

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

fuboTV MEDIA INC., and YANKA INDUSTRIES, INC.,
Petitioners,

V.

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

IPR2024-00917

U.S. PATENT NO. 10,757,156

**DECLARATION OF DR. KEVIN JEFFAY IN SUPPORT OF
PATENT OWNER RESPONSE**

Table of Contents

I.	Introduction	1
	A. Engagement.....	1
	B. Background and Qualifications.....	2
II.	Relevant Legal Principles	8
	A. Prior Art	9
	B. Anticipation.....	10
	C. Obviousness	11
III.	The Challenged Patents.....	14
	A. Downloading Rather Than Streaming	15
	B. Conventional Streaming.....	17
	C. Progressive Download	19
	D. Continuous Playback	20
IV.	Overview of Cited Art.....	23
	A. Leaning	23
	B. Gamble.....	28
	C. Klemets	29
	D. Ala-Honkola.....	29
V.	Level of Ordinary Skill in the Art.....	30
VI.	Claim Construction	31
VII.	The Petition and Dr. Houh Fail to Demonstrate that Any Claims of the Challenged Patents are Invalid	31
	A. Leaning Fails to Disclose or Suggest “Continuous Playback”	32
	1. Dr. Houh Fails to Address the “Continuous [Playback/Presentation]” Limitation.....	36

2.	Leaning Fails to Disclose “Continuous [Playback/Presentation]” Because It Intentionally Interrupts Playback at Each and Every Rate Shift.	41
3.	Dr. Houh Fails to Show that Leaning’s Teachings Regarding Delivering <i>Audio</i> Can be Applied to the Claimed Requirements for Delivering <i>Video</i>	50
4.	A POSITA Would Not Discard Leaning’s “Clear Buffer” Instruction	54
5.	Leaning’s Teaching to <i>Encode Then Segment</i> Video Cannot Achieve “Continuous Playback”	56
B.	Leaning Fails to Disclose Requesting Streamlets via “multiple Transmission Control Protocol (TCP) connections” / “a plurality of different TCP connections”	61
C.	Leaning Fails to Initially Request Low Quality Files/Streamlets to Enable Instant Playback	69
D.	Leaning Fails to Disclose “Arranging” Files/Streamlets	72
E.	Leaning Fails to Disclose Upshifting “According to an Amount of Contiguously Available Files”	75
F.	Leaning Fails to “anticipate file requests and to pre-request files” ..	79
G.	Leaning and Leaning in Combination with Klemets Fail to Disclose Arranging Streamlets “using the virtual timeline”	83
H.	Leaning Fails to Disclose “independently...playable” Streamlets ..	88
I.	Gamble Fails to Cure the Deficiencies of Leaning	108
J.	Ala-Honkola Fails to Cure the Deficiencies of Leaning	111
VIII.	Dr. Houh Fails to Show the Patents are Invalid	113
IX.	Dr. Houh’s Obviousness Opinions are Foreclosed by Strong Objective Evidence of Nonobviousness Identified by the ITC	115

A. The streaming products offered by DISH (and its predecessor MOVE Networks) embody the claimed invention.	118
B. Long-felt but unmet need and failure of others	142
C. Surprising results and superior performance	147
D. Commercial Success	151
E. Copying.....	164
F. Industry praise and peer recognition.....	168
X. Conclusion	174

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 was requested by counsel for DISH to review the Petitions for *Inter Partes* Review and accompanying Declarations of Dr. Henry Houh submitted by Petitioner Aylo Freesites Ltd in IPR2024-00940 regarding U.S. Patent No. 11,991,234 (“the ’234 Patent”), and submitted by Petitioners fuboTV Media, Inc., and Yanka Industries, Inc., in IPR2024-00917 regarding U.S. Patent No. 10,757,156 (“the ’156 Patent”), IPR2024-00918 regarding U.S. Patent No. 9,407,564 (“the ’564 Patent”), and IPR2024-00919 regarding U.S. Patent No. 8,868,772 (“the ’772 Patent”) (collectively, the “Challenged Patents” or “ABR Patents”). I understand that all of the Challenged Patents claim priority to a provisional application filed on April 30, 2004. I provide my opinion regarding validity of the Challenged Patents below.

4. I have been asked by counsel for DISH to provide technical opinions in connection with the *inter partes* reviews of the Challenged Patents. I submit this combined declaration in support of Patent Owner's Response in each of the above-listed IPR's. I reserve the right to supplement my opinions, should I be requested to provide additional analysis.

5. The specifications of the Challenged Patents are the same, with the exception of the '234 Patent, which includes certain additional disclosures as a continuation-in-part application. I also understand that Dr. Houh provided substantially similar analysis across each of his declarations in these proceedings. Accordingly, in addressing certain claim elements and limitations recited across each Challenged Patent, I will cite to the '564 Patent and the Petition and Exhibits in IPR2024-00918 unless otherwise noted.

6. 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

7. I am qualified by education and experience to testify as an expert in the field of computer networking 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.

8. 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, and again in 2024, I also served as the Chairman of the Department.

9. 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.

10. 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.

11. 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.

12. 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.

13. 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.

14. 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.

15. 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.

16. 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).

17. 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.

18. 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.

19. 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.

20. 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.

21. 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.

22. 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

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

24. I understand that the Petitioner in an IPR proceeding bears the burden of demonstrating invalidity of the Challenged Claims by a preponderance of the evidence.

A. Prior Art

25. 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.

26. 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. Anticipation

27. 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.

28. 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.

C. Obviousness

29. 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.

30. 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.

31. In evaluating whether the challenged claims would have been obvious, I have also considered the following factors:

- Whether 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.

32. 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 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 challenged patent as a road map for selecting and combining prior art.

33. 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.

34. 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.

35. 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. The Challenged Patents


36. The Challenged Patents are directed to “alleviat[ing] the problems of reliability, efficiency, and latency” when “video streaming over packet switched networks such as the Internet.” EX1001, 1:18-21, 2:39-41. The Challenged Patents’ primary focus in this regard is “the viewing or listening experience of the user.” *Id.*,

2:20-21; *see also id.*, 2:27-29. As the Challenged Patents explain, “[u]sers generally ... tend to want to see or hear ... media file[s] instantaneously,” and without any “break” or “slowdown” in playback. *Id.*, 1:49-50, 2:13-15, 2:20-21. The Challenged Patents describe these “problems and needs in the art” in the context of three different content delivery techniques—downloading, conventional streaming, and progressive downloads—and teaches solutions that overcome the drawbacks of each to achieve continuous playback while streaming. *See id.*, 1:31-2:5, 2:52-55.


A. Downloading Rather Than Streaming

37. The first conventional approach to viewing digital content over a network is “downloading.” As its name indicates, downloading involves “[d]ownloading complete media files,” which can take “large amounts of time.” EX1001, 1:44-46. In addition, “downloading is by definition restricted to receiving a specific content selection in its entirety or not at all.” *Id.*, 1:44-46, 1:54-58.

38. The Challenged Patents address these drawbacks by time-dividing media files into smaller portions, or “streamlets,” and allowing playback to begin as soon as a single streamlet is downloaded, rather than waiting for the entire media file to download. Exemplary streamlets 212 are depicted below and described by the Challenged Patents as “independent media object[s]” that correspond to 2-second portions of a larger stream 210.



Id., 6:58-63, Fig. 2c. The Challenged Patents' exemplary process for requesting (Step 606), downloading (Step 618), and playing back (Steps 610 & 614) a single streamlet is depicted below.




Id., 11:35-67, Fig. 6.

39. That the Challenged Patents are directed to allowing playback to begin after only a single streamlet is received is demonstrated by the Challenged Patents' use of the singular form of "streamlet" in the figure above and in the written description of this process, as well as the Challenged Patents' explicit reference to an "alternative" embodiment in which "a plurality of streamlets" are requested for staging and playback. *See, e.g.*, EX1001, 11:40-44 ("The agent controller module 402 then requests 606 a streamlet from the streamlet cache module 404. Alternatively, the agent controller module 402 may simultaneously request 606 a plurality of streamlets from the streamlet cache module 404"). It also is demonstrated by the Challenged Patents' reference to its inventive technology as "adaptive-rate shifting of *streaming* content" and explanation that, in general, "streaming" technology "delivers content ... to a user ... *as the data is received.*" *Id.*, Title, 1:18-21, 1:31-33 (emphasis added).


B. Conventional Streaming

40. Although conventional streaming also involves "deliver[ing] content ... to a user ... *as the data is received,*" it is "vulnerable to network failures or congestion" that can result in "delays or losses" during playback. EX1001, 1:31-33, 1:60-61, 2:13-18 (emphasis added). The Challenged Patents addresses this problem by creating different copies of its streamlets at different bit rates and inserting a step into the process depicted above—specifically, Step 610—that responds to changes

in network conditions. An example of these different copies 202 of the Challenged Patents' streamlets 212 are depicted below and described by the Challenged Patents as corresponding to "a low quality stream 204, a medium quality stream 206, and a high quality stream 208."



Id., 6:48-52, Fig. 2c. An exemplary process for determining whether to perform a rate shift is depicted below.



Id., Fig. 7. As depicted, the invention of the Challenged Patents responds to network conditions by "upshift[ing] to a higher quality streamlet" when a user's network

connection improves or “downshift to a lower quality streamlet” when a user’s network connection degrades. *Id.*, 3:11-15, 3:20-2.

41. The invention of the Challenged Patents also accounts for changes in network conditions by using multiple TCP/IP¹ connections. EX1001, 9:34-35. That way, “[w]hen a TCP connection fails completely” or “a request is not being satisfied in a timely manner, a redundant request may be sent on a different connection for the same streamlet 212.” *Id.* at 9:59-63. This “achieve[s] the latency-reduction efficiency benefits of pipelining while maintaining the independence of each streamlet 212 request,” which further ensures continuous playback. *Id.*, 10:15-18.

C. Progressive Download

42. The foregoing advances also resolve the problems associated with the third type of conventional content delivery discussed in the Challenged Patents—progressive downloading. This content delivery technique is a combination of downloading and conventional streaming in which playback begins *before* the content is “*completely downloaded.*” EX1001, 1:62-2:1 (emphasis added). But it “is not a complete or fully acceptable solution to the problem of media presentation over a network” because it “*waits* to begin playback until there is enough of the file downloaded.” *Id.*, 1:64-2:5 (emphasis added). Specifically, progressive

¹ TCP refers to “Transmission Control Protocol” and “IP” refers to “Internet Protocol” here and elsewhere in this submission.

downloading “waits” until there is enough data stored in a buffer that playback can begin without “catch[ing] up” to, or outpacing, the download, which would result in disruption of playback. *Id.*, 1:64-2:1. This can happen repeatedly during playback, which is why this technique is called “progressive.” And the resulting delays “can be substantial depending on network conditions.” EX1001, 2:2-5.


43. The invention of the Challenged Patents avoids these known issues with progressive downloading because it both (i) shifts between different bit rates in response to changes in network conditions (*see supra* § III.B) **and** (ii) allows playback to begin after only a single streamlet of the new bit rate has been received (*see supra* § III.A), rather than staying at the same bit rate and waiting for a certain amount of data to be downloaded before playback begins.

D. Continuous Playback

44. During prosecution of the '772 Patent, which issued from the parent application of the '564, '156, and '234 Patents, the Patent Owner amended each independent claim to require the making of shifting determinations “to achieve continuous playback of the video using [the files] of the highest quality one of the copies determined sustainable at that time,” explaining that this requirement distinguished over the prior art methodologies. IPR2024-00919, Paper 3 ('772 Patent Petition) at 8-9 (describing “continuous playback” amendment and related applicant argument during prosecution). The Patent Owner made the same pertinent

amendment to each independent claim of the '564 Patent, noting to the Examiner that the Board had found the prior art failed to disclose the amended claims of the '772 Patent. EX1002, 222-25, 229-30. This limitation now appears in all but one of the independent claims of the Challenged Patents: independent Claim 1 of the '234 Patent.

45. The Challenged Patents' combination of improvements over the prior art discussed above avoids disruptions and achieves continuous playback. The invention of the Challenged Patents also avoids disruptions and achieves continuous playback by playing all the streamlets already arranged in its "staging module" (depicted below), even after a decision to rate shift is made.



EX1001, Fig. 4 (highlighting added). In fact, before shifting bit rates, the invention of the Challenged Patents looks directly at “the amount of contiguously available streamlets that have been sequentially arranged by the staging module” to ensure that playback will not outpace the download of additional streamlets if an upshift is performed. EX1001, 12:43-50; *see also id.*, 1:64-2:1. The Challenged Patents refer to this as a “***read ahead*** margin” and “discourage upshifting until a larger read ahead

margin may be established.” *Id.*, 12:50-53 (emphasis added). Thus, the Challenged Patents disclose adaptive-rate shifting streaming technology that not only allows the streamlets already in its staging module to play out during a rate shift, but that also relies on the time it will take to play those streamlets to determine whether a shift can be performed in the first instance. This further ensures that playback is uninterrupted in certain embodiments of the Challenged Patents.


46. The Challenged Patents disclose at least two other techniques for achieving continuous playback. One is to “pre-request” streamlets of a different bit rate based on an “anticipate[d]” need to shift bit rates, such that the streamlets are already staged for playback before the decision to rate shift occurs. EX1001, 8:46-50. And the other is to “request streamlets at any interval less than the length of the time index,” so that smaller streamlets are received. *Id.*, 8:53-57. These techniques “enable[] a ‘fast-start’ capability with no buffering wait when starting,” including when starting a stream of a different bit rate. *Id.*, 8:50-57. Thus, the Challenged Patents disclose no fewer than six techniques directed to achieving the stated goal of continuous playback: (1) switching between different bitrate streamlets without filling a buffer first, (2) playing all the streamlets already arranged in a “staging module” after a decision to rate shift is made (3) using multiple TCP connections to request streamlets, (4) anticipating and pre-requesting streamlets, and (5) requesting only portions of streamlets.

IV. Overview of Cited Art

A. Leaning

47. International Patent Application Publication WO 2002/049343 to Leaning *et al.* (“Leaning”) “is concerned with the delivery, over a telecommunications link, of digitally coded material for presentation to a user.” EX1004, 1:3-4. More specifically, Leaning is directed to delivering digitally coded material, primarily MP3 audio, using the “progressive download” technique discussed above. EX1004, 2:5–7. Indeed, every embodiment disclosed in Leaning “*waits* to begin playback until there is enough of the file downloaded,” as described in the Challenged Patents. *Id.*; EX1001, 1:64-2:5 (emphasis added). The application for Leaning was filed on December 14, 2001.

48. Leaning includes a flow chart of pseudo-code disclosing, step-by-step, its process for rate shifting with MP3 audio files. *See* EX1004, 9-11. Although Leaning discusses modifying certain steps—*e.g.*, omitting “the first of the three AND conditions” in one step (*id.*, 15:7-8)—there are several critical steps that remain unchanged throughout every disclosed embodiment. One of those steps is that beginning, or resuming, playback is *always* conditioned on a buffer being filled with enough content to avoid playback outpacing the download of content (“If Buffer Fullness > Tp”), as depicted below.



Id., 10 (annotations added). Leaning describes the threshold that must be met in this regard as the “playout time Tp,” which is “set at 10 seconds” in Leaning’s only example. *Id.*, 5:20-21. The “playout time Tp” “determines the degree of buffering that the terminal performs.” *Id.*, 5:21-22. Thus, Leaning discloses “buffering” at least 10 seconds of content before beginning or resuming playback.

49. Another step of Leaning that remains unchanged throughout its disclosure is that the buffer is cleared (*i.e.*, emptied, flushed, erased, *etc.*) ***every time*** a rate shift or “jump” is performed, as depicted below.

If UserCommand = Jump(j%)then: <u>Clear buffer;</u> Set CurrentSubfile = Integer[(LFI+1)*j/100]; go to Step J1
--

Stepup: <u>Clear Buffer</u> Set RateD = RateC Set StemD = StemC Set RateC = RateU Set StemC = StemU Set RateU = next higher rate in list or zero if none Set StemU = Stem + “/” + item in rate list corresponding to this rate Go to Step J1A

Stepdown: <u>Clear Buffer</u> Set RateU = RateC Set StemU = StemC Set RateC = RateD Set StemC = StemD Set RateD = next lower rate in list or zero if none Set StemD = Stem + “/” + item in rate list corresponding to this rate Go to Step J1A
--

EX1004, 10-11 (annotations added). With respect to rate shifts, Leaning explains that “changing [rates] at a time when the buffer is already empty is advantageous as the sound card is not playing anything and it may be necessary to reconfigure it.” *Id.*, 8:20-23. And with respect to jumps, Leaning explains that “[w]e prefer ... to suspend decoding and clear the buffer so that the new request is sent immediately.” *Id.*, 11:19-20. A POSITA would understand that “suspend[ing] decoding and clear[ing] the buffer” also would advantageously allow a “new request” to be “sent immediately” after a rate shift.

50. Critically, the “Clear Buffer” function is always performed *before* the “If Buffer Fullness > T_p ” test during a *both* a rate upshift (“Stepup”) and a downshift (“Stepdown”), as depicted below.

User	Terminal	Server
	Decode received subfile	
	Write to buffer	
	J1A: If Buffer Fullness > Tp seconds go to Step J3	
	J2: Increment Current Subfile	
	Go to Step J1	
	J3: Begin/continue playing of buffer contents via sound card	
	J4: If Buffer Fullness < Tp seconds go to Step J2	
	If BufferFullness = 0 AND MRate < CurrentRate AND BufferLow% > Td go to Stepdown	
	If MRate > NextRate AND NextRate < 0 go to Stepup	
Input of user commands	If UserCommand = Pause then: Stop reading from buffer; Loop until Usercommand = Resume; go to J3	
	If UserCommand = Jump(j%) then: Clear buffer; Set CurrentSubfile = Integer[(LFI+1)*j/100]; go to Step J1	
	Go to Step J1A	
	Stepup: Clear Buffer Set RateD = RateC Set StemD = StemC Set RateC = RateU Set StemC = StemU Set RateU = next higher rate in list or zero if none Set StemU = Stem + "/" + item in rate list corresponding to this rate Go to Step J1A	
	Stepdown: Clear Buffer Set RateU = RateC Set StemU = StemC Set RateC = RateD Set StemC = StemD Set RateD = next lower rate in list or zero if none Set StemD = Stem + "/" + item in rate list corresponding to this rate Go to Step J1A	

EX1004, 10-11 (annotations and highlighting added). And because the condition “Buffer Fullness > Tp” is a precondition that must be satisfied before playback can resume, the “Clear Buffer” function causes playback to halt until the buffer refills beyond the “playout time Tp”—*i.e.*, until there is another 10 seconds of content stored in the buffer—*every time* a rate shift is performed. *Id.*, 6:10-11 (“If, during this process, the buffer becomes empty, the player program simply *ceases playing until further data arrives.*”) (emphasis added); *see also id.*, 8:20-23 (explaining that “the sound card is *not playing anything*” when “the buffer is already empty”) (emphasis added). Thus, Leaning is directed to the conventional process of progressive downloading, which the Challenged Patents explain “is not a complete or fully acceptable solution to the problem of media presentation over a network.” EX1001, 2:2-5. The Challenged Patents were the first to provide that solution.

51. Although Leaning describes “the internet or other packet network operating in accordance with the Hypertext Transfer Protocol” as an example of such a “telecommunications link” (EX1004 2:9-11), the highest data rate Leaning describes using is 128 kilobits per second, or kbit/s (*id.*, 19:1-5). Specifically, Leaning describes using data rates of 96 and 128 kbit/s in the context of encoding video (*id.*, 17:13-15, 19:1-5), and even lower rates—specifically, “8, 16, 24, and 32 kbit/s”—in the context of encoding audio (*id.*, 6:13-15; 6:23-26, 9:4-5). Because Leaning was assigned to British Telecom, it is my opinion that these disclosures

suggest that the inventors of Leaning were considering ISDN and other dialup circuits for the transmission of their digital media. In fact, Leaning repeatedly refers to its exemplary network as “a *telecommunications* network” (*id.*, 1:3-4, 1:15-16, 2:5-7) and also discusses the use of a “public switched telephone network (PSTN)” for both audio and video feeds (*id.*, 19:11-18). Rate fluctuations were rare in these types of networks. This provides a motivation for Leaning to focus on using H.261, a video coding scheme developed by the telephony industry for transmission of video (and not audio) over dedicated telecommunications circuits.

B. Gamble

52. U.S. Patent Application Publication No. 2004/0093420 by Gamble (“Gamble”) is titled “Method and System for Transferring Large Data Files Over Parallel Connections.” EX1020. Gamble discloses a data transfer system and method. *Id.*, Abstract. Although Gamble discloses “using parallel connections,” nowhere does Gamble disclose or suggest video streaming over those parallel connections. *Id.*, ¶ 34; *see also generally id.* Instead, Gamble merely discloses transmitting large files between two applications on two computers (“systems”) by the first application reading fixed sized portions of the file (“READCHUNKs”) and transmitting these chunks (called segments) to the second application over parallel connections. *See id.* ¶¶ 50-52.

C. Klemets

53. U.S. Patent Application Publication No. 2003/0236906 by Klemets *et al.* (“Klemets”) is titled “Client-Side Caching of Streaming Media Content.” IPR2024-00940, EX1006. Klemets discloses a technique for “[c]lient-side caching of streaming media content.” *Id.* ¶ 23. Although Klemets describes that the stored streaming media “may be separated into pieces,” nowhere does Klemets disclose or suggest dividing the resulting streams into smaller, time-divided increments, or “streamlets.” *See id.* ¶ 36.

D. Ala-Honkola

54. U.S. Patent Application Publication Number 2003/0055995 by Ala-Honkola (“Ala-Honkola”) is titled “Adaptive media stream.” IPR2024-00917, EX1017. Ala-Honkola discloses “chang[ing] the source of an adaptive media stream when an available connection speed rate changes.” *Id.* ¶ 10. According to Ala-Honkola, “[w]hen the source is changed, also a new HTTP transmission for the new source is started for a new stream containing the information of the new source.” *Id.* Ala-Honkola also discloses “reading timing marks in said adaptive media stream for defining the right point to start playing the new media stream.” *Id.* ¶ 15. In Ala-Honkola, the “plurality of media streams is contained in a single file” and “pointers are preferably embedded within said adaptive media stream data.” *Id.* ¶ 25.

According to Ala-Honkola, when “all the sources” are not “in the same file,” “the use of this kind of structure is not as efficient.” *Id.* ¶ 46.

V. Level of Ordinary Skill in the Art

55. Dr. Houh contends that a 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).” EX1003 ¶ 42.

56. I disagree. Instead, it is my opinion that a POSITA at the time of the invention 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 least two years of practical experience or coursework in the design or development of distributed multimedia delivery systems.

57. 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.

VI. Claim Construction

58. 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 (“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.

59. I have not identified any other terms of the challenged claims of the Challenged Patents that require construction for purposes of rendering the opinions I express here.

VII. The Petition and Dr. Houh Fail to Demonstrate that Any Claims of the Challenged Patents are Invalid

60. Petitioner and Dr. Houh rely on Leaning as allegedly “disclosing” every limitation of every claim of the Challenged Patents except for certain limitations that require TCP connections, multiple TCP connections, HTTP requests, or requesting sub-parts of files. But Leaning fails to disclose or suggest at least the following limitations, which also are not found in any of the secondary references relied upon by Petitioner and Dr. Houh, or in any proposed combination thereof.

A. Leaning Fails to Disclose or Suggest “Continuous Playback”

61. Each of the following claim elements of the independent claims recites a “[shift/change] ... to achieve continuous [playback/presentation]”:

Challenged Claim	Element
'564 Patent Claim 1[D]	making the successive determinations to <i>shift the playback quality based on the factor to achieve continuous playback of the video</i> using the files of the highest quality one of the copies determined sustainable at that time so that the media player up shifts to a higher quality one of the different copies when the factor is greater than a first threshold and downshifts to a lower quality one of the different copies when the factor is less than a second threshold; and
'564 Patent Claim 8[C][3]	making the successive determinations to <i>shift the playback quality based on the factor to achieve continuous playback of the video</i> using the files of the highest quality one of the copies determined sustainable at that time, wherein the making the successive determinations to shift comprises upshifting to a higher quality one of the different copies when the at least one factor is greater than a first threshold and downshifting to a lower quality one of the different copies when the at least one factor is less than a second threshold; and
'156 Patent Claim 1[D]	adapting the successive determinations to <i>shift the playback quality based on the factor to achieve continuous playback of the video</i> using the streamlets of the highest quality copy of the video that is determined to be sustainable at that time; and
'156 Patent Claim 13[C]	making successive determinations by the end user device to <i>shift the playback quality based on the factor to achieve continuous playback of the video</i> using the streamlets of the highest quality copy determined sustainable at that time; and
'772 Patent Claim 1[E]	making the successive determinations to <i>shift the playback quality based on at least one of the set of factors to achieve continuous playback of the video</i> using the files of the highest quality one of the copies determined sustainable at that time; and

'234 Patent Claim 12 [D]	making the successive determinations whether to <i>change the requested bit rate based on the factor to achieve continuous presentation of the digital content</i> using the individually-requestable portions of the highest bit rate one of the copies determined receivable at that time at a rate sufficient for presenting the digital content as the individually-requestable portions of the digital content are received, wherein the making of the successive determinations to change the requested bit rate comprises:
-----------------------------	---

62. The Challenged Patents discuss at least six ways in which “continuous playback” may be achieved. The first is to begin playing content of a different bit rate even if only a single streamlet of that new bit rate is currently stored in the cache, as discussed *supra* at § III.A.² The second is to play all the streamlets already arranged in a “staging module” even after a decision to rate shift is made. *Supra* § III.D. The third is to look directly at “the amount of contiguously available streamlets that have been sequentially arranged by the staging module”—*i.e.*, the “read ahead margin”—to ensure that playback will not outpace what has already been cached if an upshift is performed. EX1001, 12:43-53, Fig. 4; *supra* § III.D. The fifth is to “pre-request” streamlets of a different bit rate based on an “anticipate[d]” need to shift bit rates. EX1001, 8:46-50; *supra* § III.D. And the sixth is to “request streamlets at any interval less than the length of the time index.” EX1001, 8:53-57; *supra* § III.D. And while each of these exemplary techniques enables claims directed

² For example, that the '564 Patent can begin streaming after retrieving only a single streamlet also allows “instant playback,” as recited in dependent Claim 7.

to a “shift ... to achieve continuous playback,” the independent claims that include this recitation are not limited to any one these techniques.³

63. Dr. Houh relies on Leaning alone as allegedly anticipating the “continuous playback” element of each of the above-recited claim limitations. EX1003, ¶¶ 81-82. But, unlike the Challenged Patents, which disclose numerous techniques for achieving continuous playback with a rate shift, Leaning fails to disclose any technique whatsoever for achieving continuous playback. Instead, Leaning consistently teaches clearing and re-filling a playback buffer as part of every shift. This is the same “progressive download” approach disparaged by the Challenged Patents.

64. Indeed, far from reading on a “shift” that “achieve[s] continuous playback,” Leaning intentionally creates the exact type of interruptions in playback that the claimed invention avoids.

65. Dr. Houh expressly relies on Leaning’s “Clear Buffer” function and “Buffer Fullness > Tp” test as allegedly disclosing “continuous playback.” EX1003 ¶ 81. These are the very features that force Leaning’s playback process to wait for its buffer to be refilled before resuming playback after a rate shift. Thus, they are the

³ For example, dependent Claim 5 of the ’564 Patent expressly recites “upshift[ing] ... according to an amount of contiguously available files stored by the media player” and dependent Claim 6 expressly recites “pre-request[ing]” streamlets of a different bit rate based on an “anticipate[d]” need to shift.

very features that *disrupt* playback. And they do so *every time* a rate shift is performed.

66. Even if, for the sake of argument, Leaning did not intentionally disrupt playback with its buffer clearing and refilling, Petitioner and Dr. Houh still fail to carry their burden to show invalidity of the Challenged Claims directed to “video” by failing to establish that the teachings of Leaning they rely upon actually apply to video. Petitioner and Dr. Houh cite almost exclusively to Leaning’s discussion of rate shifting *audio* in their invalidity arguments. Additionally, as explained above, in its minimal disclosures of video, Leaning is primarily focused on using H.261, which encodes, stores, and transmits video separately from audio.

67. The only portion of Leaning Dr. Houh cites that relates to video does little more than vaguely state that “[t]he same [unstated] principle may be applied to the delivery of video recordings.”⁴ EX1004, 16:6-7 (cited at EX1003, ¶¶ 56, 60). Things are not that simple, as Leaning admits with its explanation that “[t]he delivery of video recordings [has] *further implications* for file partitioning if switching or jumps are to be permitted.” EX1004, 16:17-18 (emphasis added). These “further

⁴ Although Dr. Houh also cites to another portion of Leaning that also discusses video, it is directed to fast forwarding and rewinding functionality, which are a distinct (and separately claimed) functionality from rate shifting. See EX1003 ¶ 90 (citing EX1004, 16:11-18:17 as allegedly “disclos[ing] functionality for fast forward and fast reverse modes”). It is therefore largely irrelevant to the argument below, which focusses primarily on the claimed rate shifting functionality.

implications” include “serious mistracking of the decoder,” which is yet another manner in which Leaning’s rate shifting interrupts playback. *Id.*, 16:30-31. Thus, even if Leaning did not inject delays into playback through its progressive download approach, Petitioner and Dr. Houh would still fail to carry their burden of showing that the same principles Leaning discusses with respect to *audio* could also achieve “continuous [playback/presentation]” when retrieving *video*, as required by each of the Challenged Claims.

1. Dr. Houh Fails to Address the “Continuous [Playback/Presentation]” Limitation

68. Dr. Houh’s invalidity opinions ignore the recitation of “continuous [playback/presentation].” For example, when addressing element 1[D] in the ’564 Patent, Dr. Houh’s only mention of the “continuous playback” requirement is to state that it is “explained above” (with respect to elements 1[C][1] and 1[C][2]) that “Leaning discloses making the successive determinations to shift the [playback/presentation] quality based on the factor to achieve *continuous [playback/presentation]* of the video using the files (sub-files) of the highest quality copy of the video that is determined to be sustainable at that time.” EX1003 ¶ 81 (emphasis added). But neither Dr. Houh’s discussion of elements 1[C][1] and 1[C][2], nor his discussion of element 1[D], otherwise address the “continuous [playback/presentation]” limitation of the Challenged Claims. *See id.*, ¶¶ 77-82. Dr.

Houh makes this same mistake in each declaration.⁵ '156 Patent EX1003, ¶¶ 81-85; '772 Patent EX1003, ¶¶ 88-92; '234 Patent EX1003, ¶¶ 71-79, 105.

69. Indeed, the word “continuous” does not appear anywhere in those discussions other than in the quoted language above. Thus, it is my opinion that Petitioner and Dr. Houh fail to carry their burden of showing that every requirement of element 1[D] is disclosed by Leaning—the only reference Petitioner and Dr. Houh rely upon as allegedly teaching this limitation.

70. Continuing with the example of element 1[D] of the '564 Patent, that limitation recites: “making the successive determinations to shift the playback quality based on the factor to achieve continuous playback of the video using the files of the highest quality one of the copies determined sustainable at that time so that the media player upshifts to a higher quality one of the different copies when the factor is greater than a first threshold and downshifts to a lower quality one of the different copies when the factor is less than a second threshold,” with highlighting added to distinguish between the different limitations of the element. Without conceding that Petitioner and Dr. Houh satisfy their burden as to any of these requirements, below is the entirety of what Dr. Houh argues with respect to each, with the above highlighting carried through for clarity.


⁵ While 1[D] in the '234 Patent does not use the word “continuous,” I note that 12[D] does, and that the analysis for 12[D] in the '234 Patent simply refers back to 1[D] for support. EX2134 at 49.

Element 1[D]	Dr. Houh's Opinion
“making the successive determinations ... using the files of the highest quality one of the copies determined sustainable at that time”	“Leaning teaches choosing the highest quality copy that is sustainable (<i>i.e.</i>, if measured rate > current rate AND measured rate > next highest rate for period of time, then shift up until no longer true—meaning shift up to highest sustainable version).” “Leaning discloses looping ‘Stepup’ until the highest sustainable rate is achieved up to the highest available rate (by setting RateU to zero).”
“to shift the playback quality based on the factor to achieve continuous playback of the video”	—
“so that the media player upshifts to a higher quality one of the different copies when the factor is greater than a first threshold”	“the player on the terminal beginning with step J1A moves through the process and if ‘MRate [measured rate] is greater than NextRate [next highest rate] AND NextRate < 0 [there exists a next higher rate] then go to Stepup; and the Stepup process shifts up by changing the RateC [current rate] to the RateU [next highest rate] and changes RateU to the next higher rate, or to ‘0’ if the current rate is the highest rate then loops back to step J1A and continues the process” “Leaning further discloses that that the media player (player program) upshifts to a higher quality one of the different copies when the factor (the ‘measured rate’ factor) is greater than a first threshold (the ‘next higher rate’ for a given period of time)”
“so that the media player ... downshifts to a lower quality	“Leaning further discloses that that the media player (player program) ... downshifts to a lower

one of the different copies when the factor is less than a second threshold”	quality one of the different copies when the factor is less than a second threshold (‘current rate’ (and when the Buffer Low Percentage more than a Step Down Threshold))”
--	--

71. Dr. Houh spends most of his opinion addressing the “upshift” requirement and none of his opinion (here or with respect to elements 1[C][1] and 1[C][2]) address the “continuous [playback/presentation]” requirement. But particularly telling in this regard is that, with respect to the ’234 Patent, petitioner Aylo Freesites Ltd. (“Aylo”) purportedly addressed element 12[D], which recites the “continuous presentation” limitation, by citing exclusively to its analysis of element 1[D] of that patent. EX2134 (’234 Patent Petition), 49. But Claim 1 of the ’234 Patent does not recite or require shifting to achieve “continuous presentation.” This confirms that petitioner simply ignores the “continuous” limitation.

72. Dr. Houh also includes a single demonstrative in support of his 1[D] opinion, which is lifted from Leaning’s pseudo-code flowchart. EX1003 ¶ 81. But like his written opinions, Dr. Houh’s demonstrative, reproduced below, also focusses primarily on the “upshift” requirement of element 1[D] while completely ignoring the “continuous [playback/presentation]” requirement.



EX1003 ¶ 81 (citing EX1004, 10) (red annotations original to Dr. Houh's declaration, yellow highlighting added). In particular, Dr. Houh ignores the highlighted portions of the demonstrative, which establish that *every time* Leaning performs an "upshift" ("Set RateU = next higher rate in list or zero if none"), it clears the buffer ("Clear Buffer") and then *waits* until the buffer is refilled ("If Buffer Fullness > Tp") before resuming playback ("Begin/continue playing"). Petitioner

and Dr. Houh’s own demonstrative therefore shows that “[playback/presentation]” in Leaning is **not** “continuous,” during a rate shift, as required by element 1[D].

73. In sum, petitioner and Dr. Houh not only fail to address the “continuous [playback/presentation]” requirement of element 1[D] of the ’564 Patent, the portions of Leaning they cite for the other requirements of element 1[D]—*e.g.*, the “upshift” requirement—establish that playback in Leaning is, in fact, **not** continuous, but rather interrupted **every time** a rate shift is performed. *Infra* § VII.A.2. Petitioner and Dr. Houh’s analysis of the “continuous playback/presentation” limitation of each other challenged patent contains the very same flaw demonstrated above with regard to the ’564 Patent.

2. Leaning Fails to Disclose “Continuous [Playback/Presentation]” Because It Intentionally Interrupts Playback at Each and Every Rate Shift.

74. Leaning’s only disclosed “rate change” process is limited to audio only and includes a single step—Step **J3**—in which [playback/presentation] occurs.

J3: Begin/continue playing of buffer contents via <u>sound card</u>
--

EX1004, 10 (underline added). And that step is only performed under one condition: if “the contents of the buffer reach or exceed the playout time T_p ,” which is determined at Step **J1A**.


J1A: If Buffer Fullness > T_p seconds go to Step J3
--

Id., 5:24-25, 10. Thus, Leaning requires that its “buffer” be filled above a threshold—*i.e.*, the “playout time T_p ”—before playback begins.

75. In Leaning’s only example, the playout time (T_p) is 10 seconds. EX1004, 5:20-21. Leaning’s exemplary embodiment therefore requires that at least 10 seconds of content be stored in the buffer before playback can begin (or resume), which, using Leaning’s “preferred” sub-file length of “four seconds” (*id.*, 3:6-7), would require sequentially downloading and decoding at least three sub-files (or 12 seconds of content) before beginning (or resuming) playback.

76. Leaning acknowledges that this buffer filling process necessarily results in “**interruptions** in the output.” EX1004, 15:3-6 (emphasis added); *see also id.*, 12:31-34 (describing how downloading large sub-files “caus[e] **extra delay** when starting play”) (emphasis added). And such “interruptions” occur not only at the start of streaming, but also **every time** Leaning makes an upward or downward rate shift during streaming. That is because Leaning clears its buffer **every time** it makes an upward or downward rate shift.

77. This can most clearly be seen in Leaning’s pseudo-code flow chart, which shows that a “Clear Buffer” function is executed as the **first step** in both the upshift (“Stepup”) and downshift (“Stepdown”) processes when they are called.




EX1004, 10 (annotations added). Also shown in the flow chart is that the next step after an upshift (“Stepup”) or downshift (“Stepdown”) is Step **J1A**, which is the step that determines whether the buffer has been refilled before continuing playback at Step **J3**. Thus, playback is interrupted *every time* Leaning performs an upshift (“Stepup”) or downshift (“Stepdown”) because Leaning clears its buffer *every time* it performs an upshift (“Stepup”) or downshift (“Stepdown”).


78. A POSITA would understand that clearing the playback buffer as taught by Leaning would result in at least two independent types of playback discontinuities. *First*, clearing the buffer would result in a *temporal discontinuity*, where the last decoded frame would likely remain “frozen” in place on the screen until the buffer is refilled and playback “continues.” *Second*, clearing the buffer

would result in a *content-based discontinuity* in which any sub-files stored in the buffer at the time the rate shift occurs will be skipped because they are removed from the buffer and not requested.

79. The temporal discontinuity occurs because the buffer was emptied as part of the rate shift and there is a waiting period while the buffer refills. And the content-based discontinuity occurs because, as depicted below, Leaning's playback process increments the sub-file number sequentially every time a sub-file is written to the buffer, without accounting for the buffer being cleared in the future.



EX1004, 10 (annotations added). This means that the next sub-file requested after a rate shift will be the sub-file *after* the last sub-file previously stored in the buffer, which may not have been played yet, as depicted below.




Id., 3:12-18 (annotations added and “100” added to initial four sub-file names to simplify the illustration). In this example, on a rate shift occurring after sub-file “000102.bin” is retrieved but before sub-files “000101.bin” and “000102.bin” are played back, sub-files “000101.bin” and “000102.bin” will be cleared from the buffer and will never be played. Playback/presentation will later resume with sub-file “000103.bin” after the buffer refills. Either one of these temporal or content-based discontinuities alone is sufficient to show that Leaning fails to teach or suggest “continuous [playback/presentation].”⁶

⁶ Examples of Leaning’s temporal and content-based discontinuities are illustrated in the demonstrative video of Leaning’s pseudocode flowchart created at my direction and filed herewith as EX2133. Specifically, EX2133 is a video that demonstrates the application of Dr. Houh and Petitioners’ obviousness theory applying the teachings of Leaning. In the video, a standard sample video is played according to the pseudocode provided in Leaning. The sample video has been broken into four-second sub-files at 128kbps, 256kbps, and 400kbps encodings of 25 frames per second. A frame index (count) has been added to the lower-left corner of each frame of the video. During normal playback of the sample video,

80. In addition, although Leaning describes an alternate embodiment in which downshifting (“Stepdown”) can occur even if the buffer is not empty when the decision to downshift is made—*i.e.*, “the first of the three AND conditions mentioned above (namely, that the buffer is empty) would be omitted” (EX1004, 15:3-8)—Leaning consistently teaches that the buffer must *still* be cleared in the very first step of the downshift (“Stepdown”) process itself, as depicted below.

the video should play without disruption, with each frame being displayed in order from the first to the last frame. In the video demonstrative (EX2133), the playback of the sample video is shown as the network bitrate changes between 150kbps, 300kbps, and 500kbps. Leaning’s pseudocode flowchart results in numerous playback discontinuities as shown. The demonstrative video also shows a representation of the fullness of the decoded frame buffer at the bottom, with the variable T_p depicted. When playback of the sample video begins, the pseudocode begins playback once the buffer reaches T_p . As the sample video is played and the network bitrate is changed, the pseudocode is highlighted to show the resultant playback discontinuities. Notably, this demonstrative shows that Leaning’s pseudocode flowchart would result in the sample video freezing (temporal discontinuity) and dropping a substantial number (*e.g.*, hundreds) of frames before resuming playback.



EX1004, 10, 15:3-8 (annotations added). Thus, the buffer will still need to be refilled until “the contents of the buffer reach or exceed the playout time T_p ” before playback can continue. This is consistent with Leaning’s teaching that it is “advantageous” to clear the buffer during a rate shift because “the sound card is not playing anything” at that time, which allows the sound card to be “reconfigured.” EX1004, 5:24-25, 8:20-23. It also is consistent with Leaning’s teaching that clearing the buffer is preferable because it allows a request for a “new” streamlet to be “sent immediately” during a rate shift. EX1004, 11:19-20.


81. Aylo Premium Ltd., who I understand to be a real party in interest in IPR2024-00940 (the ’234 Patent IPR), expressly acknowledged Leaning’s discontinuous playback in submitting arguments in a proceeding before the European Patent Office (EPO) on a related patent. EX2016 (Aylo EPO Response). In that proceeding, Aylo argued Leaning teaches that, “for each upshift, there is a

jump in the file by at least 120 seconds or 2 minutes.” *Id.*, 9. I agree with Aylo that Leaning’s disclosed pseudo-code flowchart requires clearing a buffer at each upshift, which results in removing and not playing previously buffered content. I also agree that Leaning does not disclose or suggest subsequent (*i.e.*, redundant) requests for any subfiles which were cleared from the buffer. Accordingly, I agree that Leaning discloses content discontinuities (in addition to temporal discontinuities I discuss above). I therefore agree with Aylo’s EPO Response argument to the extent it categorizes Leaning as teaching content discontinuities.

82. However, I disagree with Aylo’s EPO argument regarding the amount of content that would be “cleared” from the buffer. Specifically, Aylo contends that two minutes of content would be dumped when clearing the buffer:

when upshifting from the lowest quality to the next higher quality, the buffer will have been filled for 120 seconds with a rate higher than 16 kbit/s [] and at the same time been drained with a rate of 8 kbit/s This corresponds to at least 120 seconds of playing time ... [such that] clearing the buffer would mean that for each upshift, there is a jump in the file by at least 120 seconds or 2 minutes.

EX2016 (Aylo EPO Response), 8-9. But this overlooks Leaning’s pseudocode flowchart, which only requests a new subfile when “Buffer Fullness” is less than “Tp,” as shown below:



EX1004, 10. Thus, Leaning's buffer can, at most, fill above the playout time threshold T_p by the duration of a single sub-file. Accordingly, for Leaning's disclosed playout time T_p of 10 seconds with four second sub-files, the largest amount of content stored in the buffer would be 14 seconds. And some content would likely be played out from the buffer while downloading the next sub-file and before Leaning's rate shift determination is made. Accordingly, a POSITA would understand that clearing the buffer during Leaning's upward rate shift would result in 10-12 seconds of content discontinuity. Aylo's arguments regarding 120 seconds of buffered content being lost are not supported by Leaning's disclosures. Aylo nevertheless acknowledges that Leaning intentionally inserts a delay when switching rates. That is the part of the EPO response I agree with.

83. At bottom, Leaning’s playback process is the same as the “progressive download” process that the Challenged Patents criticize because it “*waits* to begin playback” until there are enough sub-files downloaded that playback can resume “with the hope that,” once playback does resume, it does not “catch[] up” to the buffer. EX1001, 1:64-2:1. And as the Challenged Patents explain, this wait “can be substantial,” such that Leaning’s playback process “is not a complete or fully acceptable solution to the problem of media presentation over a network.” EX1001, 2:2-5.⁷ Thus, Leaning not only fails to disclose or suggest a “shift ... to achieve continuous [playback/presentation],” it teaches away from such a shift by lauding the benefits of clearing its buffer and interrupting playback—namely, the time it provides for the system to “reconfigure” the sound card and to retrieve the next sub-file “immediately.” EX1004, 5:24-25, 8:20-23, 11:19-20.

3. Dr. Houh Fails to Show that Leaning’s Teachings Regarding Delivering *Audio* Can be Applied to the Claimed Requirements for Delivering *Video*.

⁷ Leaning also teaches discontinuous playback in the context of “errors” encountered when streaming. Specifically, Leaning teaches to simply “stop,” or terminate the streaming process, if any error is received when downloading a subfile. EX1004, 10. By contrast, the Challenged Patents acknowledge that streamlet requests may “fail,” or may “not be[] satisfied in a timely manner,” and discloses specific solutions for error handling, including sending redundant requests for the same streamlet at either the same or a different bitrate. EX1001, 9:59-67. The Challenged Patents’ error handling processes further enable achieving “continuous playback” but are entirely absent from Leaning.

84. Petitioner and Dr. Houh’s entire argument regarding Leaning’s alleged teaching of element 1[D] (and the counterparts to this limitation that appear in the other Challenged Patents) is based on Leaning’s discussion of delivering *audio*. And while Dr. Houh contends (*for an entirely different element*) that the “same [unstated in Leaning] principle may be applied to the delivery of video recordings,” the shortcomings of Leaning’s progressive downloading process are compounded when applied to the delivery of video recordings. EX1004, 16:6-7; *see also* EX1003 ¶ 60 (citing EX1004, 16:6-12). Indeed, Leaning itself acknowledges that there are “further implications” when applying this “same principle” to video recordings. *Id.*, 16:17-18. Dr. Houh fails to consider these “further implications” and, therefore, fails to show that the principles he relies upon in Leaning with respect to *audio* could also achieve “continuous playback” when delivering *video*.

85. Tellingly in this regard, Dr. Houh relies almost exclusively on Leaning’s teachings with respect to delivering audio in arguing the invalidity of the ’564 and ’772 Patents; not just with respect to “shift[ing] ... to achieve continuous playback,” but with respect to nearly every element of every challenged claim.⁸ The one exception is Claim 6 of the ’564 Patent, which is directed to “fast-forward ...

⁸ To be clear, Petitioner does not cite to anything but audio-related portions of Leaning as allegedly teaching the “shift ... to achieve continuous [playback/presentation]” limitation, nor do they cite to any other reference as allegedly disclosing this limitation.

and rewind” functionality (Claim 6). But this is entirely different functionality than rate shifting. In fact, playback interruptions are inherent in fast-forward/rewind functions and, therefore, expected. Thus, Dr. Houh’s citations to this video-related functionality are irrelevant to whether Leaning can “achieve continuous playback” while *rate shifting*.

86. The only other portions of Leaning Dr. Houh cites to that are directed to video-related functionality do little more than state that the “same [unstated] principle” discussed above “may be applied to the delivery of video recordings,” and that, if applied, the “two or more [video] recordings corresponding to different data rates” would be “selected by the control mechanism *already described*”—*i.e.*, the mechanism by which the buffer is cleared and re-filled every time a rate shift is performed. EX1004, 16:6-12 (cited in EX1003 ¶ 60). Dr. Houh conveniently ignores what appears just a few lines later—in particular, Leaning’s explanation that “the situation is more complex” and “*serious mistracking* of the decoder would occur” if the same principle were applied to video.⁹ *Id.*, 16:20-21, 16:25-31.

87. As to the ’156 Patent, Dr. Houh relies almost exclusively on Leaning’s teachings with respect to delivering audio in arguing invalidity; not just with respect to “shift[ing] ... to achieve continuous playback,” but with respect to every element

⁹ Decoder mistracking is an error that occurs when a decoder loses track of an image, or a portion of an image required for decoding the present image.

of every challenged claim. The only portions of Leaning Dr. Houh cites to that are directed to video-related functionality do little more than state that an unidentified “same principle” discussed above “may be applied to the delivery of video recordings,” and that, if applied, the “two or more [video] recordings corresponding to different data rates” would be “selected by the control mechanism *already described*”—*i.e.*, the mechanism by which the buffer is cleared and re-filled every time a rate shift is performed. EX1004, 16:6-12 (emphasis added; cited in IPR2024-00917, EX1003 ¶ 60) (emphasis added). Dr. Houh conveniently ignores what appears just a few lines later—in particular, Leaning’s explanation that “the situation is more complex” and “*serious mistracking* of the decoder would occur” if the same principle were applied to video. *Id.*, 16:20-21, 16:25-31 (emphasis added). Dr. Houh gives the same treatment (or lack thereof) to the ’234 Patent.

88. Even if, for the sake of argument, Leaning’s reference to applying the “same [unidentified] principle” to video recordings means the “principle” Leaning discusses with respect to delivering audio recordings, this “principle” still invokes the same buffer clearing function (and corresponding playback interruptions) discussed above. But Dr. Houh fails to show that this principle could actually be applied to video in the first instance because Dr. Houh fails to address the additional complexities that are introduced when Leaning’s playback process is applied to video. As to the “continuous playback” limitation, Dr. Houh fails to point to a single

video-related portion of Leaning’s disclosure as corresponding to this limitation, and by failing to address, much less acknowledge, the additional complexities that must be addressed when applying rate shifting technology to delivering video.

4. A POSITA Would Not Discard Leaning’s “Clear Buffer” Instruction

89. To achieve continuous playback, the self-contained teaching of Leaning would have to be fundamentally changed. Indeed, a POSITA would have to abandon one of the central teachings of Leaning—specifically, the clearing of the buffer when changing the data rate. Leaning explains how clearing the playback buffer is “advantageous” and “may be necessary” during a “rate shift” because it provides time to “reconfigure” the “sound card” and to retrieve the next audio sub-file “immediately.” EX1004, 8:20-23, 11:19-20. Although Leaning primarily addresses these benefits in the context of a “sound card,” a POSITA would have understood that those teachings applied equally to hardware video cards at the time of Leaning for the decompression and display of full-motion, color video, which also would have needed to be cleared and reset in order to decode video according to the specific parameters of any new incoming bitstream. EX2139 (’332 Patent) at 10:18-32 (“for the decoder to adapt to the context” of “one MPEG transport stream to another,” “undesirable audio and video discontinuities may result.” “One attempt to solve this discontinuity problem is to reset the decoder, squelch the audio and freeze the video for a certain amount of time”).

90. Leaning's teachings are entirely consistent with this technical requirement of video cards at the time of Leaning. Leaning recognizes that a video card would be required for video playback, and that "[f]or video delivery, a video card would be fitted in place of, or in addition to, the [sound] card." EX1004, 4:9-10. But Leaning does not make any special accommodation for video cards. Thus, when Leaning teaches to "clear [the] buffer," a POSITA would understand that instruction (possibly Leaning's "same principle" to also apply to the decode buffer of a video card in the same way it applies to a sound card). Indeed, Petitioner's expert, Dr. Houh, admitted during cross-examination that (i) Leaning expressly discloses clearing the sound card buffer when rate switching and (ii) Leaning does not disclose or suggest that the corresponding video card buffer would **not** be cleared for rate switching. EX2129 (Houh Tr.) 103:18-104:9.

91. Moreover, even in the embodiment where "the first of the three AND conditions" is "omitted" such that the buffer need not be empty *at the same time* the downshift decision is made (EX1004, 15:3-8), Leaning's "preferred" configuration requires the playback buffer to empty numerous times **before** downshifting. Specifically, Leaning teaches that setting the "Step Down Threshold" (*i.e.*, "Td," or the third of the three AND conditions above) to the "preferred default value" of 5% means that, "**if the buffer empties** but the measured Buffer Low Percentage is greater 5% it **will not step down.**" EX1004, 8:27-9:2 (emphasis added). Here, Leaning

expressly acknowledges that “[f]urther *buffer empties* will obviously cause this measured rate to increase and will eventually *empty the buffer again*.” *Id.*, 8:32-9:2 (emphasis added). Thus, even if “the first of the three AND conditions” is omitted, Leaning’s “preferred” embodiment results in multiple “buffer empties”—and, in turn, multiple playback discontinuities—prior to every downshift. Removing the “clear buffer” instruction has no bearing on these discontinuities.

5. Leaning’s Teaching to *Encode Then Segment Video* Cannot Achieve “Continuous Playback”


92. Dr. Houh relies solely on Leaning’s *audio partitioning* techniques as allegedly anticipating the limitations of the Challenged Claims, all of which require streaming of “video” except the claims of the ’234 Patent. EX1003 ¶ 60 (“The **same principle** [pertaining to audio recordings] may be applied to the delivery of **video recordings**.... The manner of partitioning the file into sub-files is unchanged”) (quoting EX1004, 16:9-10) (emphasis in original).

93. The audio-based content partitioning technique of Leaning relied on by Dr. Houh involves two steps. First, the same source content is encoded at several bitrates. Second, the encoded content is “partitioned” into sub-files.

In order to provide for rate switching, the person preparing the file for loading onto the server prepares several source files - by *encoding* the same PCM file several times at different rates. He *then partitions* each source file into sub-files

EX1004, 6:21-23 (emphasis added). Thus, Dr. Houh relies on Leaning’s disclosed approach to audio segmentation that involves “encoding ... *then* partition[ing]” to arrive at the claimed “files” or “streamlets.” EX1003 ¶¶ 60-61 (emphasis added).

94. But Dr. Houh overlooks that partitioning *video* files in this manner would prevent the decoder from properly decoding any partitioned files received after a rate switch, yielding discontinuities and undesirable artifacts in playback. Leaning refers to this problem as “mistracking.” EX1004, 16:30. The results of “mistracking” are dramatic, as demonstrated by the images below, where the left image shows a properly decoded frame taken from a complete media file and the right image shows the same frame taken from a portion of that same media file:



The left image was decoded sequentially from the beginning of the media file, while the right frame was decoded sequentially from the beginning of the portion of the media file, rather than the media file itself. The result is that the right image was

decoded without the benefit of the preceding intra-frame, or “I-frame,” needed to decode it, which results in a significant loss in image data.

95. The images above were generated from a video file called “akiyo.” The “akiyo” video is a well-known video for demonstrating and testing video codecs. I downloaded the uncompressed YUV4MPEG format of the akiyo video at CIF resolution (352x288) from <https://media.xiph.org/video/derf/>. I have included a copy of the akiyo video as EX2136.¹⁰ I opened the uncompressed akiyo video in the freely available VLC media player and played it on my computer. The video played back at the expected CIF resolution without any playback issues.

96. The uncompressed akiyo video was then compressed in H.261 format specifying a 128 kbit/second bitrate using the ffmpeg tool, which is a widely-used open source video codec. <https://www.ffmpeg.org/>. The command for encoding the uncompressed akiyo video follows:

```
ffmpeg -i akiyo_cif.y4m -b:v 128k -g 99999 -c:v h261  
akiyo_cif.h261.
```

97. Next, the following command was run to segment the h.261 encoded version into four second sub-files:

```
ffmpeg -i akiyo_cif.h261 -c copy -f segment -  
segment_time 4 segment_cif_%03d.h261
```

¹⁰ I was told by counsel that the PTAB’s PTACTS filing system only accepts MPEG4-encoded videos. Accordingly, the uncompressed akiyo video is encoded in MPEG4 format in order to facilitate filing. I reviewed the MPEG4 compressed version of the video and confirmed it plays as expected at the CIF resolution.

98. Executing this command created three sub-files of the akiyo video. While the first sub-file plays at CIF resolution as expected, the second and third sub-files do not. Those sub-files demonstrate the concept of “mistracking” discussed by Leaning. The image on the right side above was generated by taking a screenshot during playback of the second subfile. I have included a copy of the second sub-file at EX2135.¹¹

99. “Mistracking” arises under these circumstances for video, but not audio, because each frame of an audio file is *independently* decodable without any dependencies between frames. By contrast, video is encoded using a mixture of *independently decodable* frames (I-frames), and *dependently decodable* frames (P- or B-frames), the latter of which can only be decoded if the decoder has access to other surrounding frames. As Leaning explains, problems arise when the player has to decode a dependently decodable frame for which the decoder does not have the surrounding frames from which it depends. EX1004, 16:26-31. Leaning expressly discloses that “mistracking” will occur when rate switching between videos compressed via “inter-frame” coding techniques, which refers to all video compression standards, including the MPEG and H.261 mentioned in Leaning.

¹¹ I transcoded the second h.261 sub-file from the segmented akiyo video in MPEG4 format in order to comport with the PTACTS filing system. I reviewed the MPEG4 compressed version to confirm it displays the same mistracking effect I experienced when playing the h.261 encoded version.

100. Although Dr. Houh does not identify or rely upon any specific video encoding embodiment, the “mistracking” problem infects each video format mentioned in Leaning. As Dr. Houh explained at deposition, “the mistracking problem” occurs regardless of encoding format “when a terminal gets a predictively coded frame for which it does not have the preceding frame.” EX2129 (Houh Tr.) 115:18-21. First, “mistracking” occurs with H.261 encoded video because that standard “do[es] not provide for the frequent, regular inclusion of intra-frames [or I-frames].” EX1004, 16:24-31. Thus, as Leaning recognizes, video encoded in H.261 and then segmented at arbitrary boundaries will almost certainly result in sub-files comprising only P-frames without the prerequisite frames required for decoding.

101. Second, MPEG encodings would also suffer from “mistracking.” A POSITA would have understood that to ensure the highest video quality at a given encoded bit-rate (an obvious goal), generating multiple full MPEG-encoded videos at different bit rates would require differing GOP structures and hence produce video segments with differing I-frame locations in the corresponding segments encoded at different bit-rates. Segmenting multiple MPEG encodings at the same time boundaries would result in sub-files that begin with B- or P-frames, and that do not include the preceding reference frames required to decode those P-frames. Thus, mistracking would result when attempting to decode those leading B- or P-frames after a rate change. *See, e.g.*, EX2139 (’332 Patent) 10:23-25, 2:6-22 (describing this

problem in the context of MPEG transport stream “splicing” and decoding in set-top boxes).

102. Moreover, even if encoded MPEG files were segmented at I-frame boundaries, switching between sub-files would still result in discontinuous playback because the I-frame boundaries, and resulting sub-file boundaries, are not guaranteed to be time-aligned.

103. Critically, Dr. Houh failed to address this “serious mistracking” problem for video as expressly disclosed by Leaning, failed to cite any disclosure of video rate switching in Leaning that might resolve “mistracking,” and failed to explain how or why a POSITA would modify Leaning’s teaching to resolve “mistracking” in video. Without any cognizable theory to overcome Leaning’s disclosure of *dis*continuous in media playback due to “mistracking,” it is my opinion that Dr. Houh and Petitioner have failed to meet their burden to show unpatentability.

B. Leaning Fails to Disclose Requesting Streamlets via “multiple Transmission Control Protocol (TCP) connections” / “a plurality of different TCP connections”

104. As listed below, some of the Asserted Claims recite requesting streamlets via “multiple Transmission Control Protocol (TCP) connections” or “a plurality of different TCP connections.” For example, dependent Claim 2 of the ’564

Patent recites “the at least one TCP connection comprises multiple Transmission Control protocol (TCP) connections.”

Challenged Claim	Element
'564 Patent Claim 2	The end user station of claim 1, wherein the at least one TCP connection comprises <i>multiple Transmission Control protocol (TCP) connections</i> with the content server.
'564 Patent Claim 9	The method of claim 8, wherein the at least one TCP connection comprises <i>a plurality of different TCP connections</i> , and wherein the requesting the plurality of sequential files includes requesting sub-parts of the files over different ones of the plurality of different TCP connections, and wherein said presenting includes reassembling the files from the received sub-parts.
'564 Patent Claim 11	The method of claim 8, wherein the at least one TCP connection comprises <i>a plurality of different TCP connections</i> , and wherein the automatically requesting includes requesting sub-parts of the files over different ones of the plurality of TCP connections, and wherein said presenting includes reassembling the files from the received sub-parts, and wherein the factor is indicative of the available bandwidth of the plurality of TCP connections.
'156 Patent Claim 2	The apparatus of claim 1, wherein the apparatus is configured to establish <i>multiple Transmission Control Protocol (TCP) connections</i> with a content server, and request streamlets of varying bitrates.

'772 Patent Claim 1[B]	requesting by the media player a plurality of sequential ones of the files of one of the copies from the set of servers over <i>a plurality of Transmission Control Protocol (TCP) connections</i> based on the time indexes;
'772 Patent Claim 1[C]	automatically requesting by the media player from the set of servers over <i>the plurality of TCP connections</i> 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,
'772 Patent Claim 7	The method of claim 1, wherein said requesting the plurality of sequential ones of the files includes requesting the plurality of sequential ones of the files over <i>different ones of the plurality of TCP connections</i> .
'772 Patent Claim 9	The method of claim 1, wherein said automatically requesting includes requesting sub-parts of the files over <i>different ones of the plurality of TCP connections</i> , and wherein said presenting includes reassembling the files from the received sub-parts.
'234 Patent Claim 16	The method of claim 12, wherein the at least one TCP connection comprises <i>a plurality of different TCP connections</i> , and wherein the requesting of the plurality of sequential individually-requestable portions includes requesting sub-parts of the individually-requestable portions over <i>different ones of the plurality of different TCP</i>

	<i>connections</i> , and wherein the arranging includes reassembling the individually-requestable portions from the requested sub-parts.
--	--

105. As the Challenged Patents explain, “a break in the continuity of a TCP connection can have serious consequences when the data must be delivered in real-time,” and “delays or losses in a TCP connection” can cause breaks or slowdowns in data transmission that are “detrimental to the viewing or listening experience of the user and therefore ... not acceptable.” EX1001, 2:13-21. The Challenged Patents avoid these delays/losses and maintain continuous playback by using multiple TCP/IP connections to retrieve content. *Id.*, 9:34-35.

106. More specifically, the Challenged Patents describe making a new, redundant content request on a different TCP connection when a content request “fails completely” or “is not satisfied in a timely manner” via the original TCP connection.

When a TCP connection *fails completely*, a *new request* may be sent on a different connection for the same streamlet **212**. In a further embodiment, if a request is *not being satisfied in a timely manner*, a *redundant request* may be sent on a different connection for the same streamlet.

EX1001, 9:59-63 (emphasis added). “[E]fficiency is raised significantly and latency is virtually eliminated” when multiple TCP connections are used in this manner. *Id.*, 9:43-47.

107. Leaning, on the other hand, makes no mention of, and has no use for, multiple TCP connections. To be sure, Leaning explicitly states that it “need contain no ... provision” for “the masking of errors due to packet loss” so that “an ordinary ‘web server’” may be used to implement its playback process. EX1004, 2:27-32. A POSITA would understand that an ordinary web server would use TCP to deliver content and thus would experience the exact pathologies the Challenged Patents concern themselves with and overcome.

108. Transmission errors from a server in Leaning are an unavoidable and unavoids component of Leaning’s playback process for multiple reasons. For example, as explained above, Leaning was assigned to British Telecom and clearly related to the use of telecom circuits, which had reliable performance. Thus, loss was very unlikely to occur with respect to audio. The transmission of video, on the other hand, was likely to cause network congestion. Leaning does not account for this fact. Instead, Leaning actually builds in (*i.e.*, creates) losses as a regular part of its playback process.

109. As Leaning explains, its playback process “is repeated until a ‘file not found error’ is returned” by the server. EX1004, 6:1-2. A POSITA would readily

appreciate that “an ordinary ‘web server’” would generate a standard “HTTP 404 – File not found” error message under these circumstances. This unique use of standard “error message[s]” also is reflected in Leaning’s pseudo-code flow chart, which shows that such a message is created when a requested sub-file does not exist, and that the user terminal will “Stop” the playback process when it receives this “error message.”

Terminal	Server
Request this URL	
	If requested subfile exists, Send requested subfile; otherwise send error message
If error message received, Stop	

Id., 10 (highlighting added). Thus, rather than seeking to avoid errors, Leaning seeks to create and use them to signal the end of a content file.

110. An HTTP or TCP socket error message also may satisfy the “If error message received, Stop” condition in Leaning’s playback process. If it does not, the playback process would simply stall while waiting for the requested sub-file to be received and decoded—*i.e.*, “Decode received subfile”—as also depicted in Leaning’s pseudo-code flow chart.



Terminal	Server
Set Current Subfile = 000000.bin	
J1: Set URL = StemC + Current Subfile	
Request this URL	
	If requested subfile exists, Send requested subfile; otherwise send error message
If error message received, Stop	
Decode received subfile	
Write to buffer	

EX1004, 10 (annotation and highlighting added). This could, in turn, result in the buffer emptying while previously downloaded content continues to play out. And, “[i]f, during this process, the buffer becomes empty, the player program simply ceases playing until further data arrives.” *Id.*, 6:3-4. Leaning’s willingness to allow these types of playback interruptions to occur while it awaits sub-files to download further demonstrates Leaning’s disinterest in addressing transmission errors of any type, and its corresponding disinterest in using multiple TCP connections to ameliorate errors.

111. Furthermore, Leaning describes requesting sub-files individually, *one-at-a-time* in a predefined sequence, such that “there is *no expectation* that a sub-file will be followed by a sub-file other than that which immediately follows it in the

original sequence.” EX1004, 12:9-11 (emphasis added);¹² *see also infra* at § VII.D.

That there is “no expectation” of receiving sub-files in any order other than their “original sequence” also confirms that Leaning fails to contemplate the use of multiple TCP connections and teaches away from such use. This is further confirmed by Leaning’s lack of any discussion whatsoever of how to handle the receipt of sub-files via multiple TCP connections, such as the ’772 Patent’s disclosure of a “staging module” that is “configured to arrange the streamlets **212** in order of ascending playback time.” IPR2024-00919, EX1001 (’772 Patent), 8:41-42.

112. Notwithstanding the foregoing, Petitioner and Dr. Houh contend that a POSITA would have understood “Leaning’s teachings of a terminal connecting to an ordinary web server ... to include teaching a plurality of TCP connections ... because a POSITA would have known that ordinary web servers supported both HTTP and TCP ... [and] each HTTP request in HTTP 1.0 *must establish a new TCP connection.*” EX1003, ¶ 59. This contention is plainly wrong. First, Leaning does not disclose use of HTTP 1.0. At the time of Leaning, HTTP 1.1, standardized in January of 1997, was in widespread use. Accordingly, POSITA would have understood Leaning’s reference to HTTP to refer to HTTP 1.1. And, by default,

¹² This also eliminates the need for Leaning’s user terminal to know the filenames in advance before requesting those sub-files from the server. EX1004, 5:6-8 (“[T]his system is organised [sic] on the basis that the sub-files are named in the manner indicated above, so that the terminal does not need to be informed of the filenames.”).

HTTP 1.1 used persistent TCP connections wherein a client could make multiple requests of a server on a single TCP connection. EX2137 (RFC 2068 – HTTP/1.1 (1997)). In fact, by 2000, even in HTTP 1.0, persistent connections were possible and widely used. EX2138 (RFC 2616 – HTTP/1.1 (1999)), at 105. Thus, Leaning’s references to HTTP do not constitute a suggestion that parallel connections are used and, given the ubiquity of persistent HTTP connections at the time, actually teach away from such use.

C. Leaning Fails to Initially Request Low Quality Files/Streamlets to Enable Instant Playback

113. As shown in the chart below, some of the Asserted Claims require initially requesting low quality files and/or streamlets.

Challenged Claim	Element
'564 Patent Claim 7	The end user station of claim 1, wherein the media player is configured to <i>initially request low quality files</i> to enable <i>instant playback</i> of the video, and to subsequently upshift to a better quality copy of the video according to the performance factor.
'156 Patent Claim 10	The apparatus of claim 1, wherein the apparatus is configured to <i>initially request low quality streamlets</i> to enable <i>instant playback</i> of the content file, and subsequent 60 upshifting according to the performance factor.
'772 Patent Claim 6	The method of claim 1, wherein said requesting the plurality of sequential one of the files includes <i>requesting the files from a low quality</i> one of the different copies to enable <i>instant playback</i> of the video, and wherein said automatically requesting includes-upshifting to a higher quality one of the different copies.

114. For example, Dependent Claim 7 of the '564 Patent recites that the claimed media player is configured to “initially request low quality files to enable instant playback of the video, and to subsequently upshift to a better quality copy of the video according to the performance factor.” Claim 7 depends directly from independent Claim 1 and as Claim 1 recites, “the performance factor” of Claim 7 is “a factor indicative of the *current* ability to sustain the streaming of the video using the files from different ones of the copies, wherein the set of one or more factors relate to the *performance of the network*.” In other words, the claimed “performance factor” indicates the “current ... performance of the network.”

115. A POSITA would understand that “instant playback” is not possible if the claimed media player “initially request[s]” files of a quality higher than the “current ... performance of the network” can handle. A POSITA would therefore understand that the claimed media player “initially request[s] low quality files” to ensure that playback can begin immediately even if the current network performance is poor. This type of “instant playback” is possible because, unlike Leaning, only a single streamlet needs to be received and archived before playback can begin in the process of Challenged Patents. *See supra* § III; *see also* EX1001, Fig. 6 (depicting a single “streamlet” being delivered to the user at Step **614** after only a single “streamlet” is archived in the streamlet cache **410** at Step **618**); *id.*, 11:55-57 (“[D]elivering **614** streamlets to the end user comprises playing video or audio

streamlets on the viewer 408.”). Leaning, on the other hand, requires that its buffer be filled a predetermined amount—“Buffer Fullness > T_p ”—before playback can begin. *See supra* §VII.A. Nowhere does Leaning describe bypassing this process before beginning playback. Thus, even if the first sub-file Leaning requests is of a low quality, Leaning cannot achieve “instant playback,” as required by Claim 7 of the ’564 Patent.

116. But the first sub-file Leaning requests is *not* of a low quality. The portion of Leaning that Petitioner relies on as allegedly disclosing this limitation describes “[i]nitially” requesting a sub-file with a bit rate of 24 kbit/s. Pet., 47 (quoting EX1004, 8:6-11). This bit rate is on the *higher* end of the data rate ranges Leaning provides in this example—specifically, “8, 16, **24** and 32 kbit/s.” EX1004, 6:14-16; *see also id.*, 6:27-7:5 (depicting the same range of data rates in the “subdirectory” of the index file). And as Leaning acknowledges, requesting higher quality (*i.e.*, larger) sub-files results in “extra delay” when starting playback. EX1004, 12:31-33 (“[E]xcessively large sub-files are disadvantageous in requiring a larger buffer and in causing *extra delay* when starting play.”). Thus, Leaning describes initially requesting *high* quality files, not “*low* quality files,” which increases delays and makes “instant playback” impossible.

D. Leaning Fails to Disclose “Arranging” Files/Streamlets

117. As shown in the chart below, some of the Asserted Claims require arranging files and/or streamlets.

Challenged Claim	Element
'772 Patent Claim 8	The method of claim 1, wherein said presenting comprises sequentially <i>arranging the requested files</i> from the different copies in the order of ascending playback with the media player on the end user station.
'234 Patent Claim 1[E]	<i>arranging the streamlets</i> requested from the at least one server in order of ascending presentation time for output to a presentation device.
'234 Patent Claim 12[E]	<i>arranging the individually-requestable portions</i> requested from the at least one server in order of ascending presentation time for output to a presentation device.

118. For example, Dependent Claim 8 of the '772 Patent includes the step of “sequentially arranging the requested files from the different copies in the order of ascending playback time.” Claim 8 further recites that this is done “with the media player on the end user station.” Claims 1 and 12 of the '234 Patent include a similar “arranging” limitation that must be performed at an “end user station” or a “client device,” respectively. In contrast, Leaning does not disclose or suggest performing any active step on its user terminal, including the step of “sequentially arranging.”

119. The '772 Patent describes how, in one embodiment of the disclosed invention, “the network controller module 406 is configured to process and request

a plurality of streamlets 212 simultaneously.” IPR2024-00919, EX1001, 9:23-26.¹³ The ’772 Patent expands on this by further explaining that, in this embodiment, “the agent controller module 402 may simultaneously request 606 a plurality of streamlets from the streamlet cache module 404.” *Id.*, 11:42-44. Thus, in this embodiment, it is necessary that “the staging module 409 then arranges 612 the streamlets into the proper order.” *Id.*, 11:52-55. “The staging module 409 may be configured to arrange the streamlets 212 in order of ascending playback time.” *Id.*, 8:41 42. Leaning never discloses or suggests “simultaneously” requesting a “plurality” of sub-files and thus, does not contemplate a need to “sequentially arrang[e]” sub-files on its user terminal.

120. Instead, Leaning describes requesting sub-files individually, one-at-a-time in a predefined sequence. This creates a scenario where the sub-files are simply stored and played back in the order in which they are received. This is accomplished by way of Leaning giving its sub-files “file names which include a serial number indicative of their sequence in the original file,” as depicted below.

¹³ Citations here are to the ’772 Patent.

```
000000.bin  
000001.bin  
000002.bin  
000003.bin  
.  
.  
000134.bin
```

EX1004, 3:10-18. A person does this by “preparing the file for loading onto the web server,” thus eliminating the need to “sequentially arrang[e]” the sub-files when they are received at Leaning’s user terminal. *Id.*, 3:19-20.

121. Leaning itself notes a “preferre[nce]” for this particular approach, which it describes as being “simple to implement” while allowing the boundaries between sub-files to be located arbitrarily. *Id.*, 6:5-8, 12:9-11. According to Leaning’s own disclosure, “[i]t matters little where the boundaries between the sub-files are located” since “there is no expectation that a sub-file will be followed by a sub-file other than that which immediately follows it in the original sequence.” *Id.*, 12:9-11. Thus, by its own admission, Leaning not only fails to disclose or suggest “sequentially arranging” sub-files at its user terminal, one goal of Leaning is to eliminate the need for such activity at the user terminal, which is accomplished by ensuring that each sub-file is received only in the “original sequence.” Therefore, Leaning actually *teaches away* from sequentially arranging sub-files at its user terminal.

E. **Leaning Fails to Disclose Upshifting “According to an Amount of Contiguously Available Files”**

122. As shown in the chart below, at least Claim 5 of the ’564 Patent requires upshifting according to an amount of contiguously available files.

Challenged Claim	Element
’564 Patent Claim 5	The end user station of claim 1 wherein the media player is configured to upshift to the higher quality copy when the performance factor is greater than the first threshold and the media player <i>determines that the higher quality playback can be sustained according to an amount of contiguously available files</i> stored by the media player.

123. Dependent Claim 5 of the ’564 Patent includes a limitation requiring that “the media player is configured to upshift to the higher quality copy when ... the media player determines that the higher quality playback can be sustained according to an amount of contiguously available files stored by the media player.” The ’564 Patent’s specification describes an “amount of contiguously available streamlets” as a “read ahead margin.” It further explains how this “read ahead margin” is used “to discourage upshifting until a larger read ahead margin may be established to withstand network disruptions.” EX1001, 12:43-53. Notably, Leaning fails to disclose or suggest considering these conditions when performing an “upshift.”

124. According to Leaning, its payback process performs an upshift (“Stepup”) *only if* two conditions are met: (i) the bit rate measured for a given period

(“MRate”) exceeds the next higher bit rate (“NextRate”) and (ii) the next higher bit rate (“NextRate”) is not zero. This conditional structure is depicted in the section of Leaning’s pseudo-code flow chart reproduced below.

If MRate > NextRate AND NextRate > 0
goto Stepup


EX1004, 10; *see also id.*, 9:24-26 (“[I]f the measured rate exceeds not only the current rate but also the next higher rate for a given period of time ... increase the current rate.”). Furthermore, since the only period Leaning describes for measuring bit rates is 120 seconds, *id.*, 9:24-26, it is clear that Leaning’s description of performing an upshift (“Stepup”) denotes looking at bit rates from the preceding two minutes, not the number of files stored in its buffer.

125. The fact that Leaning does not look at the contents of its buffer before making a rate shift also follows logic because, as discussed *supra* at § VII.A, Leaning clears its buffer—*i.e.*, calls a “Clear Buffer” function—as the first step of any rate shift (up or down). This clearing of the buffer would render pointless any attempt to use buffer fullness as a factor in determining whether to upshift (Stepup).

126. The state of the current contents of the buffer is particularly important for maintaining “continuous playback” in the Challenged Patents because it is this information that enables the invention of Claim 5 “to withstand network disruptions.” EX1001, 12:43-53. In contrast, Leaning’s backward-looking

evaluation of measured bit rate is simultaneously slow to adapt to network conditions and more likely to result in playback error. Dr. Houh apparently overlooks the “Clear Buffer” function in Leaning’s upshift (“Stepup”) process, choosing instead to focus on an entirely different process—Leaning’s playback process at Step J3. According to Dr. Houh, Leaning “only allows for shifting up and/or continuing playing if ‘Buffer Fullness > Tp seconds’” at Step J1A. EX1003, ¶ 88. This erroneous reading of Leaning misses the important fact that only playback (Step J3) depends on this condition.

127. Dr. Houh’s flawed logic breaks down further when one looks at what it would mean for downshifting. Specifically, Dr. Houh’s reading would also mean that Leaning only allows for downshifting (“Stepdown”) if “Buffer Fullness > Tp seconds” at step J1A. Such a result is depicted in Dr. Houh’s own annotated portion of Leaning’s pseudo-code flow chart (reproduced below).



EX1003, ¶ 88 (green annotations added by Respondent). The problem is that a downshift (Stepdown) would *never* occur in this situation where the system is dependent on “Buffer Fullness > Tp seconds” at Step J1A since the condition “BufferFullness = 0”—one of the conditions necessary for a downshift (Stepdown) in the above example—could *never* be met under those circumstances. Furthermore, a downshift would never occur under any other circumstance under Dr. Houh’s logic, since when the condition of Step J1A is not met—*i.e.*, “Buffer Fullness < Tp seconds” at Step J4—the process *skips* the downshift (Stepdown) inquiry altogether and reverts back to Step J2, as depicted in green above. It cannot be true that Leaning is designed such that a downshift would never occur.

128. Expecting the system to be designed to allow for downshifting, it is my opinion that a POSITA would understand that the preconditions for upshift (“Stepup”) and downshift (“Stepdown”) in Leaning are tested whether or not the tests performed at Steps J1A and J4—*i.e.*, “Buffer Fullness > Tp seconds” and “Buffer Fullness < Tp seconds,” respectively—are answered in the affirmative. This way, and only this way, would both upshifts (“Stepup[s]”) and downshifts (“Stepdown[s]”) be allowed to occur in the exemplary embodiment of Leaning. Therefore, it is clear that an upshift (“Stepup”) in Leaning is dependent *entirely* on data rates, and data rates alone, not buffer fullness, at Step J1A.

129. In any case, it is my opinion that a POSITA would readily understand Leaning’s rate change inquiries to not depend on, or be influenced by, the amount of data in the buffer at Step J1A. This understanding is bolstered by the aforementioned buffer clearing in Leaning at the first step in both an up- and a downshift, which buffer clearing renders the “read ahead” contents of the buffer meaningless for either type of shift. Thus, Leaning fails to disclose or suggest that “the media player is configured to upshift to the higher quality copy when ... the media player determines that the higher quality playback can be sustained according to an amount of contiguously available files stored by the media player,” as recited in dependent Claim 5 of the ’564 Patent.

F. Leaning Fails to “anticipate file requests and to pre-request files”

130. As shown in the chart below, at least Claim 6 of the ’564 Patent requires anticipating file requests and pre-requesting files.

Challenged Claim	Element
’564 Patent Claim 6	The end user station of claim 1, wherein the media player is further configured to <i>anticipate file requests and to pre-request files</i> to enable fast-forward, skip randomly, and rewind functionality.

131. As discussed *supra* at § III, anticipating file requests and pre-requesting files is one of the ways in which the ’564 Patent is able to “achieve continuous playback,” as recited in independent Claims 1 and 8. In fact, the ’564 Patent specifically states that a user will “experience no buffering delay” if “the agent

controller module 402 ... anticipate[s] streamlet 212 requests and pre-request[s] streamlets 212” before they are actually needed, including “in anticipation of the end user 104 selecting new content to view.” EX1001, 8:46-50, 8:57-62. In contrast, Leaning lacks any disclosure or suggestion of such a concept, especially in the context of “fast-forward, skip randomly, and rewind functionality,” which are also recited in Claim 6 of the ’564 Patent.

132. Leaning describes a process far different than the ’564 Patent’s process of anticipating and pre-requesting streamlets. Specifically, Leaning describes enabling “fast forward or fast reverse” by providing a completely different set of sub-files encoded at different rates and “amending the sub-file address” to denote that difference. EX1004, 16:11-16. “On fast forward play at, for example, 5 times normal speed, one encodes only every 5th frame.” *Id.*, 16:33-17:1. The sub-files that result from this process would then be given a subdirectory address with the suffix “x5” in it, as depicted below.

Directory	Subdirectory	Filename
mpg_name	096k_x1	000000.bin
		000001.bin
		...
	096k_x5	000134.bin
		000000.bin
		...
	.	000134.bin

Id., 17:16-19. This approach to fast forwarding and rewinding is completely different than the “anticipat[ing]” and “pre-request[ing]” functionality recited in Claim 6 of the ’564 Patent.

133. Dr. Houh analogizes Leaning’s “jump” process to the claimed “fast-forward,” “rewind,” and “skip randomly” functionality. However, even this process does not involve “anticipat[ing]” and “pre-request[ing]” streamlets. EX1003, ¶ 91. In contrast, as explained *supra* at § VII.A and depicted below, this process clears the buffer as its very first step.

If UserCommand = Jump(j%)then: <u>Clear buffer;</u> Set CurrentSubfile = Integer[(LFI+1)*j/100]; go to Step J1
--

EX1004, 10 (annotation added); *see also id.*, 11:12-20. As also explained *supra* at § VII.A, Leaning emphasizes its buffer-clearing step, explaining with respect to the “jump” process that “[w]e prefer, in the case of the jump, to suspend decoding and clear the buffer so that the new request [for a sub-file] is sent ***immediately***.” *Id.*, 19-20. Importantly, the new request for a sub-file must be sent “immediately” because Leaning neither “anticipate[s]” nor “pre-request[s]” the next sub-file needed to perform a “jump.”

134. Even if, for the sake of argument, Leaning’s discussion of storing a certain number of sub-files in a buffer before beginning playback could be

considered “pre-request[ing],” Leaning “anticipate[s]” nothing in this regard. Leaning’s clearing of its buffer as the first step for an upshift (“Stepup”), downshift (“Stepdown”), or “jump” unwinds the benefits of “anticipat[ing]” sub-files because any sub-files that are “pre-request[ed] on that basis will be discarded, and not used for playback. On the contrary, all of Leaning’s considerations in this regard, opposite of being anticipatory, are entirely backward-looking or, in the case of a “jump,” are determined on the fly.

135. In determining whether to upshift (“Stepup”), Leaning’s process looks *backward* at the measured bit rate (“MRate”) for a predetermined amount of time—*e.g.*, 120 seconds—as part of making the determination. EX1004, 8:24-26, 10. In determining whether to downshift (“Stepdown”), Leaning’s process, at best, looks to the current state of the buffer—*i.e.*, “Buffer Fullness = 0.” *Id.*, 8:18-20, 10. Beyond that possibility, Leaning’s downshift (“Stepdown”) process looks backward at the same measured bit rate (“MRate”) as the upshift (“Stepup”) process, and adds to that a backwards look at “the percentage of time that the buffer contents represent[ed] less than 25% of the playout time” in order to make the determination—*i.e.*, “BufferLow% > Td.” *Id.*, 8:27-28, 10.

136. Similarly, in the case of Leaning’s “jump” process, Leaning responds to a user “moving a cursor to a desired point on a displayed bar,” which move cannot be anticipated well enough to be used to make the determination. EX1001, 11:12-

17. Therefore, Leaning fails to disclose or suggest a media player that is configured to “anticipate file requests and pre-request files,” particularly in the context of “fast-forward, skip randomly, and rewind functionality.”

G. Leaning and Leaning in Combination with Klemets Fail to Disclose Arranging Streamlets “using the virtual timeline”

137. As shown in the chart below, at least Claim 9 of the ’234 Patent requires arranging streamlets using the virtual timeline.

Challenged Claim	Element
’234 Patent Claim 9	The end user station of claim 1, wherein the client module is further configured to transmit at least one virtual timeline request to retrieve a virtual timeline corresponding to the digital content and to receive the virtual timeline from the at least one server, and <i>arrange the streamlets</i> requested from the at least one server <i>using the virtual timeline</i> .

138. As described in the ’234 Patent, a “virtual timeline” is “configured to define a playlist.” IPR2024-00940, EX1001 at 12:35-40. This playlist, in turn, can indicate the order in which the content within a content file is viewed, including that “the same content be looped over and over again.” *Id.*, 12:40-43, 12:48-51. Furthermore, the “virtual timeline” may be used to “selectively edit content” through various methods, including filtering out—*e.g.*, muting or skipping—portions of a content file that include “offensive content.” *Id.*, 12:64-13:4. Dr. Houh treats Leaning’s user terminal as an analog to the claimed client module, yet fails to point to anything in Leaning that discloses or suggests that Leaning’s own user terminal

actively “arrange[s]” content based on a “virtual timeline,” including under Petitioner’s own construction of the term “virtual timeline.”

139. Dr. Houh contends that “virtual timeline” should be construed as “a playlist of entire content files.” EX1003 ¶ 46. Relying on this construction, Dr. Houh points to Leaning’s menu page (menu.htm) as allegedly corresponding to a “virtual timeline.” *Id.*, ¶ 117. However, this fails because Leaning’s menu page simply “contain[s] a list of recordings available, with hyperlinks to the corresponding link pages.” EX1004, 4:1-3. “The user selects one of the hyperlinks within the menu page which causes the browser to retrieve from the server, and display, the link page for the desired recording - in this example the file mp3_bwv565_link.htm.” *Id.*, 4:4-26. This information, which is the **only** information retrieved from Leaning’s menu page, is not used to “arrange” sub-files at the Leaning’s user terminal.

140. The limits of Leaning are further demonstrated by Leaning’s teaching that its user terminal does not need any additional information from the menu page to request and receive every sub-file in a content file. “[T]he terminal does not need to be informed of the filenames” for the sub-files because sub-files are assigned “file names which include a serial number indicative of their sequence in the original file,” as depicted below.

```
000000.bin  
000001.bin  
000002.bin  
000003.bin  
.  
.  
000134.bin
```

EX1004, 3:10-18, 5:6-8. This assignment is done by “the person preparing the file for loading onto the web server.” *Id.*, 3:19-20. Consequently, Leaning’s user terminal needs to do nothing more to request a sub-file than add the sequential serial number to the URL received from the server “to produce a complete address for the sub-file - *i.e.* www.server1.com/mp3_bwv565/000000.bin.” *Id.*, 5:4-6. Telling in this regard is that Leaning’s discussion of its menu page is silent as to the order in which sub-files are to be arranged since Leaning’s user terminal already knows the order—*i.e.*, 000000 and the following prefixes in Leaning’s example—***before*** a user gets so far as to access the menu page.

141. This specific manner of constructing sub-file requests used by Leaning dictates other requirements. Specifically, based on this, it is known that Leaning’s user terminal necessarily requests sub-files individually, one-at-a-time in a predefined sequence, eliminating the need to “arrange” any sub-files at the user terminal. This is true because the sub-files are already arranged.

142. Leaning expresses a “prefere[nce]” for this approach because “it is simple to implement” and allows the boundaries between sub-files to be located

arbitrarily. *Id.*, 6:5-8, 12:9-11. According to Leaning, “[i]t matters little where the boundaries between the sub-files are located” because “there is no expectation that a sub-file will be followed by a sub-file other than that which immediately follows it in the original sequence.” *Id.*, 12:9-11. In other words, Leaning not only fails to disclose or suggest “arrang[ing]” sub-files at its user terminal according to a “virtual timeline” (or anything else received from the server), Leaning seeks to eliminate the need for getting such information from the server by always requesting sub-files in the “original sequence” in which they were loaded onto the server.

143. Dr. Houh’s secondary argument for this limitation fails for the same reason. Specifically, Leaning’s index file (index.htm), the basis for this secondary argument, is likewise not used to “arrange the streamlets requested from the at least one server.” Rather than list the order in which sub-files are to be organized, the “primary purpose” of that file is “to provide a list of the data rates that are available.” EX1004, 6:27-28. Leaning describes the image below as depicting a “typical index file.”

```

<html>
<head>

<!-- Odbits: Mode = "recorded" -->
<!-- Audio = "mp3-024k_11_s mp3-032k_11_s mp3-018k_11_s mp3-016k_11_m
mp3-008k_11_m" -->
<!-- LFI = "44" -->
<!-- SL = "178" -->
<!-- Title = "Toccata and Fugue in D minor" -->
<!-- Author = "J.S. Bach" -->
<!-- Created = "6 Nov 2000" -->
<!-- Copyright = "© 1999" -->
<!-- Comments = "One of Bach's many organ works." -->
</head>
<body>
An announcement to appear on the display
</body>
</html>

```

Id., 8:2-4, Fig. 3. Notably, nothing in this file can be used to determine how sub-files will be arranged at the user terminal. Consequently, neither Leaning’s menu page nor its index file is used to “arrange the streamlets requested from the at least one server,” as recited in Claim 9. Thus, they are not “virtual timelines” under *any* construction of that term.

144. Dr. Houh further challenges Claim 9 under Ground 2, which relies on the combination of Leaning in view of Klemets (EX1006). Yet, a challenge under this combination suffers from the same deficiencies, as Dr. Houh only relies on Klemets for “teach[ing] generating a playlist of entire content files” while maintaining his reliance on Leaning to disclose “arranging the sub-files for playback in sequential order.” EX1003 ¶ 117-18. Regardless of any arguments Dr. Houh did

or could have made, Klemets cannot cure the deficiencies of Leaning because Klemets fails to disclose or suggest “arranging” sub-files at a client device using a “virtual timeline.” This is because nowhere in its disclosure does Klemets disclose any “sub-files” to begin with.

H. Leaning Fails to Disclose “independently...playable” Streamlets

145. As shown in the chart below, at least dependent Claims 5 and 15 of the ’156 Patent recite that “each of the different copies is independently requestable and playable by the [apparatus/end user device].”

Challenged Claim	Element
’156 Patent Claim 5	The apparatus of claim 1 wherein <i>each of the streamlets</i> of each of the different copies is <i>independently</i> requestable and <i>playable</i> by the apparatus.
’156 Patent Claim 15	The method of claim 14 wherein <i>each of the streamlets</i> of each of the different copies is <i>independently</i> requestable and <i>playable</i> by the end user device.

146. Claim 15 depends from Claim 14, which depends from independent Claim 13. According to independent Claim 13, “the different copies” recited in Claim 15 are “different copies of the video encoded at different bit rates,” which are “stored as multiple sets of streamlets on the server.” Consequently, dependent Claim 15 requires that each “streamlet” be “video” that is “independently ... playable by the end user device.” Leaning fails to disclose or suggest such a “streamlet” in the context of video.

147. I understand that Petitioner carries the burden to demonstrate that Leaning teaches rate shifting in the context of video generally, and it is my understanding that they fail to do so for reasons discussed above. *See supra* § VII.A. Furthermore, my review of Leaning’s discussion of video, which Dr. Houh declined to even consider, makes clear that **none** of the video sub-files described in Leaning are “independently ... playable by the end user device.” It is no surprise that Leaning’s discussion of delivering video focusses on formats that involve compression since a POSITA would understand that compression is required to effectively deliver video at least at the transmission rates contemplated by Leaning and the large size of uncompressed video files. Yet, the only two formats Leaning considers for video are the H.261 and MPEG formats. EX1004, 16:16-31. A POSITA would understand that neither of those two formats supports the creation of “independently ... playable” sub-files. For example, neither coding standard includes the concept of a “sub-file.”

148. Whereas audio is a simple two-dimensional signal of an amplitude over time, video is a more complicated multi-dimensional signal: two spatial dimensions of pixels over time (a third dimension), multiplied by multiple color dimensions—such as red/green/blue or Y/U/V (luminance/chrominance) components. Thus, whereas concepts such as sampling rate in audio (the number of samples taken of the analog audio signal per unit of time) are well defined, the same is not true for

video. The number of pixels in a video image (the image size) can change, as can the frame rate and the color depth (the number of represented colors) and the means of sampling and representing color can change.


149. Video encoding, also known as “video coding” and commonly synonymous with video compression, removes the spatial and temporal redundancy to enable storage or transmission of digital video data with fewer bits while preserving the quality of the resulting video data to the maximum extent possible. Spatial redundancy refers to the fact that many adjacent blocks of pixels within a video frame will have the same or nearly the same color values and, thus, not all pixel blocks need to be separately coded. The elimination of spatial redundancy is an example of what is known as *intra*-frame coding.

150. Temporal redundancy is removed using pixel block-based motion compensation between video frames. As such, the elimination of temporal redundancy is an example of what is known as *inter*-frame coding. Motion compensation is based on a model of the “motion” (displacement) of objects in a video scene from one location in a frame to a “nearby” location in a subsequent frame.


151. For each pixel block in a current frame, a search is made to find a best match to the block in a previous or future “reference” frame that has already been coded. The best-matched block is then used as a “prediction” for the current block.

The difference between the pixel values of the current block and the pixel values of the prediction block is called residual data. The location of the best match block in the reference frame relative to the location of the corresponding block in the current frame (the displacement of the pixel block between the frames) is called a motion vector. The pixel block in the current frame can be coded (represented) more efficiently by coding the residual data and the motion vector rather than coding the pixel block by itself. For blocks that cannot be temporally predicted efficiently, for example, because a pixel block represents a new object coming to a video scene, spatial compression within the same frame is applied.

152. This concept of encoding and decoding video is shown in the following two graphics:



EX2011 (Cardiff University MPEG) at 7




Id. at 8.

153. The CCITT, the forerunner of today's International Telecommunication Union (the "ITU," a telecommunications standards body) developed a suite of coding standards beginning in 1984, including the H.261 video coding standard. EX2012 (ITU-T Recommendation H.261 1993). H.261 encoding was specifically designed for *telecommunication* applications, which at the time were understood as applications that operated over dedicated telephony circuits. Indeed, H.261 was developed for videoconferencing/videotelephone applications over ISDN telephone circuits with a baseline data rate of 64 kbit/s and scalable in increments of 64 kbits/s. EX2011 (Cardiff University MPEG) at 11.

154. H.261 coded video using intra-frames ("I-frames") and inter-frames ("P-frames") to reduce the size of a "group of pictures" or a "GOP" (a sequence of

video frames). *Id.* at 12. I-frames serve as an initial reference point for P-frames, which are predicted from a previous I- or P-frame. Thus, P-frames depend on the immediately previous frame and by transitivity, the initial I-frame in the GOP. The number of P-frames following each I-frame determines the size of the GOP and can be fixed or dynamic in H.261.




Id. at 13.

155. To address the temporal redundancy, H.261 specifies an inter-frame coding algorithm in which the data in a video frame is predicted from the data in the previous frame, and an intra-frame coding algorithm in which the data in a video frame is coded solely using the data within that frame. EX2011 (Cardiff University MPEG) at 16. The decision whether to use the **INTER** coding algorithm or use the **INTRA** coding algorithm is determined on a 16x16 pixel block basis within each video frame. *Id.* The determination is made based on a parameter or attribute of the

data block, such as whether or not there is motion in the video scene for that block. The H.261 standard does not require regular inclusion of intra-frames.

156. MPEG is a video compression standard based on the H.261 and JPEG standards. It is essentially an attempt to overcome the shortcomings of both. MPEG improves upon H.261 by using both forward and backward prediction to interpolate pixel blocks between both past and future frames. EX2013 (ISO/IEC 11172-2:1993(en), Information technology — Coding of moving pictures and associated audio for digital storage media at up to about 1,5 Mbit/s — Part 2: Video). These “bi-directional” coded frames are called “B-frames.” A B-frame incorporates predictions based not only on previous frames, but also on future frames. Bi-directional prediction is used between previous and future P-frames or between P-frames and I-frames.

157. In MPEG, a GOP is a series of frames consisting of one I-frame and a number of B-frames and P-frames. The actual pattern of I-frames, P-frames, and B-frames looks like the following:



EX2011 (Cardiff University MPEG) at 44. Also, in MPEG, there is no defined limit to the number of consecutive B frames (or P-frames) that may be used in a GOP. However, in comparison to the H.261 standard, MPEG requires regular inclusion of I-frames. As explained in U.S. Patent No. 6,002,440 ('440 Patent), which Leaning incorporates and relies on:

[a]ccording to the MPEG1 standard[,] intra-frames are included in the data stream at regular intervals to allow random access to the video signal and editing thereof. The regular intra-frames also allow interaction with the video signal i.e. a user can request fast forward or fast reverse of a video signal and this may be accomplished by presenting the I frames of the video signal to the user until normal play is requested. However, this increases the average bitrate per picture as compared to the normal play mode of the video signal.

EX2014 ('440 Patent) at 1:42-52.

158. By contrast, H.261 does not require regular inclusion of I-frames and thus can operate in a lower bandwidth than MPEG. As explained in the '440 Patent,

however, use of H.261 is not suitable for interactive “trick play” operations such as fast forwarding or reversing:

ITU standard H.261 relates to the coding of moving pictures at rates of $n \times 64$ kbits/s (where n is an integer from 1 up to 30) (for communication over an ISDN network) and ITU standard H.263 for coding moving pictures at data rates less than 64 kbits/s (for communication over a PSTN network). These encoding schemes do not require the regular inclusion of intra-frames. ... Whilst this means that the average data rate is reduced compared to MPEG1 (the data produced for intra-frames representing a significant proportion of the total data), it is not possible to provide interactive facilities such as fast forward or reverse since there are no clear access points to the data other than at the start. Thus H.261 and H.263 coding schemes are not suitable for providing video streams with which a user can interact.

Id. at 1:53-2:3.



159. MP3 is a compressed audio format developed by the MPEG standards body (the ISO Moving Picture Experts Group). EX2015 (ISO/IEC 11172-3:1993(en), Information technology — Coding of moving pictures and associated audio for digital storage media at up to about 1,5 Mbit/s — Part 3: Audio). As part of MPEG-1 (discussed below), the MPEG group released the Layer III standard (now known as MP3), which handles audio compression.

160. MP3 compression exploits characteristics and limitations of the human hearing system to remove aspects of the audio signal that the human ear cannot

normally hear or perceive. In doing so, the total number of bits required to represent the signal can be reduced because the complexity of the signal is reduced.


161. As acknowledged by Leaning itself, rate switching is “more complex” in both the H.261 and MPEG format because both use “inter-frame techniques” to perform compression. EX1004, 16:20-25. These techniques involve the use of two or three frame types—intra-frames (I-frames), inter-frames (P-frames), and/or bi-directional frames (B-frames)—to reduce the volume of data in a GOP. *See also* EX1004, 16:21-23 (“Some such systems (for example the MPEG standards) generate a mixture of independently coded frames (‘intra-frames’) and predictively coded frames [‘inter-frames’]”). I-frames and B- and P-frames work to reduce the data in a GOP by reducing the amount of data that needs to be coded for each frame.

162. Of the three frame types, an I-frame includes the most data and is used as both a starting/restarting point for a decoder and a reference point for B- and P-frames. Meanwhile, a P-frame is encoded with nothing more than the differences that are predicted from previous a P- or I-frame. As shown below in the “Encoded Scene,” this allows P-frames to include less data.


Original SceneEncoded Scene

As a consequence of this setup, each P-frame depends on one or more previous frames, including the initial I-frame.

163. This dependency on earlier frames makes rate switching more complex for multiple reasons, some of which are tied to the format being used. For example, as explained by Leaning, the H.261 format “do[es] not provide for the frequent, regular inclusion of [I-frames],” as depicted above. Thus, when they are segmented into smaller portions, each portion may not include an I-frame, as depicted below.





EX1004, 16:24-25. Consequently, rate switching would not work well using the H.261 format for sub-file starting with a P-frame, since, if one were to switch from one bit rate to another using sub-files encoded using the H.261 format, the prior reference frame which the first P-frame needs to reference would no longer be available to the decoder, as depicted in the playout of the second of the three sub-files depicted below. EX1004, 16:25-30.



As Leaning itself admits, this would result in “*serious mistracking* of the decoder.”


EX1004, 16:30-31.

164. Using the MPEG format rather than H.261 results in a similar problem that will likewise not result in mistracking. Leaning’s description of beginning each sub-file, or segment, with an I-frame when the MPEG format is used does not dictate otherwise. The first problem is that Leaning ignores the third type of inter-frame included in GOPs that are encoded using the MPEG format—specifically, B-frames. As indicated above, the “B” in “B-frame” represents the fact that this inter-frame is encoded bi-directionally based on differences that are predicted from both previous and subsequent I- or P-frames, as depicted below.

Original SceneEncoded Scene "Display Order"

Mistracking in the MPEG format is tied to these B-frames.

165. Specifically, since B-frames can have dependencies on both I-frames and P-frames, GOPs must be encoded with the B-frames *following* both the I-frames and the P-frames on which they depend. But this includes I-frames that will not be played until *after* those B-frames. EX2017 (U.S. Patent No. 6,806,909) at 4:33-34. Consequently, frames encoded in the MPEG format must be encoded, transmitted, and decoded in a different sequence from the one in which they are played back, as depicted below.



Id., 4:18-46.¹⁴ This example assumes the video sequence shown in somewhere in the middle of a longer video file.

166. Given this additional complexity associated with using the MPEG format, simply beginning each sub-file with an I-frame is not enough to create functional, individually playable sub-files using the MPEG format.¹⁵ Yet, neither

¹⁴ Note that, in the “Decoded Sequence” example above, frame B₋₁ is from the preceding GOP in the sequence and, in the “Encoded Sequence,” frame I₇ is from the following GOP in the sequence.

¹⁵ This same problem infects the “bridging sub-files” Leaning (EX1004) describes at 17:25-19:5. And neither Leaning nor the Dr. Houh attempt to explain how to

Leaning nor Dr. Houh provide any suggestion as to what else could be done to make these sub-files functional using this MPEG format.

167. Dr. Houh also does not address that MPEG frames are stored/transmitted out of playback order. Nor do they address where you would make the cuts to an MPEG bitstream to create the sub-files. According to Dr. Houh, Leaning discloses that “*after encoding into several different versions, partitioning each source file into sub-files.*” IPR2024-00940, EX1003, ¶ 65 (emphasis added). But as I explain above, the decoding of B-frames in MPEG depends on both preceding and succeeding frames. In MPEG, “the encoded pictures are not transmitted or stored in presentation order. They are provided in the order that they are required for decoding.” EX2017, 4:33-46.

168. For example, the stream order in one example of MPEG frames looks as follows:

address this problem. Regardless, Dr. Houh also does not attempt to rely on these “bridging sub-files” to establish invalidity.

In the data stream the encoded pictures are not transmitted or stored in presentation order. They are provided in the
 35 order that they are required for decoding. That is, the B-frames follow the I and P-frames on which they are dependent. The pictures in this typical scheme are provided in stream order, as follows:


40

I_1	B_{-2}	B_{-1}	P_4	B_2	B_3	P_7	B_5	B_6	I_{10}	B_n	B_o
-------	----------	----------	-------	-------	-------	-------	-------	-------	----------	-------	-------


Note that in stream order B_{-2} and B_{-1} of the preceding GOP
 45 and I_{10} of the succeeding GOP are mixed with the pictures of this typical GOP.

Id., 4:40-44. Thus, in transmission order, the B-frames of the preceding GOP and I-frames of the succeeding GOP are mixed with the pictures of the GOP structure. A POSITA would understand that this is typical of all MPEG encodings.


169. Even if, for the sake of argument, Leaning's omission of additional details necessary to create functional, individually playable sub-files using the MPEG format was unimportant because creating sub-files using the MPEG format were as easy as Leaning suggests, the results would still include mistracking at the decoder if a video is divided into sub-files with an I-frame at the beginning of each sub-file, as depicted below.



Mistracking would occur during a rate change because, when the B-frames are decoded, the P-frames on which they depend would be in a different bit rate, as depicted below.



This mistracking would appear at the beginning of a segment every time there was a shift from one bit rate to another, as depicted in the second and third segments in the middle row below.



This mistracking results from the B-frame in a segment not having the *preceding* P-frame on which it depends “at its disposal.” This is similar to what happens in the H.261 format, where a P-frame at the beginning of a segment does not have a *preceding* I-frame “at its disposal.” Cf. EX1004, 16:25-30.


170. Given that, regardless of which of the two encoding formats Leaning uses, the playback of Leaning’s sub-files depends upon frames from either preceding or succeeding sub-files, or segments, Leaning’s sub-files are not “independently ... playable” video “streamlets,” as required by dependent Claims 5 and 15 of the ’156 Patent. Thus, Leaning fails to anticipate or render obvious those claims.

171. At most, only the first sub-file in each set of sub-files in Leaning is “independently ... playable.” But this is not does not satisfy the limitations in Claims 5 and 15 of the ’156 Patent. Instead, they require that “***each*** of the streamlets of ***each*** of the different copies”—*i.e.*, ***every*** one of a plurality of sub-files—be “independently playable.” Therefore, Leaning fails to disclose this limitation. Ala-Honkola also does not disclose or suggest this limitation.

I. Gamble Fails to Cure the Deficiencies of Leaning

172. For Ground 2 of the of the IPRs directed to the ’564, 772, and ’234 Patents, Dr. Houh relies on the combination of Leaning and Gamble as allegedly rendering obvious the portion of dependent Claims 9 and 11 pertaining to requesting “sub-parts” of files via different TCP connections. EX1003, ¶¶ 116-19. This combination fails because, like Leaning, Gamble also fails to disclose or suggest “requesting the plurality of sequential files ... over different ones of the plurality of different TCP connections.”

173. Gamble is directed, instead, to a “data processing systems ... for transferring large data files over parallel connections.” EX1020, ¶[0001]. Gamble’s disclosed file transfer techniques involve “transferring [a] data file 702A at a client 104 to client 106 via a server 108.” *Id.*, ¶[0042].



EX1020, FIG. 1 (annotations added).

174. “Initially, a send buffer” is created at the transmitting client to store “segments of [a] data file.” *Id.*, ¶[0045]. The “send buffer” includes “four sub-sections B1 through B4 to store N=4 [] segments.” *Id.* Next, each “segment in sub-buffers B1 through B4 is sent on one of the N=4 parallel connections to server 108,” where the segments are stored in “stream buffers” and “transmitted” to receiving client 106 over corresponding “independent” connections, and the data file is subsequently reconstructed at the receiving client. *Id.*, ¶[0049]. Thus, Gamble teaches **transmitting** “segments” of a larger file over multiple “connections.”

However, Gamble fails to disclose or suggest “*requesting*” segments, files, or sub-files, over multiple TCP connections as claimed.

175. In Gamble, it is the *transmitting* client that establishes “parallel connections” and transmits each segment to the *receiving* client (via that intermediate server). *Id.*, ¶¶[0021], [0039]. The *transmitting* client does all this without receiving any segment requests from the *receiving* client. In fact, Gamble is entirely silent on any request at all from the *receiving* client, much less requests over multiple TCP connections. Accordingly, Gamble fails to disclose or suggest “receiving” requests over multiple TCP connections, as claimed.

176. Dr. Houh contends that a POSITA “would have been motivated to configure Leaning’s apparatus to request sub-parts ... over different TCP connections ... as taught by Gamble.” EX1003, ¶ 118. But, as discussed directly above, Gamble does not disclose or suggest any such client “requests.” Moreover, Dr. Houh entirely fails to account for Gamble’s “server 108” (the “intermediate server” above), and how or why a POSITA would have incorporated that critical element of Gamble into Leaning’s system. Dr. Houh does not articulate any problem in Leaning that is solved through the inclusion of Gamble and, in fact, there is no such problem. Accordingly, Petitioner and Dr. Houh’s combination fails.

J. Ala-Honkola Fails to Cure the Deficiencies of Leaning

177. For Ground 2 of the '156 Patent, Dr. Houh relies on the combination of Leaning and Ala-Honkola as allegedly rendering obvious the portion of dependent Claim 2 pertaining to multiple TCP connections. EX1003, ¶¶ 114-22. This combination fails because, like Leaning, Ala-Honkola also fails to disclose or suggest “request[ing] streamlets” via multiple TCP connections as claimed.

178. Ala-Honkola discloses encoding media in a plurality of streams at different resolution levels and storing them together in a single file. EX1017, ¶[0030]. The starting byte within the file of each stream is identified by a header, and “timing marks are embedded in [each] stream at various locations [to] correlate the information in the two separate resolution streams.” *Id.*, ¶[0031]. Upon request by the terminal, “the server starts to play” the stream beginning at the starting byte for the first stream and continuing in order without any additional terminal requests for sequential streamlets. *Id.*, ¶[0036]. The terminal checks whether the bandwidth is “ok for the selected stream” and requests a higher or lower resolution stream if needed. *Id.*, ¶[0037]. “When the source is changed [] a new HTTP transmission for the new source is started for a new stream.” *Id.* ¶[0015]. And “several TCP connections may be needed or used for *transmitting* consecutive streams.” *Id.*, ¶[0039] (emphasis added).

179. Ala-Honkola does not disclose or suggest “*requesting* sequential streamlets” over multiple TCP connections as claimed since Ala-Honkola requests entire streams—*not streamlets*. In particular, Ala-Honkola discloses using the HTTP “RANGE field” to specify the starting byte of the new stream. EX1017, ¶¶ [0022], [0039]. Critically, Ala-Honkola’s “RANGE” field is unbounded in that it provides a starting byte, but no ending byte. This results in the server transmitting the entire file, all the way to its end, starting at the specified byte range. In Ala-Honkola, the server does not send, and the terminal/client does request, any smaller portion of the stream, much less a time-divided portion, or “streamlet,” with the “same portion[]” encoded at a different bitrate, as recited in independent Claim 1 (from which Claim 2 depends).” And while Ala-Honkola discloses using multiple TCP connections to transmit HTTP requests, Ala-Honkola fails to disclose transmitting individual, sequential streamlet requests over those multiple TCP connections as claimed.

180. By contrast, as explained above, the ’156 Patent describes making a new, redundant content request on a different TCP connection when a content request “fails completely” or “is not satisfied in a timely manner” via the original TCP connection. *See supra* VII.B (citing EX1001, 9:59-63). Leaning describes requesting sub-files individually, “*one-at-a-time* in a predefined sequence” (*id.* (citing EX1004, 12:9-11)) such that there is “no expectation” of receiving sub-files in any order other than their “original sequence.” Like Leaning, Ala-Honkola only

describes using “several TCP connections” for “transmitting *consecutive* streams.” EX1017 ¶ 39 (emphasis added). Ala-Honkola therefore fails to contemplate the use of multiple TCP connections as claimed in Claim 2 of the ’156 Patent.

VIII. Dr. Houh Fails to Show the Patents are Invalid

181. I understand that Petitioner has the burden of proving unpatentability of the challenged claims of the Challenged Patents by proving, by a preponderance of the evidence, that each element of the challenged claims is disclosed by the prior art or would have been obvious based on the prior art. But neither Petitioner nor Dr. Houh met such a burden here. I have reviewed Dr. Houh’s declarations and conclude that in each declaration Dr. Houh fails to show any of the claims are invalid. I disagree with Dr. Houh’s conclusion that each challenged claim is allegedly invalid over Leaning, Klemets, Gamble and/or Ala-Honkola. 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 content via HTTP), the ’444 Patent impressed me as being the earliest patent I had seen with such an architecture and the fact that it disclosed a complete, workable solution.

182. 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).

183. 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.

184. Dr. Houh's declaration and the Aylo/fuboTV/Yanka IPRs fail to show the patent is invalid for numerous reasons, which I have addressed above. In rendering my opinions above, I addressed, as examples, elements of each challenged 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.

185. The fact that I do not specifically address any claim does not constitute a waiver or concession of invalidity as to any such claims.

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

186. The Leaning IPRs and Dr. Houh opine that the challenged ABR Patent is invalid as obvious. However, both the Leaning IPRs and Dr. Houh fail to address the objective evidence of nonobviousness, which I understand must be considered where present. Similarly, the Leaning IPRs and Dr. Houh fail to acknowledge or

address the International Trade Commission's finding that the strong evidence of secondary considerations associated with the ABR Patents supported a finding of non-obviousness.

187. 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.*

EX1009 (Final ID) at 213-17 (internal citations omitted, emphasis added); *see* EX2007 at 2 (adopting same).

188. The overwhelming objective evidence here of industry praise and peer recognition, surprising results and superior performance, commercial success, long-felt but unmet need, failure of others, and copying of the inventions in the ABR Patents demonstrates nonobviousness.

189. The articles, press releases, or exhibits cited herein appear authentic and to accurately reflect the story of Move Networks, Inc. (“Move Networks,” “Move,” or “MOVE”), its ABR technology, and/or the ABR Patents. I have no reason to doubt the truthfulness or authenticity of the substance discussed in the cited material.

190. Move Networks successfully developed and patented the ABR technology of the patent claims, which literally revolutionized how content is delivered and experienced over the Internet. As MOVE’s contemporaneous documents state, “[a]ny user who ha[d] viewed a streamed video clip over the Internet [in the early 2000s] almost surely ha[d] experienced the shortcoming of the

current art: several to many seconds of delay before playback begins, hiccups, video going into ‘slide-show’ mode, total freeze with ‘buffering ...’ messages dutifully reported by the player.” EX2021 (CX-1094C, Streaming Media Technology), at 2. Using MOVE’s ABR technology embodied in these claims, users experienced less or no buffering (*e.g.*, continuous playback), fast start of video, and a dramatically improved viewing experience for high-end and low-end connections. *See* EX2021 at 2-3; EX2122 (noting “Near-Immediate and Continuous Video Delivery” and “Dramatically Optimized Quality”). For the first time, viewers could experience content in a manner rivaling that of traditional cable and satellite delivery. *See, e.g.*, EX2070 (CX-1090 Silverlight Shines at International Broadcasting Conference 2008 in Amsterdam).

A. The streaming products offered by DISH (and its predecessor MOVE Networks) embody the claimed invention.

191. I understand DISH is entitled to a presumption of nexus if the objective evidence discussed below regarding Move Networks’ Adaptive Streaming system embodies the claims and is based on features that were novel and non-obvious. As I explain below, Move Networks’ system, which includes novel and non-obvious elements such as the Move Media Player, embodies the limitations present in the challenged claims. *See also* EX2071 (’156 Patent Claim 13 Chart). Similarly, I understand that the ITC has already found the claim limitations present in Move Networks’ Adaptive Streaming system. *See* EX1009 (ITC Initial Determination) at


212 (“Dr. Jeffay has demonstrated through corroborating evidence that Move’s success is tied to the ‘unique characteristics of the claimed invention.’”). Because MOVE’s system practices the claimed invention and is co-extensive with the claims, DISH is entitled to a presumption of nexus.

192. Moreover, I understand that, absent a presumption of nexus, objective indicia of non-obviousness are probative to the extent they directly relates to claimed feature(s). As I explain in Sections IX.B—F, *infra*, even absent a presumption of nexus, the ample evidence of MOVE’s commercial success relates directly to the claimed feature(s).

193. Move’s contemporaneous documents prove MOVE’s Media-Player practiced the patented claims. *See* EX2022 (Move Networks Tech Overview); EX2021 (Streaming Media Technology); EX1009 (ITC Initial Determination) at 212. In addition, my analysis of MOVE’s Media-Player shown in the representative claim chart exemplifies how Move practiced the limitations present in the challenged claims. EX2071 (’156 Patent Claim 13 Chart).

194. Specifically, Move’s system included an end user station/apparatus for adaptive rate streaming of live and on-demand digital content from a video server over a network. *See, e.g.*, EX2022; EX2021; EX2071 at 1. For example, “Move Networks provides and operates the premier video distribution platform” that “delivers pre-recorded (VOD) video content or live broadcasts over the Internet.”

EX2022 at 4; EX2071 at 1-2; EX2077 (CX-1127, Move Adaptive Stream) at 1; EX2069 (CX-1247C, Move Networks Publishing System) at 6. A high-level diagram of the Move platform is shown below.



EX2022 at 6. “The platform supports a broadcast publishing process that includes the video source acquisition, encoding and preparing the video for transmission, the storage and distribution of video content, [and] client viewing of the video content.” *Id.* at 6. The Move system “provide[s] a self-optimizing, variable bitrate viewing experience.” *Id.* at 10. More specifically, “a smooth end user experience” is provided “as the Move Media Player up-shifts and down-shifts in response to network and client CPU ability ... to provide a consistency of service at different bitrates.” *Id.* at 12. “For instance, if the Quantum Media Manager senses that the current environment has poor connectivity, the Quantum Delivery Manager will be directed

to only request streamlets from low-bandwidth profiles. If the connection speed or quality of the connection suddenly increases, the Quantum Delivery Manager starts downloading higher quality streamlets. The Quantum Media Manager continues to test the environment to provide the best viewing experience currently possible. This ability to ‘upshift’ and ‘downshift’ to multiple viewing profile bit rates on a live stream is called Adaptive Streaming.” *Id.* at 23.

195. In the Move system, the live or on-demand video is encoded using Move Simulcode into 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, and at least one of the different quality streams was encoded at a bitrate of no less than 600 kbps. *See, e.g.,* EX2022 at 12, 19-21; EX2021 at 3-4.

196. More specifically as to the different viewing profiles and bitrates, “Move Networks encode[d] to either high definition (16x9) or standard definition (3x4) aspect ratios, using one of the following ‘best practice’ profile sets, based on an aspect ratio (16x9) frames-per-second. Profiles up to 2.4 Mbps are used for video streams intended for distribution over high-bandwidth networks. The low-end profiles include a ‘thumbnail’ profile (Stream Type = T), which is used in multi-player mosaic Web pages, and an ‘emergency profile (Stream Type = E) which provides a minimal experience for poor connectivity or low-bandwidth connections. The remaining profiles are ‘normal’ and are situated at bit rates that offer a smooth

end user experience as the Move Media Player up-shifts and down-shifts in response to the network and CPU availability.” EX2022, at 19.


197. Further, each of the different bitrate/quality streams are stored on and accessed by the server as groups or sets of streamlets/files that correspond to portions of the video, and each group is encoded at the same respective one of the different bitrates. *See, e.g.*, EX2022; EX2021; EX2077 at 2. For example, “[t]he process of atomizing and encoding video for transmission is called Simulcoding. The first step in this process of preparing video for transmission decomposes the captured video content into small streamlets for encoding and transmission The process of Simulcoding refers to the process of producing media that matches multiple viewing profiles (based on device and performance profiles) in a single process. Live content streamlets are encoded ... to provide a self-optimizing, variable bitrate viewing experience. EX2022 at 10. The figure below provides an overview of this aspect of the Move System.



Id. at 11. “Move Networks deploys a cluster server ... for the process of encoding each video to multiple profiles The process of breaking the video into smaller


streamlets or Web transmission and encoding multiple bit rates for various viewing profiles simultaneously is called Simulcoding.” *Id.* at 11. “These streamlets are sent to Web servers hosted within a content distribution network (CDN).” *Id.* at 21. The “servers propagate the streamlets to caching servers at the CDN for distribution through the Internet infrastructure to the Move Networks Quantum Media Player.” *Id.*

198. In the Move system, the streamlets across the groups include a time index, are time aligned, and corresponding streamlets across the groups have equal durations. *See e.g.*, EX2022; EX2021. For example, “Move Networks provides an integrated Content Management System (CMS) for entering the metadata associated with video programming, such as titles, descriptions, topics, program thumbnails, and time offsets (beginning and end points of a video segment). Time offsets can also be automatically detected in some cases based on time indexes and sudden changes in video keyframes.” EX2022 at 25. Further, the figure below provides an exemplary illustration of the duration and alignment of streamlets/files of varying bitrates in the Move system.



Id. at 24.

199. In the Move system, the Move Media Player automatically places a streamlet request using a HTTP GET request for each streamlet over one or more network connections to the stream to retrieve a streamlet/file for the selected stream, and the HTTP GET requests identify the selected streamlet, and the requested streamlets are received from the server. *See e.g.*, EX2022; EX2021. To playback the video normally, the Move Media Player requests sequential files/streamlets from one of the copies from the video server based on the time indexes. *See e.g.*, EX2022; EX2021. For example, the Move system provides for “[s]imple HTTP protocol transfer of media files from standard Web servers rather than deployment of expensive media servers.” EX2022 at 22. This is further illustrated in the figure below.




Id. at 22.

200. The stream selection and automatic streamlet/file requests by the Move Media Player are based upon successive local determinations to shift the playback quality to a higher or lower quality one of the copies and includes repeatedly generating a factor relating to the performance of the network that is indicative of the media player's current ability to sustain streaming the video using the streamlets/files from different ones of the copies. *See e.g.*, EX2022; EX2021. The Move Media Player makes such successive determinations to shift to achieve continuous playback of the video using the highest quality determined sustainable such that the player upshifts to a higher quality when the factor is greater than a first threshold and downshifts when the factor is less than a second threshold. *See e.g.*, EX2022; EX2021. In this context, the Move client evaluates bandwidth available and pulls streamlets accordingly, where such "pull" would be made with an HTTP GET request. *See id.*


201. For example, the Move system "provide[s] a self-optimizing, variable bitrate viewing experience." EX2022 at 10; EX2069 at 19 ("MOVE Adaptive

Stream works better because it can almost always seamlessly adapt to changing Internet conditions and the others often cannot.”). These processes produced several benefits including “continuous playback—no noticeable buffering” and “[q]uality of service optimization based on hardware profiles and speed and quality of the connection to the Internet.” EX2022 at 11. The Move system was configured to “offer a smooth end user experience as the Move Media Player up-shifts and down-shifts in response to network and client CPU availability.” *Id.* at 19. More specifically, the “Move Media Player, the Move Networks client software, delivers the video experience to the end consumer. The client decodes and renders the video streamlets to provide a continuous, optimal stream of video. A major function of the client is managing the viewing experience. The Quantum Delivery Manager ensures the player is constantly receiving the streamlets required to deliver a continuous stream of video. The Quantum Media Manager deduces the current environment to tell the Delivery Manager which streamlets to pull into the player. For instance, if the Quantum Media Manager senses that the current environment has poor connectivity, the Quantum Delivery manager will be directed to only request streamlets from low-bandwidth profiles. If the connection speed or quality of the connection suddenly increases, the Quantum Delivery Manager starts downloading higher quality streamlets. The Quantum Media Manager continues to test the environment to provide the best viewing experience currently possible. This ability

to ‘upshift’ and ‘downshift’ to multiple viewing profile bit rates ... is called Adaptive Streaming.” *Id.* at 23. The diagram below illustrates how Move Adaptive Streaming starts instantly, doesn’t stall, and adapts the streaming.



Id. at 24. Importantly, “the intelligence that controls the quality of service resid[es] in the client.” *Id.* at 22. See also EX2077 at 2 (below):



202. The client of the Move system “is always on the watch to use an appropriate stream encoding rate selection so that the available bandwidth is not

exceeded by downloading of the component streamlets.” EX2021 at 8. In this way, the stream selection and automatic streamlet/file requests by the Move Media Player are based upon successive local determinations to shift the playback quality to a higher or lower quality one of the copies and includes repeatedly generating a factor relating to the performance of the network that is indicative of the media player’s current ability to sustain streaming the video using the streamlets/files from different ones of the copies. *See e.g.*, EX2022 at 23; EX2021 at 8; EX2071. The basic question is not, ‘what is my available bandwidth?’ but ‘*am I receiving the current stream fast enough to present it in real time?*’.” EX2021 at 8. Because “*the client* [of the Move system] *can select and retrieve directly from an array of streamlets of different encode rates*, there is no need to inform any media server and ask it to change the data it is sending.” *Id.* Thus, “[i]f the client observes that over some modest interval of time (say 20 seconds or so) the player is, on average, consuming streamlets more quickly than they arrive, then a stream of lower quality must be selected.” *Id.* In view of this, the Move Media Player makes successive determinations as to whether or not to shift to achieve continuous playback of the video using the highest quality determined sustainable. *See, e.g.*, EX2022 at 23.

203. As part of the streaming, the Move Media Player also requested and received a virtual timeline from the server based on the selected one of the streams. *See e.g.*, EX2022 at 26 (referencing “QVT”).

204. Further, as has been described above, the Move Media Player renders and presents the requested and received streamlets on the end user station/apparatus in ascending order of playback time to playback the video normally. *See e.g.*, EX2022; EX2021. For example, “Move Media Player, the Move Networks client software, delivers the video experience to the end consumer. The client decodes and renders the video streamlets to provide a continuous, optimal stream of video.” EX2022 at 23. The figure below further illustrates the components of the Move Media Player that render the video.



EX2022 at 25. Additionally, as noted above, inventor Robert Drew Major testified that “the system described in the first patent was used to encode and then deliver” the live stream of the October 1, 2005 LDS General Conference. EX2020 at ¶23; *see also* EX2127 at ¶4; EX2100 at ¶ 17; EX2064 at 1; EX2120.

205. The Move system therefore is covered by the challenged claims and, resultingly, exhibits the benefits of the inventions claimed therein. Moreover, it is

the claimed features of the Move system that resulted in the benefits described below, and there are no critical features of Move's Media Player that are unclaimed or otherwise would impact the co-extensiveness of the Media Player with the challenged claims. As explained further below, the inventions of the ABR Patents are a verifiable example that some of the best inventions involve simple but elegant ways to solve multiple problems.

206. For example, the Move Networks Simulcoding process described above, which is responsible for preparing the different bitrate/quality streams and streamlets, “produce[s] several benefits for the video customer, including the following: 1. Quality of service benefits for continuous playback—no noticeable buffering; 2. Quality of service optimization, based on hardware profiles and the speed and quality of the connection to the Internet; 3. The ability to modify the video timeline, such as replacing a program, inserting a commercial, or switching between live and prerecorded programming; and 4. Encoding process enables scalable delivery infrastructure.” EX2022 at 11. And “the small video streamlets” are what “allow for the rich functionality available in delivery and viewing of the video. These streamlets enable the Move Networks system to operate efficiently and effectively.” *Id.* In an elegant solution, by aligning the streamlets, Move's solution allowed for seamless retrieval of content while changing bitrates during a session so that the video could continue even though the network bandwidth was inconsistent and not

stable. And even though the playout of the video may decrease in quality (in ways that may or may not be discernable to the viewer), the delivery and hence the playout was continuous, and viewers accepted this tradeoff. *See infra* §IX.C; *see also* EX2036, 2 (“more than 50 minutes”); EX2035, 1 (“more than 50 minutes”); *see also* EX2034 (“22 minutes a session, compared with 3 or 4 minutes” previously); EX2048 (“It also means people watch longer.”). This aspect of the inventions is fundamentally different from the prior art in which viewers suffered from discontinuous playback, constant buffering, packet losses, and latency. *Cf.* EX2077. Thus, the inventions of the ABR Patents do not suffer from the same limitations to functionality and performance of prior systems plagued by common streaming solutions at that time.

207. Further, because the video is encoded in this way, simple HTTP can be used to stream, 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. EX2022 at 22. It also enables rich caching, which results in the following benefits: “1. Scalability to support hundreds or thousands or millions of viewers with little impact on managed infrastructure by utilizing CDN caching servers; 2. Rapid delivery of video to viewer clients; 3. Allows quick switching between live and pre-recorded content; 4. Supports ad insertion and other ‘stream assembly’ functionality without compromising the viewer experience; 5. The using

of caching proxies to filter requests allows for rights management and access control; and 6. The ability to quickly add capacity through CDN peering (expanding the capacity of an existing CDN by connecting it to another CDN through a CDN peering gateway).” *Id.* What the inventors did was invent how to prepare streaming content so it could be stored and served from anywhere without requiring proprietary or custom server hardware and software. Some of the largest technology companies in the world were trying to push streaming forward using proprietary hardware and software and RTSP technology. EX2020 at ¶ 6. Such systems, however, are not as scalable as the inventions of the ABR Patents and thus, had such technology been fully adopted by the market, streaming video over the Internet may never have become as ubiquitous as it is today.

208. Further, “[t]his advanced video storage and distribution is leveraged in the player environment on the client once the content is delivered to the client. The benefits of this distribution model include the intelligence that controls the quality of service residing in the client and the low cost of infrastructure when compared with hardware running media server software that requires a constant session connection with the client.” EX2022 at 22. In addition, giving the “client control” provides the benefit of more “rapid adaptive management” for a better end user experience. EX2022 at 22; *see also* EX2077 at 1-2; EX2069 at 19. That is, the inventors moved the intelligence into the client so that the server did not have to

track what every client was doing or needed and thereby enable the server to serve more clients.

209. Therefore, the Move Networks system is covered by the inventions disclosed and claimed in the challenged claims. Moreover, the features claimed in the challenged claims that impact the performance of the streaming system, and that led to the inventive aspects of the challenged claims, are present in the Move Networks system. *See, e.g.*, EX2071; EX2020, ¶¶[24-29].

210. As will be apparent from the discussion above regarding the elements of the challenged claims and the particular benefits that arise therefrom, the objective evidence discussed in the sections below directly ties to the novel elements of the challenged claims. In other words, the objective indicia of long-felt need, failure of others, unexpected results/skepticism, industry praise, commercial success, and copying were present because of the patented features of the challenged claims, and not present in the prior art. *See, e.g.*, EX2071. And while the objective evidence is tied directly to the novel claim elements that were practiced by Move and are practiced by Sling TV (*see, e.g.*, EX1009 at 101-112; EX2020, ¶ 24; EX2023; EX2100 ¶[21]; EX2127 ¶[21]), the objective evidence also directly relates to claimed features and benefits that the claims as a whole provide. *See* EX2071.

211. For example, regarding a nexus between the objective evidence of long-felt need and the challenged claims, the objective evidence shows that the claimed

features regarding streamlets, offering playback at multiple bitrates, distributing the streamlets based on end-user demand via traditional servers, putting the end-user environment in control, and “ensur[ing] the fluidity of playback” “makes it possible for example, to see HD when the capacity is there but lower quality video when someone else taps into your shared wireless connection. It also means people watch longer.” EX2048; EX2031 (“Move’s trick ... is to deliver video as a standard web object, cached in small bits. ... Move doesn’t use special media servers and adjusts the video quality based on available bandwidth. The streaming is made possible using a small applet. ... Benefits to the viewer: no buffering, no stalling ... high quality video”); EX2039 (“Streaming enables users to click and start watching immediately without the delay associated with a download”).

212. Similarly, regarding a nexus between the objective evidence of failure of others and the challenged claims, the objective evidence shows that the claimed features regarding using standard asynchronous HTTP communication requests were not expected to support streaming video content over the internet. Further, the evidence of commercial success discusses other attempts to develop internet streaming solutions attained so few viewers that, for industry leaders such as AOL, “the company didn’t even track it.” EX2047. These attempts were incompatible with certain computers and browsers, forcing the companies leading these other attempts to “giv[e] up on their respective attempts at delivering HD quality video online.”

EX2047. The evidence of commercial success also discusses that existing hardware and software solutions did not support dynamically changing audio or video formats mid-stream and, therefore, could not support adaptive bitrate streaming. EX2043 at 1-4.

213. The objective evidence of surprising results/superior performance also relates to claimed features of the challenged claims, such as the ability to load quickly (fast start) and breaking the video up into streamlets. EX2075 (“ABC.com Readies Online HD Shows as ‘Streamlettes’ from Move”). As another example, the objective evidence relating to commercial success also relates to the claimed features, such as segmenting the stream into “small stream loads” (*i.e.*, streamlets/files) to maximize video flow and image quality (especially for live content), the recognized difference over other technologies is that Move’s streaming protocol is controlled by the downloaded media player rather than being controlled by a media server (thereby enabling standard HTTP web servers to be used), using HTTP requests to retrieve streamlets, all of which contributes to having much lower communication overhead that gets more bits delivered. EX2050 (“ABC.com to Stream HD Video”). Further, the evidence of commercial success also discusses how Move’s technology automatically upshifts/downshifts to enable the best streaming experience possible based on the end-user’s local computer environment. EX2050; EX2034 (touting the ability to retrieve video at the best possible rate based on the

network traffic and unmatched scalability using standard internet protocols). Further, the evidence of commercial success also discusses how Move’s technology delivered a “full-screen HD-quality” broadcast that was “shockingly good” and that the writer “was a bit shocked” “to come home one night and find Lost streaming across a 19” flat-screen computer monitor in crisp, clear glory” where the “quality of the picture was better than you got on a traditional TV set just five years ago.” EX2065 (Apple and Move NDA Agreement).

214. The objective evidence of commercial success also relates to claimed features of the challenged claims and not any parallel technical advances, such as improved bandwidth capacities. *See, e.g.*, EX2041 (“There is enough bandwidth, but they cannot get the same results you do.”); *cf.* EX2048 (“Move Networks and the Quality Solution”); EX2099 ¶¶ 8–11; EX2100 ¶ 29. Indeed, the objective evidence of commercial success discusses that Move Networks was handling “more than a million full episode streams a week and viewership [w]as doubl[ing] every month” and was delivering roughly as much data as YouTube “but more than the Web sites of CNN, NBC, and CBS.” EX2024. Additionally, the objective evidence of commercial success discusses that “the full-screen version offered by Move was ‘incredibly crisp and clean compared to Flash.’” EX2026 at 2. Praising Move’s product as superior because it could “stream videos at an adaptive rate” continuously, many touted Move’s ability to adapt the video quality during the

stream without sacrificing playback continuity. *See* EX2051. The objective evidence of commercial success also discusses Move Networks' successful implementations of adaptive bitrate streaming for ABC, Fox, The CW, CBS, the Democratic National Convention, Oprah Book Club, Televisa, CW, Comcast, ProSieben, and ESPN. EX2024; EX2070. Further, the objective evidence of commercial success discusses that Move Networks' technology provided benefits to both users (no buffering, no stalling, and backwards and forwards navigation within high-quality and live streams) and content providers (cheaper, detailed reporting, and high-quality video), resulting in longer Internet video streaming sessions. EX2021; EX2029 ("Streaming enables users to click and start watching immediately without the delay associated with a download"); EX2026. Other objective evidence of commercial success discusses Move Networks was, at one point, adding new viewers at a rate of 100,000 per day. Further objective evidence of commercial success discusses that third parties, such as Level 3, adopted Move Networks' technology, because it could be sure that "viewers of their content will always be able to watch the video stream on whatever computer they happen to be using" and because it enabled Level 3 "to optimize for a more centralized and cost-effective content delivery platform." EX2029. Yet other objective evidence of commercial success discusses how Microsoft was interested and excited about the partnership with Move Networks. *See generally* EX2060 (Move and Microsoft Term Sheet). Yet more objective

evidence of commercial success discusses demonstrations, including one before Microsoft vice president Scott Guthrie, the-then Corporate Vice President of the Developer Division at Microsoft, and the then-CEO of Disney, Bob Iger, showing Move Networks' plug-in "could shift between different bitrate streams without pausing or freezing" and produced "noticeable changes" in the quality of the video in response to the throttle of bandwidth." EX2020 at ¶¶24-26. Additional objective evidence of commercial success discusses demonstrations at the 2008 National Association of Broadcasters (NAB), MIX08, and Democratic National Convention in August 2008, where the "live, online webcast of the Democratic National Convention was the first large demonstration of the integration of Move Networks and Silverlight," "was a technological success and tremendously popular," and the "highest-quality video experience [Microsoft's customers] have ever seen online." EX2058. Additional objective evidence of commercial success discusses multiple investments in Move Networks, based on its innovative ABR technology, including \$11.3 million investment in its first institutional round of venture funding in February 2007, another \$46 million in a round C of venture financing, which included investments from Steamboat Ventures, the venture capital arm of The Walt Disney Company; Hummer Winblad Venture Partners; Comercio Mas, an affiliate of Televisa with whom Move Networks had worked; Comcast; Level 3; and Cisco; and an additional \$8 million, including \$5 million from Microsoft. EX2048 ("Here

Comes More HD Video On the Web – Move Networks Raises \$46 Million in Round C”); EX2058.

215. As another example, the objective evidence of copying relates to claimed features of the challenged claims, such as “the encoding and use of multiple bit rate ‘streamlets’ with the same time index; such that the aggregate set of streamlets, stored on standard web servers, independently yield playback of the original full length video. This is accomplished through intelligent requests of successive optimal streamlets sent by a client device.” EX2024. “Following [Move’s] market introduction in 2006, practically every major online video provider – including Akami, Apple, Adobe, Limelight, Microsoft, Netflix, Widevine and several others – has deployed an adaptive bit rate architecture inspired by [Move’s] invention. EX2024. The objective evidence of copying also discusses that Move Networks brought adaptive bitrate streaming “to the fore” and that “many companies” were “turning toward adaptive bitrate streaming, a technique of detecting a watcher’s bandwidth and CPU capabilities in real time and then adjusting the quality of a video stream[t]hat requires encoding a single video at multiple bitrates and switching to the most appropriate one on a moment-by-moment basis. The result: very little buffering, fast start time and a good experience for both high-end and low-end connections. EX2046. Additionally, the evidence of copying states that “Akamai, Apple, Adobe, Limelight, Microsoft, Netflix and Widevine all [] use

adaptive bitrate streaming that was ‘inspired’ by Move’s invention” and that Move Networks disclosed its technology to at least Apple and Microsoft under non-disclosure agreements. EX2065; EX2066.

216. Likewise, the objective evidence relating to industry praise and peer recognition frequently discusses and relates to the claimed features of the challenged claims, such as seamlessly adjusting the video quality to the bandwidth available (e.g., ensuring continuous playback), allowing greater amounts of data to be transferred by streaming the data from standard HTTP servers, delivering the video as a standard web object that’s cached in small bits, all of which provide a “technology advantage” and a “dramatically better viewing experience.” EX2026. For example, industry participants praised the fluidity with which playback occurred, commenting that while “the Verizons of the world [were] installing fiber to deliver IPTV over a closed network,” MOVE was focused on “improving what [it] call[ed] ‘*quality of experience*’ for online media delivery over public networks.” EX2048. “*The key*,” one article writes, “*is in [MOVE’s] plug-in*...Because it puts the end-user environment in control, it can determine the best end-user experience *intelligently...ensur[ing] fluidity of playback*.”

217. Additionally, the objective evidence relating to industry praise and peer recognition discusses the superiority of Move’s technology to its peers and notes multiple awards received, for example, for its “innovation, market potential,

customer adoption, media buzz and investor value creation,” including the OnHollywood 100, OnMedia 100 Winner, Red Herring 100 Winner, Red Herring Global 100, and a TechCrunchies Finalist. EX2026; EX2027; EX2028; EX2029; EX2030. Each of these awards recognized the innovation of Move’s Media Player, which was praised for delivering the highest quality content and a preeminent viewing experience of online television. Further, the objective evidence relating to industry praise and peer recognition discusses Move Networks’ technology as “a significant technology enhancement – a true step forward – in the delivery of HD content over the Internet” providing “the highest-quality online video viewing experience we’ve seen, as demonstrated on ABC.com. Level 3 will be integrating Move’s technology and reselling it.” EX2039 (“Level 3 Collaborates With Move Networks to Provide ‘True’ HD Streaming”) at 3-4; EX2044. More recently, Sling TV (which I understand the ITC found to practice claims of the challenged ’564 and ’156 Patents, *see* EX1009 at 101-12) was recognized as a “pioneer in adaptive bit rate streaming” and received “the Technology & Engineering Emmy Award in the category of “Pioneering Development of Large Scale, Cloud Served, Broadcast Quality, Linear Channel Transmission” at the 71st annual Technology & Engineering Emmy Awards on April 19, 2020 at the National Association of Broadcasters annual conference in Las Vegas.” EX2023.

B. Long-felt but unmet need and failure of others

218. The claimed invention solved a long-felt need for streaming live and on-demand content over the Internet. I understand that evidence of a long-felt but unresolved need weighs in favor of the non-obviousness of an invention because it is reasonable to infer the need would not have persisted had the solution been obvious. I also understand that a long-felt need coupled with the failure of others to adequately address the need is particularly probative of non-obviousness.

219. Prior to the introduction of adaptive bitrate streaming technology, “the quality, fluidity and reliability of most broadband streaming [] rarely ever matched that of typical cable, DBS or terrestrial television.” EX2040 (“The Sport Page ESPN 360”). There were “well-known problems associated with streaming at high bit rates over TCP networks,” including but not limited to high latency, packet loss, and congestion, which were manifested to the viewer as buffering, delays, stalling, and unresponsiveness. EX2029; EX2066; EX2059. Additionally, non-streaming video content had delays as content was downloaded, unlike streaming video which enables users to click and start watching immediately without the delay. *See, e.g.*, EX2021 at 2-3. Typical broadband streaming experiences were “hesitant” and had “freeze-up tendencies,” meaning they were unable to provide continuous playback. EX2040. Dial-up users using contemporary technology would not “typically attempt video streaming, and broadband users often give up after prolonged waits or

unresponsiveness due to regular rebuffering.” EX2067. Put simply, pre-existing streaming solutions were “too choppy and too fuzzy, with a viewing window far too small to enjoy.” EX2103 (“Baseball Tests Online Broadcasts” – NY Times 2002). Others, like Adobe’s Flash, were seen as “broken” even with “enough bandwidth.” EX2041 (internal emails at Move Networks from 2007).

220. Move’s ABR technology directly addressed these issues and actually changed viewership patterns. In the absence of Move Networks’ ABR technology which provided continuous playback of content, many Internet users limited their Internet streaming to only a few minutes or “chunks” of content. EX2024; EX2025. Even in late 2007, “short clips under three minutes still ma[d]e up the vast majority of what people watch[ed]” and industry estimates “placed average video viewing times across the Internet close[] to three minutes.” EX2025; EX2026 at 2. One study examined more than 2,400 different Internet video streams showing that “51% of stream viewers were frustrated with the experience; 17% of streams have an immediate fatal error; and 25% take between 5 and 10 seconds to start playing, and stop to rebuffer before the video is finished.” EX2067. These experiences, including “jerky, five-minute productions on YouTube,” failed to satisfy “consumer’s real, pent-up demand for HD content.” EX2024; EX2027.

221. Before the development of Move Networks’ ABR technology, it was felt that “HD is not yet ready for primetime, because it consumes a lot more

bandwidth and therefore costs more to provide,” and due to the high costs of “high end media servers used for traditional Flash and Windows Media encoding.” EX2038; EX2029. However, “whoever can solve or bypass these problems [with network bottlenecks and slow connection speeds at people’s homes] will become adopted by more video sites as the demand for HD video rises.” EX2048.

222. Following the introduction of Move’s technology, users began “spending more than an hour per session watching content Move provides.” EX2064 (LDS General Conference 2005 Transcript). Move’s technology allowed users to enjoy the high-quality long-form content at the high-quality levels users expected. *Id.* Move’s “technology is good enough” that some thought that Move could do what “many in Hollywood and Silicon Valley have so far been unable to do: take Internet video from its YouTube origins to a full-fledged television service with dozens of channels.” EX2028.

223. Move Networks was successfully developing its patented ABR technology, which provided continuous content playback, while others failed. EX2026, at 2. For example, other technologies, such as progressive download, peer-to-peer, and proprietary RTP/RTSP streaming systems requiring specialized servers and dedicated circuits were proposed as alternatives, but could not achieve the same quality of experience as that provided by the patented ABR technology and were ultimately disfavored in comparison to the patented ABR technology. *See* EX2026

at 2. Indeed, several entities developing other means to deliver high-quality long form content over the Internet ultimately abandoned their proprietary systems after years of development and adopted the patented ABR technology designed and developed by Move Networks. In November 2007, both AOL and CBS gave up their attempts to deliver HD quality video online. “Fred McIntyre, SVP of AOL Video, said the viewership for its two-year-old Hi-Q service was so low, the company didn’t even track it.” EX2037. CBS Interactive’s president, Quincy Smith, stated an issue with its player included limited compatibility with all computers and browsers. *Id.*

224. As shown in EX2037 and EX2024, other companies were offering clearly inferior alternatives or “giving up on their respective attempts at delivering HD quality video online” in view of the technological hurdles that Move ultimately overcame. YouTube offered “jerky, five-minute productions” and other Flash-based implementations offered only “grainy videos,” as explained in EX2024 and EX2041. Neither NBC, using a Flash-based solution, nor CBS’s Real Player-based solution “provided great visual quality,” as stated in EX2043. Ultimately, ABC.com abandoned its use of Flash in favor of Move Networks’ adaptive bitrate streaming technology, as shown in EX2034.

225. I also note that, despite British Telecom’s possession of Leaning and incentive to satisfy the industry’s long-felt need, British Telecom never developed or commercialized the video teachings of Leaning. *See* EX2126. As British

Telcom's Chief IP Lawyer explains, "BT engages with network equipment manufacturers, and other third parties, when it needs to test, trial or procure ('commercialize') products or services." *Id.* Further, "BT is one of the world's oldest telecommunications companies and ... due to BT's significance in world history, BT has built an archive of records extensively documenting its research, development, history, and commercialization. Known as the 'BT Archives,' the archive includes records such as minutes, advertising material, photographs, and major collections of the company's history." EX2126 ¶5. BT has no records of commercializing the video teachings of Leaning, as confirmed by British Telecom's chief IP lawyer and the direct manager of the Leaning inventors. EX2126 ¶¶10-11.

226. Move's patented ABR technology, on the other hand, was so transformative that the president of Level 3's Content Markets Group called Move's Media Player "a significant technology enhancement – a true step forward – in the delivery of HD content over the Internet" because of its ability to "provide a better live video experience for fans than that currently offered by any other service." EX2039; EX2040.

227. Indeed, after others failed to develop their own competing solutions, practically every major online video provider has shifted to deploy an adaptive bitrate architecture that was inspired by the patented ABR technology. *See, e.g.*, EX2024; EX2020.

228. For example, Move Networks began sharing its technology with Microsoft in April 2007. At that time, Microsoft's system, including its Silverlight browser plug-in, did not support dynamically changing audio or video formats mid-stream and, therefore, could not support adaptive bitrate streaming. EX2042.

229. Move frequently faced skepticism that its technology, which included asynchronous communication protocols, such as standard HTTP communication requests, could support high-definition video streaming when existing solutions, including for example, progressive download or server-side technologies were not working well. *See, e.g.*, EX2041 at 1-2 (“[W]hy are the current streaming protocols broken? There is enough bandwidth but they [competitors] cannot get the results you [Move] do.”).

C. Surprising results and superior performance

230. I understand that evidence of surprising results and superior performance weighs in favor of non-obviousness of an invention because it is reasonable to infer that there would be no surprise regarding the invention if it had been obvious. As I explain below, it was the Move Media Player's ability to provide continuous playback of content that resulted in a high-quality playback experience, prompting substantial viewer and industry surprise.

231. The introduction of Move Networks' ABR technology led some to suggest that “someday you'll be able to ditch your cable or satellite provider and get

your content from cheaper sources that also happen to be legal” and make “online streaming HD downloads a legitimate competitor to traditional media.” EX2032; EX2037.

232. The commentary surrounding products embodying Move’s ABR technology indicate a substantial amount of surprise. *See, e.g.*, EX2053 (“Kudos to ABC for being a frontrunner in providing quality content to their customers.”). ABC.com’s use of its Move-powered streaming player gave viewers a continuous stream of “full-screen ‘HD-quality’” broadcast that was “shockingly good.” EX2032; *see also* EX2054 (“Move continues to offer the highest quality online video viewing experience we’ve seen, as demonstrated on ABC.com”). ABC.com’s player was deemed one “of the most spectacular online video players” because it “shows remarkable, full-screen high quality” content and “load[ed] quickly.” EX2075; EX2082; EX2083; EX2084.

233. In February 2008, one user, who previously “thought we would be waiting years to see viable streaming HD technology,” stated he “was a bit shocked” “to come home one night and find Lost streaming across a 19” flat-screen computer monitor in crisp, clear glory” and that the “quality of the picture was better than you got on a traditional TV set just five years ago.” EX2037. He further stated he couldn’t “remember the last time that technology actually surprised [him] by

arriving before [he] expected it” but “[j]udging from the ABC site, the Move Networks stuff really works.” *Id.*

234. Additionally, ABC’s online player “clearly deliver[ed] a better, sharper picture” using Move Networks’ technology than NBC’s player, which “relie[d] on Adobe’s Flash to stream their video” and “clearly doesn’t come close to the quality that ABC is offering.” EX2049. Additionally, unlike Adobe Flash, where the brains of the operations were located at the server, the key to Move Networks’ technology “is in the plugin,” which “can determine the best end-user experience intelligently” to provide continuous playback of content. EX2048. “Move’s adaptive protocol means that the plugin can determine the best encoding profile for a user based on available bandwidth.” *Id.* “This ensures fluidity of playback and makes it possible for example, to see HD when the capacity is there but lower quality video when someone else taps into your shared wireless connection. It also means people watch longer.” *Id.* Additionally, unlike “BitTorrent and Joost, Move delivers its video as a standard web object that’s cached in small bits—making its service infinitely more ISP-friendly than traditional streaming technology. Move also doesn’t use special media servers, and adjusts the video quality based on available bandwidth. This means no buffering, no stalling, and backwards and forwards navigation within a live stream while you’re watching.” EX2036 at 2.

235. Move's technology was also regarded as superior to previous streaming solutions. For example, it was noted that Move's streams "looked better than the flash based full-screen grainy videos" and "Move Networks player stream videos at an adaptive rate, which means the quality of the video depends on your CPU and bandwidth, as opposed to Flash videos that look the same no matter how good your resources are." EX2051 ("ABC Adds Full-Screen Mode to Online Shows"). "CBS' InnerTube system uses Real Player, while NBC's video content is flash based" and "[n]either provides great visual quality. ABC uses a player by Move Networks, which offers a much clearer picture quality" – "it's clear that ABC has the superior product." EX2053 ("ABC Offering Full Screen Streaming HD") at 1.

236. Others also noticed the Move Networks player on ABC.com, noting that "[o]ne of the most spectacular online video players" and "had attracted quite a bit of attention. The player shows remarkable, full-screen high quality. And, it loads quickly." EX2084. After a demonstration of Move's technology to Level 3, a senior vice president at Level 3 asked Move personnel "why are the current streaming protocols broken? There is enough bandwidth but they cannot get the results you [Move] do." EX2041.

237. Adopters of Move's technology noted that "[n]ot only is Move's technology superior to others we have used, but their ability to manage large scale setup and operations with flawless delivery is unmatched." EX2052.

D. Commercial Success

238. There is a substantial amount of real-world evidence that shows, in my opinion, that Move’s product that embodied the Asserted Claims was commercially successful from a number of perspectives.

239. Move’s initial demonstrations were for the Church of Jesus Christ of Latter-Day Saints. In October 2005, MOVE collaborated with Brigham Young Television (“BYU.tv”) to broadcast the Church’s General Conference. EX2064 (October General Conference Notes, at 1); EX2127 ¶5. Move’s former Vice President of Operations, David Ericson, stated that “the event was tremendously successful, and BYU.tv has recognized that its 2005 use of MOVE’s adaptive bitrate technology made it ‘the first television network to stream its signal uninterrupted and free of charge 24/7 on the Internet.’” EX2100 ¶17 (quoting EX2120). In fact, BYU has dedicated a plaque commemorating the BYU.tv and Move collaboration, which states “BYU TV launches live stream using Move Networks’ adaptive bitrate technology, essentially making BYU TV the first television network to stream its signal uninterrupted and free of charge 24/7 on the internet.” EX2120; EX2027 ¶5; EX2128, 28:9-29:13. Mark Mitchell, the Manager of Business Intelligence for BYU Broadcasting, states that “the ABR technology provided by MOVE and its Media Player allowed to provide continuous streaming experiences to viewers—even during times of network congestion—by shifting the quality, ‘to prevent the stream

from buffering.’” EX2127 ¶5 (quoting (EX2128, 29:14-30:7). “At the same time,” according to Mr. Mitchell, “BYU TV also offered a Windows Media Stream, but the stream using the MOVE Media Player ‘was definitely better than our Windows Media stream by leaps and bounds.’” *Id.* ¶5 (quoting EX2128 at 90:17-91:6).

240. Relying on its successful early demonstrations (most notably, its streaming of the October 2005 General Conference for the Church of Jesus Christ of Latter-Day Saints), the objective evidence shows that Move worked to ensure that tens of thousands of people from all over the world could pull the same files at the same time to watch the same event without freezing, buffering, or extensive latency, which plagued nearly every commercially available solution at that time. EX2064 (October General Conference Notes, at 1). MOVE’s ABR streaming technology was subsequently used in the October 2006 and October 2007 General Conferences for the Church of Jesus Christ of Latter-Day Saints. EX2100 ¶17 (citing EX2073); *see also* EX2061.

241. In EX2034, it was recognized that Move Networks, and its ABR technology, were “at the forefront” of Internet streaming. This article explains that “Move does for video what voice over Internet networks did for telephone calls: it breaks up the video into bits and efficiently reorganizes them over the network so there’s no need for the special computer servers and dedicated transmission lines required on streams using Flash.” The commercial result of Move’s technology, as

this article reflects, is that Move Networks was handling “more than a million full episode streams a week and viewership [w]as doubl[ing] every month.” EX2034. In May 2007, Move Networks was delivering roughly as much data as YouTube “but more than the Web sites of CNN, NBC, and CBS.” EX2034.

242. In April 2007, within months of Move’s commercial launch, ABC turned away from Adobe’s Flash, which was then used by many large media companies, and launched on ABC.com a new broadband video player using Move Networks’ technology that permitted users to watch shows such as “Lost,” “Dancing with the Stars,” “Grey’s Anatomy” and “Desperate Housewives” in a full-screen format. EX2083. Notwithstanding that Move was a new company, “ABC was astounded by the video quality we were able to deliver online and said we need to have this for our customers,” as is stated in EX2044. In fact, it was reported that Move was handling “more than a mullion full episode streams a week,” with a sustained viewership that doubled every month. EX2034. By March 2008, viewers were averaging over 50 minutes or more per stream (*compare with* EX2025 (viewers previously averaged roughly three minutes)).

243. ABC’s employees said publicly, such as ABC’s vice president of digital media, Alexis Rapo, stated in EX2044 and EX2036 (Move Over Joost) at 2 that “the full-screen version offered by Move was ‘incredibly crisp and clean compared to Flash’ and only had minimal buffering at the beginning of the stream,” which she

hoped would “lead to a better viewing experience for users, which in turn, should boost the effectiveness of the commercials that advertising are placing within the streams.” Rapo also stated that “Move’s type of streaming ‘allows the network to dramatically improve the user experience and operational efficiency’ on the newly available player for streaming ‘long form’ (full episodes) of TV series.” It appears ABC was right. Bloggers “raved about ABC’s new player,” as shown in EX2044. And as shown in EX2045 (Disney ABC Throws Down), others saw “the recent launch of streamed HD shows on the ABC.com website” as evidence that Move Networks was removing “technical barriers to delivering online HDTV.” In addition, after moving away from Adobe’s Flash, and adopting Move Networks’ technology, ABC’s website, ABC.com, delivered “clear and smooth streamed images” whose quality was “good, no question.” EX2034; EX2050.

244. In my opinion, the success of Move Networks’ technology was attributable to its unique design. Objective evidence shows that people were adopting Move’s product was because of the technology, for example, Move Networks’ ABR technology “dramatically improve[d] the viewing experience.” EX2050; EX2036 at 2. Early experiences with Move’s Media Player proved that the quality of its streaming delivery was “most excellent” and “actually watchable.” EX2031. In fact, some said that Move did “what many in Hollywood and Silicon

Valley have so far been unable to do: take Internet video from its YouTube origins to a full-fledged television service with dozens of channels.” EX2038.

245. Regarding Move’s unique design, “[r]ather than streaming being controlled by a media server, in [Move’s] solution the server is just a standard http Web server [because Move’s] delivery protocol is based on http—the same that’s used to serve up a page of any other online content.” EX2050. ABC also believed that Move Networks’ streaming solution enabling long-form content would “lead to increased viewing by users and to improved economics for [ABC’s] streaming product.” EX2050 at 2.

246. Move’s commercial success is also attributable to its Media Player’s ensuring that each end user device “adjust the streaming at its fullest potential—depending on a variety of variable conditions that will be different with each desired viewing (*i.e.*, online traffic, broadband connection, computer speed, type of content, etc.).” Additionally, EX2050 also recognizes that Move Networks’ software also optimized “the video for playback,” by cutting it “up into small stream loads” or streamlets and “dealing with much smaller files to maximize video flow and image quality.” *Id.* Because Move’s technology “automatically upshift[ed] and downshift[ed]” while still providing seamless playback, ABC commented that Move’s “streaming software [did] not require any significant additional technical commitment on ABC’s part, nor [did] it cause any stress on [the] Web site’s

infrastructure.” *Id.* at 2. Another article summarized the impact of Move’s technology by stating that “Move’s success in attracting viewers since its launch proves that if you offer Internet viewers high quality, full-screen video content, they just might sit still and watch an hourlong program.” EX2036 at 2.

247. Move Networks’ technology provided benefits to both viewers (no buffering, no stalling, and backwards and forwards navigation within high-quality and live streams) and content providers (cheaper, detailed reporting, and high-quality video). EX2031; EX2036 at 2; and EX2039. For example, “Fox, which tested Move Networks’ technology on the Web sites of 24 of its affiliates last year, plan[ned] to expand to 200 affiliates” in the summer of 2007. EX2034. This article also notes that “[a]nother broadcaster, which was paying \$1 million a month to deliver jerky versions of its shows, expects those costs to drop up to 85% with Move.” *Id.* Content providers, like Fox, saw the benefit of Move Networks’ technology as enabling increased advertising revenues and supporting consumption of additional content. Ronald Berryman, senior vice president at Fox Interactive Media noted that its customers, who were now enjoying the benefits of Move Networks’ ABR technology, were now watching “22 minutes a session, compared with 3 or 4 minutes,” so we “can increase the advertising and deliver long-form content,” which was reported in EX2034.

248. Additionally, EX2034 states that “[u]nlike YouTube-style Flash video, which fetches streamed bits in a series of requests sent from a set of servers at the sender’s end, Move gets bits from the closest storage cache ... and brings them back to the screen at the best streaming rate based on the network’s traffic load. It uses standard internet protocols, which means it can take advantage of the many server farms around the world that offer up Web pages. Both ABC and Fox say there is nothing else that can stream at this scale.”

249. Moreover, it is recognized that “Move helped paved the way for HD video” in EX2086. By November 2007, Move Networks was providing “high-quality streaming video platforms to media companies like ABC” such that “more than twice as many people are using its platform than in August” of 2007. EX2085; EX2087. “In absolute terms, more than six million people watched long-form video using Move in October [2007] and Move [was] adding new viewers at a rate of 100,000 per day.” EX2085.

250. Around the same time, Level 3 announced a collaboration with Move Networks. EX2039. This article states that “Level 3’s Move-enabled streams will dynamically adapt to changing network and end-user computer conditions. If bandwidth availability is constrained for any reason, or if computing power is limited, the Level 3 stream will reduce its bit rate automatically, without buffering. ‘It is a graceful transition – up and down – that most users will not even notice,’ said

Level 3's Rafuse. 'With the Move Media Player, customers know in advance that viewers of their content will always be able to watch the video stream on whatever computer they happen to be using.' This technology will enable Level 3 to optimize for a more centralized and cost-effective content delivery platform. ... Utilizing Move's technology, Level 3 will be able to provide a very large amount of capacity, immediately, for its customers who may have very large audiences." *Id.*

251. In October 2008, ESPN re-launched its "360 online video player, and its really pretty groundbreaking." EX2040. "[P]owered by Move Networks, the player offer[ed] the first instance of live streaming using Move's technology, promising a better live video experience for fans than that currently offered by any other service," also reported in that article. And even "on a slow wireless connection," a user could "get a relatively good looking picture and more importantly, an entirely fluid stream." *Id.*

252. Move Networks demonstrated the benefits of its ABR technology to Microsoft and was working with Microsoft to integrate that technology into Microsoft's Silverlight software. Following Move's demonstrations to Microsoft in 2007, Microsoft had internal debates as to "the merits of investing [in] or acquiring" Move Networks, Microsoft invited John Edwards, the then-President and CEO of Move Networks, and Greg Smith, the then-CTO of Move Networks to travel to Redmond, Washington to meet with Microsoft personnel from a number of business

units, such as Media Center and Xbox, to discuss how Move Networks' technology may benefit each group in early June 2007, as shown on page 4 of EX2042. Following that meeting, Greg Smith of Move met with Chris Carper, who oversaw the business development of Microsoft's Silverlight browser plug-in and subsequently worked closely with Move Networks to integrate Move Networks ABR technology into Silverlight to address the lack of adaptive streaming support in Silverlight, which is confirmed by EX2043, discussing different functionality that is included in Silverlight and not included. EX2043 at 2.

253. Before its integration with Move, Microsoft's software did not support resolution changes or switching between streams with different bitrates during playback, both of which were key to implementing Move Networks' ABR technology, as shown in EX2043 at 1-4. The emails between Move and Microsoft personnel show that computer hardware at this time was such that displaying an HD stream required hardware acceleration. Given these limitations, Move Networks quickly found a solution for Microsoft in which a Silverlight window could be filled with video received and processed by Move Networks' client, which supported both DRM and hardware acceleration. *See, e.g.*, EX2081.

254. Move's business continued to grow in late 2007 and early 2008. By March 2008, users watching content provided by Move Networks were averaging 50 minutes or more per session and numbered more than 34 million, substantially

more than the 20 million unique viewers it had achieved in November 2007, as reported in EX2074; EX2036 at 2. By December 2008, Move Networks had powered seven of the top ten online broadcast TV network entertainment programs. *See* EX2114.

255. Microsoft also requested Move Networks support a demonstration of the “full featured player” for the 2008 National Association of Broadcasters (NAB) conference, which is shown in EX2088. Further, EX2055 and EX2089 show that in March 2008 at MIX08, Microsoft’s web developer and designer conference for showcasing upcoming web technologies, Microsoft also discussed announcing its relationship with and investment in Move Networks, and then invited Move Networks to attend Advance, Microsoft’s internal tradeshow directed to advertising, business, entertainment, and communication, in mid-May 2008.

256. Move and Microsoft collaborated to support streaming of the Democratic National Convention in August 2008, with additional assistance from Level 3, which Microsoft executives praised as “a technological success and tremendously popular, with visitors to the official convention Web site viewing more than 350,000 hours of video.” EX2070 at 44. *See also* EX2060 at 49-51; EX2058. For example, as shown in EX2091, several praised the HD “online video feed” at DemConvention.com, which used Move Networks’ plugin, labeling it as “the most awesome (I have probably never used that word in seven years of this site) online

video feed.” In one article, Microsoft’s Scott Guthrie explained that the “live, online webcast of the Democratic National Convention was the first large demonstration of the integration of Move Networks and Silverlight” and that viewers “experience[d] the convention online like never before.” He also reported it was the “highest-quality video experience [Microsoft’s customers] have ever seen online.” EX2070 at 5.

257. Microsoft also discussed Microsoft’s “strategic investment in Move Networks Inc., as the two companies continue to work together to deliver high-quality, skip-free Internet television programming over HTTP-based networks” in EX2070. Microsoft explained that “Move Networks’ technologies allows users to experience high quality, skip-free Internet television programming” in EX2070 at 5. After calling Move “a leader in HTTP-based video delivery,” Microsoft highlighted that Move had “developed an adaptive streaming technology that seamlessly adjusts image quality to fit available bandwidth on each viewer’s Internet connection, creating a high-quality Internet television viewing experience. [Its] investment builds on a strategic partnership [Microsoft] announced with Move Networks in March of 2008.” EX2070 at 5; *see also* EX2092 at 4-5. With this investment, Microsoft not only recognized Move as the “leader in HTTP-based video delivery”, but it also saw so much value in the success of the patented ABR technology that it

worked with Move to integrate its Silverlight product with the Move Media Player. EX2079.

258. The success, and successful adoption, of Move Networks' ABR technology was recognized by investors as well as other media companies. Move Networks became "one of the preferred video streaming partners for many major media sites, including ABC.com, Discovery.com, ESPN.com, and Fox.com" led to investments in a multiple series of venture funding. EX2058. For example, Move Networks received a \$11.3 million investment in its first institutional round of venture funding in February 2007. EX2057. Move Networks also received \$46 million in a round C of venture funding in April 2008. EX2058; EX2068. Investors in this round of funding included major players like Steamboat Ventures, the venture capital arm of The Walt Disney Company; Hummer Winblad Venture Partners; Comercio Mas, a Televisa affiliate; Comcast; Level 3; and Cisco; and it later included an additional \$8 million in investments, including \$5 million from Microsoft. *See* EX2058; EX2068.

259. Content delivery networks were also interested in Move Networks' technology. Brady Rafuse, president of Level 3's Content Markets Group, stated that "Level 3 believes that Move's service is a significant technology enhancement -- a true step forward -- in the delivery of HD content over the Internet." EX2039 at 3-4. Level 3 explained that "higher bitrates, and subsequently higher viewing quality,

encourage people to watch longer. Clearly, longer session times increase potential advertising revenue when compared to lower-quality streaming solutions. We believe that this raises the potential for our customers to actually make more money from HD than at the lower bit rates.” *Id.*

260. By September 2010, Move Networks’ customers also included Oprah Book Club, Televisa, CW, Comcast, and ProSieben, as shown in EX2024.

261. DISH’s licenses for the Challenged Patents with ViacomCBS and Univision Communications mentioned in EX2063 further evidence the commercial success of the Challenged Patents. EX2063, 23 n.100. DISH has also licensed the ABR technology in the Challenged Patents to Peloton for \$75 Million, to lululemon, and to iFit. EX2093; EX2094; EX2095. *See also* Peloton Will Pay Dish \$75M to Avoid Import Ban, Law360 (<https://www.law360.com/articles/1604881/peloton-will-pay-dish-75m-to-avoid-import-ban>) (May 4, 2023).

262. I understand that in 2010, Microsoft solicited a license to MOVE’s ABR technology covered by the ’783 patent application, which is disclosed in the challenged patents. EX2125 at ¶4. MOVE declined to grant the license. *Id.*

263. Ultimately, because of the impact of the 2008 financial collapse, Move’s business was put up for sale and acquired by EchoStar. *See* EX2020, ¶¶18; EX2090; EX2100 ¶¶11, 28, 30. Move Networks was acquired by EchoStar, and its

technology was used to create and build Sling TV, a successful Internet streaming television platform. *See, e.g.*, EX2142; EX2020, ¶¶[19-21]; EX2090.

E. Copying

264. I understand that copying, particularly by competitors, can indicate that the inventions in a patent are non-obvious. I also understand that copying can be established by showing that an accused copier had access to the patented invention and that the accused system shared substantial similarities with the patented invention. I also understand that copying by competitors can indicate non-obviousness. It is my opinion that the objective evidence suggests that competitors copied the patented ABR technology.

265. By summer 2009, after “adaptive bitrate streaming [was] brought to the fore by Move Networks,” many companies turned “toward adaptive bitrate streaming, a technique of detecting a watcher’s bandwidth and CPU capabilities in real time and then adjusting the quality of a video stream. That requires encoding a single video at multiple bitrates and switching to the most appropriate one on a moment-by-moment basis. The result: very little buffering, fast start time and a good experience for both high-end and low-end connections,” as stated in EX2046 and shown in EX2036 at 2.

266. In September 2010, Move Networks announced the approval of its application for “a fundamental patent on adaptive streaming.” EX2024; EX2096.

The industry recognized that “Move Networks just got a lot more valuable on the intellectual property front, as it has just been granted a patent for its HTTP-based adaptive streaming technology. The awarding of the patent also means that it could pose a threat to Adobe, Microsoft, Apple and many other companies that have rolled out adaptive streaming technology of their own.” EX2062. This article also provides the context that “Move Networks was an early proponent of HTTP-based adaptive streaming technology for delivery of video streams that adjust to the amount of bandwidth that is available to the end user. Since the introduction of Move’s technology and the filing of the patent in question, a number of video providers have followed its example with similar technologies of their own.” *Id.* It was reported that “Akamai, Apple, Adobe, Limelight, Microsoft, Netflix and Widevine all [] use adaptive bitrate streaming that was ‘inspired’ by Move’s invention.” *Id.*

267. Similarly, after Move Networks disclosed its technology to Apple under a non-disclosure agreement shown in EX2055 and EX2056, Apple “embraced adaptive bitrate streaming over HTTP at its developer conference” in June 2009, announcing that it would “be available with the new iPhone OS 3.0 on June 17 and in September with the new version of QuickTime for Mac OS X Snow Leopard.” EX2036. Like Move Networks’ technology, “Apple’s new HTTP streaming,” which is known as “HLS” would support “live and on-demand” streaming video content. EX2036 at 2. As Move’s former VP of Operations observes, HLS, as a proposed

standard is “very similar to what Move had developed.” EX2100 ¶26. *See also* EX2100 ¶27; EX2102 at 97:1-4, 118:18-21, 157:4-5. Indeed, the HLS Standard published by Apple states in its introduction that “HTTP Live Streaming provides a reliable, cost-effective means of delivering *continuous* and long-form video over the Internet” and “[u]sing this protocol, a client can receive a *continuous* stream of media.” EX2131, 4.

268. The objective evidence also indicates that Microsoft was very interested in Move’s ABR technology. As I discussed above, in 2010, Microsoft solicited a license to MOVE’s ABR technology covered by the ’783 patent application, which is disclosed in the challenged patents. EX2125 (Declaration Regarding Microsoft Document Search) at ¶4. After Move declined Microsoft’s offer, Move and Microsoft subsequently worked closely together. *See, e.g.*, EX2099 ¶¶10-16; EX2070; EX2112; EX2110; EX2113; EX2109; EX2055; EX2089; EX2092; EX2060; EX2108. For example, a meeting between Move Networks and Microsoft occurred in June 2007, when Microsoft invited John Edwards, who at the time was the President and CEO of Move, and Greg Smith to meet with Microsoft personnel to discuss how the adaptive bitrate streaming technology could benefit each group “from both a business and technology” perspective. EX2042. Following that meeting, Greg Smith worked with Microsoft on a plug-in for Microsoft’s Silverlight. EX2042. Following an investment by Microsoft into Move (EX2050),

Microsoft used the patented ABR technology in collaborations including the streaming of the 2008 Democratic National Convention. EX2060. However, despite the successful collaborations, Microsoft decided to implement its own version of adaptive bitrate streaming technology, Smooth Streaming, as “a new extension” to enable “adaptive streaming of media to Silverlight client via HTTP.” EX2067. *See also* EX2100 ¶¶ 21, 24 (quoting EX2102 at 97:1-4). Microsoft pitched its Silverlight technology as providing nearly the exactly same benefits that Move’s technology provides: enabling “existing HTTP caches and caching networks to greatly increase the scalability of a media experience and make true HD (720p+) a reality;” offering “an easy to manage and deploy approach for content delivery;” “minimal buffering and quick start-up time;” adapting “the quality of the video stream in real-time based upon the consumer's changing bandwidth and CPU conditions;” and “using true HD (720p+) video and uninterrupted delivery to extend viewing times.” *Id.*

269. Thus, the objective evidence demonstrates that the adaptive bitrate streaming solution “brought to the fore by Move Networks” was influential in others’ adoption of “substantially similar” ABR technology. EX2024; EX2026; EX2100 ¶26.

F. Industry praise and peer recognition

270. Move and its technology received a substantial amount of praise and recognition in the industry at all levels from users, content distributors, content providers, and industry-wide organizations.

271. Move Networks was recognized as “especially impressive” given that “the quality of content Move delivers is higher than that offered on platforms like Hulu or any of the standard players like Flash.” EX2033. This praise was directed towards Move’s Media Player’s ability to “actively adjust[] its video quality to the bandwidth available. *Id.* That article also states that “Move Networks blends into the inconsistent landscape of users’ different download speeds by actively adjusting its video quality to the bandwidth available. It delivers pieces of content progressively via a piece of software downloaded and installed through the browser — importantly, users don’t need to open any separate program.” *Id.*

272. Users recognized that Move Networks’ technology took advantage of “your CPU and bandwidth, as opposed to Flash videos that look the same no matter how good your resources are” to provide an “improved high-quality player” that “looked better than the flash based full-screen grainy videos” and is “well worth your time.” EX2041. They also compared it to other technologies. EX2043. This article notes that “CBS’ InnerTube system uses Real Player, while NBC’s video content is flash based. Neither provides great visual quality. ABC uses a player by

Move Networks, which offers a much clearer picture quality” and “looks awfully nice.” *Id.* To the industry, that article states that “[I]t’s clear that ABC has the superior product.” *Id.* In addition, the industry focused on why Move was superior, noting that in contrast with Flash, which required specialized Flash Media Servers, “Move’s system allow[ed] for greater amounts of data to be transferred by streaming the data from standard HTTP servers” and enabled ABC to be a “frontrunner in providing quality content to their consumers over numerous devices.” *Id.* As noted above, in 2007, the broadcast delivered by ABC’s full-screen HD-quality video player was declared “shockingly good.” EX2022.

273. The industry also praised Move’s patented approach of using an intelligent client which “determine[d] the best end-user experience intelligently.” EX2038 (explaining that, with “Move’s adaptive protocol ... the plugin can determine the best encoding profile for a user based on available bandwidth”). When users praised the Move Media Player as being “especially impressive,” they tied this praise to the combination of claimed features, which smoothly shifted between bitrates and provided continuous playback. EX2033. Those in the industry also recognized that “MOVE delivers its video as a standard web object that’s cached in small bits—making its service *infinitely more ISP-friendly* than traditional streaming technology. Others also praised Move Networks’ technology by identifying “Move’s bottom-line technology advantage” as “its service is profoundly

more efficient and less expensive, and it provides a dramatically better viewing experience. Move's technology also lets advertisers target users and offers excellent analytics and reporting." EX2026. Or to put it less formally, users commented that "Boy, does this rock! It's sharp and clear – does full-screen video. Lost is stunningly beautiful on it. Demonstrates that on YouTube or iTunes we sure aren't getting anywhere close to the full experience that is possible today." EX2070.

274. At the content distributor level of the industry, Level 3's Rafuse also praised Move Networks' technology. "Level 3 believes that Move's service is a significant technology enhancement – a true step forward – in the delivery of HD content over the Internet." EX2029; EX2044. Level 3 praised the ability of Move's ABR technology to "automatically upshift[] and downshift[] [video content] to provide continuous video playback at the highest-quality profile possible." EX2130 (Email Regarding non-binding indication of interest).

275. At the content provider level of the industry, ABC and ESPN and other content providers praised Move's technology, as I stated above. Others also praised Move Networks' technology thusly: "Move's bottom-line technology advantage: its service is profoundly more efficient and less expensive, and it provides a dramatically better viewing experience. Move's technology also lets advertisers target users and offers excellent analytics and reporting." EX2026 at 2. Or to put it less formally: "Boy, does this rock! It's sharp and clear – does full-screen video.

Lost is stunningly beautiful on it. Demonstrates that on YouTube or iTunes we sure aren't getting anywhere close to the full experience that is possible today." EX2070.

276. Competitors similarly praised Move Networks. In addition to previously attempting to license the technology, *see* EX2125 ¶4, Microsoft praised Move's patented ABR technology as providing, "the highest-quality viewing experience available online." EX2055 at 4. Again, the praise was tied to Move's "adaptive streaming technology that seamlessly adjust[ed] image quality to fit available bandwidth on each viewer's Internet connection." EX2058 at 1.

277. Outside of the U.S., Move Networks received similar praise. Tony Rallo, the then-vice president of Technology Strategy for Televisa, stated that "Move has proven to be an excellent technology partner for Televisa time and again. ... Not only is Move's technology superior to others we have used, but their ability to manage large scale setup and operations with flawless delivery is unmatched." EX2042.

278. Move's ABR technology also earned multiple industry-wide awards. For example, Move Networks was recognized in June 2008, over 1000 other peer-nominated entries identified by leading venture capitalists, investment bankers, and industry analysts, in AlwaysOn's OnHollywood 100 for its commercial success based on "innovation, market potential, customer adoption, media buzz and investor value creation," as reported in EX2026 (OnHollywood 100) and EX2027

(OnHollywood 100 Winner). EX2026 also recognized that “Move Networks Internet television platform delivers the highest quality content and a preeminent viewing experience of online television while providing both publishers and advertisers with the ability to meet the demands of scalability and profits at the same time. By combining scalability with meaningful reporting and analytics, Move Networks has created the industry’s first online television business model that focuses on revenue generation for its content providers while distributing their content to millions of viewers simultaneously worldwide. Currently, Move Networks is used to deliver HD video content online by leading media companies such as ABC, FOX, ESPN, HARPO and many others.” *Id.*

279. In addition, Move Networks was also an OnMedia 100 Winner, as reported in EX2028 (OnMedia 100 Winner). Move was also a Red Herring 100 Winner, an award given to the top 100 private technology companies based in North America, in May 2007 and a Red Herring Global 100 in December 2007 based on evaluations “made on both quantitative and qualitative criteria such as financial performance, innovation, management, global strategy, and ecosystem integration, because it gave “content publishers the ability to deliver broadcast-quality streaming video to viewers. EX2029. This high-quality experience, coupled with comprehensive viewer metrics, and the delivery of highly relevant ads, enable

content publishers to attract and capture the attention of today's viewer that demands access to premium content anytime, anywhere." EX2030; EX2029.

280. In 2007, Move Networks was also a TechCrunchies Finalist for Best Technology Innovation, selected as one of five companies from 82,000 nominations, as reported in EX2030 (Move Networks TechCrunchies Finalist). More recently, Sling, which the ITC also found to practice the claims of the challenged '156 and '564 Patents (*see* EX1009 at 101-12), was recognized as a "pioneer in adaptive bit rate streaming" and received "the Technology & Engineering Emmy Award in the category of "Pioneering Development of Large Scale, Cloud Served, Broadcast Quality, Linear Channel Transmission" at the 71st annual Technology & Engineering Emmy Awards on April 19, 2020 at the National Association of Broadcasters annual conference in Las Vegas", which was reported in EX2023 (Sling TV Granted Prestigious Tech).

281. In sum, the objective indicia of non-obviousness or secondary considerations discussed above further support my opinions that the Challenged Claims are not invalid as obvious; Leaning expressly teaches away from the claimed invention; and the proposed combinations with Klemets, Gamble, and/or Ala-Honkola fail to address Leaning's central flaw or otherwise disclose or suggest the limitations they are relied upon to teach.

X. Conclusion


282. 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.

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

284. I reserve the right to supplement and/or amend my analyses and opinions expressed herein in response to any new positions taken by Petitioner's 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.

285. 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', is written over a horizontal line.

Executed On

February 28, 2025 at Chapel Hill, NC

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, 2024.

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
Five Years)**

Corrigent Corporation, v. Cisco Systems, Inc., United States District Court for the Western District of Texas, Waco Division, Case No. 6:22-cv-369-ADA-DTG.

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

Koninklijke KPN N.V. v. Cisco Systems, Inc., Arbitration Pursuant to the Rules of Arbitration of the International Chamber of Commerce, ICC Case No. 28424/PDP.

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

CosmoKey Solutions GMBH & Co. KG v. Duo Security, Inc. n/k/a Duo Security LLC and Cisco Systems, Inc., United States District Court for the District of Delaware, Case No. 18-1477-CFC.

Provided deposition testimony on behalf of defendants Duo Security LLC and Cisco Systems, Inc.

In re: Inter Partes review of U.S. Patent Nos. 9,407,564, 10,469,554, 10,951,680, 11,470,138, 11,677,798.

Provided deposition testimony on behalf of Patent Owner DISH Technologies LLC.

IBM Corporation, v. Zynga, Inc., United States District Court for the District of Delaware, Case No. 1:22-cv-00590-GBW.

Provided deposition and trial testimony on behalf of plaintiff IBM Corp.

Netsocket Inc., v. Cisco Systems, Inc., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:22-00172-JRG.

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

DataCore Software Corporation, v. Scale Computing, Inc., United States District Court for the District of Delaware, Case No. 22-535-GBW.

Provided deposition and trial testimony on behalf of defendant Scale Computing Systems.

Touchstream Technologies Inc., v. Comcast Cable Communications, LLC, d/b/a XFINITY, et al., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:23-cv-00062-JRG.

Provided deposition testimony on behalf of defendant Comcast Systems *et al.*

Koninklijke Philips N.V. and Philips North America LLC, v. Hewlett Packard, Inc., United States District Court for the District of Delaware, Case No. 2:20-cv-1241-CFC.

Provided deposition and hearing testimony on behalf of defendant Hewlett Packard, 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 and trial testimony on behalf of defendant Nokia Corporation.

Lionra Technologies Ltd., v. Cisco Systems, Inc., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:22-cv-00305-JRG-RSP.

Provided deposition testimony on behalf of defendant Cisco Systems.

In re: Inter Partes review of U.S. Patent No. 7,895,641.

Provided deposition testimony on behalf of Petitioner Palo Alto Networks, Inc.

Correct Transmission, LLC, v. Nokia Corporation, et al., United States District Court for the Eastern District of Texas, Marshall Division, Case No. 2:22-cv-0343-JRG-RSP.

Provided deposition testimony on behalf of defendant Nokia Corporation.

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.

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 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.

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 ACM/SIGCOMM Internet Measurement Conference Steering Committee, 2005–2009.

Committees 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 11th International Workshop on Quality-of-Service (Co-Chair), Monterey, CA, June 2003.

Program Chair 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* ACM SIGCOMM Internet Measurement Conference, San Diego, CA, October 2007.
- General Chair* 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.

Conference Finance Chair, IEEE Real-Time Systems Symposium, Miami, FL, December 2005.

Organizing Finance Chair, IEEE Real-Time Systems Symposium, Lisbon, Portugal, December 2004.

Committees 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, 14th IEEE Symposium on Real-Time Systems, Durham, NC, December 1993.

Refereed Publications

Technical Papers *Explicit Topology Management for Continental-Scale Airborne Networks*, B. Newton, J. Aikat, K. Jeffay, IEEE INFOCOM Workshop on Wireless Communications and Networking in Extreme Environments,

Atlanta, GA, May 2017, pages 72-77.

Scientific Training in the Era of Big Data: A New Pedagogy for Graduate Education, 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.

Geographic Routing in Extreme-Scale Highly-Dynamic Mobile Ad hoc Networks, 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.

Simulating Large-Scale Airborne Networks with ns-3, 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. Monroe, 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 Lectures *The Synthetic Traffic Generation Problem: The least sexy problem in computer networking*
University of Minnesota, Minneapolis, MN, January 2008.

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.