, NJ, October 2009.
 9th Passive and Active Measurement Conference 2008, Cleveland, OH, April 2008.
 ACM SIGMETRICS 2008, Annapolis, MA, June 2008.
 IEEE INFOCOM 2008, Phoenix, AZ, April 2008.
 6th ACM SIGCOMM Workshop on Hot Topics in Networks, Atlanta, GA, November 2007.
 15th IEEE International Conference on Network Protocols, Beijing, China, October 2007.
 5th IEEE Workshop on Embedded Systems for Real-Time Multimedia, Salzburg, Austria, October 2007.
 Workshop on Experimental Computer Science, ACM FCRC, San Diego, CA, June 2007.
 15th International Workshop on Quality-of-Service, Chicago, IL, June 2007.
 17th International Workshop on Network and Operating System Support for Digital Audio and Video, Urbana-Champaign, IL, June 2007.
 10th IEEE Global Internet Symposium 2007, Anchorage, AK, May 2007.
 ACM SIGMETRICS 2007, San Diego, CA, June 2007.
 14th IEEE International Conference on Network Protocols, Santa Barbara, CA, October 2006.
 ACM Internet Measurement Conference 2006, Rio de Janeiro, Brazil, October 2006.
 Second ACM SIGCOMM Conference on Future Networking Technologies, Lisbon, Portugal, December 2006.
 16th International Workshop on Network and Operating System Support for Digital Audio and Video, Newport, RI, June 2006.
 IEEE Workshop on Research Directions for Security and Networking in Critical Real-Time and Embedded Systems, San Jose, CA, April 2006.
 IEEE INFOCOM 2006, Barcelona, Spain, April 2006.
 13th IEEE International Conference on Network Protocols, Boston, MA, November 2005.
 15th International Workshop on Network and Operating System Support for Digital Audio and Video, Skamania, WA, June 2005.
 IEEE INFOCOM 2005, Miami, FL, March 2005.
 SPIE/ACM Multimedia Computing and Networking 2005, San Jose, CA, January 2005.
 ACM Multimedia 2004, New York, NY, October 2004.
 ACM Internet Measurement Conference 2004, Taormina, Italy, October 2004.
 16th EUROMICRO Conference on Real-Time Systems, Catania, Italy, June-July 2004.
 14th International Workshop on Network and Operating System Support for Digital Audio and Video, Cork, Ireland, June 2004.
 12th International Workshop on Quality-of-Service, Montreal, Canada, June 2004.
 10th IEEE Real-time and Embedded Technology and Applications Symposium, Toronto, Canada, May 2004.
 IEEE INFOCOM 2004, Hong Kong, March 2004.
 24nd IEEE Real-Time Systems Symposium, Cancun, Mexico, December 2003.

19th ACM Symposium on Operating Systems Principles, Lake George, New York, October 2003.

15th EUROMICRO Conference on Real-Time Systems, Porto, Portugal, July 2003.

Ninth IEEE Real-Time and Embedded Technology and Applications Symposium, Toronto, CA, May 2003.

SPIE/ACM Multimedia Computing and Networking 2003, San Jose, CA, January 2003.

23rd IEEE Real-Time Systems Symposium, Austin, TX, November 2002.

10th International Conference on Network Protocols, Paris, France, November 2002.

Second Workshop on Embedded Software, Grenoble, France, October 2002.

22nd IFIP International Symposium on Computer Performance Modeling, Measurement and Evaluation (Performance 2002), Rome, Italy, September 2002.

Eighth IEEE Real-Time and Embedded Technology and Applications Symposium, San Jose, CA, September 2002.

10th International Workshop on Quality-of-Service, Miami, FL, June 2002.

Seventh Global Internet Symposium (held in conjunction with Globecom 2002), Taipei, Taiwan, November 2002.

14th EUROMICRO Conference on Real-Time Systems, Vienna, Austria, June 2002.

ACM SIGMETRICS 2002, Marina del Rey, CA, June 2002.

SPIE/ACM Multimedia Computing and Networking 2002, San Jose, CA, January 2002.

22nd IEEE Real-Time Systems Symposium, London, UK, December 2001.

Sixth Global Internet Symposium (held in conjunction with Globecom 2001), San Antonio, TX, November 2001.

27th EUROMICRO Conference, Warsaw, Poland, September 2001.

Sixth IEEE Symposium on Computers and Communications, Hammamet, Tunisia, July 2001.

Seventh IEEE Real-Time Technology and Applications Symposium, Taipei, Taiwan, June 2001.

Ninth IFIP International Workshop on Quality of Service, Karlsruhe, Germany, June 2001.

SPIE/ACM Multimedia Computing and Networking 2001, San Jose, CA, January 2001.

Fifth Global Internet Mini-Conference (held in conjunction with Globecom 2000), San Francisco, CA, November 2000.

21st IEEE Real-Time Systems Symposium, Orlando, FL, December 2000.

Tenth International Workshop on Network and Operating System Support for Digital Audio and Video, Chapel Hill, NC, June 2000.

Sixth IEEE Real-Time Technology and Applications Symposium, Washington, D.C., June 2000.

SPIE/ACM Multimedia Computing and Networking 2000, San Jose, CA, January 2000.

19th IEEE Real-Time Systems Symposium, Phoenix, AZ, December 1999.

Fourth Global Internet Mini-Conference (held in conjunction with Globecom '99), Rio de Janeiro, Brazil, November 1999.

Fifth International Workshop on Multimedia Information Systems, Palm Springs, CA, October 1999.

Ninth International Workshop on Network and Operating System Support for Digital Audio and Video, Basking Ridge, NJ, June 1999.

IEEE Workshop on QoS Support for Real-Time Internet Applications, Vancouver, Canada, June, 1999.

Second ACM Workshop on Internet Server Performance (held in conjunction with SIGMETRICS '99),

Atlanta, GA, May 1999.

Seventh IEEE International Workshop on Parallel and Distributed Real-Time Systems, San Juan, Puerto Rico, April 1999.

SPIE/ACM Multimedia Computing and Networking 1998, San Jose, CA, January 1999.

19th IEEE Real-Time Systems Symposium, Madrid, Spain, December 1998.

Third Global Internet Mini-Conference (held in conjunction with Globecom '98), Sydney, Australia November 1998.

Sixth ACM International Conference on Multimedia, Bristol, UK, September 1998.

Third IEEE/ACM International Workshop on Multimedia Database Management Systems, Dayton, OH, August 1998.

Eighth International Workshop on Network and Operating System Support for Digital Audio and Video, Cambridge, UK, July 1998.

SPIE Interactive Multimedia Services and Equipment, Zurich, Switzerland, May 1998.

IEEE Workshop on Dependable and Real-Time E-Commerce Systems, Denver, CO, June 1998.

Fourth IEEE Real-Time Technology and Applications Symposium, Denver, CO, June 1998.

IEEE International Conference on Multimedia Computing Systems, Austin, TX, June 1998.

SPIE/ACM Multimedia Computing and Networking 1998, San Jose, CA, January 1998.

16th IEEE Symposium on Reliable Distributed Systems, Durham, NC, October 1997.

Third IEEE Real-Time Technology and Applications Symposium, Montreal, Canada, June 1997.

IEEE Real-Time Education Workshop, Montreal, Canada, June 1997.

Seventh International Workshop on Network and Operating System Support for Digital Audio and Video, St. Louis, MO, May 1997.

17th International Conference on Distributed Computing Systems, Distributed Real-Time Systems Track, Performance of Distributed Systems Track, Distributed Multimedia Track, Baltimore, MD, May 1997.

Fifth IFIP International Workshop on Quality of Service, New York, NY, May 1997.

Third International Conference on Computer Science & Informatics, Raleigh-Durham, NC, March 1997.

SPIE/ACM Multimedia Computing and Networking 1997, San Jose, CA, February 1997.

Fourth ACM International Conference on Multimedia, Boston, MA, November 1996.

Second USENIX Symposium on Operating Systems Design and Implementation, Seattle, WA, October/November 1996.

ACM SIGSOFT '96: Fourth Symposium on the Foundations of Software Engineering, San Francisco, CA, October 1996.

Second IEEE International Workshop on Multimedia Database Management Systems, Blue Mountain Lake, NY, August 1996.

Second IEEE Real-Time Technology and Applications Symposium, Boston, MA, June 1996.

Sixth International Workshop on Network and Operating System Support for Digital Audio and Video, Zushi, Japan, April 1996.

First International Workshop on Real-Time Databases: Issues and Applications, Newport Beach, CA, March 1996.

SPIE Multimedia Computing and Networking 1996, San Jose, CA, February 1996.

16th IEEE Real-Time Systems Symposium, Pisa, Italy, December 1995.

Third ACM International Conference on Multimedia, Technical Paper Program, San Francisco, CA,

November 1995.

Third ACM International Conference on Multimedia, Video Program, San Francisco, CA, November 1995.

1995 ACM Conference on Organizational Computing Systems, Technical Track, Milpitas, CA, August 1995.

Fifth IEEE Workshop on Hot Topics in Operating Systems, Orcas Island, WA, May 1995.

15th International Conference on Distributed Computing Systems, Distributed Real-Time Systems Track and CSCW track, Vancouver, British Columbia, Canada, May 1995.

Fifth International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, NH, April 1995.

IEEE Workshop on the Role of Real-Time in Multimedia, Research Triangle Park, NC, December 1993.

14th IEEE Real-Time Systems Symposium, Research Triangle Park, NC, December 1993.

13th International Conference on Distributed Computing Systems, Real-Time Issues Track, Pittsburgh, PA, May 1993.

12th IEEE Real-Time Systems Symposium, San Antonio, TX, December 1991.

IEEE Conference on Communication Software: Communications for Distributed Applications and Systems, Chapel Hill, NC, April 1991.

<i>Conference Organizing Committees</i>	Finance Chair, IEEE Real-Time Systems Symposium, Miami, FL, December 2005.
	Finance Chair, IEEE Real-Time Systems Symposium, Lisbon, Portugal, December 2004.
	Finance Chair, IEEE Real-Time Systems Symposium, Cancun, Mexico, December 2003.
	Poster Session Chair, ACM Symposium on Operating System Principles, Bolton Landing, NY, October 2003.
	Demonstrations Chair, ACM Conference on Computer-Supported Cooperative Work, Research Triangle Park, NC, November 1994.
	Wine Steward, ACM Symposium on Operating System Principles, Asheville, NC, December 1993.
	Local Arrangements Chair, 14 th IEEE Symposium on Real-Time Systems, Durham, NC, December 1993.

Refereed Publications

<i>Technical Papers</i>	<i>Explicit Topology Management for Continental-Scale Airborne Networks</i>, B. Newton, J. Aikat, K. Jeffay, IEEE INFOCOM Workshop on Wireless Communications and Networking in Extreme Environments, Atlanta, GA, May 2017, pages 72-77. <i>Scientific Training in the Era of Big Data: A New Pedagogy for Graduate Education</i>, J. Aikat, T. Carsey, K. Fecho, K. Jeffay, A. Krishnamurthy, P.J. Mucha, A. Rajasekar, S. Ahalt, Big Data, Volume 5, Number 1, 2017, pages 12-18. <i>Geographic Routing in Extreme-Scale Highly-Dynamic Mobile Ad hoc Networks</i>, B. Newton, J. Aikat, K. Jeffay, Proceedings of the 24th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), London, UK, September 2016, pages 205-210. <i>Simulating Large-Scale Airborne Networks with ns-3</i>, B. Newton, J. Aikat, K. Jeffay, Proceedings of the 2015 ACM Workshop on ns-3, Barcelona, Spain, May 2015, pages 32-39.
-------------------------	--

Explicit Topology Control for Airborne Networks, B. Newton, K. Jeffay, J. Aikat, Proceedings of the 22nd IEEE International Conference on Network Protocols, Research Triangle, NC, October 2014, pages 368-373.

The Continued Evolution of Web Traffic, B. Newton, K. Jeffay, J. Aikat, Proceedings of the 21st IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), San Francisco, CA, August 2013, pages 80-89.

Gini in a Bottle: A Case Study of Pareto's Principle in the Wild, A. Blate, K. Jeffay, International Journal of Computer Networks and Communications Security, Vol. 1, No. 1, June 2013, pages 30-39.

Experiment Replication Using ProtoGENI Nodes, D. O'Neill, J. Aikat, K. Jeffay, Proceedings of the Second GENI Research and Educational Experiment Workshop, Salt Lake City, UT, March 2013, pages 9-15.

Towards Traffic Benchmarks for Empirical Networking Research: The Role of Connection Structure in Traffic Workload Modeling, J. Aikat, S. Hasan, K. Jeffay, F.D. Smith, Proceedings of the 20th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Arlington, VA, August 2012, pages 78-86.

Generating Realistic Synthetic TCP Application Workloads, J. Aikat, K. Jeffay, F.D. Smith, Proceedings, GENI Infrastructure Workshop, Princeton, NJ, June 2010, 8 pages.

Correlations of Size, Rate, and Duration in TCP Connections: The Case Against, C. Park, F. Hernández-Campos, J.S. Marron, K. Jeffay, F.D. Smith, Annals of Applied Statistics, Volume 4, Number 1, May 2010, pages 26-52.

Passive, Streaming Inference of TCP Connection Structure for Network Server Management, J. Terrell, K. Jeffay, F.D. Smith, J. Gogan, J. Keller, IFIP International Workshop on Traffic Monitoring and Analysis, Aachen, Germany, May 2009, in, Lecture Notes in Computer Science, Volume 5537, Springer, Berlin, Germany, pages 42-53.

Exposing Server Performance to Network Managers Through Passive Network Measurements, J. Terrell, K. Jeffay, F.D. Smith, J. Gogan, J. Keller, IEEE Internet Network Management Workshop 2008, Orlando, FL, October 2008, pages 1-6.

Multi-Resolution Anomaly Detection for the Internet, L. Zhang, Z. Zhu, K. Jeffay, J.S. Marron, F.D. Smith, IEEE Workshop on Network Management, Phoenix, AZ, April 2008, 6 pages.

The Effects of Active Queue Management and Explicit Congestion Notification on Web Performance, L. Le, J. Aikat, K. Jeffay, F.D. Smith, IEEE/ACM Transactions on Networking, Volume 15, Number 6, December 2007, pages 1217-1230.

Co-Scheduling Variable Execution Time Requirement Real-time Tasks and Non-Real-Time Tasks, A. Singh, K. Jeffay, Proceedings of the 19th Euromicro Conference on Real-Time Systems, Pisa, Italy, July 2007, pages 191-200.

Quantifying the Effects of Recent Protocol Improvements to Standards-Track TCP: Impact on Web Performance, M.C. Weigle, K. Jeffay, F.D. Smith, Computer Communications, Volume 29, Number 15, September 2006, pages 2853-2866.

Tmix: A Tool for Generating Realistic TCP Application Workloads in ns-2, M.C. Weigle, P. Adurthi, F. Hernández-Campos, K. Jeffay, F.D. Smith, ACM Computer Communications Review, Volume 36, Number 3, July 2006, pages 67-76.

A Loss and Queuing-Delay Controller for Router Buffer Management, L. Le, K. Jeffay, F.D. Smith, Proceedings of the 26th IEEE International Conference on Distributed Computing Systems, Lisbon,

Portugal, July 2006, 10 pages.

Understanding Patterns of TCP Connection Usage with Statistical Clustering, F. Hernández-Campos, A.B. Nobel, F.D. Smith, K. Jeffay, 13th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Atlanta, GA, September 2005, pages 35-44.

Delay-Based Early Congestion Detection and Adaptation: Impact on Web Performance, M.C. Weigle, K. Jeffay, F.D. Smith, Computer Communications, Volume 8, Number 8, May 2005, pages 837-850.

Extremal Dependence: Internet Traffic Applications, F. Hernández-Campos, K. Jeffay, C. Park, J.S. Marron, S.I. Resnick, Stochastic Models, Volume 21, Number 1, 2005, pages 1-35.

Generating Realistic TCP Workloads, F. Hernández-Campos, F.D. Smith, K. Jeffay, Proceedings of the Computer Measurement Group's 2004 International Conference, Las Vegas, NV, December 2004, pages 273-284.

Differential Congestion Notification: Taming the Elephants, L. Le, J. Aikat, K. Jeffay, F.D. Smith, Proceedings of the 12th IEEE International Conference on Network Protocols, Berlin, Germany, October 2004, pages 118-128.

Stochastic Models for Generating Synthetic HTTP Source Traffic, J. Cao, W.S. Cleveland, Y. Gao, K. Jeffay, F.D. Smith, M.C. Weigle, Proceedings of IEEE INFOCOM 2004, Hong Kong, March 2004, Volume 3, pages 1546-1557.

Variability in TCP Roundtrip Times, J. Aikat, J. Kaur, D. Smith, K. Jeffay, Proceedings of the 2003 ACM SIGCOMM Internet Measurement Conference, Miami Beach, FL, October 2003, pages 279-284.

Tracking the Evolution of Web Traffic: 1995-2003, F. Hernández-Campos, K. Jeffay, F.D. Smith, Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Orlando, FL, October 2003, pages 16-25.

The Effects of Active Queue Management on Web Performance, L. Le, J. Aikat, K. Jeffay, F.D. Smith, Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany, August 2003, pages 265-276.

Managing Latency and Buffer Requirements in Processing Graph Chains, S.M. Goddard, K. Jeffay, The Computer Journal, Volume 44, Number 6, 2001, (special issue on high-assurance systems), pages 486-503.

Rate-Based Resource Allocation Methods for Embedded Systems, K. Jeffay, S.M. Goddard, in, *Embedded Software*, Proceedings of the First International Workshop on Embedded Software (EMSOFT 2001), Tahoe City, CA, October 2001, Lecture Notes in Computer Science, Volume 2211, T. Henzinger, C. Kirsch, editors, Springer-Verlag, Berlin, Germany, 2001, pages 204-222.

Beyond Audio and Video: Multimedia Networking Support for Distributed, Immersive Virtual Environments, K. Jeffay, T. Hudson, M. Parris, Proceedings of the 27th EUROMICRO Conference, Warsaw, Poland, September 2001, pages 300-307.

Tuning RED for Web Traffic, M. Christiansen, K. Jeffay, D. Ott, F.D. Smith, IEEE/ACM Transactions on Networking, Volume 9, Number 3, (June 2001), pages 249-264.

What TCP/IP Protocol Headers Can Tell Us About the Web, F.D. Smith, F. Hernández Campos, K. Jeffay, D. Ott, Proceedings of ACM SIGMETRICS 2001/Performance 2001, Cambridge, MA, June 2001, pages 245-256.

Experiments in Best-Effort Multimedia Networking for a Distributed Virtual Environment, T. Hudson, M.C. Weigle, K. Jeffay, R.M. Taylor II, in *Multimedia Computing and Networking 2001*, Proceedings,

SPIE Proceedings Series, Volume 4312, San Jose, CA, January 2001, pages 88-98.

Analyzing the Real-Time Properties of a U.S. Navy Signal Processing System, S.M. Goddard, K. Jeffay, International Journal of Reliability, Quality and Safety Engineering, Volume 8, Number 4, December 2001, (special issue of HASE '99 best papers) pages 301-322.

A Comparative Study of the Realization of Rate-Based Computing Services in General Purpose Operating Systems, K. Jeffay, G. Lamastra, Proceedings of the Seventh IEEE International Conference on Real-Time Computing Systems and Applications, Cheju Island, South Korea, December 2000, pages 81-90.

The Synthesis of Real-Time Systems from Processing Graphs, S.M. Goddard, K. Jeffay, Proceedings of the Fifth IEEE International Symposium on High Assurance Systems Engineering, Albuquerque, NM, November 2000, pages 177-186.

Tuning RED for Web Traffic, M. Christiansen, K. Jeffay, D. Ott, F.D. Smith, Proceedings of ACM SIGCOMM 2000, Stockholm, Sweden, August-September 2000, pages 139-150.

Towards a Better-Than-Best-Effort Forwarding Service for Multimedia Flows, K. Jeffay, IEEE Multimedia, Volume 6, Number 4, October-December 1999, pages 84-88.

A Theory of Rate-Based Execution, K. Jeffay, S.M. Goddard, Proceedings of the 20th IEEE Real-Time Systems Symposium, Phoenix, AZ, December 1999, pages 304-314.

Parallel Switching in Connection-Oriented Networks, S. Baruah, K. Jeffay, J. Anderson, Proceedings of the 20th IEEE Real-Time Systems Symposium, Phoenix, AZ, December 1999, pages 200-209.

Analyzing the Real-Time Properties of a U.S. Navy Signal Processing System, S.M. Goddard, K. Jeffay, Proceedings of the Fourth IEEE International Symposium on High Assurance Systems Engineering, Washington, DC, November 1999, pages 141-150.

Application-Level Measurements of Performance on the vBNS, M. Clark, K. Jeffay, Proceedings of the IEEE International Conference on Multimedia Computing and Systems, Volume 2, Florence, Italy, June 1999, pages 362-366.

Lightweight Active Router-Queue Management for Multimedia Networking, M. Parris, K. Jeffay, F.D. Smith, in *Multimedia Computing and Networking 1999*, Proceedings, SPIE Proceedings Series, Volume 3654, San Jose, CA, January 1999, pages 162-174.

Proportional Share Scheduling of Operating System Services for Real-Time Applications, K. Jeffay, F.D. Smith, A. Moorthy, J.H. Anderson, Proceedings of the 19th IEEE Real-Time Systems Symposium, Madrid, Spain, December 1998, pages 480-491.

Efficient Object Sharing in Quantum-Based Real-Time Systems, J.H. Anderson, R. Jain, K. Jeffay, Proceedings of the 19th IEEE Real-Time Systems Symposium, Madrid, Spain, December 1998, pages 346-355.

A Better-Than-Best-Effort Service for Continuous Media UDP Flows, M. Parris, K. Jeffay, F.D. Smith, J. Borgersen, Proceedings of the Eighth International Workshop on Network and Operating System Support for Digital Audio and Video, Cambridge, UK, July 1998, pages 193-197.

Managing Memory Requirements in the Synthesis of Real-Time Systems from Processing Graphs, S. Goddard, K. Jeffay, Proceedings of the Fourth IEEE Real-Time Technology and Applications Symposium, Denver, CO, June 1998, pages 59-70.

Fair On-Line Scheduling of a Dynamic Set of Tasks on a Single Resource, S.K. Baruah, J.E. Gehrke, C.G. Plaxton, I. Stoica, H. Abdel-Wahab, K. Jeffay, *Information Processing Letters*, Volume 64, Number

1, October 1997, pages 43-51.

A Two-Dimensional Audio Scaling Enhancement to an Internet Videoconferencing System, P. Nee, K. Jeffay, M. Clark, G. Danneels, Proceedings of the International Workshop on Audio-Visual Services over Packet Networks, Aberdeen, Scotland, UK, September 1997, pages 201-206.

Feasibility Concerns in PGM Graphs With Bounded Buffers, S. Baruah, S. Goddard, K. Jeffay, Proceedings of the Third IEEE International Conference on Engineering of Complex Computer Systems, Como, Italy, September 1997, pages 130-139.

Analyzing the Real-Time Properties of a Dataflow Execution Paradigm Using a Synthetic Aperture Radar Application, S. Goddard, K. Jeffay, Proceedings of the Third IEEE Real-Time Technology and Applications Symposium, Montreal, Canada, June 1997, pages 60-71.

Real-Time Computing with Lock-Free Shared Objects, J.H. Anderson, S. Ramamurthy, K. Jeffay, ACM Transactions on Computing Systems, Volume 15, Number 2, May 1997, pages 134-165.

The Performance of Two-Dimensional Media Scaling for Internet Videoconferencing, P. Nee, K. Jeffay, G. Danneels, Proceedings of the Seventh International Workshop on Network and Operating System Support for Digital Audio and Video, St. Louis, MO, May 1997, pages 237-248. Republished in "Readings in Multimedia Computing," K. Jeffay, H.J. Zhang, editors, Morgan Kaufmann, 2002.

On the Duality between Resource Reservation and Proportional Share Resource Allocation, I. Stoica, H. Abdel-Wahab, K. Jeffay, in *Multimedia Computing and Networking 1997*, Proceedings, SPIE Proceedings Series, Volume 3020, San Jose, CA, February 1997, pages 207-214.

Strategic Directions in Real-Time and Embedded Systems, J.A. Stankovic, A. Burns, K. Jeffay, M. Jones, G. Koob, I. Lee, J. Lehoczky, J. Liu, A. Mok, K. Ramamritham, J. Ready, L. Sha, A. van Tilborg, *ACM Computing Surveys*, Volume 28, Number 4, December 1996, pages 751-763.

A Proportional Share Resource Allocation Algorithm For Real-Time, Time-Shared Systems, I. Stoica, H. Abdel-Wahab, K. Jeffay, S.K. Baruah, J.E. Gehrke, C.G. Plaxton, Proceedings of the 17th IEEE Real-Time Systems Symposium, Washington, DC, December 1996, pages 288-299.

A General Framework For Continuous Media Transmission Control, T.M. Talley, K. Jeffay, Proceedings of the 21st IEEE Conference on Local Computer Networks, Minneapolis, MN, October 1996, pages 374-383.

A Router-Based Congestion Control Scheme For Real-Time Continuous Media, K. Jeffay, M. Parris, T. Talley, F.D. Smith, Proceedings of the Sixth International Workshop on Network and Operating System Support for Digital Audio and Video, Zushi, Japan, April 1996, pages 79-86.

Lock-Free Transactions for Real-time Systems, J.H. Anderson, S. Ramamurthy, M. Moir, K. Jeffay, Proceedings of the First Workshop on Real-Time Databases: Issues and Applications, Newport Beach, CA, March 1996, pages 107-114.

Real-Time Computing with Lock-Free Shared Objects, J.H. Anderson, S. Ramamurthy, K. Jeffay, Proceedings of the 16th IEEE Real-Time Systems Symposium, Pisa, Italy, December 1995, pages 28-37.

Early Experience with the Repository for Patterned Injury Data, D. Stotts, J.B. Smith, K. Jeffay, P. Dewan, D.K. Smith, W. Oliver, Proceedings of the SPIE International Symposium on Investigative and Trial Image Processing, San Diego, CA, July 1995, SPIE Volume 2567, 1995, pages 249-260.

Early Prototypes of the Repository for Patterned Injury Data, P. Dewan, K. Jeffay, J. Smith, D. Stotts, W. Oliver, Proceedings of Digital Libraries '95, The Second Annual Conference on the Theory and Practice of Digital Libraries, Austin, TX, June 1995, pages 123-130.

Support For Real-Time Computing Within General Purpose Operating Systems: Supporting co-resident operating systems, G. Bollella, K. Jeffay, Proceedings of the IEEE Real-Time Technology and Applications Symposium, Chicago, IL, May 1995, pages 4-14.

A Rate-Based Execution Abstraction For Multimedia Computing, K. Jeffay, D. Bennett, Proceedings of the Fifth International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, NH, April 1995, published in *Lecture Notes in Computer Science*, T.D.C. Little, R. Gusella, editors, Volume 1018, pages 64-75, Springer-Verlag, Heidelberg, Germany, 1995.

An Empirical Study of Delay Jitter Management Policies, D.L. Stone, K. Jeffay, *ACM Multimedia Systems*, Volume 2, Number 6, January 1995, pages 267-279. Republished in "Readings in Multimedia Computing," K. Jeffay, H.J. Zhang, editors, Morgan Kaufmann, 2002.

On the Partitioning of Function in Distributed Synchronous Collaboration Systems, J. Menges, K. Jeffay, ACM CSCW '94, Proceedings of the Workshop on Distributed Systems, Multimedia, and Infrastructure Support in CSCW, Research Triangle Park, NC, October 1994, SIGOIS Bulletin, Volume 15, Number 2, December 1994, pages 34-37.

Two-Dimensional Scaling Techniques For Adaptive, Rate-Based Transmission Control of Live Audio and Video Streams, T.M. Talley, K. Jeffay, Proceedings of the Second ACM International Conference on Multimedia, San Francisco, CA, October 1994, pages 247-254.

Transport and Display Mechanisms For Multimedia Conferencing Across Packet-Switched Networks, K. Jeffay, D.L. Stone, F.D. Smith, *Computer Networks and ISDN Systems*, Volume 26, Number 10, July 1994, pages 1281-1304. Republished in "A Guided Tour of Multimedia Systems and Applications," B. Furht, M. Milenkovic, editors, IEEE Computer Society Press, 1995.

A Patterned Injury Digital Library for Collaborative Forensic Medicine, D. Stotts, J.B. Smith, P. Dewan, K. Jeffay, F.D. Smith, D. Smith, S. Weiss, J. Coggins, W. Oliver, Proceedings of Digital Libraries '94, The First Annual Conference on the Theory and Practice of Digital Libraries, College Station, TX, June 1994, pages 25-33.

The Artifact-Based Collaboration System: An infrastructure for supporting and studying collaboration, K. Jeffay, J.B. Smith, F.D. Smith, D.E. Shackelford, J. Menges, Proceedings of the 15th Interdisciplinary Workshop on Informatics and Psychology, Schärding, Austria, May 1994, 27 pages.

On Latency Management in Time-Shared Operating Systems, K. Jeffay, Proceedings of the 11th IEEE Workshop on Real-Time Operating Systems and Software, Seattle, WA, May 1994, pages 86-90.

Inverting X: An Architecture for a Shared Distributed Window System, J. Menges, K. Jeffay, Proceedings of the Third Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises, Morgantown, WV, April 1994, IEEE Computer Society Press, pages 53-64.

Issues, Problems, and Solutions in Sharing X Clients on Multiple Displays, H. Abdel-Wahab, K. Jeffay, *Internetworking — Research and Practice*, Volume 5, Number 1, March 1994, pages 1-15.

Dynamic Participation in a Computer-based Conferencing System, G. Chung, K. Jeffay, H. Abdel-Wahab, *Computer Communications*, Volume 17, Number 1, January 1994, pages 7-16.

Accounting for Interrupt Handling Costs in Dynamic Priority Task Systems, K. Jeffay, D.L. Stone, Proceedings of the 14th IEEE Real-Time Systems Symposium, Raleigh-Durham, NC, December 1993, pages 212-221.

Queue Monitoring: A Delay Jitter Management Policy, D.L. Stone, K. Jeffay, Proceedings of the Fourth International Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, UK, November 1993, published in *Lecture Notes in Computer Science*, D. Shepherd, G. Blair, G. Coulson, N. Davies, F. Garcia, editors, Volume 846, pages 149-160, Springer-Verlag, Heidelberg,

Germany, 1994.

The Real-Time Producer/Consumer Paradigm: A paradigm for the construction of efficient, predictable real-time systems, K. Jeffay, Proceedings of the 1993 ACM/SIGAPP Symposium on Applied Computing, Indianapolis, IN, ACM Press, February 1993, pages 796-804.

Accommodating Late-Comers in Shared Window Systems, G. Chung, K. Jeffay, H. Abdel-Wahab, *IEEE Computer*, Volume 26, Number 1, January 1993, pages 72-74.

Scheduling Sporadic Tasks with Shared Resources in Hard-Real-Time Systems, K. Jeffay, Proceedings of the 13th IEEE Real-Time Systems Symposium, Phoenix, AZ, December 1992, pages 89-99.

Adaptive, Best-Effort, Delivery of Audio and Video Data Across Packet-Switched Networks, K. Jeffay, D.L. Stone, T. Talley, F.D. Smith, Proceedings of the Third International Workshop on Network and Operating System Support for Digital Audio and Video, La Jolla, CA, November 1992, published in *Lecture Notes in Computer Science*, V. Rangan, editor, Volume 712, pages 3-14, Springer-Verlag, Heidelberg, Germany, 1993.

Architecture of the Artifact-Based Collaboration System Matrix, K. Jeffay, J.K. Lin, J. Menges, F.D. Smith, J.B. Smith, ACM CSCW '92, Proceedings of the Conference on Computer-Supported Cooperative Work, Toronto, Canada, ACM Press, November 1992, pages 195-202.

Client/Server Protocol Filtering: An infrastructure for supporting group collaborations, K. Jeffay, J. Menges, J.-K. Lin, ACM CSCW '92, Proceedings of the Workshop on Tools and Technologies, Toronto, Canada, November 1992, pages 96-99.

Kernel Support for Live Digital Audio and Video, K. Jeffay, D.L. Stone, F.D. Smith, *Computer Communications*, Volume 15, Number 6, July/August 1992, pages 388-395.

On Kernel Support for Real-Time Multimedia Applications, K. Jeffay, Proceedings of the Third IEEE Workshop on Workstation Operating Systems, Key Biscayne, FL, April 1992, pages 39-46.

On Non-Preemptive Scheduling of Periodic and Sporadic Tasks, K. Jeffay, D.F. Stanat, C.U. Martel, Proceedings of the Twelfth IEEE Real-Time Systems Symposium, San Antonio, TX, December 1991, pages 129-139.

Kernel Support for Live Digital Audio and Video, K. Jeffay, D.L. Stone, F.D. Smith, Proceedings of the Second International Workshop on Network and Operating System Support for Digital Audio and Video, Heidelberg, Germany, November 1991, published in *Lecture Notes in Computer Science*, R.G. Herrtwich, editor, Volume 614, pages 10-21, Springer-Verlag, Heidelberg, Germany, 1992.

UNC Collaboratory Project Overview, J.B. Smith, F.D. Smith, P. Calingaert, J.R. Hayes, D. Holland, K. Jeffay, M. Lansman, Proceedings of the 1991 Symposium on Command and Control Research, National Defense University, Washington, D.C., June 1991, pages 341-391.

YARTOS: Kernel support for efficient, predictable real-time systems, K. Jeffay, D. Stone, D. Poirier, Proceedings of the Joint Eighth IEEE Workshop on Real-Time Operating Systems and Software and IFAC/IFIP Workshop on Real-Time Programming, Atlanta, GA, May 1991, in *Real-Time Systems Newsletter*, Volume 7, Number 4, Fall 1991, pages 8-13. Republished in "Real-Time Programming," W. Halang, K. Ramamritham, editors, 1992.

System Design for Workstation-Based Conferencing With Digital Audio and Video, K. Jeffay, F.D. Smith, Proceedings of the IEEE Conference on Communication Software: Communications for Distributed Applications and Systems, Chapel Hill, NC, April 1991, pages 169-180.

Designing a Workstation-Based Conferencing System Using the Real-Time Producer/Consumer Paradigm, K. Jeffay, F.D. Smith, Proceedings of the First International Workshop on Network and

Operating System Support for Digital Audio and Video, International Computer Science Institute, Berkeley, CA, November 1990, pages 40-55.

Analysis of a Synchronization and Scheduling Discipline for Real-Time Tasks with Preemption Constraints, K. Jeffay, Proceedings of the Tenth IEEE Real-Time Systems Symposium, Santa Monica, CA, December 1989, pages 295-305.

Corset & Lace: Adapting Ada Runtime Support to Real-Time Systems, T.P. Baker, K. Jeffay, Proceedings of the Eighth IEEE Real-Time Systems Symposium, San Jose, CA, December 1987, pages 158-167.

Research in Real-Time Systems, A.C. Shaw, C. Binding, W.L. Hu, K. Jeffay, Proceedings of the Third IEEE Workshop on Real-Time Operating Systems, Boston, MA, February 1986, pages 121-131.

Book Chapters *Rate-Based Resource Allocation Methods*, K. Jeffay, in, "Handbook of Real-Time and Embedded Systems," I. Lee, J. Y-T. Leung, S.H. Son, editors, Chapman & Hall/CRC Press, Boca Raton, FL, 2008, pages 4-1 – 4-15.

Visualization and Natural Control Systems for Microscopy, R.M. Taylor II, D. Borland, F.P. Brooks Jr., M. Falvo, M. Guthold, T. Hudson, K. Jeffay, G. Jones, D. Marshburn, S.J. Papadakis, L.-C. Qin, A. Seeger, F.D. Smith, D.H. Sonnenwald, R. Superfine, S. Washburn, C. Weigle, M.C. Whitton, P. Williams, L. Vicci, W. Robinett, in "Visualization Handbook," C. Johnson, C. Hansen, editors, Harcourt Academic Press, 2005, pages 875-900.

An Empirical Study of Delay Jitter Management Policies, D.L. Stone, K. Jeffay, in "Readings in Multimedia Computing and Networking," K. Jeffay, H.J. Zhang, editors, Morgan Kaufmann, San Francisco, CA, 2002, pages 525-537.

The Performance of Two-Dimensional Media Scaling for Internet Videoconferencing, P. Nee, K. Jeffay, G. Danneels, in "Readings in Multimedia Computing and Networking," K. Jeffay, H.J. Zhang, editors, Morgan Kaufmann, San Francisco, CA, 2002, pages 581-592.

Lock-Free Transactions for Real-time Systems, J.H. Anderson, S. Ramamurthy, M. Moir, K. Jeffay, in "Real-Time Database Systems: Issues and Applications," A. Bestavros, K.J. Lin, S.H. Son, editors, Kluwer Academic Publishers, Norwell, MA, 1997, pages 215-234.

Transport and Display Mechanisms For Multimedia Conferencing Across Packet-Switched Networks, K. Jeffay, D.L. Stone, F.D. Smith, in "A Guided Tour of Multimedia Systems and Applications," B. Furht, M. Milenkovic, editors, IEEE Computer Society Press, Los Alamitos, CA, 1995, pages 439-461.

Contributor to: *R&D for the NII: Technical Challenges*, M.K. Vernon, E.D. Lazowska, S.D. Personick, editors, Interuniversity Communications Council (EDUCOM), 1994.

YARTOS: Kernel support for efficient, predictable real-time systems, K. Jeffay, D. Stone, D. Poirier, in "Real-Time Programming," W. Halang, K. Ramamritham, editors, Pergamon Press, Oxford, UK, 1992, pages 7-12.

Videotapes *Adaptive, Best-Effort Delivery of Live Audio and Video Across Packet-Switched Networks*, K. Jeffay, D.L. Stone, ACM Multimedia '94 Video Proceedings, San Francisco, CA, October 1994, 6 minutes. Excerpts also appear on the CD-ROM version of the conference proceedings.

Abstracts & Short Papers *Quantifying the Effects of Recent Protocol Improvements to Standards-Track TCP*, M.C. Weigle, K. Jeffay, and F.D. Smith, Proceedings of the 11th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS), Orlando, FL, October 2003, pages 226-229.

Engineering of Rate Based Services for Real-Time Computing, G. Lamastra, K. Jeffay, 20th IEEE Real-

Time Systems Symposium Work in Progress Proceedings, Phoenix, AZ, December 1999, pages 51-55.

Lightweight Active Router-Queue Management for Multimedia Networking, M. Parris, K. Jeffay, F.D. Smith, IS&T/SPIE International Symposium on Electronic Imaging: Science and Technology, San Jose, CA, January 1999, pages 280-281.

Support For Real-Time Computing in Windows NT, K. Jeffay, Usage Abstracts, USENIX Windows NT Workshop, Seattle, WA, August 1997, page 84.

On the Duality between Resource Reservation and Proportional Share Resource Allocation, I. Stoica, H. Abdel-Wahab, K. Jeffay, IS&T/SPIE Ninth Symposium on Electronic Imaging: Science and Technology 1997, San Jose, CA, February 1997, pages 135-136.

Technical and Educational Challenges For Real-Time Computing, K. Jeffay, *ACM Computing Surveys*, Volume 28A, Number 4(es), December 1996, URL <http://www.acm.org/pubs/contents/journals/surveys/1996-28/#4es>, 3 pages.

Distributed Real-Time Dataflow: An Execution Paradigm for Image Processing and Anti-Submarine Warfare Applications, S. Goddard, K. Jeffay, 19th IEEE Real-Time Systems Symposium Work in Progress Proceedings, Washington, DC, December 1996, pages 55-58.

Why Using the Request Abstraction in Proportional Share Allocation Systems is Useful, I. Stoica, H. Zhang, K. Jeffay, Proceedings of the IEEE Real-Time Systems Symposium Workshop on Resource Allocation Problems in Multimedia Systems, Washington, DC, December 1996, 4 pages.

Future distributed embedded and real-time applications will be adaptive: Meanings, challenges, and research paradigms, A.K. Mok, C.L. Heitmeyer, K. Jeffay, M.B. Jones, C.D. Locke, R. Rajkumar, 15th Proceedings of the 15th IEEE International Conference on Distributed Computing Systems, Vancouver, British Columbia, Canada, May 1995, pages 182-184.

Adaptive, Best-Effort Delivery of Live Audio and Video Across Packet-Switched Networks, K. Jeffay, D.L. Stone, Proceedings of the Second ACM International Conference on Multimedia, San Francisco, CA, October 1994, pages 487-488.

Adaptive Rate-Based Flow and Latency Management of Audio and Video Streams, T.-M. Chen, T. Talley, K. Jeffay, Abstracts of the Fourth Workshop on Network and Operating System Support for Digital Audio and Video, Lancaster, UK, November 1993, pages 17-20.

Client/Server Protocol Filtering: An infrastructure for supporting group collaborations, K. Jeffay, J. Menges, J.-K. Lin, ACM CSCW '92, Proceedings of the Workshop on Tools and Technologies, Toronto, Canada, November 1992, pages 96-99.

Network and Operating Systems Support for Digital Audio and Video, K. Jeffay, in 13th ACM Symposium on Operating System Principles "Work in Progress" Abstracts, E.D. Lazowska, editor, Operating Systems Review, Volume 26, Number 2, April 1992, pages 7-31.

Posters *Education Modules for Networking, Cloud Computing, and Security in Systems Courses*, J. Aikat, M.K. Reiter, K. Jeffay, ACM Special Interest Group on Computer Science Education Technical Symposium, Memphis, TN, March 2016.

Efficient Management of a High-Capacity Airborne Network of Commercial Aircraft, B. Newton, J. Aikat, K. Jeffay, NSF Cyber-physical Systems Young Professional Workshop, Washington, DC, March 2014.

Variability of TCP Round-trip Times (RTTs) Over the Years, W. Yau, J. Aikat, K. Jeffay, Grace Hopper Celebration of Women in Computing (GHC), Baltimore, MD, October 2012.

Characterizing Modern Web Traffic using TCP Headers, S. Zolayvar, J. Aikat, K. Jeffay, Grace Hopper Celebration of Women in Computing (GHC), Baltimore, MD, October 2012.

Multivariate SVD Analyses For Network Anomaly Detection, J. Terrell, L. Zhang, K. Jeffay, A. Nobel, H. Shen, F.D. Smith, Z. Zhu, ACM SIGCOMM 2005 Poster Session, Philadelphia, PA, August 2005.

How Real Can Synthetic Traffic Be?, F. Hernández-Campos, K. Jeffay, F.D. Smith, ACM SIGCOMM 2004 Poster Session, Portland, OR, August 2004.

A Non-Parametric Approach to Generation and Validation of Synthetic Network Traffic, F. Hernández-Campos, A. Nobel, F.D. Smith, K. Jeffay, IMA Workshop on Measurement, Modeling, and Analysis of the Internet, Minneapolis, MN, January 2004.

Sync-TCP: Using GPS Synchronized Clocks for Early Congestion Detection in TCP, M.C. Weigle, K. Jeffay, F.D. Smith, ACM SIGCOMM 2000 Poster Session, Stockholm, Sweden, August 2000, page 6.

Reviews *Advance Reservation Systems*, K. Jeffay, Proceedings of the Fifth International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, NH, April 1995, published in *Lecture Notes in Computer Science*, T.D.C. Little, R. Gusella, editors, Volume 1018, pages 1-2, Springer-Verlag, Heidelberg, Germany, 1995.

Books and Proceedings

The Effects of Traffic Structure on Application and Network Performance, J. Aikat, K. Jeffay, and F. D. Smith, Springer Publishing, New York, NY, 2012, 289 pages.

The Effects of Active Queue Management on TCP Application Performance, An experimental performance evaluation, L. Le, K. Jeffay, F.D. Smith, Lambert Academic Publishing, Berlin, Saarbrücken, Germany, 2009, 222 pages.

Quality-of-Service — IWQoS 2003, K. Jeffay, I. Stoica, K. Wehrle, editors, Lecture Notes in Computer Science, Springer-Verlag, Berlin, Heidelberg, Germany, Volume 2707, 2003, ISBN 3-540-40281-0, 517 pages.

Readings in Multimedia Computing and Networking, K. Jeffay, H.J. Zhang, editors, Morgan Kaufman, San Francisco, CA, 2002, ISBN 1-55860-651-3, 863 pages.

Proceedings, The 22nd IEEE Real-Time Systems Symposium, K. Jeffay, G. Buttazzo, editors, IEEE Computer Society Press, Los Alamitos, CA, 2001, ISBN 0-7695-1420-0, 321 pages.

Proceedings, The Sixth IEEE International Symposium on Computers and Communications, K. Jeffay, R. Steinmetz, editors, IEEE Computer Society Press, Los Alamitos, CA, 2001, ISBN 0-7695-1177-5, 754 pages.

Proceedings, The 21st IEEE Real-Time Systems Symposium, K. Jeffay, W. Zhao, editors, IEEE Computer Society Press, Los Alamitos, CA, 2000, ISBN 0-7695-0900-2, 311 pages.

Proceedings, The 10th International Workshop on Network and Operating System Support for Digital Audio and Video, K. Jeffay and H. Vin, editors, Technical Report, Department of Computer Science, University of North Carolina at Chapel Hill, 2000, 320 pages. Proceedings also published on-line at <http://www.cs.unc.edu/nossdav2000>.

ACM Multimedia '99, D. Bulterman, K. Jeffay, H.J. Zhang, editors, ACM Press, Los Angeles, CA, 1999, ISBN 0-8194-3125-7, 500 pages.

Multimedia Computing and Networking 1999, D.D. Kandlur, K. Jeffay, T. Roscoe, editors, Proceedings of SPIE, Volume 3654, SPIE, Bellingham, WA, 1998, ISBN 0-8194-3125-7, 328 pages.

Multimedia Computing and Networking 1998, K. Jeffay, D.D. Kandlur, T. Roscoe, editors, Proceedings of SPIE, Volume 3310, SPIE, Bellingham, WA, 1997, ISBN 0-8194-250-0, 266 pages.

Proceedings, 1997 IEEE Real-Time Technology and Applications Symposium, R. Rajkumar, K. Jeffay, editors, IEEE Computer Society Press, Los Alamitos, CA, 1997, ISBN 0-8186-8016-4, 269 pages.

Proceedings, IEEE Real-Time Systems Symposium Workshop on Resource Allocation Problems in Multimedia Systems, K. Jeffay, editor, <http://www.cs.unc.edu/~jeffay/meetings/mm-wrkshp96/prog.html>, 1996, 250 pages.

Proceedings, 1996 IEEE Real-Time Technology and Applications Symposium, K. Jeffay, W. Zhao, editors, IEEE Computer Society Press, Los Alamitos, CA, 1996, ISBN 0-8186-7448-2, 264 pages.

Unrefereed Publications

Invited Papers *Modeling and Generation of TCP Application Workloads*, F. Hernández-Campos, K. Jeffay, F.D. Smith, Proceedings of the 4th IEEE International Conference on Broadband Communications, Networks, and Systems, Raleigh, NC, September 2007, 10 pages.

Statistical Clustering of Internet Communication Patterns, F. Hernández-Campos, A.B. Nobel, F.D. Smith, K. Jeffay, Proceedings of the 35th Symposium on the Interface of Computing Science and Statistics, Salt Lake City, UT, July 2003, Computing Science and Statistics, Volume 35, 2004.

Rate-Based Resource Allocation Methods for Embedded Systems, K. Jeffay, S.M. Goddard, in, *Embedded Software*, Proceedings of the First International Workshop on Embedded Software (*EMSOFT 2001*), Tahoe City, CA, October 2001, Lecture Notes in Computer Science, Volume 2211, T. Henzinger, C. Kirsch, editors, Springer Verlag, Berlin, Germany, 2001, pages 204-222.

Beyond Audio and Video: Multimedia Networking Support for Distributed, Immersive Virtual Environments, K. Jeffay, T. Hudson, M. Parris, Proceedings of the 27th EUROMICRO Conference, Workshop on Multimedia and Telecommunication, Warsaw, Poland, September 2001, pages 300-307.

Experiments in Best-Effort Multimedia Networking for a Distributed Virtual Environment, T. Hudson, M.C. Weigle, K. Jeffay, R.M. Taylor II, in *Multimedia Computing and Networking 2001*, Proceedings, SPIE Proceedings Series, Volume 4312, San Jose, CA, January 2001, pages 88-98.

Towards a Better-Than-Best-Effort Forwarding Service for Multimedia Flows, K. Jeffay, IEEE Multimedia, Volume 6, Number 4, October-December 1999, pages 84-88.

Network Support For Distributed, Immersive, Virtual Laboratories K. Jeffay, Proceedings of the NSF Workshop on Automated Control of Distributed Instrumentation, Beckman Institute for Advanced Science and Technology, University of Illinois at Urbana-Champaign, Urbana, IL, April 1999, pages 73-78.

Storage Requirements For Distributed Virtual Laboratories, K. Jeffay, Proceedings of the Internet2 Distributed Storage Infrastructure Application Workshop, University of North Carolina at Chapel Hill, Chapel Hill, NC, March 1999, pages 57-59.

Efficient Kernel Support for Continuous Time Media Systems, K. Jeffay, Proceedings of the ITC Workshop on Continuous Time Media, Carnegie Mellon University, Pittsburgh, PA, June 1991, 4 pages.

White Papers Contributor to: *NSF Report on Network Research Testbeds*, B. Braden, M. Gerla, J. Kurose, J. Lepreau, R. Rao, J. Turner, editors, December 2002. A white paper commission by NSF that led to the creation of a \$10 million funding program for network research testbeds.

Contributor to: *R&D for the National Information Infrastructure: Technical Challenges*, M.K. Vernon,

E.D. Lazowska, S.D. Personick, editors, Interuniversity Communications Council (EDUCOM), 1994.

Technical Reports *Beyond Window Sharing Hacks: Support for First-Class Window Sharing*, J. Menges, K. Jeffay, Report TR97-021, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, April 1997.

A Proportional Share Resource Allocation Algorithm For Real-Time, Time-Shared Systems, I. Stoica, H. Abdel-Wahab, K. Jeffay, S.K. Baruah, J.E. Gehrke, C.G. Plaxton, Report TR96-038, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, September 1996.

Predicting Worst Case Execution Times on a Pipelined RISC Processor, S.J. Bharrat, K. Jeffay, Report TR94-072, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, April 1994.

The Design, Implementation, and Use of a Sporadic Tasking Model, K. Jeffay, D. Becker, D. Bennett, S. Bharrat, T. Gramling, M. Housel, Report TR94-073, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, April, 1994.

UNC Collaboratory Project Overview, J.B. Smith, F.D. Smith, P. Calingaert, J.R. Hayes, D. Holland, K. Jeffay, M. Lansman, Report 90-042, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, November 1990.

An Implementation and Application of the Real-Time Producer/Consumer Paradigm, D. Poirier, K. Jeffay, Report 90-038, University of North Carolina at Chapel Hill, Department of Computer Science, Chapel Hill, NC, October 1990.

The Real-Time Producer/Consumer Paradigm: Towards Verifiable Real-Time Computations, K. Jeffay, Report 89-09-15, University of Washington, Department of Computer Science, Seattle, WA, September 1989, (Ph.D. Dissertation).

On Optimal, Non-Preemptive Scheduling of Periodic and Sporadic Tasks, K. Jeffay, R. Anderson, Report 88-11-06, University of Washington, Department of Computer Science, Seattle, WA, November 1988.

On Optimal, Non-Preemptive Scheduling of Periodic Tasks, K. Jeffay, Report 88-10-03, University of Washington Department of Computer Science, Seattle, WA, October 1988.

Software Engineering of Real-Time Operating Systems, A.C. Shaw, K. Jeffay, Report 88-01-01, University of Washington Department of Computer Science, Seattle, WA, January 1988.

Concurrent Programming with Time, A thesis proposal, K. Jeffay, Report 87-10-03, University of Washington Department of Computer Science, Seattle, WA, October 1987.

A Lace for Ada's Corset, T.P. Baker, K. Jeffay, Report 86-09-06, University of Washington Department of Computer Science, Seattle, WA, September 1986.

Software Distributions

GENI tmix — A synthetic TCP workload generator for GENI, J. Aikat, D. O'Neil, B. Newton, K. Jeffay, June 2012.

Distributed to the GENI networking research community.

tmix — A synthetic TCP workload generator, F. Hernández Campos, K. Jeffay, F.D. Smith, June 2005.

Distributed to other academic and industry networking research groups via the Web.

UNC VNC — A version of the public domain workspace sharing system VNC to support secure sharing of workspace regions and windows of individual applications. J. Branscomb, L. Fowler, K. Jeffay, August 2004.

Distributed to other research groups and the public via the Web and SourceForge.

thttp 2003 — An HTTP/v1.0 and v1.1 traffic generation program used to simulate the traffic generated by a collection of geographically distributed web browsers and servers. L. Le, A. Van Osdol, F.D. Smith, K. Jeffay, January 2004.

Distributed to other academic and industry networking research groups via the Web.

thttp 2001 — An HTTP/v1.0 and v1.1 traffic generation program used to simulate the traffic generated by a collection of geographically distributed web browsers and servers. F. Hernández Campos, F.D. Smith, K. Jeffay, August 2002.

Distributed to other academic and industry networking research groups via the Web.

Packmime Implementation in ns — An implementation of the Lucent Bell Labs (Bill Cleaveland) HTTP traffic model in ns. M.C. Weigle, K. Jeffay, F.D. Smith, September 2001.

Distributed as part of the Network Simulator (**ns-2**) distribution.

UNC Campus Network Trace Data and Empirical Distributions — A collection of network traces and empirical distributions illustrating the nature of web browsing and the structure of web pages. F. Hernández Campos, K. Jeffay, F.D. Smith, June 2001.

Distributed to other academic and industry networking research groups via the Web.

HTTP Analysis Tools — A set of tools for analyzing http connections. F. Hernández Campos, K. Jeffay, F.D. Smith, June 2001.

Distributed as part of the Network Simulator (**ns-2**) distribution.

thttp — An HTTP traffic generation program used to simulate the traffic generated by a collection of geographically distributed web browsers and servers. M. Christiansen, F.D. Smith, K. Jeffay, January 2000.

Distributed to other academic and industry networking research groups via the Web.

FlowGen — A general purpose, programmable network flow generator. M. Clark, K. Jeffay, F.D. Smith, December 1999.

Distributed to Advanced Networking Systems Inc. for use in performance evaluation studies of Internet 2 and Abilene. Also available via anonymous *ftp* from <http://www.cs.unc.edu/Research/Dirt>.

TruTime Device Drivers for FreeBSD — A set of FreeBSD device drivers for the TruTime GPS boards, A. Moorthy, K. Jeffay, and F.D. Smith, December 1998.

Distributed to Advanced Networking Systems Inc. for use in the IETF Internet Performance Measurement Initiative (IPMI) Surveyor network.

YARTOS (Yet Another Real-Time Operating System) — A real-time operating system kernel for Intel x86 platforms that uses a novel task implementation based on a single, shared run-time stack, an earliest-deadline-first processor scheduling algorithm developed at UNC, and integrates processor scheduling and inter-process communication. K. Jeffay and D. Stone, January 1996.

Distributed to several universities for use in operating system research and teaching. Also adopted by researchers at IBM's Networking Software Division, RTP, NC, for use in constructing a network protocol evaluation testbed.

XPEL (The X Protocol Engine Library) — A C++ library for constructing modular X Window System pseudo-servers, J. Menges and K. Jeffay, November 1993.

Distributed on the Internet via anonymous *ftp* from <ftp.cs.unc.edu>.

XTV (X Terminal View) — An X Window System pseudo-server that supports window sharing and conferencing across the Internet, H. Abdel-Wahab (Old Dominion University) and K. Jeffay, August 1991.

Distributed by MIT as part of the contributed software portion of the X Window System distribution (release 11, version 5). Last revised January 1993.

Also adopted by the NEC Corporation of Japan as the basis for their *Xprotodist* (X protocol distributor) product.

Patents

Methods, Systems, and Computer Readable Media For Topology Management and Geographic Routing in Mobile Ad-Hoc Networks, U.S. Patent Number 9,832,705, B. Newton, K. Jeffay, J. Aikat, November 2017.

High-Performance Virtualized Applications Platform, U.S. Patent Application No. 421-343. A. Blate, K. Jeffay, May 2014.

Methods, Systems, and Computer Program Products For Network Server Performance Anomaly Detection, U.S. Patent No. 8,938,532, J.S. Terrell. K. Jeffay, F.D. Smith, R. Broadhurst, January 2015.

Method for Understanding the Use of TCP/IP Networks by Users, and Non-Parametric Generation of Synthetic Internet Traffic, U.S. Patent Number 7,447,209, F. Hernández-Campos, K. Jeffay, F.D. Smith, A. Nobel, November 2008.

User Controlled Adaptive Flow Control for Packet Networks, U.S. Patent Number 5,892,754, V. Kompella, F.D. Smith, J.P. Gray, and K. Jeffay, April 1999.

Public Demonstrations

Education Modules for Networking, Cloud Computing and Security in Systems Courses, J. Aikat, K. Jeffay, Demo session at the ACM SIGCSE, Special Interest Group for Computer Science Education, Memphis, TN, March 2016.

Exploring the Internet's Future with GENI (Global Environment for Network Innovations), J. Aikat, K. Jeffay, NSF Higher Education Technology Conference 2013, Morgantown, WV, October 2013.

Your first experiment on GENI: a Hands-on Tutorial, J. Aikat, K. Jeffay, NSF Higher Education Technology Conference 2013, Morgantown, WV, October 2013.

NSF Workshop on GENI in Education (one day), J. Aikat, K. Jeffay, New York University - Polytechnic, NY, October 2013.

Using the GENI Networking Testbed, J. Aikat, K. Jeffay, ACM SIGCSE, Special Interest Group for Computer Science Education, Denver, Colorado, March 2013.

Run your Systems/Networking Experiments using (free) GENI Resources, Lunchtime Table Topics, S. Zolayvar, W. Yau, J. Aikat, K. Jeffay, Grace Hopper Celebration of Women in Computing, Baltimore, Maryland, October 2012.

NSF Workshop on Designing Tools and Curricula for Undergraduate Courses in Distributed Systems, J. Aikat, K. Jeffay, Boston, MA, July 2012.

Introduction to Network Experiments using the GENI CyberInfrastructure, J. Aikat, K. Jeffay, ACM SIGMETRICS/Performance Joint International Conference on Measurement and Modeling of Computer Systems, London, UK, June 2012.

The nanoManipulator "reverse field trip" to Orange County High-School, G. Jones, R. Superfine, R.M. Taylor, M. Whitton, T. Lovelace, R. Parameswaran, K. Jeffay, F.D. Smith, Orange County High School,

Hillsboro, NC, November 1999.

The Distributed NanoManipulator Project, R.M. Taylor, K. Jeffay, T. Hudson, M. Clark, an invited demonstration for the Internet 2 Spring 1999 Members Meeting, Washington, DC, April 1999.

A Demonstration of Internet Access to UNC-CH Advanced Microscopy Facilities for Science Education Outreach, K. Jeffay, F.D. Smith, R. Superfine, G. Jones, R.M. Taylor, T. Hudson, M. Clark, Orange County High School, Hillsboro, NC, June 1998.

The Artifact-Based Collaboration (ABC) System, J.B. Smith, K. Jeffay, D. Shackelford, J. Menges, J. Hilgedic, B. Ladd, M. Parris, E. Kupstas, T. Ellis, and T. Hudson, CSCW '94, ACM Conference on Computer-Supported Cooperative Work, RTP, NC, October 1994.

CONCUR: A Window System Supporting Window Sharing, J. Menges and K. Jeffay, CSCW '94, ACM Conference on Computer-Supported Cooperative Work, RTP, NC, October 1994.

Adaptive, Best-Effort Delivery of Live Audio and Video Across Packet-Switched Networks, K. Jeffay, Second MCNC/NSF Packet Video Workshop, Research Triangle Park, NC, December 1992.

Grants and Awards

GENI in the classroom: Course Modules for Teaching Networking Concepts, J. Aikat, K. Jeffay, National Science Foundation (GENI Project: Raytheon Corporation), October 2013, 3 years, \$178,978.

II-New: Seeing the Future: Ubiquitous Computing in EyeGlasses, J.-M. Frahm, *et al.*, National Science Foundation, November 2013, 3 years, \$613,623.

GENI Tutorial at ACM SIGCSE, J. Aikat, K. Jeffay, National Science Foundation, grant number CNS 13-32340, March 2013, 1 year, \$13,388.

Development of Education and Training Resources for GENI Experimenters, K. Jeffay, J. Aikat, Global Environment for Network Innovations/National Science Foundation, December 2011, 3 years, \$388,590.

Exploring Privacy Breaches in Encrypted VoIP Communications, F. Monrose, K. Jeffay, National Science Foundation, grant number CNS 10-17318, August 2010, 3 years, \$496,492.

Synthetic Traffic Generation Tools and Resource: A Community Resource for Experimental Networking Research, K. Jeffay, F.D. Smith, UNC, M. Weigle, Old Dominion University, A. Vahdat, University of California San Diego, P. Barford, University of Wisconsin, National Science Foundation, grant number CRI 07-09081, August 2007, 3 years, \$799,745.

Modeling and Testing of Application Workloads on Corporate Enterprise Networks, K. Jeffay, F.D. Smith, IBM Global Services Faculty Award, September 2006, 2 years, \$70,000 (exempt from indirect costs).

Tera-pixels: Using High-resolution Pervasive Displays to Transform Collaboration and Teaching, K. Jeffay, A. Lastra, K. Mayer-Patel, L. McMillan, F.D. Smith, National Science Foundation (CISE RI Program), grant number EIA 03-03590, August 2003, 5 years, \$962,902.

Extracting and Using Semantic Information in Network Workloads, K. Jeffay, F.D. Smith, IBM Global Services Faculty Award, September 2006, 1 year, \$40,000 (exempt from indirect costs).

Modeling and Testing of Application Workloads on Corporate Enterprise Networks, K. Jeffay, F.D. Smith, IBM Global Services Faculty Award, September 2006, 1 year, \$40,000 (exempt from indirect costs).

Generation and Validation of Synthetic Internet Traffic, K. Jeffay, F.D. Smith, A.B. Noble, National Science Foundation, grant number ANI 03-23648, September 2003, 3 years, \$470,000.

Rate-Based Resource Allocation Methods for Real-Time Embedded Systems, K. Jeffay, F.D. Smith, National Science Foundation, grant number CCR 02-08924, August 2002, 3 years, \$179,990.

Empirical Workload Characterizations for Advanced Networks, K. Jeffay, F.D. Smith, Cisco Systems, November 2000, 3 years, \$212,500.

Support for Active Queue Management, Distributed XP Programming, and a Teaching Laboratory for Web Services, F.D. Smith, K. Jeffay, J.B. Smith, P.D. Stotts, IBM Shared University Research Grant, August 2002, 1 year, \$83,000.

Internet Traffic Measurement and Analysis and A Teaching Laboratory for Enterprise Java Computing, F.D. Smith, K. Jeffay, J.B. Smith, P.D. Stotts, IBM Shared University Research Grant, April 2001, 1 year, \$83,000.

An NS Implementation of the HTTP Connection Model, F.D. Smith, K. Jeffay, M. Clark Weigle, Lucent Technologies, April 2001, \$10,000.

Multimedia Networking Research in Support of Virtual Field Trips, K. Jeffay, F.D. Smith, R.M. Taylor II, Dell Computer Corporation Strategic Technology and Research (STAR) program, December 2000, 1 year, \$25,000.

Rate-based Scheduling Technology for Latency-Sensitive Graphics Applications, J. Anderson, S. Baruah, K. Jeffay, R. Taylor, National Science Foundation, grant number ITR 00-82866, September 2000, 3.5 years, \$350,000.

Active Router Queue Management for Congestion Control and Quality-of-Service, K. Jeffay, F.D. Smith, National Science Foundation, grant number ITR 00-82870, September 2000, 3 years, \$451,903.

Internet Measurements and Analysis to Support the nanoManipulator and Tele-Immersion, F.D. Smith, K. Jeffay, P. Jones, IBM Shared University Research Grant, September 2000, 1 year, \$250,000.

The Performance of Differentiated Services Implementations for Supporting Distributed Virtual Environment Applications, K. Jeffay, North Carolina Network Initiative, September 2000, 1 year, \$25,000.

Empirical Application-Workload Characterizations for Advanced Networks, K. Jeffay, F.D. Smith, Sun Microsystems, September 2000, \$79,000.

Empirical Application-Workload Characterizations for Advanced Networks, F.D. Smith, K. Jeffay, MCNC, September 2000, 1 year, \$50,000.

A Quality-of-Service Network for the nanoManipulator, K. Jeffay, Cabletron Inc. and the North Carolina Network Initiative, September 1999, \$25,000.

Interactive Graphics for Molecular Studies and Microscopy — Supplement for Collaborative Science, F.P. Brooks, D. Erie (Chemistry), K. Jeffay, J. Samulski (Gene Therapy), F.D. Smith, D. Sonnenwald (Information and Library Science), R. Superfine (Physics and Astronomy), R. Taylor, National Institute of Health, October 1998, 4 years, \$1,902,544.

Computing Power for Collaborative Science, S.R. Aylward, G. Bishop, D.W. Brenner, E. Bullitt, E.L. Chaney, V.L. Chi, B.J. Dempsey, N. England, A.G. Gash, D. Fritsch, H. Fuchs, B. Hemminger, J. Hermans, K. Jeffay, K. Keller, A. Lastra, M. Lin, D. Manocha, L.S. Nyland, S.M. Pizer, J.W. Poulton, J.F. Prins, J. Rosenman, F.D. Smith, R. Superfine, R.M. Taylor, S. Tell, G. Tracton, A. Tropsha, S. Washburn, G. Welch, Intel Corporation, August 1998, 3 years, \$2,858,747. PI on *Multimedia Networking*

and *A Distributed Teaching Laboratory for Networking and Internet Technologies* sections (F.D. Smith, B.J. Dempsey, Co-PIs), \$557,669.

Empirical Application-Workload Characterizations for Advanced Networks, F.D. Smith, K. Jeffay, North Carolina Network Initiative, September 1998, \$40,000.

Congestion control for high-speed networks, F.D. Smith, K. Jeffay, IBM Corporation, 1998, \$20,000.

A Communications Middleware For Immersive Distributed Virtual Environments, K. Jeffay, North Carolina Network Initiative, 1998, \$15,000.

Technology for Real-Time Services in Protocol Stack Implementations, K. Jeffay, IBM Corporation, 1997, \$80,000.

Internet Access to UNC-CH Advanced Microscopy Facilities for Science Education Outreach, K. Jeffay, R. Superfine (Physics and Astronomy), G. Jones (Education), University of North Carolina at Chapel Hill Instructional Technology Award, 1997, \$29,750.

Support for the Real-Time Technology and Applications Symposium, K. Jeffay, Office of Naval Research, 1997, \$7,000.

Research Experience For Undergraduates supplement to *Object Sharing Technology For Real-Time Systems*, J. Anderson, K. Jeffay, National Science Foundation, 1996, \$10,000.

Infrastructure for Research in Collaborative Systems, S.F. Weiss, J.B. Smith, K. Jeffay, P. Dewan, P. D. Stotts, D.K. Smith, F.D. Smith, W. Oliver (U.S. Armed Forces Institute of Pathology), National Science Foundation, grant number CDA-9624662, August 1996, 5 years, \$1,260,830.

Collaboration Bus: An Infrastructure for Supporting Interoperating Collaborative Systems, P. Dewan, K. Jeffay, H. Abdel-Wahab (Old Dominion University), P. D. Stotts, L. Nyland, J.B. Smith, J. Mchugh (Portland State University), J. Menges (Hewlett Packard), Advanced Research Projects Agency, grant number 96-06580 (High Performance Distributed Services Technology), 1996, \$973,334.

Object Sharing Technology For Real-Time Systems, J. Anderson, K. Jeffay, National Science Foundation, grant number CCR 95-10156, 1996, \$209,442.

Flexible Shared Windows, P. Dewan, K. Jeffay, National Science Foundation, grant number IRIS 95-08514, 1995, \$343,811.

An ATM Testbed For Multimedia Networking and Computer-Supported Cooperative Work, K. Jeffay, IBM Corporation, 1995, \$400,000.

An Examination of Flow and Congestion Control Mechanisms for Media Transmission in Collaborative Systems, K. Jeffay, IBM Corporation, 1995, \$125,242.

Software Infrastructure for the Rapid Development of Interactive and Collaborative Educational Simulations, J.F. Prins, K. Jeffay, P.D. Stotts, L.S. Nyland, Advanced Research Projects Agency, grant number 95-36871 (Computer Aided Education and Training Initiative), 1995, \$200,000.

A Proposal For Network Routers, K. Jeffay, IBM Corporation, Research Triangle Park, NC, 1994, \$40,000.

System Support For Video Teleconferencing Across Local Area Networks, K. Jeffay, Intel Corporation, 1993, \$253,385.

The Integration and Use of Digital Audio and Video in Desktop Computing Environments, K. Jeffay, IBM Corporation, 1993, \$100,000.

Construction of a UNC-Tektronix MBONE Link, K. Jeffay, Tektronix Corporation, 1993, \$10,000.

An Empirical Determination of the Limits of Human Perception of Properties of Digital Audio and Video Streams, K. Jeffay, University of North Carolina Research Council, 1993, \$2,000.

Processor and Resource Allocation Problems in Hard-Real-Time Systems: Theory and Practice, K. Jeffay, National Science Foundation, RIA Award, 1991, \$59,400.

Accommodating Continuous Media in a Local Area Network: An exercise in distributed real-time computing, K. Jeffay, University of North Carolina Junior Faculty Development Award, 1990, \$3,000.

Building and Using a Collaboratory: A Foundation for Supporting and Studying Group Collaborations, J.B. Smith, F.D. Smith, P. Calingaert, K. Jeffay, J.R. Hayes, D. Holland, M. Lansman, National Science Foundation, grant number ICI-9015443, 1990, \$946,000.

Enhanced program of research and teaching in communications software and distributed systems, K. Jeffay, IBM Corporation, 1989, \$429,000.

The Design and Construction of Predictable Real-Time Systems, K. Jeffay, Digital Faculty Program Award, Digital Equipment Corporation, 1989, \$180,000.

Invited Presentations

Keynote Addresses *Network Neutrality Considered Harmful*,
PearlHacks, University of North Carolina at Chapel Hill, Chapel Hill, NC, March 2014.
Carolina Science Café, University of North Carolina at Chapel Hill, Chapel Hill, NC,
 May 2015.

Network Neutrality Considered Harmful, Revisiting the Quality-of-Service Morass
 ACM International Workshop on Network and Operating System Support for Digital
 Audio and Video, Newport, RI, May 2006.

Rate-Based Resource Allocation Methods for Multimedia Computing
 SPIE Multimedia Computing and Networking 2003, Santa Clara, CA, January 2003.

Network Support For Distributed Virtual Environments: The Tele-nanoManipulator
 NCNI Advanced Networking Symposium, Research Triangle Park, NC, May 1999.
 The Internet 2 Spring Members Meeting (presented jointly with R.M. Taylor),
 Washington, DC, April 1999.

Network Support For Distributed, Immersive, Virtual Environments
 Workshop on Automated Control of Distributed Instrumentation, Beckman Institute for
 Advanced Science and Technology, University of Illinois at Urbana-Champaign,
 Urbana, IL, April 1999.

The Future of Networking: The Networking Revolution Has Yet to Begin
 North Carolina Governor's Board of Science and Technology Retreat, Banner Elk, NC,
 September 1998.

Distinguished *The Synthetic Traffic Generation Problem: The least sexy problem in computer networking*
 University of Minnesota, Minneapolis, MN, January 2008.

Lectures

The Effect of Active Queue Management on Web Performance: The Good, the Bad, and the Ugly

University of Nebraska, Lincoln, NE, November 2003.

University of Pennsylvania, April 2004.

The Evolution of Quality-of-Service on the Internet

Mälardalen University, Sweden, February 2001.

Colloquia Efficient Topology Management and Geographic Routing in High-Capacity Continental- Scale Airborne Networks

University of Virginia, November 2017.

The Evolution of Quality-of-Service on the Internet

IEEE Computer Society Seattle Chapter, Seattle, WA, September 2005.

Modeling and Generating TCP Application Workloads

Georgia Institute of Technology, Atlanta, GA, September 2005.

Microsoft Research, Redmond, WA, September 2005.

Worcester Polytechnic Institute, Worcester, MA, July 2005.

Cisco Systems, RTP, NC, May 2005 (presently jointly with F. Hernandez-Campos and F.D. Smith).

A Rate-Based Execution Abstraction For Embedded Real-Time Systems

University of Pennsylvania, April 2004.

How “Real” Can Synthetic Network Traffic Be?

University of Virginia, March 2004.

Non-Parametric Approach to Generation and Validation of Synthetic Network Traffic

Columbia University, New York, NY, January 2004.

Is Explicit Congestion Notification (ECN) Worthwhile?

Cisco Systems, San Jose, CA, October 2003.

Intel Architecture Labs, Hillsboro, OR, October 2003.

Tuning RED for Web Traffic: RED Considered Harmful?

Sprint Advanced Technology Research Laboratories, San Francisco, CA, March 2000.

University of Illinois at Urbana-Champaign, Urbana, IL, March 2000.

Internet Traffic: Measurement & Generation

Ganymede Software, Research Triangle Park, NC, June 1999. (Presented jointly with F.D. Smith.)

Lightweight Active Router-Queue Management for Multimedia Networking

University of Toronto, October 2000.

HP Laboratories, Palo Alto, CA, January 1999.

Sprint Advanced Technology Research Laboratories, San Francisco, CA, January 1999.

A Better-Than-Best-Effort Service For UDP: Lightweight Active Router-Queue Management for Multimedia Networking

University of Virginia, January 1999.

Congestion Control Mechanisms for Real-Time Communications on the Internet

IBM Networking Systems, Research Triangle Park, NC, November 1998.

A Better Than Best-Effort Forwarding Service For UDP

Microsoft Research, Redmond, WA, March 1998.

IBM Networking Systems, Research Triangle Park, NC, March 1998.

Multimedia Networking Research at UNC Chapel Hill

IBM Networking Systems, Research Triangle Park, NC, February 1998.

Lucent Technologies, Bell Laboratories, Holmdel, NJ, June 1997.

Design Principles for Distributed, Interactive, Virtual Laboratories

Mitsubishi Electric Research Laboratory, Cambridge, MA, January 1998.

Lucent Technologies, Bell Laboratories, Holmdel, NJ, June 1997.

The Performance of Two-Dimensional Media Scaling for Internet Videoconferencing

Mitsubishi Electric Research Laboratory, Cambridge, MA, January 1998.

Carnegie Mellon University, Pittsburgh, PA, March 1997.

Some Approaches to Enabling Real-Time Computation on Desktop Operating Systems

Honeywell Technology Center, Minneapolis, MN, October 1996.

Two-Dimensional Scaling Techniques for Adaptive, Rate-Based Transmission Control of Live Audio and Video Streams

Oregon State University, Corvallis, OR, May 1996.

A Router-Based Congestion Control Scheme For Real-Time Continuous Media

Intel Architecture Development Laboratory, Hillsboro, OR, May 1996.

NetEdge Inc., Research Triangle Park, NC, March 1996.

North Carolina State University, Raleigh, NC, March 1996.

A Hybrid Reservation-Based/Best-Effort Transmission Scheme For Real-Time Transmission of Multimedia Data

Intel Research Council, Hillsboro, OR, November 1995.

Operating System Support For Multimedia Computing

IBM Networking Systems, Research Triangle Park, NC, July 1995.

Network Support For Multimedia Computing

IBM Networking Systems, Research Triangle Park, NC, July 1995.

A Rate-Based Execution Abstraction For Multimedia Computing

Tektronix Computer Research Laboratory, Beaverton, OR, May 1996.
 INRIA, Rocquencourt, France, December 1995.
 AT&T Bell Laboratories, Murray Hill, NJ, July 1995.

A Theory of Rate-Based Scheduling

Carnegie Mellon University, Pittsburgh, PA, August 1994.

Demonstration of Videoconferencing Over Packet-Switched Internetworks

Global Lecture Hall of the U.S.-Russian Electronic Distance Education System/TELE-TEACHING '93, Trondheim, Norway (via satellite from Chapel Hill), August 1993.

Transport and Display Mechanisms For Multimedia Conferencing Across Packet-Switched Networks

York University, York, England, November 1993.
 Purdue University, West Lafayette, IN, February 1993.
 Intel Architecture Development Laboratory, Hillsboro, OR, February 1993.
 Oregon Graduate Institute, Beaverton, OR, February 1993.
 Second MCNC/NSF Packet Video Workshop, Research Triangle Park, NC, December 1992.
 Carnegie Mellon University, Pittsburgh, PA, November 1992.

Multimedia and Conferencing Research at UNC-CH

IBM Multimedia Networking, Research Triangle Park, NC, September 1992.

System Support for Synchronous Collaboration

North Carolina Artificial Intelligence and Advanced Computing Symposium, Raleigh, NC, March 1992.

Software Architectures for Predictable Real-Time Computer Systems

Carnegie Mellon University, Pittsburgh, PA, March 1992.

Operating System Requirements for Digital Audio and Video

MCNC/NSF Packet Video Videoconferencing Workshop, Research Triangle Park, NC, December 1991.

Some Experiments With Live Digital Audio and Video on a Local-Area Network

IBM European Networking Center, Heidelberg, Germany, November 1991.

Network and Operating System Support for Digital Audio and Video

Duke University, Durham, NC, November 1994.
 Carnegie Mellon University, Pittsburgh, PA, April 1994.
 Intel Corporation, Hillsboro, OR, December 1993.
 IBM Networking Systems, Research Triangle Park, NC, June 1992.
 IBM TJ Watson Research Center, Yorktown Heights, NY, March 1992.
 University of Washington, Seattle, WA, February 1992.
 Old Dominion University, Norfolk, VA, October 1991.

Systems Research in the UNC Collaboration Project

Intel Architecture Development Laboratory, Hillsboro, OR, May 1992.

University of Colorado, Boulder, CO, March 1991.

Some Deterministic Resource Allocation Problems in Real-Time Computer Systems

Research Triangle Institute, Research Triangle Park, NC, June 1990.

University of North Carolina at Chapel Hill, Department of Operations Research, Chapel Hill, NC, October 1989.

The Real-Time Producer/Consumer Paradigm: Towards Verifiable Real-Time Computations

Carnegie-Mellon University, Pittsburgh, March 1991.

Duke University, Durham, NC, June 1990.

Technical University of Denmark, Lyngby, Denmark, September 1989.

University of California at Davis, Davis, CA, April 1989.

University of Minnesota, Minneapolis, MN, April 1989.

Purdue University, West Lafayette, IN, April 1989.

University of North Carolina at Chapel Hill, Chapel Hill, NC, April 1989.

University of British Columbia, Vancouver, BC, Canada, April 1989.

Rice University, Houston, TX, March 1989.

University of Arizona, Tucson, AZ, March 1989.

Oregon Graduate Center, Beaverton, OR, March 1989.

Tektronix Computer Research Laboratory, Beaverton, OR, March 1989.

University of Colorado, Boulder, CO, March 1989.

IBM TJ Watson Research Center, Yorktown Heights, NY, March 1989.

IBM Systems Integration Division, Owego, NY, March 1989.

New York University, New York, NY, March 1989.

Bell Communications Research, Morristown, NJ, March 1989.

Concurrent Programming With Time

Olivetti Research Center, Menlo Park, CA, December 1987.

IBM TJ Watson Research Center, Yorktown Heights, NY, September 1987.

Columbia University, New York, NY, September 1987.

Short Courses Trends in Congestion Control and Quality-of-Service: Active queue management on the Internet of the future

Mälardalen University, Sweden, February 2001.

Rate-Based Execution Models For Real-Time Multimedia Computing

Scuola Superiore Santa Anna, Pisa, Italy, September 1997.

Tutorials Issues in Multimedia Delivery Over Today's Internet

IEEE International Conference on Multimedia Computing Systems, Austin, TX, June 1998.

Systems Issues in the Design and Realization of Desktop Videoconferencing Systems,

Second ACM International Conference on Multimedia, San Francisco, CA, October 1994.