IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re *Inter Partes* Review of:)
U.S. Patent No. 8,230,101)
Issued: Jul. 24, 2012)
Application No.: 12/527,777)
Filing Date: August 19, 2009)

For: **Server Device for Media, Method for Controlling Server for Media, and Program**

DECLARATION OF MARK CROVELLA IN SUPPORT OF PETITION FOR *INTER PARTES* REVIEW OF U.S. PATENT NO. 8,230,101

Google
Exhibit 1003

TABLE OF CONTENTS

I. Introduction and Qualifications 1

II. Summary of Materials Reviewed and Considered 5

III. Understanding of Legal Standards 7

A. Claim Construction 7
B. Obviousness 8
C. Person of Ordinary Skill in the Art ("POSITA") 12

IV. The '101 Patent 13

A. Summary of the '101 Patent (Ex. 1001) 13
B. Overview of the File History of the '101 Patent 17
C. Claim Construction 18
1. "transfer" 18
2. Means-Plus-Function Terms 20

V. Identification of Grounds 21

VI. Ground 1: Claims 1-3, 5, and 7 are obvious under § 103 over Sloss in view of Lamkin 21

A. Summary of Sloss (Ex.1004) 21
B. Summary of Lamkin (Ex.1005) 23
C. Motivation to Combine 26
D. Independent Claim 1 31
[1.0] *A server device for media, the server device for media comprising:* 31
[1.1] *an internal storage device for storing digital contents, wherein the server device for media responds to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;* 35
[1.2.1] *a transfer control unit adapted to transfer and store part of held digital contents in the internal storage device to a network storage device,* 40
[1.2.2] *wherein the network storage device is connected to the network and is capable of storing data,* 43

[1.2.3] *and wherein said transfer control unit does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device;* 44

[1.3.1] *a list information transmission unit adapted to respond to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,*.................47

[1.3.2] *wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,*...51

[1.3.3] *and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;* 53

[1.4] *a search unit adapted to respond to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and*...56

[1.5] *a digital contents data transmission processing unit adapted to allow the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,*..57

[1.6] *wherein the server device for media is a media player.*.............60

E. Dependent Claim 2..61

[2.1] *The server device for media according to claim 1, wherein said digital contents data transmission processing unit causes the network storage device to transmit the corresponding data to the server device for media, and then transmits the corresponding data received from the network storage device from the server device for media to the network player.* 61

F. Dependent Claim 3..63

[3.1] *The server device for media according to claim 1, wherein said digital contents data transmission processing unit transmits the corresponding data and information for identifying the network storage device to the network player, and causes the network storage device to directly transmit the corresponding data to the network player*. 63

G. Dependent Claim 5..64

[5.1] *The server device for media according to claim 1, wherein said list information transmission unit makes the list information to be transmitted to the network player* 64

[5.2] *include information for identifying whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player*...65

H. Independent Claim 7 67

[7.0] *A method for controlling a server device for media which is equipped with an internal storage device for storing digital contents, the method comprising the steps of:*..............67

[7.1] *responding to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;* 67

[7.2.1] *transferring and storing part of held digital contents in the internal storage device to a network storage device,* 67

[7.2.2] *wherein the network storage device is connected to the network and is capable of storing data,*....................................67

[7.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device;* 67

[7.3.1] *responding to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,*..67

[7.3.2] *wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,* 68

[7.3.3] *and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;* 68

[7.4] *responding to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and* 68

[7.5] *allowing the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,* 68

[7.6] *wherein the server device for media is a media player.* 68

VII. Ground 2: Claim 4 is obvious under § 103 over Sloss in view of Lamkin, and further in view of Rou 68

A. Summary of Rou (Ex.1015) 68

B. Motivation to Combine 70

C. Dependent Claim 4 71

[4.1] *The server device for media according to claim 1, further comprising a return control unit adapted to cause the digital contents corresponding to a predetermined condition among the digital contents which have been transferred to the network storage device to be returned from the network storage device to the internal storage device.* 71

VIII. Ground 3: Claims 6, 8, 9, 11, and 12 are obvious under § 103 over Sloss in view of Lamkin, and in further view of Chamberlain 73

A. Summary of Chamberlain (Ex.1010) 73

B. Motivation to Combine 74

C. Independent Claim 6 77

Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are identical to Claims [1.0]-[1.2.2], [1.3.1]-[1.6] 77

[6.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;* 77
D. Dependent Claims 8, 9, and 11 78
Claims 8, 9, and 11 are identical to Claims 2, 3, and 5, respectively........... 78
E. Independent Claim 12 79
Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are identical to Claims [7.0]-[7.2.2], [7.3.1]-[7.6]... 79
[12.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;* 79

IX. Ground 4: Claim 10 is obvious under 35 U.S.C. § 103 over Sloss and Lamkin in further view of Chamberlain and Rou. 79

A. Claim 10 is identical to Claim 4 79

X. Ground 5: Claims 1-4 and 7 are obvious under § 103 over Roden in view of Van Hoff, Ito, and Rathbone 79

A. Summary of Roden (Ex.1006) 79
B. Summary of Ito (Ex.1008) 83
C. Summary of Van Hoff (Ex.1007)... 85
D. Summary of Rathbone (Ex.1009) 89
E. Motivation to Combine 90
F. Independent Claim 1 94
[1.0] *A server device for media, the server device for media comprising:* 94
[1.1] *an internal storage device for storing digital contents, wherein the server device for media responds to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;*.................................... 95
[1.2.1] *a transfer control unit adapted to transfer and store part of held digital contents in the internal storage device to a network storage device,* 99
[1.2.2] *wherein the network storage device is connected to the network and is capable of storing data,*................................. 101

[1.2.3] *and wherein said transfer control unit does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device;* 102
[1.3.1] *a list information transmission unit adapted to respond to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,*...............106
[1.3.2] *wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,*..110
[1.3.3] *and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;* 111
[1.4] *a search unit adapted to respond to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and*..115
[1.5] *a digital contents data transmission processing unit adapted to allow the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,*..116
[1.6] *wherein the server device for media is a media player.*...........117
G. Dependent Claim 2..118
[2.1] *The server device for media according to claim 1, wherein said digital contents data transmission processing unit causes the network storage device to transmit the corresponding data to the server device for media, and then transmits the corresponding data received from the network storage device from the server device for media to the network player.* 118
H. Dependent Claim 3..120

[3.1] *The server device for media according to claim 1, wherein said digital contents data transmission processing unit transmits the corresponding data and information for identifying the network storage device to the network player, and causes the network storage device to directly transmit the corresponding data to the network player.* 120

I. Dependent Claim 4 121

[4.1] *The server device for media according to claim 1, further comprising a return control unit adapted to cause the digital contents corresponding to a predetermined condition among the digital contents which have been transferred to the network storage device to be returned from the network storage device to the internal storage device.* 121

J. Independent Claim 7 122

[7.0] *A method for controlling a server device for media which is equipped with an internal storage device for storing digital contents, the method comprising the steps of:* 122

[7.1] *responding to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;* 122

[7.2.1] *transferring and storing part of held digital contents in the internal storage device to a network storage device,* 122

[7.2.2] *wherein the network storage device is connected to the network and is capable of storing data,* 123

[7.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device;* 123

[7.3.1] *responding to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,* 123

[7.3.2] *wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,* 123

[7.3.3] *and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;* 123

[7.4] *responding to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and* 123

[7.5] *allowing the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,* 123

[7.6] *wherein the server device for media is a media player.* 124

XI. Ground 6: Claim 5 is obvious under § 103 over Roden in view of Van Hoff, Ito, and Rathbone, and in further view of Lamkin 124

A. Motivation to Combine 124

B. Dependent Claim 5 128

[5.1] *The server device for media according to claim 1, wherein said list information transmission unit makes the list information to be transmitted to the network player* 128

[5.2] *include information for identifying whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player.* 128

XII. Ground 7: Claims 6, 8-10, and 12 are obvious under § 103 over Roden in view of Van Hoff, Ito, Rathbone, and in further view of Harris 131

A. Summary of Harris (Ex.1011) 131

B. Motivation to Combine 132

C. Independent Claim 6 134

Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are identical to Claims [1.0]-[1.2.2], [1.3.1]-[1.6] 134

[6.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;* 134
D. Dependent Claims 8-10 135
Claims 8, 9, 10 are identical to Claims 2, 3, 4, respectively 135
E. Independent Claim 12 135
Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are identical to Claims [7.0]-[7.2.2], [7.3.1]-[7.6] 135
[12.2.3] *and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;* 135

XIII. Ground 8: Claim 11 is obvious under 35 U.S.C. § 103 over Roden, Van Hoff, Ito, and Rathbone in further view of Harris and Lamkin 136

XIV. Conclusion 137

I. Introduction and Qualifications

1. I, Mark Crovella, have been retained as an expert witness on behalf of Google LLC ("Petitioner") to provide my opinion concerning the validity of U.S. Patent No. 8,230,101 (attached to the accompanying Petition as Ex. 1001 and henceforth referred to as "the '101 patent") in support of this Petition for *Inter Partes* Review.

2. I am being compensated for my time in connection with this IPR at my rate of $700 per hour. My compensation is not dependent in any way upon the outcome of this matter.

3. I am a Professor in the Department of Computer Science at Boston University. From 2013 to 2018 I served as Chair of the Department. Since 2020 I am also a Professor in the Faculty of Computing and Data Sciences, and I serve as the Chair of Academic Affairs there. A brief summary of my background is as follows.

4. I earned an undergraduate degree in Biology from Cornell University, which I received in 1982. I received a master's degree in Computer Science from the University of Buffalo in 1989. My master's degree project involved the development of code using the PVM system running on an Intel iPSC hypercube (parallel) computer.

5. From 1982 to 1984, I worked as a computer programmer for the State of Colorado. From 1984 to 1994, I was employed at Calspan Corporation, a research and development firm in Buffalo, NY, where I rose to the level of Senior Computer Scientist. My work at Calspan focused on development of experimental software and large-scale simulation software in support of contracts between Calspan and the US Department of Defense.

6. I received a Ph.D. in Computer Science from the University of Rochester in 1994. My Ph.D. research concerned the measurement and analysis of parallel programs – software in which work is distributed over multiple processors or multiple independent computers.

7. In 1994, I joined the faculty of Boston University, initially as an Assistant Professor of Computer Science. I was promoted to the rank of Associate Professor in 2000 and became a full Professor in 2006. I served as Chair of the Department from 2013 to 2018. Since 2020 I have served as Chair of Academic Affairs for the Faculty of Computing and Data Sciences. In 2024 I was recognized as a Duan Family Fellow at Boston University.

8. While at Boston University, much of my teaching has focused on parallel computing, distributed computing, the Web, and the Internet. I have taught the Department's course in Computer Networking for many years, a course in which we teach students how the communication protocols underlying

distributed systems work, and my students in that course learn to write distributed programs that run on multiple computers and communicate using protocols such as UDP, TCP, and via sockets.

9. I have been pursuing research in the area of computer networking since 1994. I have conducted research in a variety of areas related to the Internet and the World Wide Web. Among other areas, I have studied the efficient design of Web servers and content distribution systems; I have studied the statistical properties of Internet traffic; and I have made extensive measurements of Internet infrastructure and the behavior of Internet protocols. From 2007 to 2009, I was the Chair of ACM SIGCOMM (the Special Interest Group in Computer Communication), the main professional organization for scientists in the field of computer networking. On the strength of my work on parallel and distributed systems, I have been elected a Fellow of the IEEE and also of the ACM.

10. I began my research at Boston University in 1994, as the Web was beginning to be widely used. At the time, there were many important and unanswered questions about how factors such as Web protocol design, Web server design, Internet structure and topology, and user behavior affected the performance (speed and reliability) of the Web. These questions were my principal research focus during the period from 1994 to roughly 2002.

11. One of the topics that I studied in detail during that period concerned how Web browsers presented information to users, and how user requests for information affected the overall traffic on the Internet. During that period my group produced a number of highly cited papers exposing how the statistical properties of Web pages, and user requests for Web pages, affected the performance of the Web.

12. I am co-author of *Internet Measurement: Infrastructure, Traffic, and Applications* (Wiley Press, 2006), which is the first book written on the subject. I am the author of over two hundred research papers; according to Google Scholar, my work has been cited by over 30,000 other publications. I hold nine U.S. Patents derived from my Internet related research. I am a past editor of principal journals in the field of networking: *Computer Communication Review*, *IEEE/ACM Transactions on Networking*, *Computer Networks*, and *IEEE Transactions on Computers.*

13. I also have considerable experience working with Web protocols, distributed systems, and content delivery systems in industry. I co-founded a company that developed protocols for content delivery in the Web, which through a chain of acquisitions became part of Network Appliance, Inc. At Network Appliance I served for two years as a Technical Director working on content distribution and Web caching systems.

14. I have also served as Chief Science for Guavus, Inc., a company founded by one of my PhD students, that commercialized work that we did analyzing the properties of Internet traffic. Guavus was subsequently acquired and is now a division of Thales, Inc.

15. My detailed employment background, professional experience, and list of technical papers and books are contained in the attached Curriculum Vitae, which I understand is attached as Exhibit A.

16. I am very familiar with the subject matter of this case. I consider myself an expert in distributed systems and in computer networking, which includes network protocols, applications, and architecture.

II. Summary of Materials Reviewed and Considered

17. In preparing this Declaration, I have reviewed the '101 patent and considered the documents identified below in light of the general knowledge in the relevant art. In forming my opinions, I relied on my education, knowledge, and experience and considered the level of ordinary skill in the art as discussed below.

Ex.	Description
1001	U.S. Patent No. 8,230,101
1002	Prosecution History of U.S. Patent No. 8,230,101
1004	WO 2001/076192 to Sloss et al.
1005	U.S. 2006/0161635 to Lamkin et al.
1006	U.S. 2006/0101489 to Roden et al.
1007	U.S. 7,895,633 to Van Hoff et al.
1008	WO 2006/073040 to Ito et al.

1009	"TiVo for Dummies," Andy Rathbone, Wiley Publishing, Inc. (2004)
1010	U.S. 2002/0026563 to Chamberlain et al.
1011	U.S. 5,835,698 to Harris et al.
1012	Plaintiff's Fourth Amended Disclosure of Asserted Claims and Infringement Contentions, Appendix E-1 - Claim Chart for U.S. Patent No. 8,230,101 Against Google Cloud Content Delivery Network (CDN) and Google Products Utilizing CDN
1013	Plaintiff's Fourth Amended Disclosure of Asserted Claims and Infringement Contentions, Appendix E-2 - Claim Chart for U.S. Patent No. 8,230,101 Against YouTube and Google Products Utilizing YouTube
1014	Plaintiff's Fourth Amended Disclosure of Asserted Claims and Infringement Contentions, Appendix E-3 - Claim Chart for U.S. Patent No. 8,230,101 Against Google Home App and Compatible Google Products
1015	"Online File Storage System" by T.K. Rou, Proceedings of the 2002 Student Conference on Research and Development, IEEE (2002)
1017	Microsoft Computer Dictionary, 5th ed.
1018	"File Structures," Folk et al., 2nd ed.
1019	2005 Service Update Guide, TiVo Inc. (2005)
1020	"Torvalds on TiVo," Forbes, March 9, 2006, *available at* https://www.forbes.com/2006/03/09/torvalds-linux-licensing-cz_dl_0309torvalds2.html

18. My opinions are additionally guided by my appreciation of how a person of ordinary skill in the art ("POSITA") would have understood the claims of the '101 patent at the time of the invention, which I have been asked to assume is March 2, 2007 for all claims.

III. Understanding of Legal Standards

19. I am not an attorney, and I do not opine on matters of law. However, I have been informed by counsel of certain legal standards, which are set forth below. I have applied those standards in my analysis in this declaration. The material in this section has been supplied to me by counsel.

20. I have been informed by Counsel that the validity analysis is a two-step process. First, the patent claims are construed to ascertain their proper scope. Second, the construed claims are compared to the identified prior art to determine if the claims are valid over the prior art.

A. Claim Construction

21. I have been informed by Counsel that the claims of a patent define the limits of the patentee's exclusive rights. I have been informed by Counsel that to determine the scope of the claimed invention, courts typically construe claim terms, the meaning of which the parties may dispute. I have been informed by Counsel that claim terms should generally be given their ordinary and customary meaning as understood by one of ordinary skill in the art at the time of filing of the patent application, after reading the patent and its prosecution history. I also have been informed by Counsel that a basic tenet of claim construction presumes that different words in a claim have different meanings, unless there is evidence to the contrary.

22. I have been informed by Counsel that claims must be construed in light of, and consistent with, the intrinsic evidence. In this context, I have been informed by Counsel that intrinsic evidence includes the claims themselves, the written disclosure in the specification, and the patent's prosecution history, including prior art that was considered by the United States Patent and Trademark Office ("USPTO"). I have been informed by Counsel that the specification is always highly relevant to the claim construction analysis and often is the single best guide to the meaning of a disputed term. I have been informed by Counsel that extrinsic evidence may also be considered when construing claims and may include, for example, technical dictionaries, technical publications and books, treatises, and expert testimony.

B. Obviousness

23. I have been informed by Counsel that patent claims can be deemed invalid if the differences between the claimed subject matter and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the pertinent art. While such conclusions are often based on more than one piece of prior art, only one is required. It is not sufficient that a collection of prior art references merely recites the various elements of a challenged patent claim. Rather, as I have been informed by Counsel, the prior art must present the elements in a manner that is consistent

with their arrangement or use in the challenged claims. I have been informed by Counsel that obviousness cannot be based on the hindsight combination of components selectively culled from the prior art.

24. I have been informed by Counsel that a person having ordinary skill in the art (POSITA) is a hypothetical person who is presumed to have known the relevant art at the time the invention was made. I have been informed by Counsel that the condition "at the time the invention was made" is imposed to rule out impermissible hindsight. I also have been informed by Counsel that an expert is to analyze the prior art from the perspective of a person of ordinary skill in the art at the time the invention was made, and not simply to provide his/her own personal conclusions.

25. I also have been informed by Counsel that an obviousness determination includes several factual inquiries, including (1) determining the scope and content of the prior art; (2) ascertaining the differences between the claimed invention and the prior art; (3) resolving the level of ordinary skill in the pertinent art; and (4) taking into consideration any objective indicia of nonobviousness.

26. I have been informed by Counsel that obviousness must be determined as of the date of the invention (i.e., the effective filing date or priority date of the patent). Thus, in considering the reason or motivation to combine

references, it is essential to avoid using hindsight. For example, the problem examined when considering obviousness is the general problem that confronted the inventor before the invention was made, not the specific problem solved by the invention. Defining the problem in terms of its solution reveals improper hindsight in the selection of the prior art relevant to obviousness. Further, an overly narrow statement of the problem can represent a form of hindsight, because often the inventive contribution lies in defining the problem in a new revelatory way. Similarly, an assertion that a person of ordinary skill could combine the references, rather than that they would have been motivated to do so, is an impermissible form of hindsight. Moreover, knowledge of a problem and motivation to solve it are entirely different from motivation to combine particular references to reach the particular claimed invention.

27. I have been informed by Counsel that a motivation to solve a problem (or issue) is not the same as a motivation to have combined references. I have been informed by Counsel that whether a skilled artisan would be motivated to make a combination includes whether he/she would select particular references in order to combine their elements.

28. I have been informed by Counsel that the analysis of the motivation to combine should be explicit. Additionally, it must include articulated reasoning with rational underpinnings to support the conclusion of obviousness. Providing a

full explanation of the motivation to combine the references and the reasonable expectation of success is a necessary component of the obviousness inquiry, as inventions in most, if not all, instances rely upon building blocks long since uncovered, and claimed discoveries almost of necessity will be combinations of what, in some sense, is already known. However, conclusory statements fail to adequately explain why a person of ordinary skill would have a motivation to combine the potential prior art references.

29. I have been informed by Counsel that common sense, common wisdom, and common knowledge may be used to support a motivation to combine, so long as the use of common sense in the analysis is explained with sufficient reasoning. But, common sense can only be used to supply a limitation missing from the prior art where the technology is unusually simple and straightforward.

30. I also have been informed by Counsel that evidence suggesting reasons to combine cannot be viewed apart from evidence suggesting reasons not to combine. In analyzing motivation, both advantages and disadvantages must be considered. Teaching away, for example, is a statement in the prior art that either (i) discourages one from following the path of the claimed invention, or (ii) encourages one to follow a path that diverges from the path of the claimed invention. The concept of teaching away bears directly on whether there was a reason to combine prior art. Known disadvantages in prior art technology or

devices which would naturally discourage a search for new inventions may also be taken into account in determining obviousness.

C. Person of Ordinary Skill in the Art ("POSITA")

31. I understand that a person of ordinary skill in the art ("POSITA") is a hypothetical person who is presumed to be aware of all pertinent art, possesses conventional wisdom in the art, is a person of ordinary creativity, and has common sense. I understand that this hypothetical person is considered to have the normal skills and knowledge of a person in a certain technical field (including knowledge of known problems and desired features in the field).

32. I have been asked to focus my analysis on the '101 patent, and prior art relating thereto, from the perspective of such a person at the time of the alleged invention. I understand that the '101 patent issued from U.S. Application 12/527,777 ("'777 application"), which was a U.S. National Stage application filed under 35 U.S.C. §371 on August 19, 2009. The '101 patent claims priority to PCT Application PCT/JP2007/054603 ("'603 application"), which was filed March 2, 2007.

33. It is my opinion that a person of ordinary skill in the art in March 2007 would have been someone knowledgeable and familiar with the field of networked media devices that is pertinent to the '101 patent. That person would have a bachelor's degree in electrical engineering, computer engineering, computer

science, or equivalent training, and two years of experience working in the field of media device networking. Lack of work experience can be remedied by additional education, and vice versa.

34. As of March 2, 2007, I would have qualified as at least a POSITA, and my opinions herein are informed by my own knowledge based on my personal experiences and observing others of various skill levels (including those above and below the level of a POSITA).

35. My opinions below are not restricted to the precise definition of a POSITA above. The claims of the ’101 patent are directed to media device networking concepts that were well-known in the art and taught by numerous prior art references, including the references discussed below. Thus, my opinions below would apply under any reasonable definition of a POSA.

IV. The ’101 Patent

A. Summary of the ’101 Patent (Ex. 1001)

36. The ’101 patent is titled “Server Device for Media, Method for Controlling Server for Media, and Program.” The ’101 patent purports to “provide a server device for media … capable of maintaining the convenience of playback in a network player, while properly dealing with the large total size of held digital content.” Ex.1001 [’101 patent], 2:6-11.

37. In pertinent part, the ’101 patent intends to allow for digital content (e.g., music, photographs, and movies) to be shared and played across different network players, with digital content stored across internal storage (i.e., storage within a “server device for media” such as Figure 1’s “HDD portable player 10”) and network storage (e.g., a “network HDD 35” which stores content, for example, when internal storage gets full). One example of the ’101 patent’s system is shown in Figure 3’s embodiment, which includes a network 26 coupling a HDD portable player 10, network player 27, and network HDD 35, among other components. Ex.1001 [’101 patent], 4:13-25. The network 26 “is a DLNA network and built in the home 21” but the claims of the ’101 patent are not limited to “DLNA” or any specific form of network. The ’101 patent explains the “HDD portable player 10 functions as a server for delivering music pieces, pictures and other contents.” Ex.1001 [’101 patent], 4:13-25.

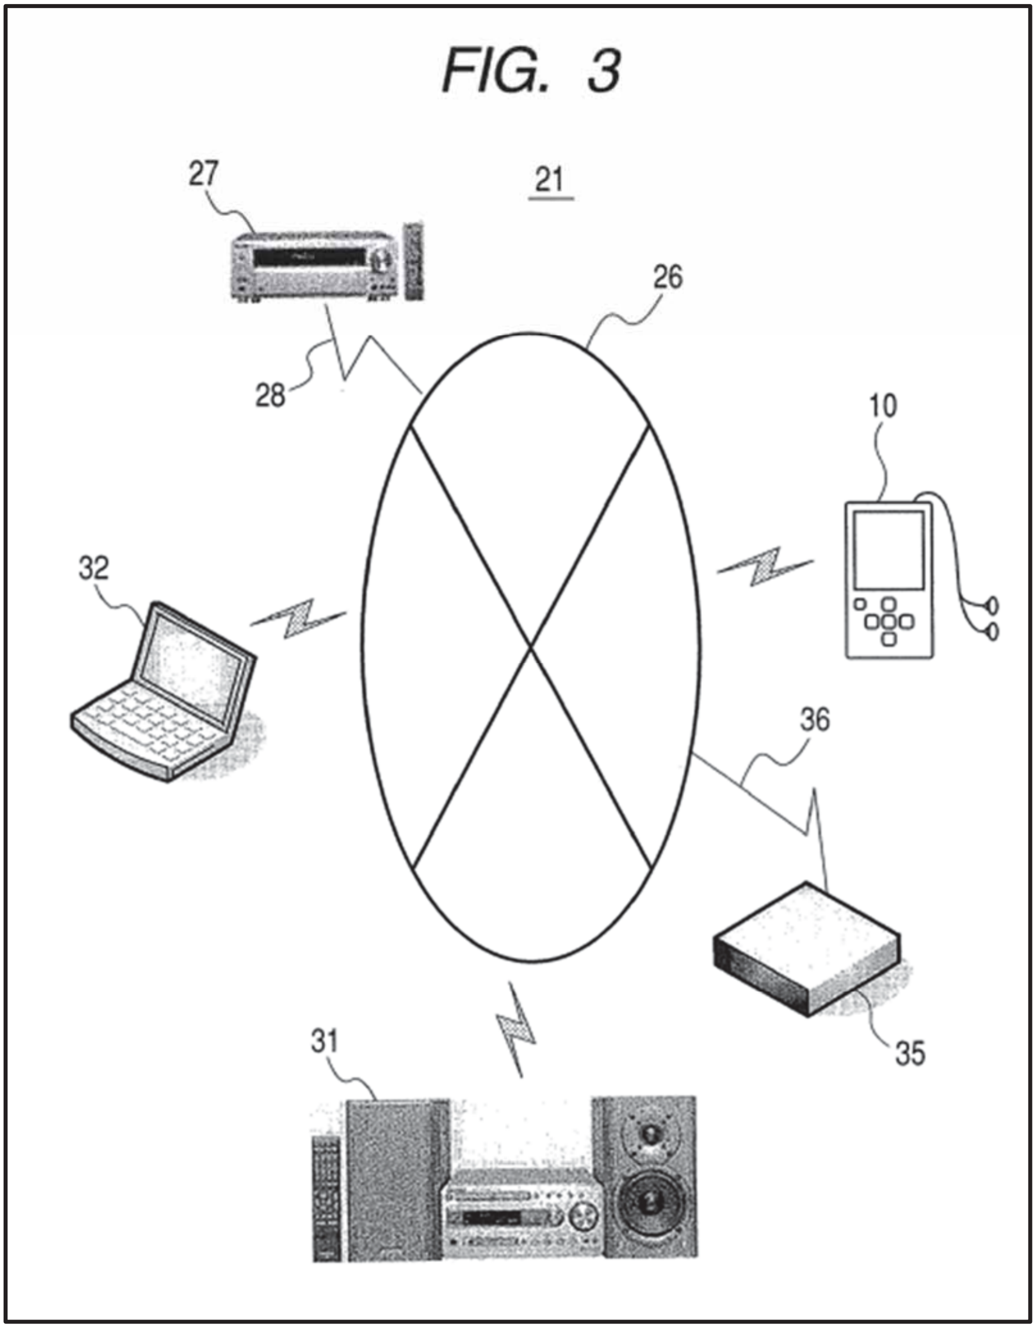

Ex.1001 [’101 patent], FIG. 3

38. In one example, the ’101 patent describes “transferring and storing part of held digital contents in the internal storage device to a network storage device.” Ex.1001 [’101 patent], 2:23-26. In other words, the ’101 patent describes

the well-known technique of offloading (e.g., moving or copying) content from a device with a comparatively small amount of storage (e.g., HDD portable player 10) to a device with a comparatively larger amount of storage (e.g., network HDD 35) but still making that offloaded content accessible. This includes means for providing a list to a network player that identifies digital content that is left in the internal storage device of the server device, and digital content that has been transferred from internal storage device to the network storage device. Ex.1001 [’101 patent], 2:27-35. The ’101 patent also discloses responding to a request to identify where the digital content is stored by “searching for a location where the held digital contents are currently stored,” (Ex.1001 [’101 patent], 2:36-40), and that corresponding data can be “stream-delivered from the network storage device to the network player, if the result of the search shows the [digital contents are stored on the] network storage device.” Ex.1001 [’101 patent], 2:41-44.

39. The ’101 patent also purports to address a problem that, due to “copyright protection, some digital contents are adapted not to permit the overlapped part of the same content exists in a plurality of storage devices for more than several seconds in the playback time.” Ex.1001 [’101 patent], 9:26-29. According to the ’101 patent, as a result of this copyright limitation, if the “transferring operation of such a digital content from the internal storage device 51 to the network storage device 57 is interrupted by a failure in the network 55 or the

like, the original digital content cannot be recovered neither in the internal storage device 51 nor the network storage device 57, which causes a great damage to the user." Ex.1001 ['101 patent], 9:26-35. The '101 patent attempts to provide a solution in this case, where, if digital content cannot be recovered if a network failure occurs during transfer, that digital content will either (1) not be transferred, or (2) will be transferred after obtaining permission from the user. '101 patent 9:36-46.

40. Analyzed according to its priority date of March 2, 2007, the features recited in the challenged claims were all well-known, as shown below, by Sloss, Roden, and other references and would have been obvious to a POSITA. Thus, the challenged claims should not have been allowed.

B. Overview of the File History of the '101 Patent

41. The '101 patent issued from U.S. Application 12/527,777 ("'777 application"), which was a U.S. National Stage application filed under 35 U.S.C. §371 on August 19, 2009, which claims priority to PCT Application PCT/JP2007/054603 ("'603 application"), filed March 2, 2007. The '101 patent's priority date is March 2, 2007.

42. During prosecution, the claims were rejected as obvious multiple times (*see* Ex.1002, 432-451, 487-506, 522-538) over combinations of references not at issue here. After two office actions and responses (*see id.* 515-518, 547-554),

the examiner conducted an interview with the applicant, and proposed "merging subject matter in the preamble of the independent claims into the body of the claims" and other amendments, which ultimately led to allowance. *Id.*, 567-575.

43. However, as illustrated in this declaration, the challenged claims' limitations were already known in the art by the '101 patent's priority date.

C. Claim Construction

44. Aside from the terms construed below, in my analysis, I have applied the plain and ordinary meanings of the claim terms of the '101 patent as a POSITA would have understood them in the context of the patent at the time of its priority date.

45. The opinions that I provide in this declaration would remain true under any reasonable construction of the claim terms in the '101 patent.

1. "transfer"

46. I understand that the claim term "transfer," as used in the limitation "transfer control unit does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device" ('101 patent, Claim 1) is not defined within the '101 patent. In my opinion, a POSITA would have recognized this term

to include two potential meanings consistent with the claim language and specification.

47. First, a POSITA would understand that "transfer" could mean "copying" or "caching." In this case, a POSITA would have recognized that this interpretation would be consistent with common usage in the art; for example, a user may transfer content from a first device to a second device to make that content available on both devices, in which case, the content is copied from the first device to the second device, but not necessarily deleted from the first device. I understand this interpretation would be consistent with Patent Owner's interpretation of the claim language in its district court filings. For example, I have reviewed Patent Owner's infringement allegations relative to the '101 patent, in which Patent Owner states: "While delivering the data from cloud server to various APIs, the data is only cached to the cloud CDN storage if it is cacheable (i.e., wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device[])." *See* Ex.1012 [E-1 chart], 10-11, Ex.1013 [E-2 chart], 11-12 (similar). That is, Patent Owner takes the interpretation that caching data by making it available on one or more different devices, which includes copying that data to the one or more different devices, fits within its interpretation of "transfer." A POSITA would have recognized that a cache is a

memory storing copies of data. *See* Ex.1017 [Microsoft Computer Dictionary 5th ed.], p. 81 ("cache *n*. A special memory subsystem in which frequently used data values are **duplicated** for quick access.") (emphasis added). Thus, a POSITA would have recognized that, under Patent Owner's interpretations, a "transfer" of content would include "copying" that content.

48. Second, a POSITA would understand that "transfer" could also mean "moving data from one location to another," indicating just one instance of the data, being moved from one place to another place, with no "copying" involved. This interpretation is consistent with technical dictionary definitions, for example, the Microsoft Computer Dictionary 5th edition referenced above defines the noun "transfer" as "movement of data from one location to another" and the verb "transfer" as "[t]o move data from one place to another." Ex.1017 [Microsoft Computer Dictionary 5th ed.], p. 527.

49. As stated above, it is unnecessary to construe this term as the '101 patent claims are unpatentable under either construction of the word "transfer."

2. Means-Plus-Function Terms

50. I have also been informed by Counsel that Claims 1 and 6 recite various "unit" terms that are means-plus-function terms because they do not "'recite sufficiently definite structure' or else recite[] 'function without reciting sufficient structure for performing that function.'".

51. I have also been informed by Counsel that the Petitioner has identified these terms as indefinite means-plus-function terms in the parties' co-pending litigation, and, that Petitioner has identified function and corresponding structure in its Petition.

V. Identification of Grounds

Ground	Statutory Basis	Claims	Prior Art
1	§103	1-3, 5, 7	Sloss, Lamkin
2	§103	4	Sloss, Lamkin, Rou
3	§103	6, 8, 9, 11, 12	Sloss, Lamkin, Chamberlain
4	§103	10	Sloss, Lamkin, Chamberlain, Rou
5	§103	1-4, 7	Roden, Van Hoff, Ito, Rathbone
6	§103	5	Roden, Van Hoff, Ito, Rathbone, Lamkin
7	§103	6, 8-10, 12	Roden, Van Hoff, Ito, Rathbone, Harris
8	§103	11	Roden, Van Hoff, Ito, Rathbone, Harris, Lamkin

VI. Ground 1: Claims 1-3, 5, and 7 are obvious under § 103 over Sloss in view of Lamkin

A. Summary of Sloss (Ex.1004)

52. The reference that I refer to as "Sloss" is PCT Application Publication WO2001076192A2. It was filed on March 14, 2001 and published on October 11, 2001. Ex.1004 [Sloss], Cover. Reed J. Sloss is the first-named inventor. Sloss

Cover. On its face, it is assigned to Intel Corporation. Ex.1004 [Sloss], Cover. I understand that Sloss is prior art under § 102(b).

53. Sloss is titled "Distributed Edge Network Architecture." Ex.1004 [Sloss], Cover. Sloss teaches an "invention related generally to the field of network services" and "an improved architecture for network data distribution." Ex.1004 [Sloss], 1:5-7. To summarize, Sloss teaches a multi-level caching technique that aims to provide requested data quickly upon a user's request, by locating that data closer to the end user using intermediate and edge caches. Thus, Sloss's techniques use a "multi-tiered networking architecture" which includes "one or more data centers 220-222" where each data center is comprised of "groups of network servers on which various types of network content may be stored and transmitted to end users 250, including … live & on-demand multimedia streaming files." Ex.1004 [Sloss], 5:3-5, 11-14. In one example of Sloss, requested content is transmitted "directly from" a source (e.g., a data center) to an end user or client. *See* Ex.1004 [Sloss], 1:21-2:3. In Sloss's distribution technique, after content is uploaded to a server at a data center 220-222, it may be "automatically distributed from the data center 220-222 to one or more of the intermediate POPs 230-234, and/or edge POPs 240-245 …". Ex.1004 [Sloss], 8:13-16. Sloss explains that not all data is distributed, as it teaches that only cacheable data will be transferred to

the local proxy cache, whereas non-cacheable data is not transmitted from the data centers to the POPs. *See* Ex.1004 [Sloss], 1:18-2:3.

B. Summary of Lamkin (Ex.1005)

54. The reference that I refer to as "Lamkin" is U.S. Publication No. 2006/0161635. It was filed December 16, 2005 and published on July 20, 2006. Ex.1005 [Lamkin] Cover. Allan B. Lamkin is the first-named inventor. Lamkin Cover. On its face, it is assigned to Sonic Solutions. Lamkin Cover. I understand that Lamkin is prior art under § 102(a).

55. Lamkin is titled "Methods and System for Use in Network Management of Content." Ex.1005 [Lamkin], Cover. Lamkin teaches "methods, systems and apparatuses for use in managing content" on a network. Ex.1005 [Lamkin], Abstract. Lamkin indicates that its methods and systems manage content "on at least a local network" (Ex.1005 [Lamkin], Abstract) but also that its teachings are applicable to "a wide area network (WAN) and/or other distributed networks." Ex.1005 [Lamkin], ¶161. In one aspect of Lamkin's content management teachings, Lamkin discloses that "a user 224 at the client device 124 is provided a user interface that accesses and/or shows content accessible, for example, through a CDS [content directory service] from various devices and/or local content, and allows the user to select the content to be pulled to the client device to be accessed and/or recorded." Ex.1005 [Lamkin], ¶56. Lamkin's user

interface also includes a listing of content available on various network-connected devices: "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87.

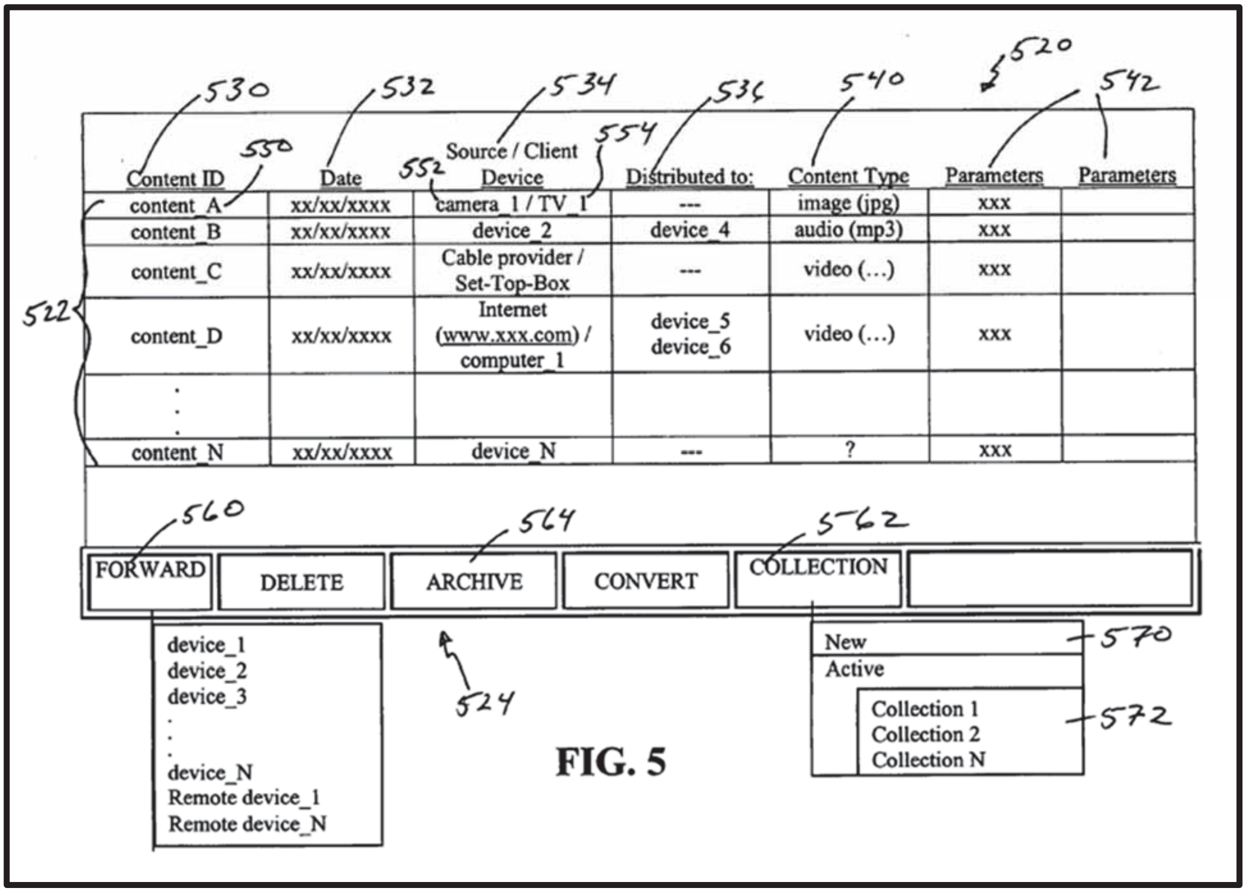

Ex.1005 [Lamkin], FIG. 5.

56. Lamkin teaches a "media in-box user interface" that "can identify content and in some implementations present the content according to an organized structure, similar to a file structure" like a hierarchical tree structure. Ex.1005 [Lamkin], ¶95, FIG. 6.

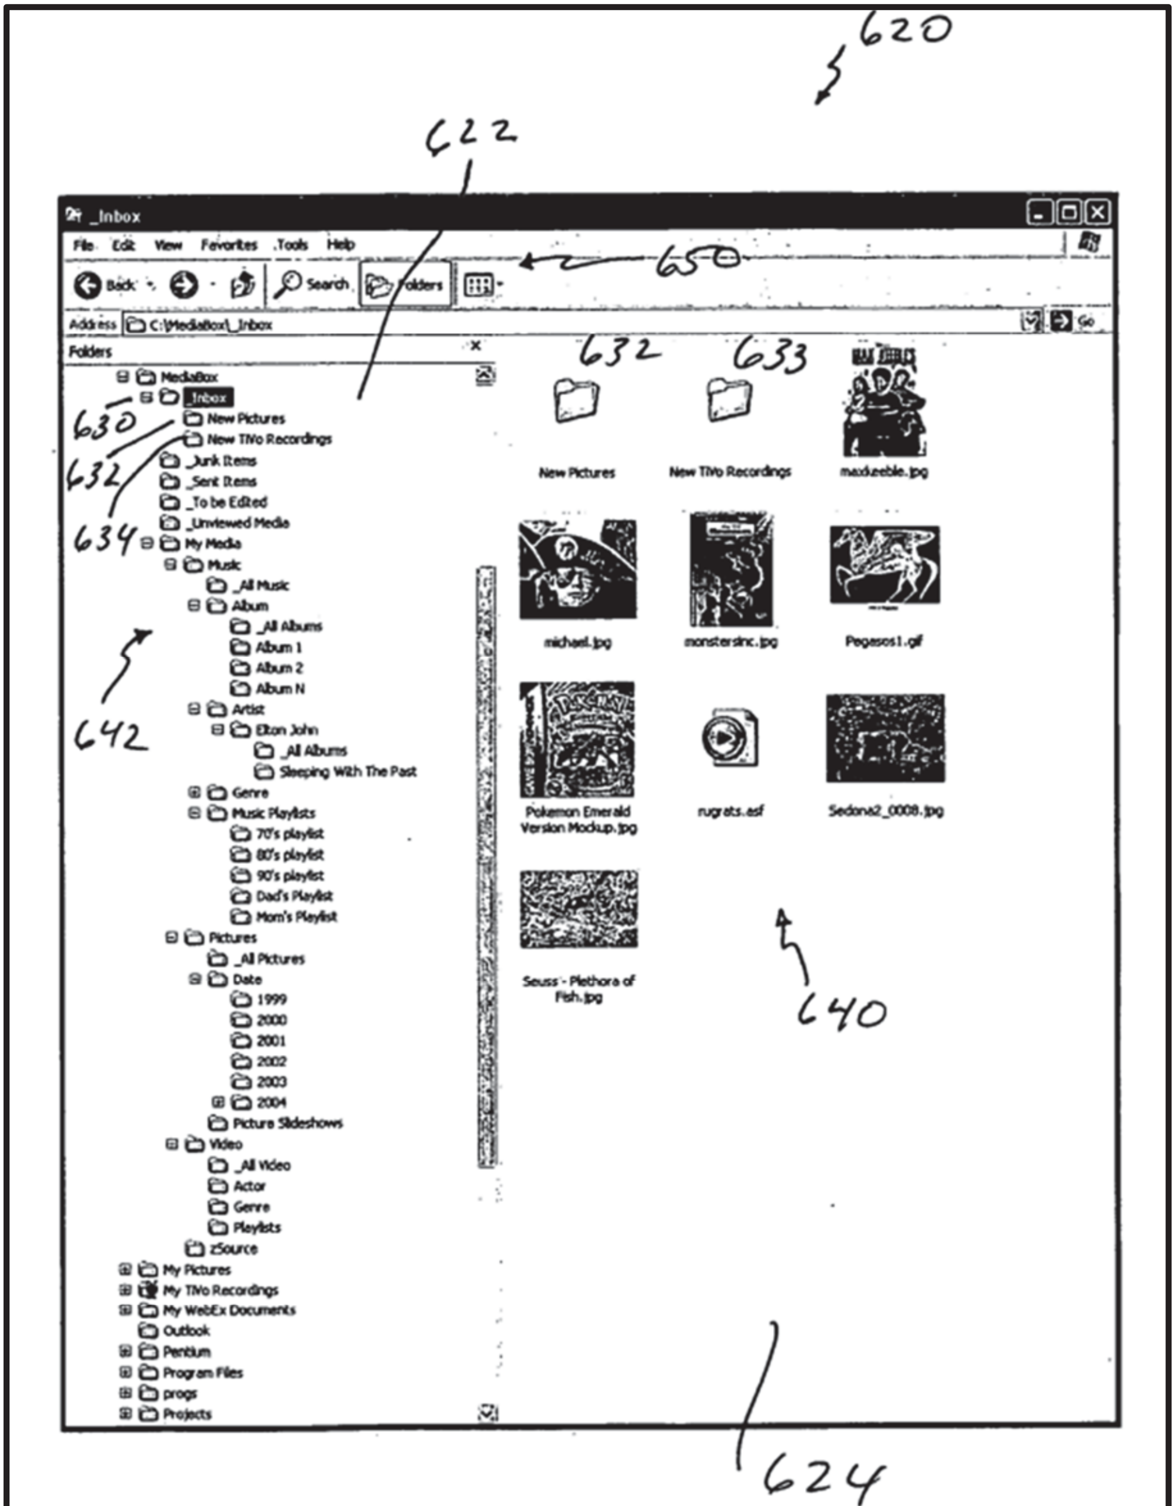

FIG. 6

Ex.1005 [Lamkin], FIG. 6

C. Motivation to Combine

57. In my opinion, a POSITA would have been motivated to combine the teachings of Sloss with those of Lamkin. It would have been obvious, beneficial, and predictable to apply Lamkin's teaching of a list showing where digital contents are stored, including whether digital contents are stored in internal storage or network storage, to Sloss's disclosures of network data distribution. This combination would achieve the benefits disclosed by Lamkin, of a "simplified … user interface" in its methods of "managing content" over a "distribution network," as disclosed by Sloss and Lamkin. Ex.1005 [Lamkin], Abstract.

58. A POSITA, when considering the teachings of Sloss, would have also considered the teachings of Lamkin. Both Sloss and Lamkin relate to managing digital contents over a network. Ex.1004 [Sloss], Abstract; Ex.1005 [Lamkin], Abstract. As such, it is my opinion that both Sloss and Lamkin are analogous art to the '101 patent, as both are within the field of endeavor of the '101 patent. *See*, *e.g.*, Ex.1001, 1:7-12.

59. Sloss discloses "one or more data centers 220-222" where each data center is comprised of "groups of network servers on which various types of network content may be stored and transmitted to end users 250, including … live & on-demand multimedia streaming files." Ex.1004 [Sloss], 5:3-5, 11-14. Sloss indicates that when files are "copied to or deleted" from various nodes in the

network, its "content distribution subsystem creates or removes a 'FileLocation' database record in the central content database 530." Ex.1004 [Sloss], 17:1-5. Sloss also states that the "database 530 in one embodiment keeps track of exactly where content has been distributed throughout the system." Ex.1004 [Sloss], 21:11-14. However, Sloss does not depict the structure of the database, or explicitly disclose that the contents' locations are specified in provided in a "list" or that such a list can be displayed.

60. Showing the contents of a storage device in the computing context, including in a list, however, would have been well-known to a POSITA at the time of the '101 patent. For example, the Microsoft Computer Dictionary provides a definition of the well-known "dir" command in MS-DOS (and Windows) which "instructs a computer to display **a list of files** and subdirectories in the current directory or folder." Ex.1017 [Microsoft Computer Dictionary], p. 161. Unix-based operating systems contained a similar command known as "ls" which similarly provided "**a list of files**." Ex.1017 [Microsoft Computer Dictionary], p. 320. And, as the Microsoft Computer Dictionary confirms, these commands could be used for networked locations; for example, the Microsoft Computer Dictionary indicates "ls" can be used on "FTP sites." Ex.1017 [Microsoft Computer Dictionary], p. 320. Consistent with this knowledge in the art, Lamkin discloses "a user interface that accesses and/or shows content accessible, for example,

through a CDS from various devices and/or local content." Ex.1005 [Lamkin], ¶56. This is also consistent with Sloss's teachings of the content database, and provides a well-known structure and depiction of how certain fields within Sloss's database may appear. Further, Lamkin explicitly states that a "database, listing, or other methods can be employed to locate desired content." Ex.1005 [Lamkin], ¶54.

61. A POSITA would have recognized at the time of the '101 patent that lists were a routine way to keep track of available digital contents that could be provided to end users. For example, it was well known in the field that playlists were a routine way to keep track of media to present to end users in an organized manner. For example, Lamkin explicitly teaches "a simplified example of a user interface 520 according to some embodiments that identifies content added and/or altered since a previous network user access, within a time period, content not previously distributed, and other criteria or combinations of criteria. The user interface 520 can include a listing of content 522." Ex.1005 [Lamkin], ¶87. Lamkin further discloses that the "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87. And, Lamkin explicitly discloses a flexible user interface that includes a listing of content that is available on various network-connected devices. Lamkin also provides an explicit motivation for doing providing such a list, as the list is part of its "simplified example of a user interface"

and Lamkin characterizes the invention as "methods, systems and apparatuses for use in managing content" on a network, which would have motivated a POSITA to use its list teachings to manage the digital contents available to a user to provide an easy method of managing such content in other networks and systems, like the system of Sloss. Ex.1005 [Lamkin], Abstract.

62. A POSITA would have recognized additional benefits of utilizing a list to manage digital contents, as Lamkin teaches, within a networked media system (e.g., in the context of Sloss's teachings). For example, in Lamkin's background section, it states that "[t]here has been a drastic increase in the number of consumer electronic devices capable of communicating with one or more computers or other consumer electronic devices." Ex.1005 [Lamkin], ¶4. As such, a consumer will need a way to "manage [that] media content over [said] network." Ex.1005 [Lamkin], ¶2. Likewise, a POSITA would have recognized that nearly all operating systems at the time of the invention kept records of files accessible within the computer on which the operating system ran, as well as records for files accessible over a network from the computer. And, such records were kept in a hierarchical file system in which a tree structure was used to represent the location of such files. This was very common even at a time much earlier than the invention. A hierarchical file system where directories could be nested within other directories, creating a tree-like structure was first introduced as part of the Multics

operating system in the 1960s, several decades before the priority date. Since that time, this organization system has become widely adopted, so that most modern computers use hierarchical file systems in a tree structure to organize and present the user with a visual representation of where files were stored. *See also* Ex.1017 [Microsoft Computer Dictionary], p. 252 ("hierarchical file system *n.* A system for organizing files on a disk in which files are contained in directories or folders, each of which can contain other directories as well as files."). Thus, Lamkin's teachings would have been well-known to a POSITA by the time of the '101 patent.

63. A POSITA would have therefore recognized that combining the teachings of Sloss and Lamkin would have been nothing more than the combination of known prior art elements (e.g., Sloss's database indicating where content is stored with Lamkin's teachings of the structure and display of such a database and a tree structure) according to known methods (e.g., configuring Sloss's database with the fields shown in Lamkin). Further, a POSITA would have found it obvious to apply Lamkin's teachings to Sloss' disclosure of streaming media over a network because there was only a finite number of ways to keep track of and display where media is stored in a networked media distribution system. It would have been obvious to a POSITA to use a list and tree organizational structure to inform the user where each piece of media was stored, in internal storage or

network storage. In addition, the combination would have been obvious because it represents known potential options with a reasonable expectation of success.

64. A POSITA would have also had a reasonable expectation of success in combining these teachings. Specifically, given the close overlap in subject matter between Sloss and Lamkin, a POSITA would have expected success implementing Lamkin's list feature of managing digital contents in Sloss, which likewise deals with distributing digital content across a networked media system. Further, contemporaneous research confirms that it was possible, predictable, and well known to provide a list of digital contents to client devices in a streaming media context. Indeed, before the '101 patent's earliest priority date, Netflix, which was a typical content provider consistent with Sloss's disclosure, provided lists of available digital content to end users (for example, in a list of available movies for an end-user to stream on their device).

65. Accordingly, it is my opinion that a POSITA would have been motivated to combine the teachings of Sloss and Lamkin.

D. Independent Claim 1

[1.0] ***A server device for media, the server device for media comprising:***

66. Sloss discloses or renders obvious the preamble of claim 1. Specifically, Sloss teaches: "[a] distributed edge network architecture is described in which a **data center** serves as a primary repository for content uploaded by

content providers. From the data center, the content is replicated at all, or a selected group of, geographically dispersed 'intermediate' point of presence ('POP') sites. A plurality of edge POP sites communicate with the intermediate POP sites and serve as network caches, storing content as it is requested by end users." Sloss, Abstract. Sloss also discloses an embodiment "of the system and method for distributing and streaming multimedia files." Ex.1004 [Sloss], 24:4-26:20.

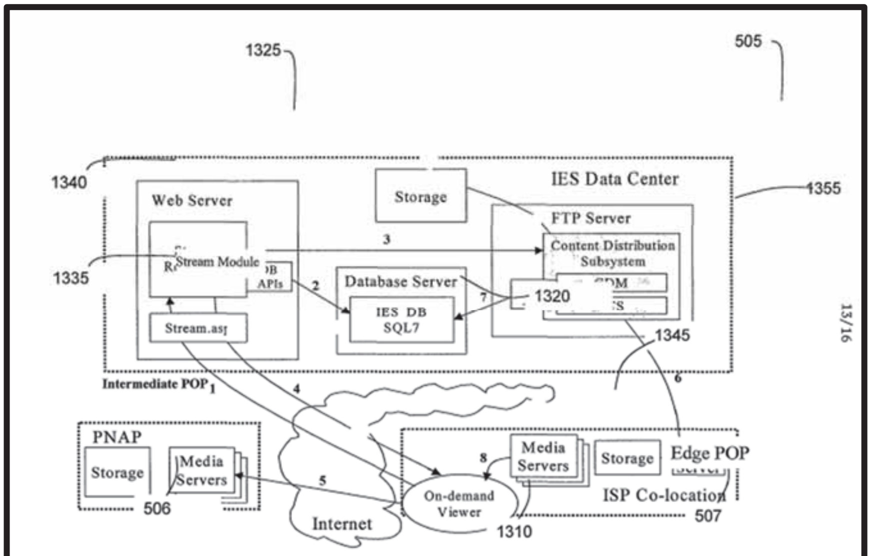

FIG. 13

Ex.1004 [Sloss], FIG. 13

67. Figure 2 of Sloss teaches a "multi-tiered networking architecture" which includes "one or more data centers 220-222" and "edge POPs 240-245" where each data center and edge POP is comprised of "groups of **network servers** on which various types of network content may be stored and transmitted to end users 250, including … live & on-demand multimedia streaming files." Ex.1004

[Sloss], 5:3-6:15.[1] The set of servers (data centers 220-222 and edge POPs 240-245) are mapped in this analysis to the claimed "*server device for media*."

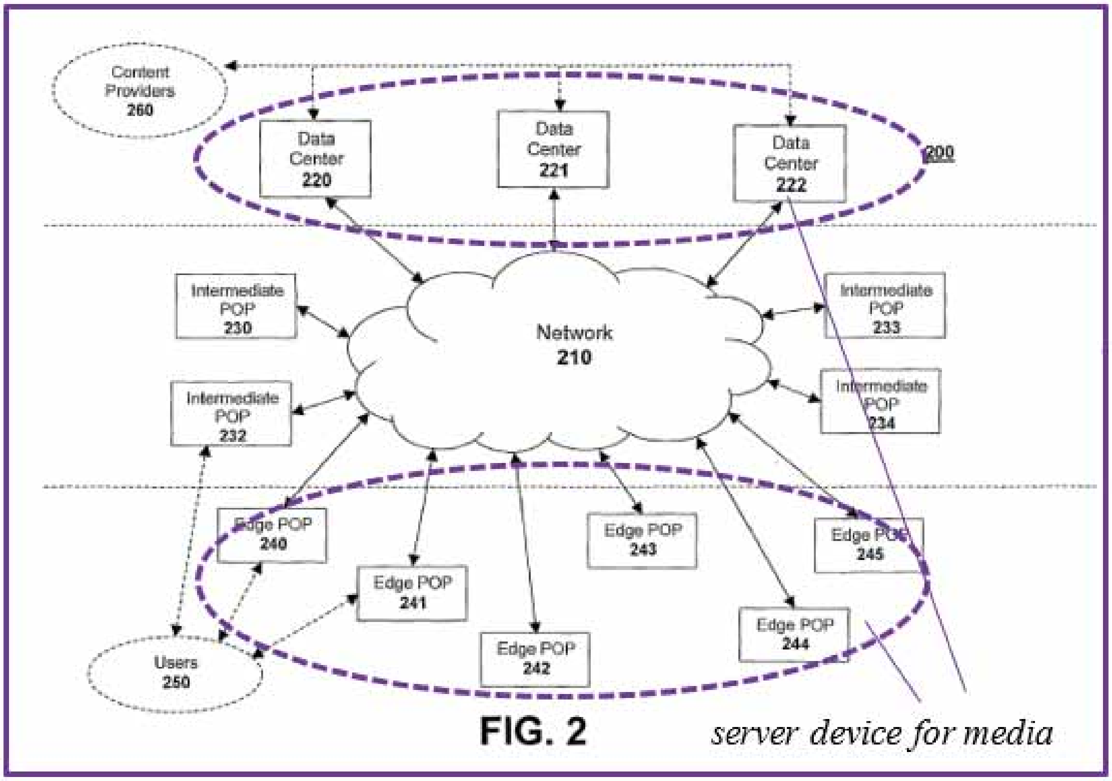

Ex.1004 [Sloss], FIG. 2 (annotated)

68. Sloss also explains, by way of background, a "traditional network caching system" in Figure 1, including "remote Internet server 120." Ex.1004 [Sloss], 1:9-2:17. A POSITA would have understood that Sloss's description of techniques in the traditional network caching system's remote Internet server 120 would also apply to the set of servers at data centers 220-222, as those servers provide the same function (e.g., both sets of servers are the origin server at which

---

[1] Emphasis in quoted material is added throughout.

content originates or is first stored). Thus, in the analysis below, I have applied the teachings of Sloss's remote Internet server 120 to inform my understanding of the functions of Sloss's data centers 220-222.

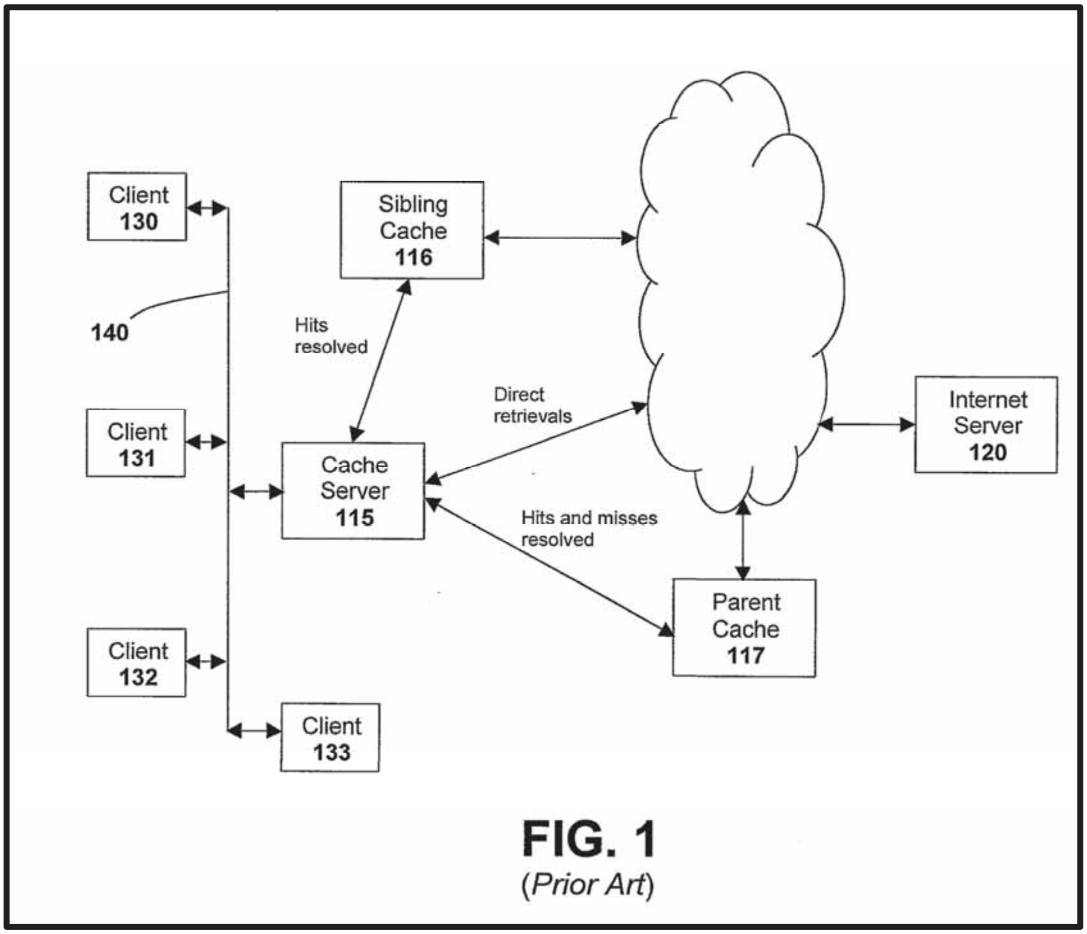

Ex.1004 [Sloss], FIG. 1

69. Thus, Sloss's set of servers at data centers 220-222 and edge POPs 240-245 disclose or render obvious a "*server device for media*" as recited.

[1.1] ***an internal storage device for storing digital contents, wherein the server device for media responds to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;***

70. In my opinion, Sloss discloses or renders obvious this limitation.

71. First, Sloss teaches "*an internal storage device for storing digital contents*." In Sloss, "data centers 220-222 serve as the primary initial repositories for network content." Ex.1004 [Sloss], 8:9-9:6. Sloss also teaches that "data centers 220-222 must be capable of **storing** and transmitting vast amounts of **content provider 260 data**" and are thus "**equipped with disk arrays** capable of storing hundreds of terabytes of data." Ex.1004 [Sloss], 8:19-21.

72. Additionally, Sloss teaches that the data center is "comprised of a plurality of servers for storing digital content." Ex.1004 [Sloss], 38:4 (Claim 1). Sloss also describes "a computer system 300 representing exemplary … servers for implementing elements of the present invention" which includes "**data storage device 327** such as a magnetic disk … **for storing information**." Ex.1004 [Sloss], 6:17-7:21. Further, Sloss teaches that edge POPs have "storage capacity." Ex.1004 [Sloss], 22:12-17. Thus, Sloss teaches "*an internal storage device for storing digital contents*."

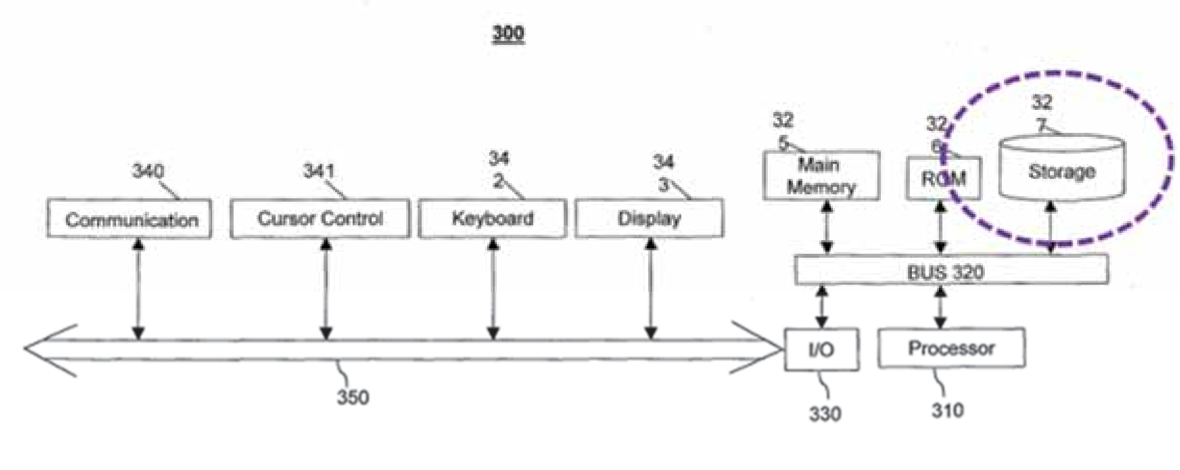

FIG. 3

Ex.1004 [Sloss], FIG. 3 (annotated)

73. Second, Sloss teaches that the set of servers at the data centers and the edge POPs "*respond[] to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network.*" For example, in describing the prior art "network caching system," Sloss teaches that clients 130-133 communicate "over a local area network and/or a larger network 110 (e.g., the Internet)" and "run a browser application." Ex.1004 [Sloss], 1:9-2:17. The teachings of clients 130-133 also apply to Sloss's "end users 250" in Figure 2, which receive "Web pages, network news data, e-mail data, File Transfer Protocol ('FTP') files, and live & on-demand multimedia streaming files." Ex.1004 [Sloss], 5:9-14. Thus, a POSITA would have recognized that the "end users 250"

are clients which execute Web browsers to request and receive information. The clients 130-133 and end users 250 therefore correspond to the recited "*network player*."

74. Sloss further explains that "[t]he browser on each client 130-133 may be configured so that all requests for information (e.g., Web pages) are transmitted through a local cache server 115, commonly referred to as a 'proxy cache.'" Ex.1004 [Sloss], 1:16-18. Then, "[w]hen a client 130 requests information from a remote Internet server 120," (which teachings are applicable, e.g., to servers at data center 220-222), "the local proxy cache 115 examines the request and initially determines whether the requested content is 'cacheable.'" Ex.1004 [Sloss], 1:18-21. "If the local proxy cache 115 detects a non-cacheable request, it forwards the request directly to the content source (e. g., Internet server 120)." Ex.1004 [Sloss], 1:21-2:3.

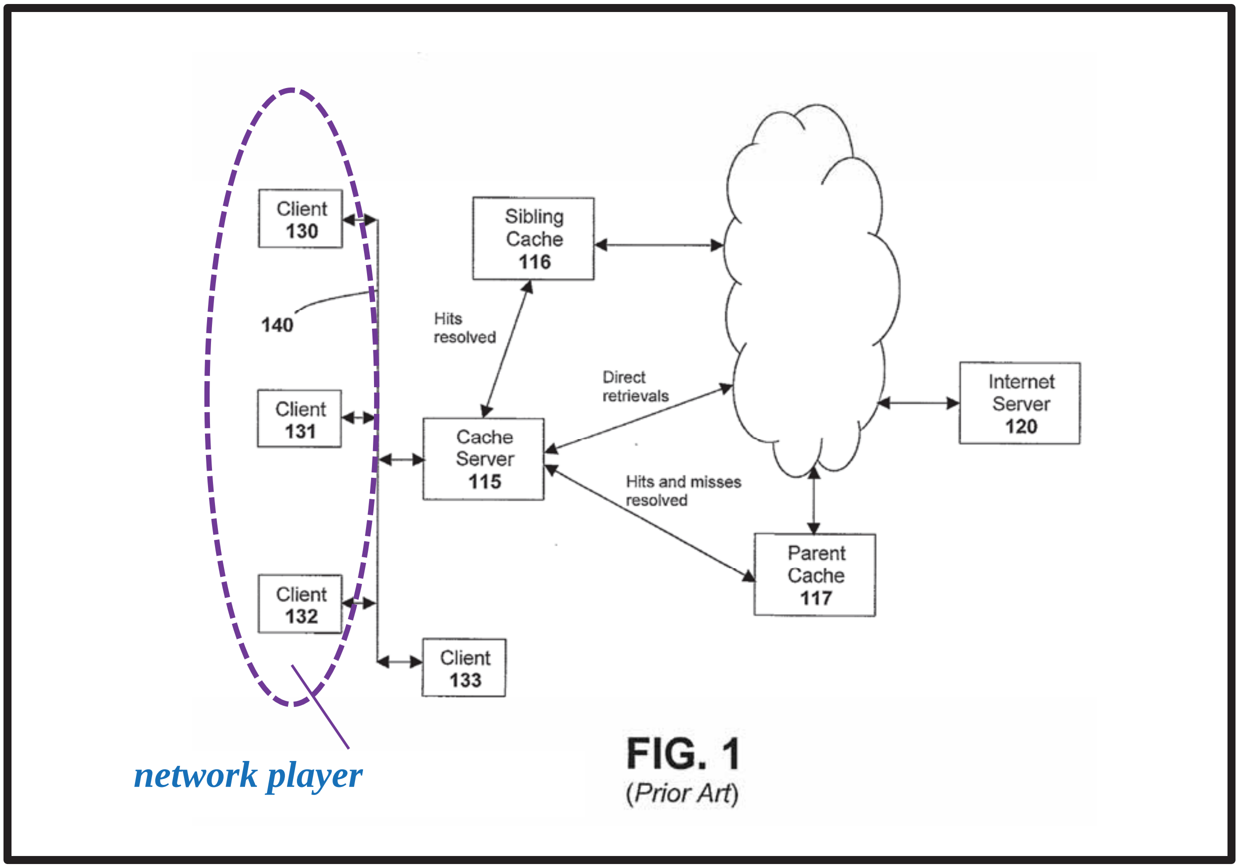

Ex.1004 [Sloss], FIG. 1 (annotated)

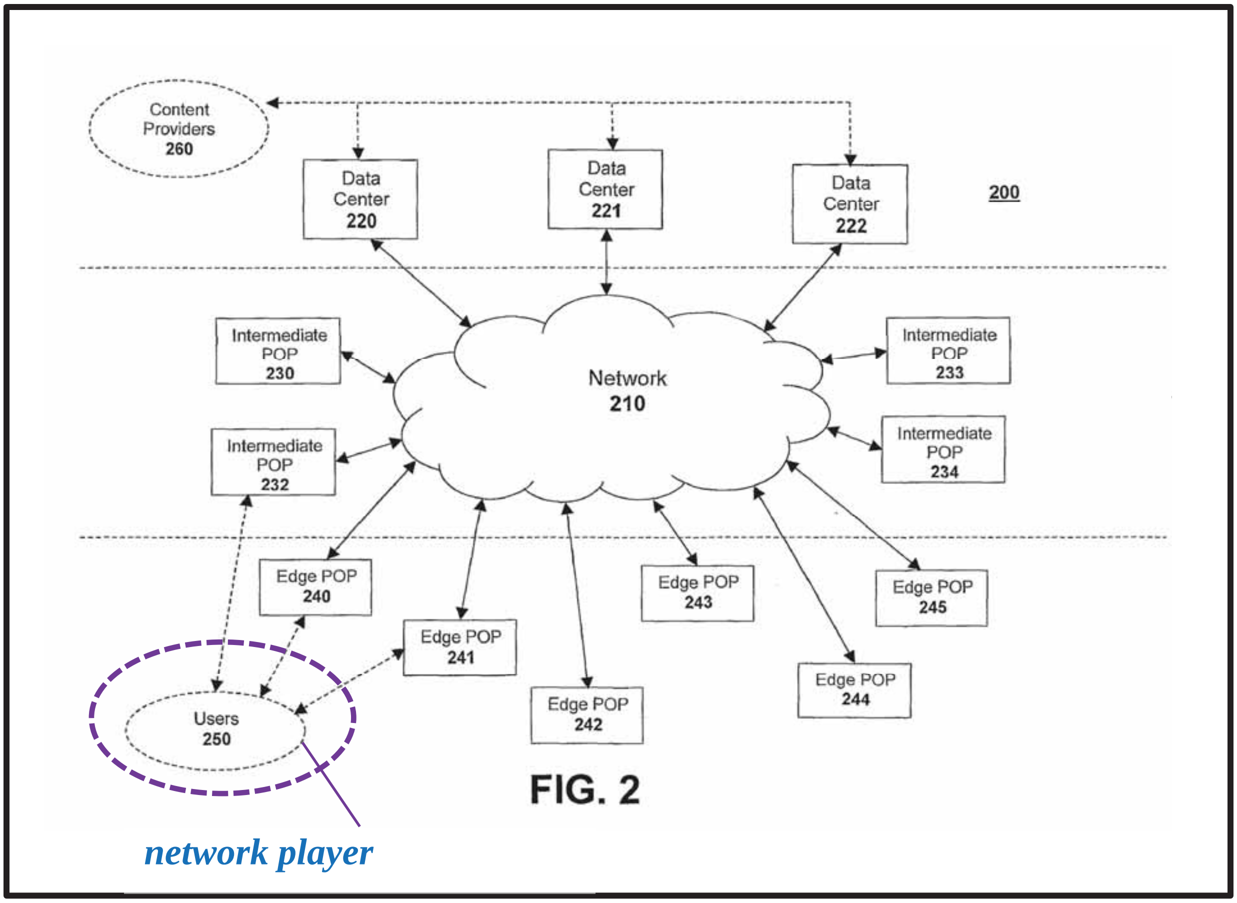

Ex.1004 [Sloss], FIG. 2 (annotated)

75. In other words, a POSITA would understand that Sloss teaches that a network player (e.g., a browser at a client device or an end user 250) transmits, to a local proxy cache, a request for content, and when that request corresponds to non-cacheable content, the request is forwarded to the internet server (or a server at data centers 220-222). Thus, the data centers 220-222 receive "*a data transmission request from a network player*." Responsive to this request, Sloss teaches that "requested content is then transmitted directly from the source 120 to the client 130 and is not stored locally on the proxy cache 115," streaming directly

from the remote internet server (e.g., servers at data centers 220-222) to the network player (e.g., clients 130-133 and end users 250). Ex.1004 [Sloss], 1:21-2:3. Sloss explains that this is "*stream-delivering*" as claimed because data centers transmit "on-demand multimedia streaming files." Ex.1004 [Sloss], 5:12-14; *see also* 8:9-18. Sloss also teaches that content cached at an edge POP may be "transmitted to the user." Ex.1004 [Sloss], 23:2-5. A POSITA would understand that transfer of data to be "stream-delivering" as well.

76. Thus, Sloss teaches that "*the server device for media*" (set of servers at data centers 220-222 and edge POPs) "*responds to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network*" (the servers transmit streaming data to the client device's browser over the Internet from the storage of the servers).

77. Accordingly, Sloss discloses or renders obvious this limitation.

[1.2.1] ***a transfer control unit adapted to transfer and store part of held digital contents in the internal storage device to a network storage device,***

78. In my opinion, Sloss discloses or renders obvious this limitation.

79. First, Sloss discloses multiple "*network storage device[s]*" in its teachings of Intermediate POPs in Figure 2. As with the data centers and edge POPs, Sloss teaches that "intermediate POPs 230-234 … are comprised of groups

of network servers on which various types of network content may be stored and transmitted." Ex.1004 [Sloss], 5:9-16, 8:9-18.

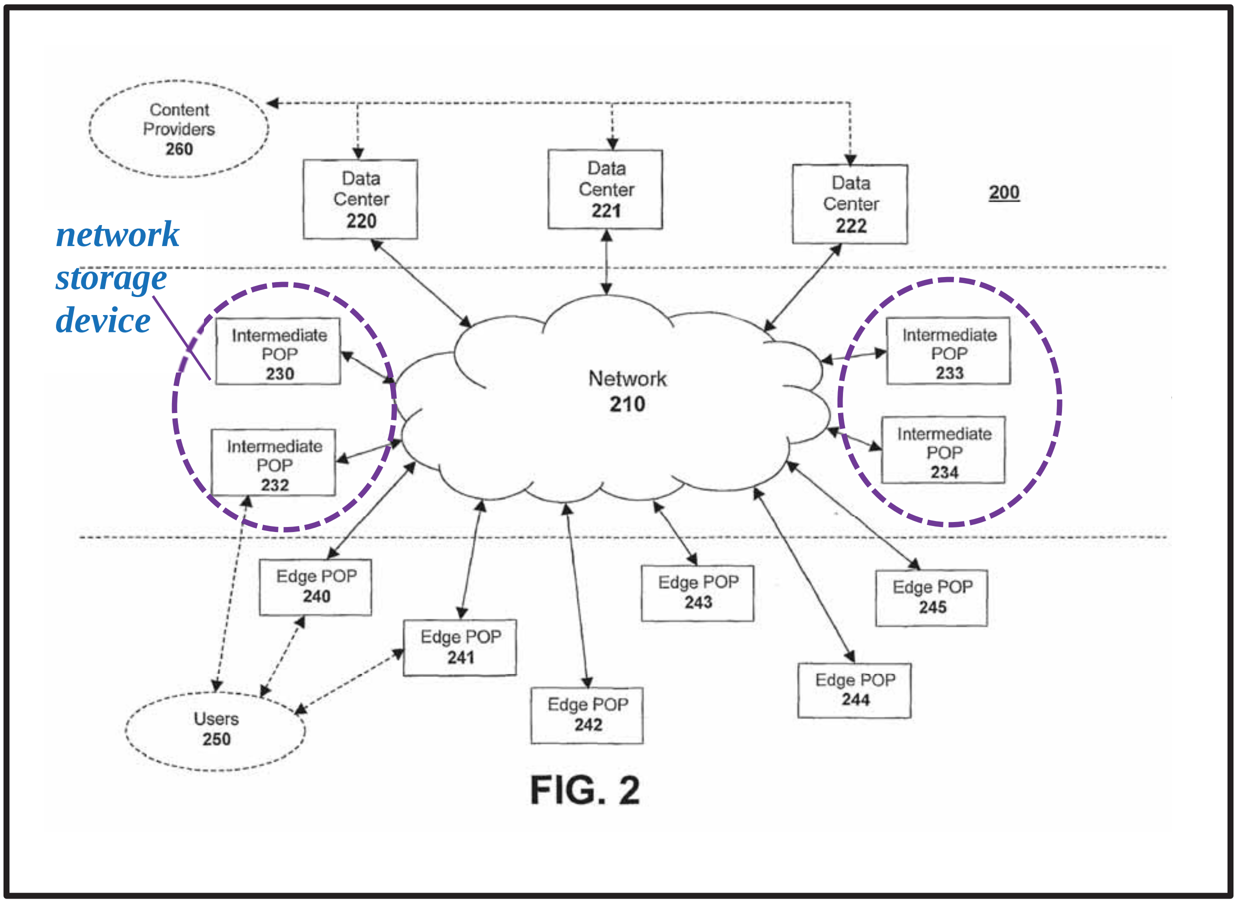

Ex.1004 [Sloss], FIG. 2 (annotated)

80. Sloss also discloses data being transferred and stored from the data centers to the intermediate POPs. For example, Sloss teaches that after content is uploaded to "a streaming server located at a data center 220-222 … [t]he file will then be automatically distributed from the data center 220-222 to one or more of the intermediate POPs 230-234 …" Ex.1004 [Sloss], 8:9-18, claim 1. Thus, Sloss teaches that the servers at the data centers "*transfer and store part of held digital*

contents in the internal storage device to a network storage device" (i.e., an intermediate POP).

81. This is consistent with the teachings of Sloss in Figure 1, which details the internet server distributing content to a local proxy cache: "[w]hen a client 130 requests information from a remote Internet server 120, the local proxy cache 115 examines the request and initially determines whether the requested content is 'cacheable' (a significant amount of Internet content is 'non-cacheable')….If a cache 'miss' occurs, however, the content is retrieved from the source Internet server 120, transmitted to the client 130 and a copy is stored locally on the proxy cache 115." Ex.1004 [Sloss], 1:18-2:17. Thus, Sloss teaches transmitting data from the data centers 220 to the intermediate POPs to store that content, which corresponds to "*transfer and store part of digital contents in the internal storage device to a network storage device*" as recited.

82. Third, Sloss teaches "*a transfer control unit*" to perform this functionality. For example, Sloss teaches that the computer system of Figure 3, which is representative of the servers at the data centers, include a "processor 310" and "instructions to be executed by processor 310." Ex.1004 [Sloss], 7:1-6.

83. Thus, Sloss discloses or renders obvious this limitation.

[1.2.2] ***wherein the network storage device is connected to the network and is capable of storing data,***

84. In my opinion, Sloss discloses or renders obvious this limitation. As shown in Figure 2, the intermediate POPs are connected to a network.

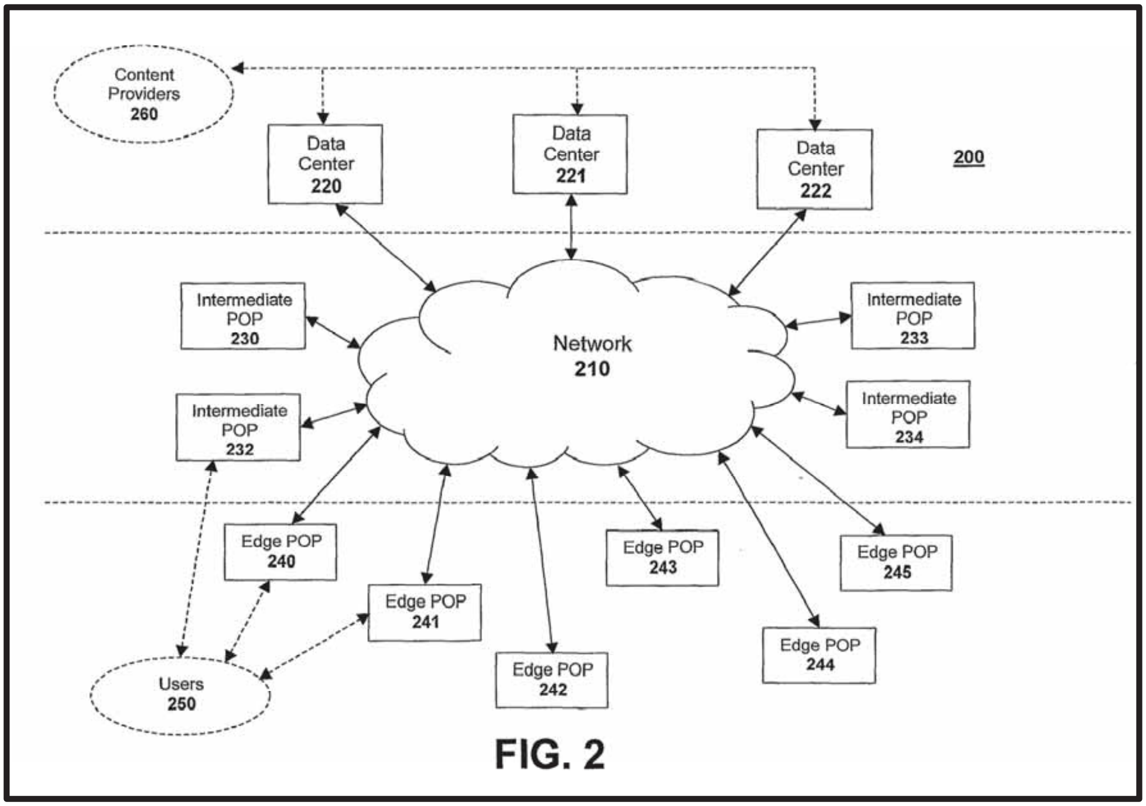

Ex.1004 [Sloss], FIG. 2

See also Ex.1004 [Sloss], FIG. 4.

85. Additionally, as detailed in the analysis of limitation [1.2.1], the intermediate POPs "are comprised of groups of network servers on which various types of network content may be stored and transmitted." Ex.1004 [Sloss], 5:9-14. Thus, the intermediate POPs are "*connected to the network*" and "*capable of storing data*" as recited.

[1.2.3] ***and wherein said transfer control unit does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device;***

86. In my opinion, Sloss discloses or at least renders obvious this limitation.

87. In describing the caching procedures with respect to Figure 1, Sloss teaches that only cacheable data will be transferred to the local proxy cache. Specifically, Sloss states,"[w]hen a client 130 requests information from a remote Internet server 120, the local proxy cache 115 examines the request and initially determines whether the requested content is 'cacheable' (a significant amount of Internet content is 'non-cacheable'). If the local proxy cache 115 detects a non-cacheable request, it forwards the request directly to the content source (e. g., Internet server 120). The requested content is then transmitted directly from the source 120 to the client 130 **and is not stored locally on the proxy cache 115**." Ex.1004 [Sloss], 1:18-2:3; *see also* Ex.1004 [Sloss], 2:4-10 (describing functionality when content is cacheable); *see also* Ex.1004 [Sloss], 22:12-18 (detailing a caching policy that considers "whether the requested file is non-cacheable"). Indeed, a POSITA would have considered it basic knowledge that only cacheable data would be transferred to a local proxy cache (as implied by the term "cacheable") and in fact, this basic concept is fundamental in my teachings to students.

88. My understanding of the Patent Owner's interpretation of this claim language in its district court filings is that, in its allegations of infringement, the Patent Owner has stated: "While delivering the data from cloud server to various APIs, the data is only cached to the cloud CDN storage if it is cacheable (i.e., wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device[])." *See* Ex.1012 [E-1 chart], 10-11, Ex.1013 [E-2 chart], 11-12 (similar).[2] In my view, this means that, when data is cached or not cached on the basis of a determination that the data is "cacheable" or not, then such contingent caching meets the language of this claim limitation.

89. Sloss teaches that it does not store non-cacheable data on the proxy cache, which is just like data only being cached "to the cloud CDN storage if it is cacheable" in Patent Owner's view, and therefore Sloss discloses or renders obvious this claim language.[3] In particular, Sloss teaches that the source 120 (which teachings are applicable, e.g., to the data centers 220-222) "*does not transfer, from the internal storage device to the network storage device*" (from the

---

[2] I have been informed by counsel that a patent owner's own infringement contentions may be considered by the Board when interpreting the claims.

[3] Although I set forth my understanding of Patent Owner's allegations in this paragraph, I am not expressing an agreement with Patent Owner's position that this concept is encompassed by the claim language, or that Patent Owner's allegations of infringement are correct in any sense.

data centers to the intermediate POPs) "*the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device*" (the requested content is not stored on the equivalent of the proxy cache, e.g., the intermediate POPs, if the content is not cacheable).

90. To the extent the Board finds that the claim scope only allows for "transfer" to mean "move" and not "copy" or "cache," it would have been obvious to a POSITA to use Sloss's disclosure of *caching* copies closer to the end user, and apply that to *moving* copies of the data closer to where the end user is requesting it. For example, as I have explained above, moving and copying data were well-known concepts at the time of the '101 patent (indeed, fundamental concepts within computing), and moving a content element (e.g., a movie) instead of copying/caching that movie would have reduced the amount of duplicate copies of the movie, and reduced the total amount of storage occupied by duplicate copies of the movie, which a POSITA would have recognized as desirable given the cost of storage at the time of the '101 patent. *See also, e.g.*, Ex.1017 [Microsoft Computer Dictionary], p. 349.

> move *n*. A command or an instruction to transfer information from one location to another. Depending on the operation involved, a move can affect data in a computer's memory or it can affect text or a graphical image in a data file. In programming, for example, a move

> instruction might transfer a single value from one memory location to another. In applications, on the other hand, a move command might relocate a paragraph of text or all or part of a graphic from one place in a document to another. Unlike a copy procedure, which duplicates information, a move indicates that information either is or can be deleted from its original location.

As the above definition from the Microsoft Computer Dictionary explains, move and copy are well-known alternatives to each other, where a “move” is a “command or an instruction to transfer information from one location to another” while a copy “duplicates information,” and a move “indicates that information either is or can be deleted from its original location.” Given this well-known concept, and for example, within the context of Sloss, it would have been nothing more than an obvious variation of Sloss’s techniques to **move** data (e.g., a movie) from one of the data centers 220-222 to an intermediate POP, such that only one copy of that movie exists, and not **copy** that movie, which would have the desirable benefit of saving costly storage space within the Figure 2 system.

91. Thus, Sloss renders obvious this limitation.

[1.3.1] ***a list information transmission unit adapted to respond to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,***

92. In my opinion, the combined teachings of Sloss and Lamkin render obvious this limitation.

93. Sloss teaches a "database for storing information relating to distributed network content" (Ex.1004 [Sloss], 6:10-15, 15:17-16:10) but does not explicitly teach that the data center servers can provide a list of digital contents to the client devices; however this would have been well-known to a POSITA in the streaming media context of Sloss and prior to the invention. For example, Netflix's streaming service launched in January 2007, prior to the '101 patent's earliest effective filing date, and a POSITA would have known that Netflix would have been a typical content provider consistent with the disclosure of Sloss, and one which provided lists of available digital content to end users (e.g., a list of available movies for an end-user to stream on their device).

94. Consistent with this knowledge of a POSITA, Lamkin teaches responding to "*a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player*." For example, Lamkin teaches "methods, systems and apparatuses for use in managing content" on a network and teaches that "a user 224 at the client device 124 is provided a user interface that accesses and/or shows content accessible, for example, through a CDS [content directory service] from various devices and/or local content, and allows the user to select the content to be pulled to the client device to be accessed and/or recorded." Ex.1005 [Lamkin], Abstract, ¶56; *see also* ¶¶70-77, 128, Fig. 4. A POSITA would have understood that

Lamkin's user interface is provided responsive to a request for that user interface, for example, Lamkin teaches "[u]sers can access a content user interface, for example, by selecting the content button or option" and be provided the "user interface 520 of FIG. 5" and thus Lamkin at least renders obvious "*respond[ing] to a list presentation request*." Ex.1005 [Lamkin], ¶204. A POSITA would have recognized that selecting a content button or option to be a common method of users sending requests to various systems, such as those disclosed in Lamkin.

95. Lamkin depicts, in Figure 5, "a simplified example of a user interface 520 according to some embodiments that identifies content added and/or altered since a previous network user access, within a time period, content not previously distributed, and other criteria or combinations of criteria. The user interface 520 can include a listing of content 522." Ex.1005 [Lamkin], ¶87; *see also* ¶¶68 ("listing or database"), 54 ("database, listing or other methods"), 113.

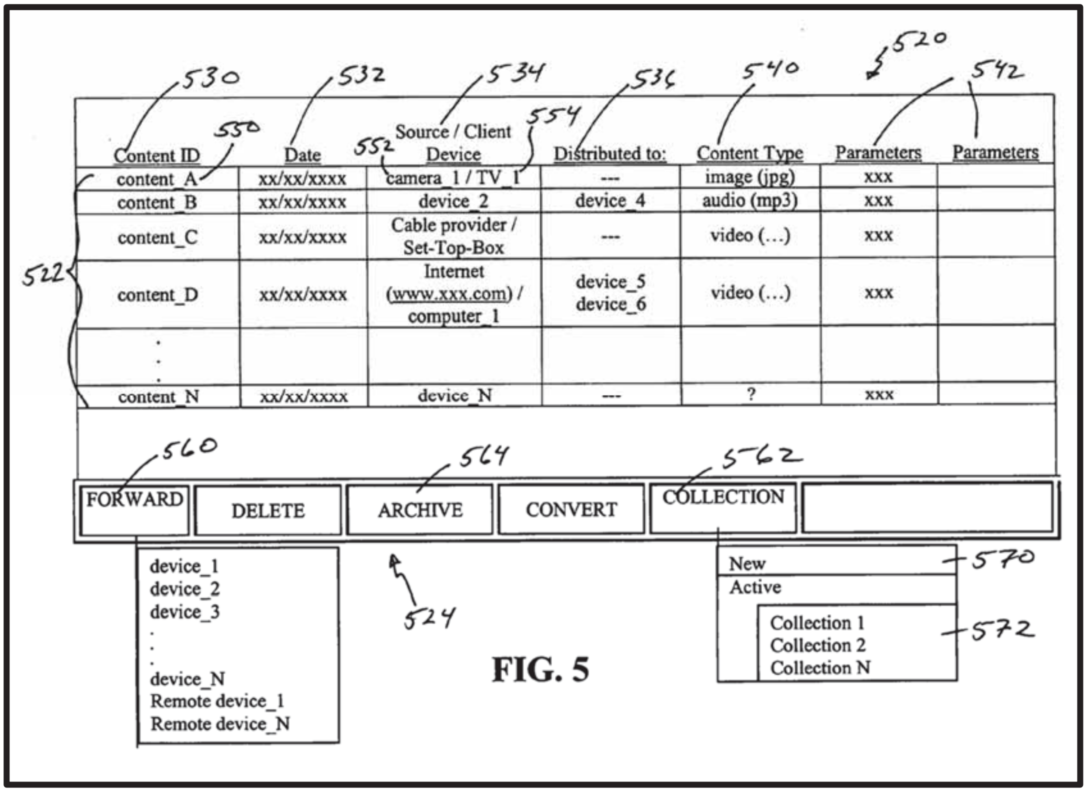

Ex.1005 [Lamkin], FIG. 5

96. As shown in Figure 5, the "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87. Thus, Lamkin's provision of a user interface corresponds to "*transmitting list information to the network player*." In the combination with Sloss, this functionality of Lamkin would be implemented by Sloss's servers at data centers 220-222, which therefore correspond to the recited "*list information transmission unit*."

97. Accordingly, the combined teachings of Sloss and Lamkin render obvious "*a list information transmission unit adapted to respond to a list presentation request for the held digital contents of the server device for media*

from the network player" (Lamkin's providing of a user interface with a listing with Sloss's teachings of data centers 220-222).

[1.3.2] *wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,*

98. In my opinion, the combined teachings of Sloss and Lamkin render obvious this limitation.

99. Sloss teaches that "the database 530 in one embodiment keeps track of exactly where content has been distributed throughout the system." Ex.1004 [Sloss], 21:11-14, 15:17-16:10, claim 9. Thus, a POSITA would have recognized that Sloss teaches keeping track of whether data is "*in the internal storage device*" (e.g., at a data center 220-222 or edge POP) or has been "*transferred from the internal storage device to the network storage device and stored in the network storage device*" (e.g., within an intermediate POP).

100. This is confirmed by Lamkin's teachings. As shown in Figure 5 and detailed above, Lamkin's user interface includes a listing of content that is available on various network-connected devices. Specifically, Lamkin's "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87. In more detail, Figure 5 shows **content_C** which is available on client device set-top-box (which may be, in one example, a personal video recorder). This content

item corresponds to one example of "*the digital contents left in the internal storage device*" which, in the combination, would be content stored at Sloss's data centers or edge POPs. Lamkin's Figure 5 also shows **content_D** which has been distributed to device_5 and device_6, one of which may be, as Lamkin explains "network attached storage (NAS) 132" and in the combination, would be content stored at Sloss's intermediate POPs. Ex.1005 [Lamkin], ¶85. This content item thus corresponds to one example of "*the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device*" because a POSITA would have recognized network attached storage as a backup device consistent with Lamkin's disclosure of "remote storage of content" done to "archive content." Such a backup device would be similar to the intermediate POPs of Sloss, i.e., the "*network storage device*." Systems that allow data to be stored in internal storage as well as remote storage, and that provide a list showing where each data item was stored, were well known since the mid-1990s. For example, the "archie" system and the "gopher" system were both mid-1990s systems that had such functionality.

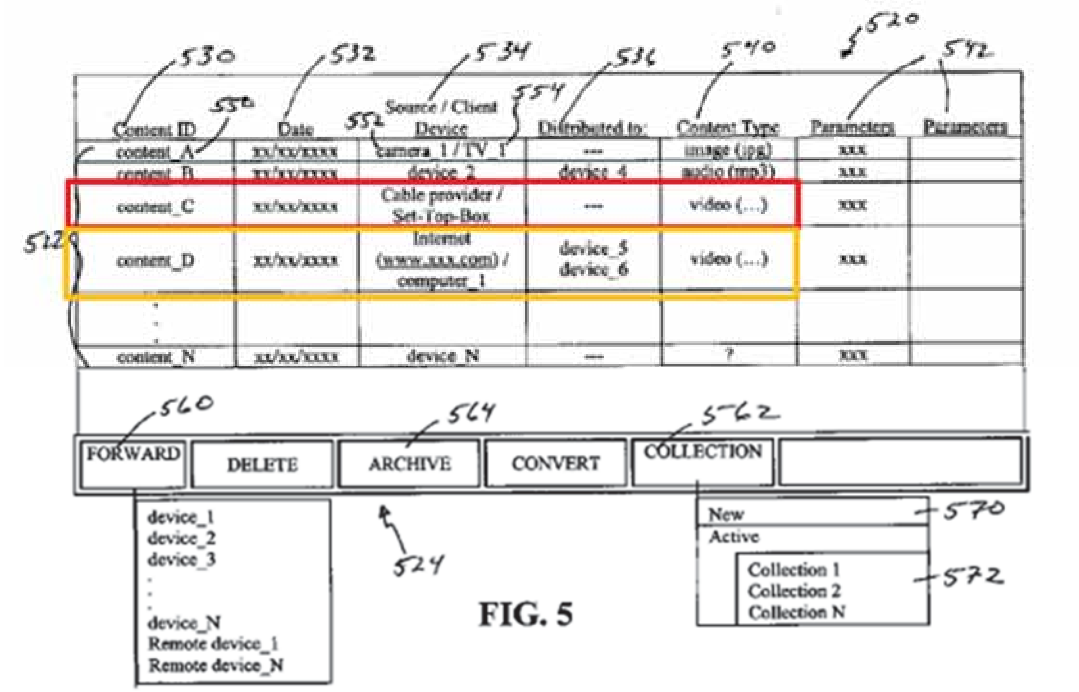

Ex.1005 [Lamkin], FIG. 5 (annotated)

101. Thus, the combined teachings of Sloss with Lamkin's user interface in Figure 5 render obvious "*list information*" that "*lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device*" as recited.

[1.3.3] ***and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;***

102. It is my opinion that Lamkin discloses or at least renders obvious this limitation. Lamkin teaches that in some implementations the "user interface 520 includes … a media in-box user interface." Ex.1005 [Lamkin], ¶94. Lamkin's

"media in-box user interface" is shown in Figure 6, which depicts a simplified example of a media in-box 620 according to some embodiments. The media in-box 620 can identify content and in some implementations present the content according to an **organized structure, similar to a file structure**." Ex.1005 [Lamkin], ¶95. As shown in Figure 6, the file structure appears as a hierarchical tree:

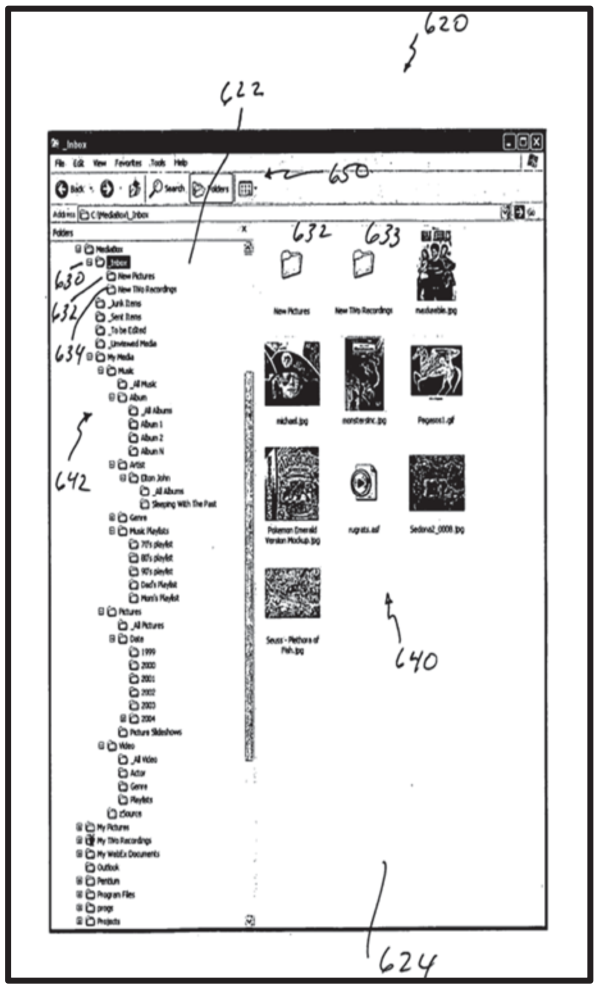

Ex.1005 [Lamkin], FIG. 6

103. The tree structure depicted in Figure 6 would have been recognized by a POSITA as common to many file systems at the time of the ’101 patent. For example, a POSITA would have recognized that both Windows and Unix-based operating systems stored files in a hierarchical file system, in which directories/folders and files were represented by a tree structure, tracing all the way back to the first instance in the 1960s. *See* Ex.1018 (File Structures), p. 22 (“The UNIX filesystem is a tree-structured organization….”); Ex.1017 [Microsoft Computer Dictionary], pp. 457 (“root directory *n*. The point of entry into the directory **tree** in a disk-based hierarchical directory structure.”), 529 (“tree view *n*. A hierarchical representation of the folders, files, disk drives, and other resources connected to a computer or network. For example, **Windows Explorer uses a tree view** to display the resources that are attached to a computer or a network.”). Confirmed by Lamkin’s presentation of such a tree-structured file system (e.g., a Windows file system), therefore, a POSITA would have recognized that Lamkin’s devices would have “*maintain[ed] a tree structure of the digital contents in the internal storage device*” at all times, including “*before transferring the digital contents to the network storage device*” as recited. *See also* Ex.1005 [Lamkin], ¶¶264, 296-297 (teaching a tree structure).

104. Thus, Lamkin teaches that "*the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device*."

[1.4] *a search unit adapted to respond to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and*

105. It is my opinion that Sloss discloses or at least renders obvious this limitation.

106. Sloss teaches, in one example, "when a user 705 requests content stored on a particular Internet site ... the request is received by a load balancer module ('LBM') 710." Ex.1004 [Sloss], 20:21-21:3. Sloss explains that the "LBM in one embodiment is a module which resides at a data center (e.g., running on a Web server)" and thus the LBM corresponds to a "*search unit*" which is part of the recited "*server device*" and which receives "*a data transmission request for the held digital contents from the network player*" (e.g., the LBM receives a request for content from a user at a client device). Ex.1004 [Sloss], 21:4-19.

107. Sloss further teaches that the LBM, or "*search unit*," is "*adapted to respond to a data transmission request … by searching for a location where the held digital contents are currently stored*" as it teaches that "[i]n one embodiment, the LBM 710 finds the most appropriate edge POP 507 and determines whether the content is available at the edge POP 507 by querying the central database 530

(i.e., the database 530 in one embodiment keeps track of exactly where content has been distributed throughout the system). If the requested content is available at the edge POP 507, it is transmitted to the user 705. If, however, the content is not available at the edge POP 507, then the LBM 710 redirects the request to the second most appropriate POP, (e.g., intermediate POP 506 in the illustrated embodiment), which then transmits the content to the user 705." Ex.1004 [Sloss], 21:11-19, Fig. 7, Fig. 8. Thus, the LBM searches for where the requested content is stored, which corresponds to this claim language. A POSITA would know that this basic search feature was already well known in the art at the time of the '101 patent.

108. Accordingly, Sloss's teachings of an LBM at the data center discloses or renders obvious this limitation.

[1.5] ***a digital contents data transmission processing unit adapted to allow the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,***

109. It is my opinion that Sloss discloses or renders obvious this limitation.

110. Sloss teaches that, in one example, "a client 130 requests information from a remote Internet server 120" and when the client's request is cacheable "it searches for a copy of the content locally" and may transmit the content from a cache. Ex.1004 [Sloss], 1:9-2:10. As detailed above, the intermediate POPs in Sloss are caches and correspond to the recited "*network storage device*" and thus at least this example of Sloss renders obvious "*allow the corresponding data in*

held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device." *See also* Ex.1004 [Sloss], 21:16-19 (describing intermediate POP transmitting content to user), Fig. 8 (step 850).

111. Sloss also teaches a streaming module through which content is transmitted to the user:

> "[o]ne particular embodiment of the system and method for distributing and **streaming multimedia files** will now be described with respect to Figure 13. A viewer 1310 connected to the Internet through an edge POP 507 in this example, makes a request to stream an on-demand file. The file is referenced in the IES database 1320 by a "FileInfo" record with the ID to the record embedded as a parameter in the URL the viewer clicked on to access a Web server 1325 at the data center 505. The web server 1325 in this embodiment brings up a **streaming module** (e. g., a Web page; "stream. asp" for Windows 98TM) 1335 to process the request."

Ex.1004 [Sloss], 24:4-15.

112. Sloss explains that the streaming module within the web server 1325, which corresponds to the "*digital contents data transmission processing unit*," calls the stream redirector to determine the path of the content (where the content is being stored, it will find a match and identify an intermediate POP server (e.g., a "*network storage device*")) and then return a streaming path that redirects the viewer (a "*network player*") to a media server:

> "[t]he Stream Redirector 1340 in one embodiment is an out-of-proc COM server running on the Web server 1325. When called by the **streaming module** 1335 to create the streaming server path to the on-demand file, the redirector 1340 first checks the viewer's 1310 IP address against a list of site IP masks collected earlier from the database 1320. In the illustrated embodiment, the redirector 1340 finds a match and correctly identifies the edge POP site 507 the viewer 1310 is connecting from. It checks the database 1320 (e. g., using database API's) to determine if the desired file exists at the viewer's edge POP site 507. If it finds a FileLocation record matching this site 507 using the FileInfo ID from the URL, it returns a streaming path that redirects the viewer to a media server 1345 colocated at the edge POP site 507. **If it doesn't find the file there (i.e., resulting in a cache 'miss'), it instead generates a path redirecting the viewer to one of the intermediate POP sites 506 where the file is known to be located**…. The redirector 1340 returns the intermediate POP redirection path to the streaming module 1335 where it is inserted into the metafile and returned to the viewer's 1310 browser."

Ex.1004 [Sloss], 24:16-25:7, 25:17-19, *see also* 23:6-14 ("intermediate POP server transmits the content to the user at 850"). Thus, Sloss teaches allowing the file to be "*stream-delivered from the network storage device*" (from the intermediate POP server) "*to the network player, if the result of search shows the network storage device*."

113. Thus, Sloss renders obvious this limitation.

[1.6] ***wherein the server device for media is a media player.***

114. Sloss discloses that the "data centers 220-222" and "edge POPs 240-245" (i.e., the "*server device for media*") include servers "on which various types of network content may be stored and transmitted to end users 250, including, for example … live & on-demand multimedia streaming files." Sloss explains, in one example, that a redirector returns a file path to the streaming module, where the file is returned to the viewer's browser, which then hands it over to the streaming media player: "The viewer's 1310 browser receives the metafile and hands it over to the **streaming player (e. g., RealPlayer®, Windows® . . . etc)**. The player parses the metafile for the redirection path." Ex.1004 [Sloss], 25:17-26:2.

115. This interpretation is consistent with my understanding of Patent Owner's infringement allegations in district court. I understand that Patent Owner has argued that "Google Cloud CDN servers act as media players by hosting and streaming content via services like YouTube …. The Google Cloud servers with integration to its services like YouTube act as a media player." Ex.1012 [E-1], 32. Thus, consistent with Patent Owner's allegations that servers hosting content that is delivered to users constitute "*a media player*," Sloss's servers at data centers 220-222 and edge POPs 240-245 ("*the server device for media*") correspond to "*a media player*" as recited.

E. Dependent Claim 2

[2.1] ***The server device for media according to claim 1, wherein said digital contents data transmission processing unit causes the network storage device to transmit the corresponding data to the server device for media, and then transmits the corresponding data received from the network storage device from the server device for media to the network player.***

116. In my opinion, Sloss discloses or renders obvious claim 2. Specifically, as discussed in Claim [1.0], Sloss teaches a set of servers located at data centers 220-222 and edge POPs which are mapped to the claimed "*server device for media*."

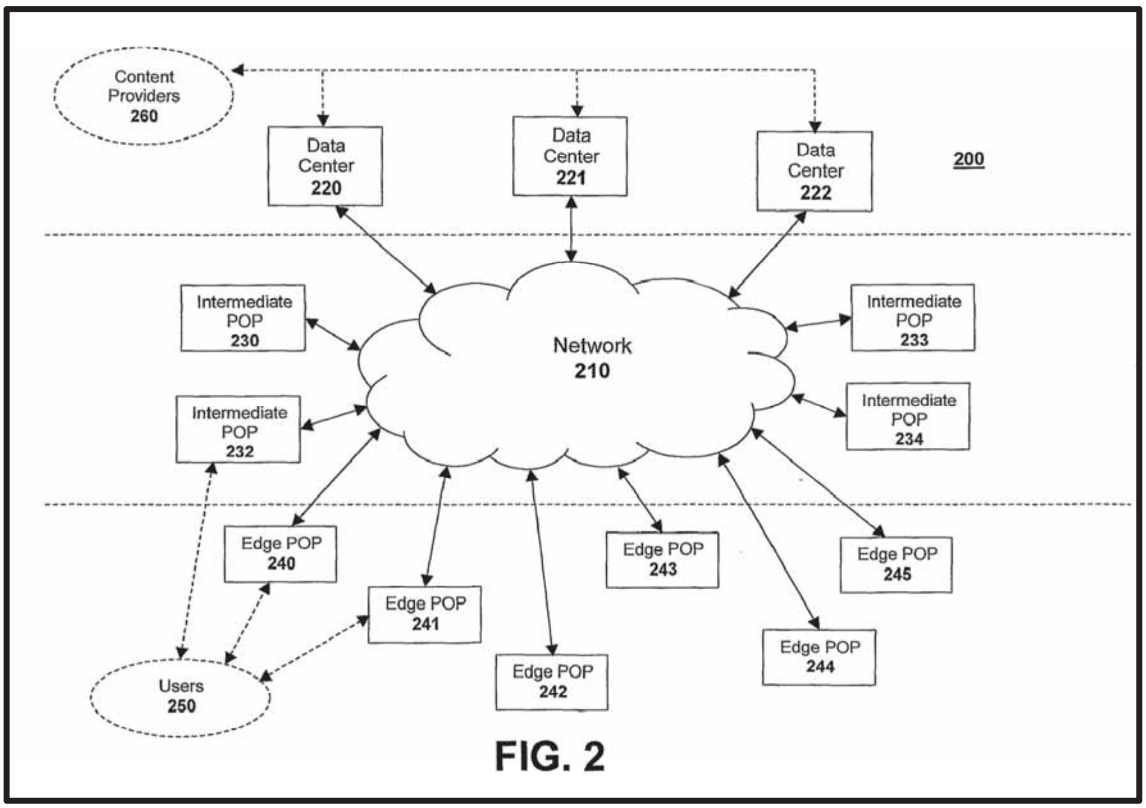

Ex.1004 [Sloss], FIG. 2

117. As detailed above, the intermediate POP in Sloss is a cache and corresponds to the recited "*network storage device*." Sloss teaches, in some

examples, that its method may "determine[] whether a copy of the requested content should be stored locally at the edge POP" and if so, "then the content is transferred to the edge POP site" from an "intermediate POP server" and thus Sloss renders obvious "*said digital contents data transmission processing unit causes the network storage device*" (Sloss's transfer from the intermediate POP server) "*to transmit the corresponding data to the server device for media*" (content is transferred to the edge POP).

118. As discussed in Claim [1.5] above, Sloss also teaches a streaming module through which content is transmitted to the user at the network player from the edge POP (i.e., "*from the server device for media to the network player*"): "[o]ne particular embodiment of the system and method for distributing and **streaming multimedia files** will now be described with respect to Figure 13. A viewer 1310 connected to the Internet through an **edge POP 507** in this example, makes a request to stream an on-demand file…. The web server 1325 in this embodiment brings up a **streaming module** (e. g., a Web page; "stream. Asp" for Windows 98TM) 1335 to process the request." Ex.1004 [Sloss], 24:4-15. It was also well known to a POSITA that a user could stream content from the edge POP to the network player (where the user would view the content). Thus, Sloss renders obvious this limitation.

F. Dependent Claim 3

[3.1] ***The server device for media according to claim 1, wherein said digital contents data transmission processing unit transmits the corresponding data and information for identifying the network storage device to the network player, and causes the network storage device to directly transmit the corresponding data to the network player.***

119. In my opinion, the combined teachings of Sloss and Lamkin render obvious this limitation.

120. In the teachings of the Sloss-Lamkin combination, as detailed above with respect to claim 1, Lamkin's list information transmitted to the client device includes a location of the content (e.g., the source/client device and the devices to which the content is distributed to). Thus, the combination teaches transmitting "*information for identifying the network storage device to the network player*" as recited.

121. The combination further teaches transmitting "*the corresponding data … to the network player … and causes the network storage device to directly transmit the corresponding data to the network player*" as recited. For example, as explained with respect to claim 1, in Sloss, if content is available at an intermediate POP (a "*network storage device*"), the content may be transferred from the POP to the client device ("*to the network player*"). *See also* Ex.1004 [Sloss], 23:6-14 ("intermediate POP server transmits the content to the user at 850").

122. Thus, the combined teachings of Sloss and Lamkin render obvious this limitation.

G. Dependent Claim 5

[5.1] ***The server device for media according to claim 1, wherein said list information transmission unit makes the list information to be transmitted to the network player***

123. In my opinion, the combined teachings of Sloss and Lamkin render obvious this limitation.

124. Again, Sloss does not explicitly teach that the data center servers or edge POPs can provide a list of digital contents to the client devices, but this would have been well-known to a POSITA in the streaming media context of Sloss prior to the invention. *See* [1.3.1].

125. Consistent with this knowledge of a POSITA, Lamkin teaches "*list information to be transmitted to the network player*." For example, Lamkin teaches "methods, systems and apparatuses for use in managing content" on a network and teaches that "a user 224 at the client device 124 is provided a user interface that accesses and/or shows content accessible, for example, through a CDS from various devices and/or local content, and allows the user to select the content to be pulled to the client device to be accessed and/or recorded." Ex.1005 [Lamkin], Abstract, ¶0056. And again, a POSITA would have found it obvious to combine Lamkin because providing lists of content to client devices was well-known at the

time of the '101 patent (e.g., streaming providers provided lists of available movie content to client devices for an end-user to stream), consistent with the disclosure of Sloss,.

126. Accordingly, Lamkin renders obvious "*said list information transmission unit makes the list information to be transmitted to the network player*" (Lamkin's providing of a user interface on the network player with a listing of digital contents).

[5.2] ***include information for identifying whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player.***

127. In my opinion, Lamkin renders obvious this limitation. As shown in Figure 5 and detailed above, Lamkin's user interface includes a listing of content that is available on various network-connected devices. Specifically, Lamkin's "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87. *See* [1.3.2].

128. As Figure 5 shows, the listing 522 includes "*information for identifying whether each digital content is currently stored in the internal storage device*" (e.g., content that is stored in a **source/client device**) "*or the network storage device*" (e.g., content that is **distributed to and available from another**

device) "*in the display list of the network player*" (e.g., at a client device's browser in Sloss):

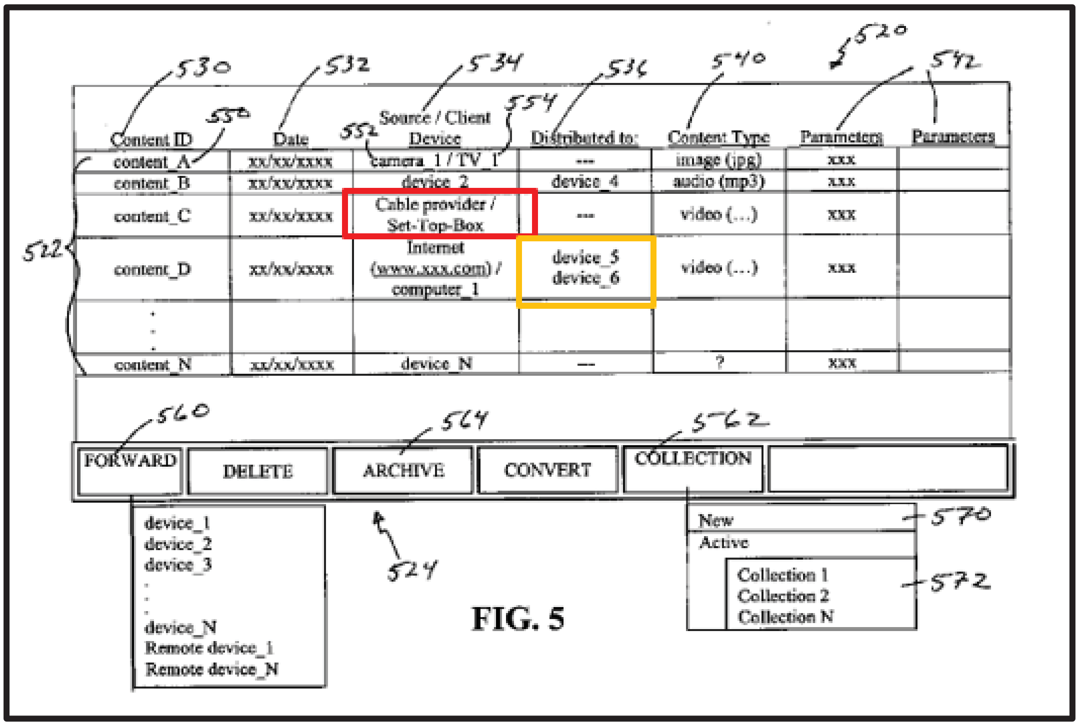

Ex.1005 [Lamkin], FIG. 5 (annotated)

129. Thus, Lamkin's user interface in Figure 5 renders obvious "*whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player*" (again, Lamkin's providing of a user interface on the network player with a listing of digital contents, noting which network-connected device each content is located on). A POSITA would be familiar with this functionality because the "archie" system and the "gopher" system were both mid-1990s systems that allowed data to be stored in internal as

well as remote storage and provided a list showing where each data item was stored.

H. Independent Claim 7

[7.0] ***A method for controlling a server device for media which is equipped with an internal storage device for storing digital contents, the method comprising the steps of:***

130. *See* [1.0].

[7.1] ***responding to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;***

131. *See* [1.1].

[7.2.1] ***transferring and storing part of held digital contents in the internal storage device to a network storage device,***

132. *See* [1.2.1].

[7.2.2] ***wherein the network storage device is connected to the network and is capable of storing data,***

133. *See* [1.2.2].

[7.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device;***

134. *See* [1.2.3].

[7.3.1] ***responding to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,***

135. *See* [1.3.1].

[7.3.2] ***wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,***

136. *See* [1.3.2].

[7.3.3] ***and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;***

137. *See* [1.3.3].

[7.4] ***responding to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and***

138. *See* [1.4].

[7.5] ***allowing the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,***

139. *See* [1.5].

[7.6] ***wherein the server device for media is a media player.***

140. *See* [1.6].

VII. Ground 2: Claim 4 is obvious under § 103 over Sloss in view of Lamkin, and further in view of Rou

A. Summary of Rou (Ex.1015)

141. The reference that I refer to as "Rou," is a paper published by IEEE in 2002. Ex.1015 (Rou). Kai Rou is the first-named author. Ex.1015, 1. IEEE, the Institute of Electrical and Electronics Engineers, is a professional organization that is well-known for publishing timely research on various topics, including topics related to streaming media, the Internet, and others. IEEE papers are regularly

referred to by academics and industry members to gain an understanding of current techniques and technologies, and are readily accessible by members of the public, including interested POSITAs. I am familiar with IEEE's publication of research papers and its prominence based at least on my status as an IEEE fellow, and the fact that I have had over 40 papers published by the IEEE. I understand that Rou is prior art under § 102(b).

142. Rou discloses an "Online File Storage System." The authors of Rou developed "an online file storage system called UniFile" to which is described as "leverag[ing] the use of the Internet to save [] files in [a] remote server and to use the system to access their files." Ex.1015 [Rou], 83. In one aspect, Rou details an ability to upload and download contents to and from an online photo gallery (i.e., the ability to upload and download digital content from a local device to network storage): "[t]he online photo gallery facilitates the uploading of pictures from digital cameras or previously scanned image to their own sites." Ex.1015 [Rou], 83. Rou also discloses tools to perform simple yet powerful functions, "by assisting user in managing and organizing their remote online files with a few simple mouse clicks . . . Upload, Download. . . CopyTo, MoveTo" allowing for photos to be uploaded to an online photo gallery (e.g., a network storage device) and then downloaded back to the PC (e.g., an internal storage device). Ex.1015 [Rou], 84.

A POSITA would have recognized the techniques described in Rou as common to network-based storage technologies.

B. Motivation to Combine

143. For the reasons set forth below, a POSITA would have been motivated to combine the teachings of Rou with the teachings of Sloss and Lamkin. It would have been obvious, beneficial, and predictable to apply Rou's teachings regarding downloading data from network storage back to internal storage to the Sloss and Lamkin combination disclosing uploading data to network storage from internal storage.

144. A POSITA, when considering the teachings of Sloss and Lamkin, would have also considered the teachings of Rou. Rou is also in the field of managing digital content in a networked system, and as such, Rou is analogous art to the '101 patent.

145. A POSITA would have been specifically motivated to incorporate Rou's teachings because doing so would have provided details on how to implement functionality disclosed in Sloss and Lamkin. Specifically, Rou's teachings provide detail on how the data saved on the network storage device can be returned back to the internal storage, for example, if that data is to be deleted from the network storage device. This would have been consistent, for example, with Sloss's teachings of copying data from an intermediate POP (i.e., network

storage device) to an edge POP based on a caching policy. *See* Ex.1004 [Sloss], 21:11-22:11.

146. A POSITA would have also had reasonable expectation of success in making the combination. The combination does not change the intended functionality of Sloss, but rather, Rou's teachings merely provide implementation details for Sloss. Accordingly, a POSITA would have found it obvious to apply Rou's teachings to Sloss because the combination merely amounts to the combination of prior art elements according to known methods.

C. Dependent Claim 4

[4.1] ***The server device for media according to claim 1, further comprising a return control unit adapted to cause the digital contents corresponding to a predetermined condition among the digital contents which have been transferred to the network storage device to be returned from the network storage device to the internal storage device.***

147. In my opinion, the combined teachings of Sloss and Rou render obvious this limitation.

148. Sloss teaches various "Storage Space Management" techniques implemented at the POPs (including the intermediate POPs, i.e., implemented at the "*network storage device*" of the claims). These techniques include, for example, "file removal operations" for cached content files that have passed their expiration dates (e.g., "*digital contents corresponding to a predetermined condition among the digital contents*"). Ex.1004 [Sloss], 27:11-16, 16:11-21.

149. Although Sloss does not explicitly teach those "*digital contents which have been transferred to the network storage*" are "*returned from the network storage to the internal storage device*," when those files have passed their expiration date, this technique would have been well-known and desirable to a POSITA in the streaming media context of Sloss and prior to the invention. For example, a POSITA would have recognized that, if content was moved (and not copied) to the intermediate POP from the data center, it would need to be moved back to the data center to ensure that an instance of the content still existed.

150. Rou discloses "tools" that "perform simple yet powerful functions, by assisting user in managing and organizing their remote online files with a few simple mouse clicks. They are 16 functions altogether, namely: New Folder, New File, **Upload**, **Download**, Zip Download, Cut, Copy, Paste**, CopyTo, MoveTo**, Properties, Rename, Preview, Edit, Delete and Zip." Ex.1015 [Rou], 84. Rou's tools apply to a user's "remote folder." Thus, Rou teaches returning (Download), content from a remote folder ("*cause the digital contents … which have been transferred to the network storage*") to the user's local PC device ("*to be returned from the network storage device to the internal storage device*") as recited, and in combination with Sloss renders obvious this limitation. A POSITA would have known that there were applications at the time of the '101 patent, that would require a user to return data offloaded from internal storage to network storage, back to the

internal storage. In fact, moving data between internal and remote storage was commonly used in distributed file systems since at least the mid-1980s. For example, in the Andrew File System, an offloaded file would be downloaded from the remote server back to the user's local machine when it is edited, and then moved back to the remote server when the edit is completed.

VIII. Ground 3: Claims 6, 8, 9, 11, and 12 are obvious under § 103 over Sloss in view of Lamkin, and in further view of Chamberlain

A. Summary of Chamberlain (Ex.1010)

151. The reference that I refer to as "Chamberlain" is U.S. Patent Publication US 2002/0026563. It was filed on January 25, 1999 and published on February 28, 2022. Ex.1010 [Chamberlain] Cover. John T. Chamberlain is the first-named inventor. Ex.1010 [Chamberlain], Cover. I understand that Chamberlain is § 102(b) prior art.

152. Chamberlain is titled, "Cache Override Control in an Apparatus for Caching Dynamic Content," and teaches a "caching system and method [] that allow for the caching of web pages that have dynamic content. . . . The caching system only caches those responses having dynamic content that are deemed cacheable." Ex.1010 [Chamberlain], Abstract.

153. Chamberlain also discloses that "the automatic caching system can be overridden by the information author, the page creator or the system designer," and "the caching system of the present invention further comprises a system for

overriding the automatic analysis performed by the system." Ex.1010 [Chamberlain], Abstract, ¶18.

B. Motivation to Combine

154. For the reasons set forth below, a POSITA would have been motivated to combine the teachings of Chamberlain with the teachings of Sloss and Lamkin. It would have been obvious, beneficial, and predictable to apply Chamberlain's teachings regarding a manual override feature of an automatic caching system to the Sloss and Lamkin combination---for example, to provide a more customizable system.

155. For example, in Chamberlain, a "validity analyzer . . . automatically invalidates [web]pages based upon whether the candidate cached response is stale. . . . Because of [the cacheability analyzer's and attribute analyzer's] conservative nature, they err on the side of indicating that the response is not cacheable in order to guarantee the freshness of the page returned. The cacheability analyzer [is] configured to allow the page designer to decide that certain pages ought to be cached where the caching strategy generator, following the conventions discussed above, determines that they cannot be. The designer can utilize the following controls to override the cache behavior where appropriate." Ex.1010 [Chamberlain], ¶131. As such, a POSITA would have known that this feature to override an automatic caching protocol to create a more customizable experience

was well-known in the art at the time of the ’101 patent. For example, the HTTP and DNS protocols both support user-configurable settings that allow the content creator to control whether and how long a content item can be held in a cache. The HTTP protocol allows the page creator to put a “Cache-control:” header which overrides the caching behavior of proxy caches, and the DNS protocol allows the record creator to add a “TTL” (time to live) field to the record which overrides the caching behavior of DNS resolvers. These features have both been present in these protocols since the mid-1990s.

156. A POSITA, when considering the teachings of Sloss and Lamkin, would have also considered the teachings of Chamberlain. Chamberlain, Sloss, and Lamkin all relate to managing content over a network or the Internet and caching data for optimal efficiency. As such, Chamberlain is analogous art to the ’101 patent.

157. A POSITA would have been specifically motivated to incorporate Chamberlain’s teachings because doing so would have allowed the user more control and specific tailoring of the system when it came to specifying the digital content the user would like to cache. A few reasons why a POSITA would want to tailor a system to specifically cache or not cache certain digital content are: to ensure that content that is subject to frequent change is not held overly long in caches, or to allow the creator to make changes to the content. For example, if a

Web page contains a weather report, the page author may use a "Cache-control:" header to ensure that the page is not held for a period longer than the weather report is valid. In the case of DNS, the record creator may put a short TLL on a DNS record if they record creator intends to change the record on a frequent basis. As disclosed in Chamberlain, "the automatic caching system can be overridden by the information author, the page creator or the system designer," and "the caching system of the present invention further comprises a system for overriding the automatic analysis performed by the system" meaning that, even if (like in Sloss) a content item is non-cacheable, that policy can be overwritten, if the user so wishes. Ex.1010 [Chamberlain], Abstract, ¶18, 142 (noting that this capability "allows maximum flexibility for web site design and implementation"). Combining Chamberlain's teachings would have been nothing more than the combination of prior art elements according to known methods (e.g., combining the teachings of Chamberlain of overriding cache settings within the caching system of Sloss and Lamkin).

158. A POSITA would have also had a reasonable expectation of success in the combination. The combination does not change the intended functionality of Sloss, but rather adds an additional feature of user tailoring with regards to what the system does automatically, like caching. Any modifications to Sloss's methods and software to accommodate Chamberlain's additional information would have

been relatively straightforward to a POSITA, as doing so would have been well-within the level of skill in the art. Accordingly, a POSITA would have found it obvious to apply Chamberlain's teachings to Sloss's methods because the combination merely amounts to applying a known technique to a method ready for improvement.

C. Independent Claim 6

Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are identical to Claims [1.0]-[1.2.2], [1.3.1]-[1.6]

159. *See* Ground 1 above. Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are obvious for the same reasons discussed above for identical Claims [1.0]-[1.2.2], [1.3.1]-[1.6].

[6.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;***

160. This limitation is the only limitation in Claim 6 that differs from Claim 1. The combined teachings of Sloss and Chamberlain render obvious this limitation.

161. Sloss does not explicitly teach that non-cacheable data can be cached upon user permission, i.e., that data that cannot be recovered if a network failure occurs during the transfer from internal storage to network storage, ***can*** be transferred after obtaining permission from a user, but this would have been well-known to a POSITA in the context of Sloss. For example, a POSITA would have

known that many automatic systems can be overridden by user permissions. For example, Meriam Webster defines "override" as "an ability or allowance to correct, change, supplement, or suspend the operation of an otherwise automatic mechanism, system, etc." A POSITA would have known that an automatic system, here, determining whether some data should be cached or not, could be overridden by a user.

162. Chamberlain explicitly teaches this, as it discloses "the automatic caching system can be overridden by the information author, the page creator or the system designer," and "the caching system of the present invention further comprises a system for overriding the automatic analysis performed by the system" meaning that, even if (in Sloss) a content item is non-cacheable, that policy can be overwritten/overridden. Ex.1010 [Chamberlain], Abstract, ¶18.

163. Accordingly, Chamberlain renders obvious "*after obtaining permission from a user*" (Chamberlain's providing of a user overriding the automatic system policy), and resulting in "*the digital contents . . . [being] transferred.*"

D. Dependent Claims 8, 9, and 11

Claims 8, 9, and 11 are identical to Claims 2, 3, and 5, respectively

164. *See* Ground 1 above. Claims 8, 9, and 11 are obvious for the same reasons discussed above for identical Claims 2, 3, and 5.

E. Independent Claim 12

Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are identical to Claims [7.0]-[7.2.2], [7.3.1]-[7.6]

165. *See* Ground 1 above. Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are obvious for the same reasons discussed above for identical Claims [7.0]-[7.2.2], [7.3.1]-[7.6].

[12.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents <u>from the internal storage device to the network storage device is transferred after obtaining permission from a user</u>;***

166. The additional underlined portion is the only limitation in Claim 12 that differs from Claim 7. *See* analysis for [6.2.3].

IX. Ground 4: Claim 10 is obvious under 35 U.S.C. § 103 over Sloss and Lamkin in further view of Chamberlain and Rou.

A. Claim 10 is identical to Claim 4

167. *See* Grounds 2 above. Claim 10 is obvious for the same reasons discussed above for identical Claim 4.

X. Ground 5: Claims 1-4 and 7 are obvious under § 103 over Roden in view of Van Hoff, Ito, and Rathbone

A. Summary of Roden (Ex.1006)

168. The reference that I refer to as “Roden” is U.S. Patent Publication US 2006/0101489. It was filed on November 9, 2004 and published on May 11, 2006. Ex.1006 [Roden], Cover. Barbara J. Roden is the first-named inventor. Ex.1006 [Roden], Cover. I understand that Roden is prior art under § 102(a).

169. Roden teaches "managing content storage and selection services over a network." Ex.1006 [Roden], ¶1.

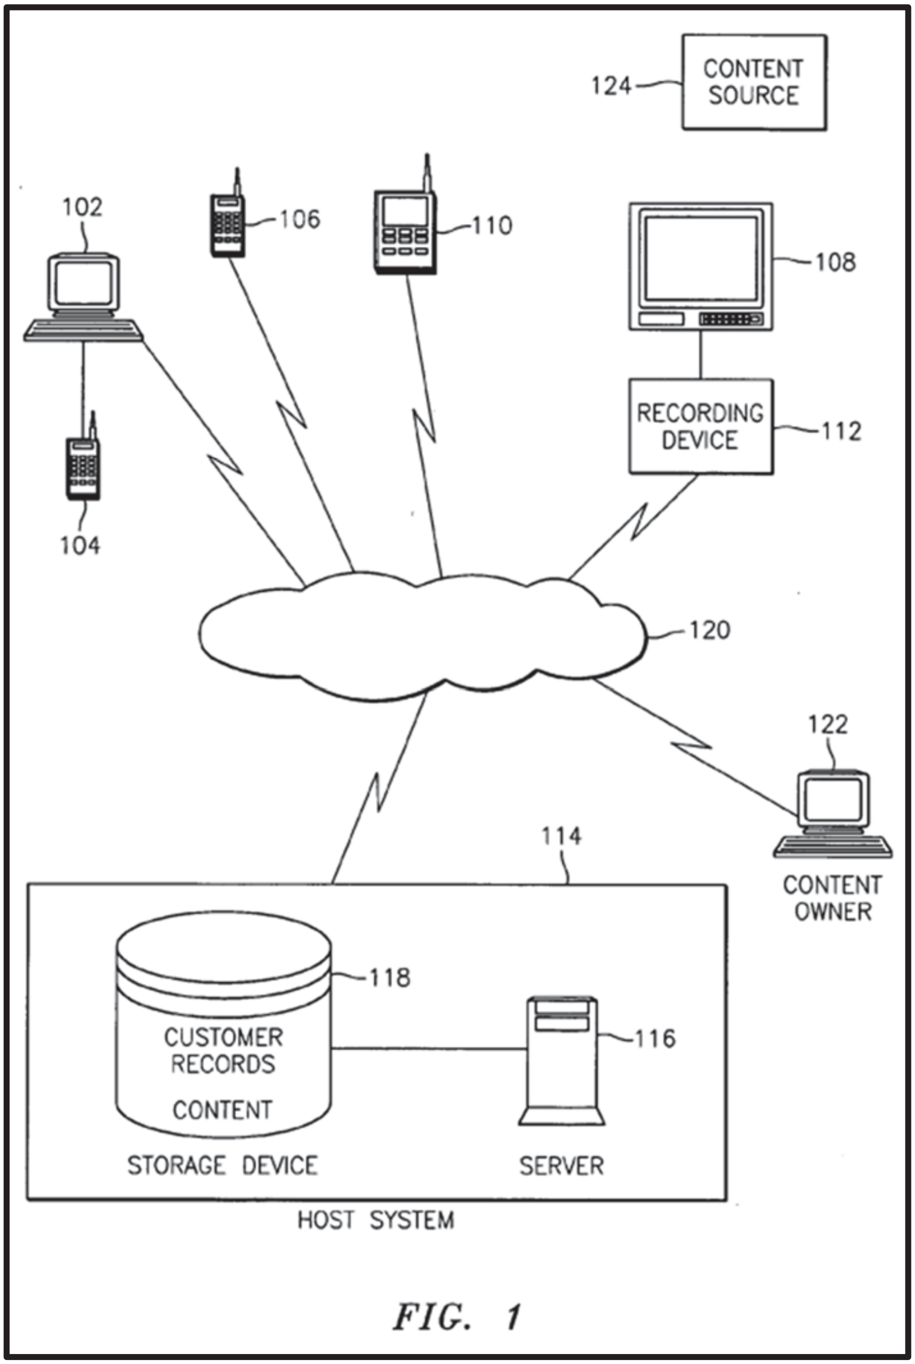

Ex.1006 [Roden], FIG. 1

Roden focuses on the ability for content and storage selection applications to "remove[] selected content from a content device" and "manage[] the off-site network storage system" Ex.1006 [Roden], ¶24. Roden discloses that "[t]he content storage and selection application includes logic that performs functions such as determining what content will be stored locally on a user's content device and what content should be stored on the network (i.e., storage device 118)." *Id*.

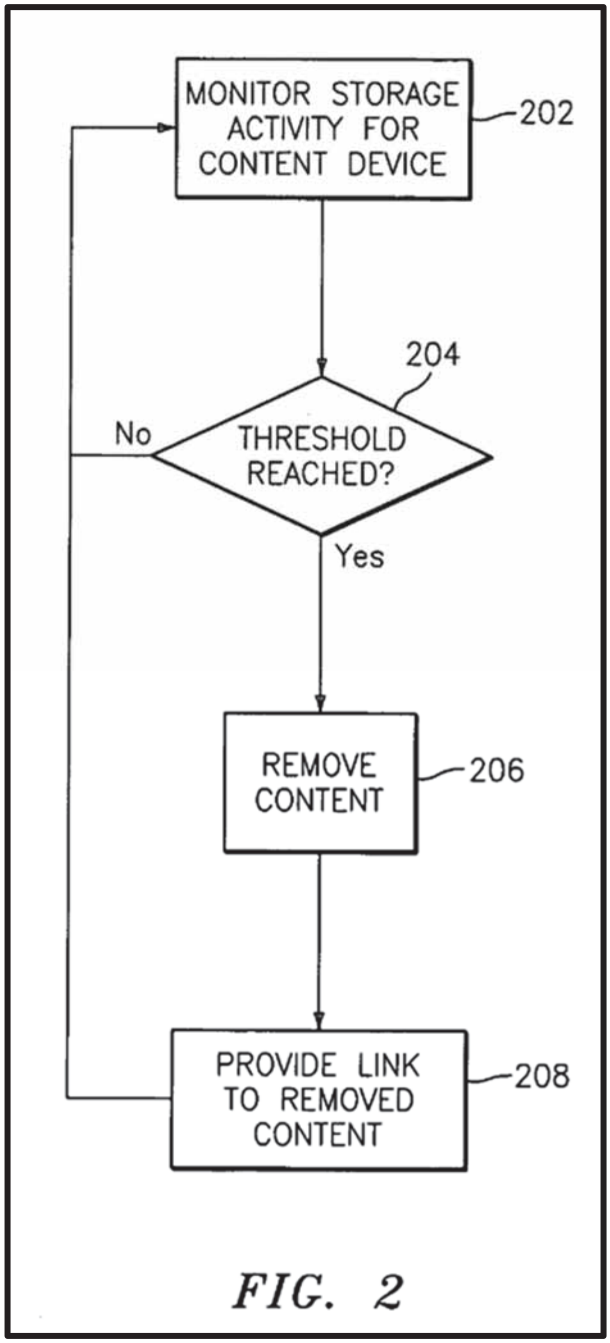

Ex.1006 [Roden], FIG. 2

170. Roden also discloses a user interface that may "present a list of all content in the user's library on a display screen of the content device" and the "listing of content stored locally on the device and the content stored on the

network may be integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user." Ex.1006 [Roden], ¶22. Roden also teaches that in response to "a request to access content stored in storage device 118," "the content storage and selection application grants access to the content." Ex.1006 [Roden], ¶¶35-36. Roden also teaches that customers "may continue to access" media content removed to network storage "e.g., via a download over the network 120) at step 208." Ex.1006 [Roden], ¶34.

B. Summary of Ito (Ex.1008)

171. The reference that I refer to as "Ito" is PCT Application Publication WO2006/073040A1. Ex.1008 (Ito). It was filed on December 12, 2005 and published on July 13, 2006 in Japanese; accordingly, I have reviewed a certified English translation of Ito. Ex.1008 [Ito], Cover. Yoshikatsu Ito is the first-named inventor. Ex.1008 [Ito], Cover. I understand that Ito is prior art under § 102(a).

172. Ito relates to "generating a backup of digital content." Ex.1008 [Ito], ¶1. Like Roden and the '101 patent, Ito describes "transmitting a content to a backup device." Ex.1008 [Ito], Abstract. In Ito's backup process, an HD recorder 100 (similar to Roden's recording device 112) includes a main control unit 118 that performs the backup process. Ex.1008 [Ito], FIGs. 18 and 19.

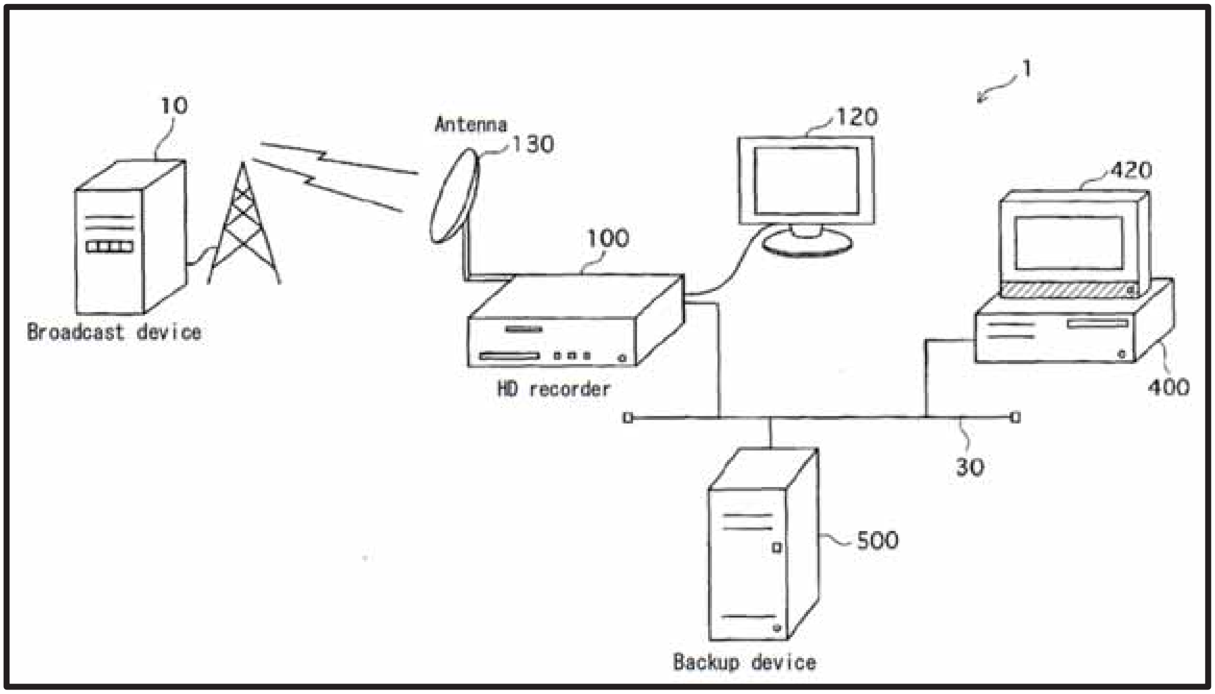

Ex.1008 [Ito], FIG. 1

In one technique of Ito, a backup request is sent from the HD recorder to a backup device 500 (which is analogous to Roden's storage device 118) connected over a network. Ito explains that "main control unit 118 transmits a backup request instructing a backup…via the transmitting and receiving unit 101," but in some instances, "main control unit 118 receives an error notification" because "the backup request cannot be accepted." Ex.1008 [Ito], ¶¶120-121. Subsequently, the backup process ends, and the content is not sent from an HD recorder to a backup device. Ito teaches that the HD recorder 100 may fail to communicate with the backup device "due to a malfunction of the LAN 30" (i.e., a network failure occurs).

C. Summary of Van Hoff (Ex.1007)

173. The reference that I refer to as "Van Hoff" is U.S. Patent 7,895,633. It was filed on November 21, 2005 and issued on February 22, 2011. Ex.1007 (Van Hoff) Cover. Arthur Van Hoff is the first-named inventor. Ex.1007 [Van Hoff], Cover. On its face, it is assigned to TiVo Inc. Ex.1007 [Van Hoff], Cover. I understand that Van Hoff is prior art under at least § 102(e).

174. Van Hoff teaches a "method and apparatus for secure transfer and playback of multimedia content" which "enables the secure transfer of multimedia content from a digital video recorder (DVR) to a personal computer (PC) and further to a handheld device." Ex.1007 [Van Hoff], Abstract. Van Hoff also teaches "securely transferring multimedia content between devices in a computer network." Ex.1007 [Van Hoff], 1:15-17.

175. Van Hoff teaches the PC 101 sends a content request to the DVR 102A at step 310B, and that the "DVR 102A sends the **program stream** to PC 101." Ex.1007 [Van Hoff], 11:3-4, FIG. 1A.

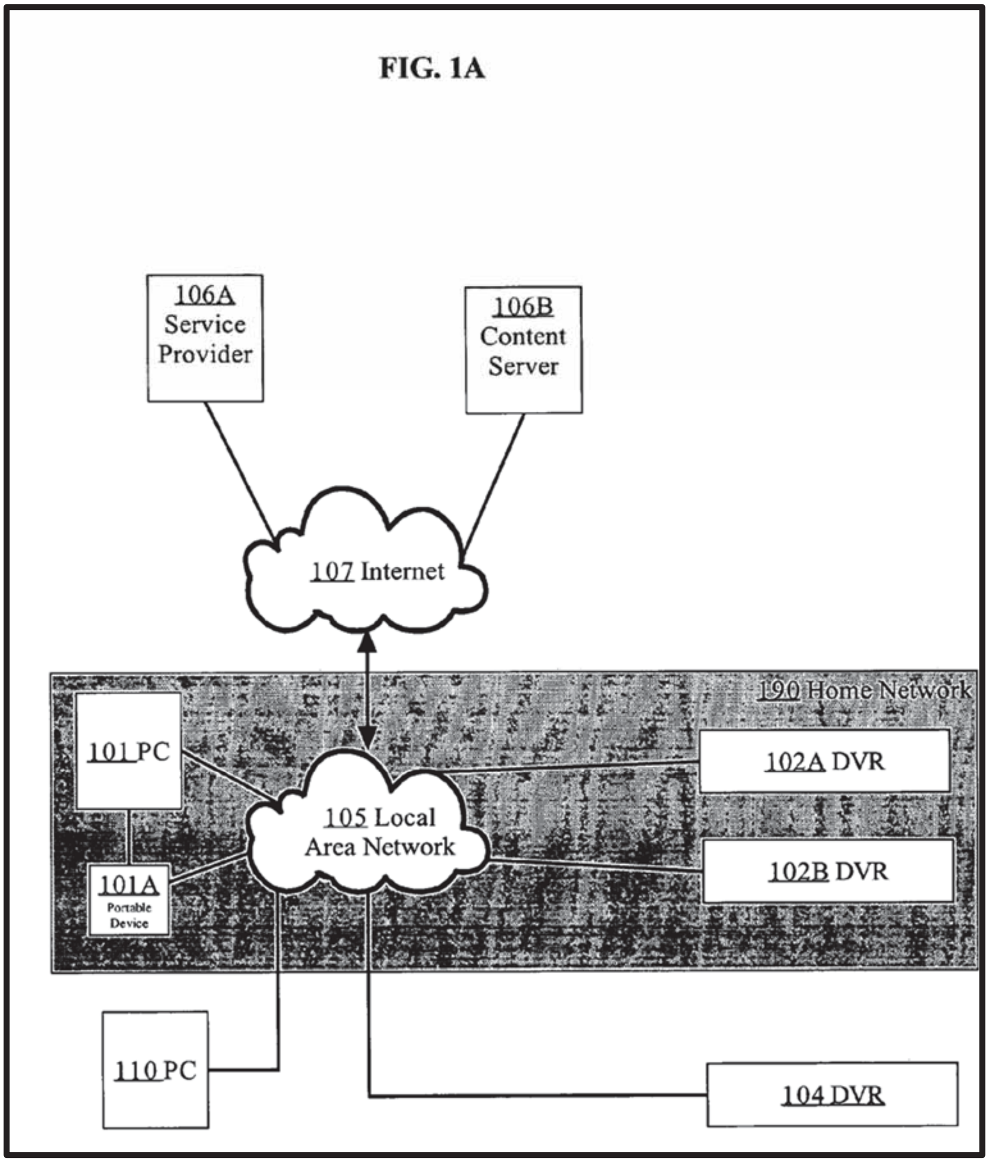

Ex.1007 [Van Hoff], FIG. 1A

176. Van Hoff also teaches "securely transferring multimedia content between a DVR and a PC." Ex.1007 [Van Hoff], 2:20-22. Van Hoff discloses that "at step 310A, using the media application program, the user may select multimedia from the list of available content 358 on DVR 102A. . . The DVR 102A, at step 311, receives the request and locates the content, which may be stored

in the DVR's memory or other storage mechanism accessible by the DVR 102A. Once the DVR 102A has located the requested multimedia content, the DVR 102A begins the process of preparing the content for transfer to PC 101." Ex.1007 [Van Hoff], at 9:35-47. Once the content is converted to a MPEG-2 stream, "[a]t step 319, DVR 102A sends the **program stream** to PC 101. . . After PC 101 has decrypted the program stream, PC 101 may store the program stream onto the PC's local or removable storage device." Ex.1007 [Van Hoff], 11:3-10, FIG. 3A.

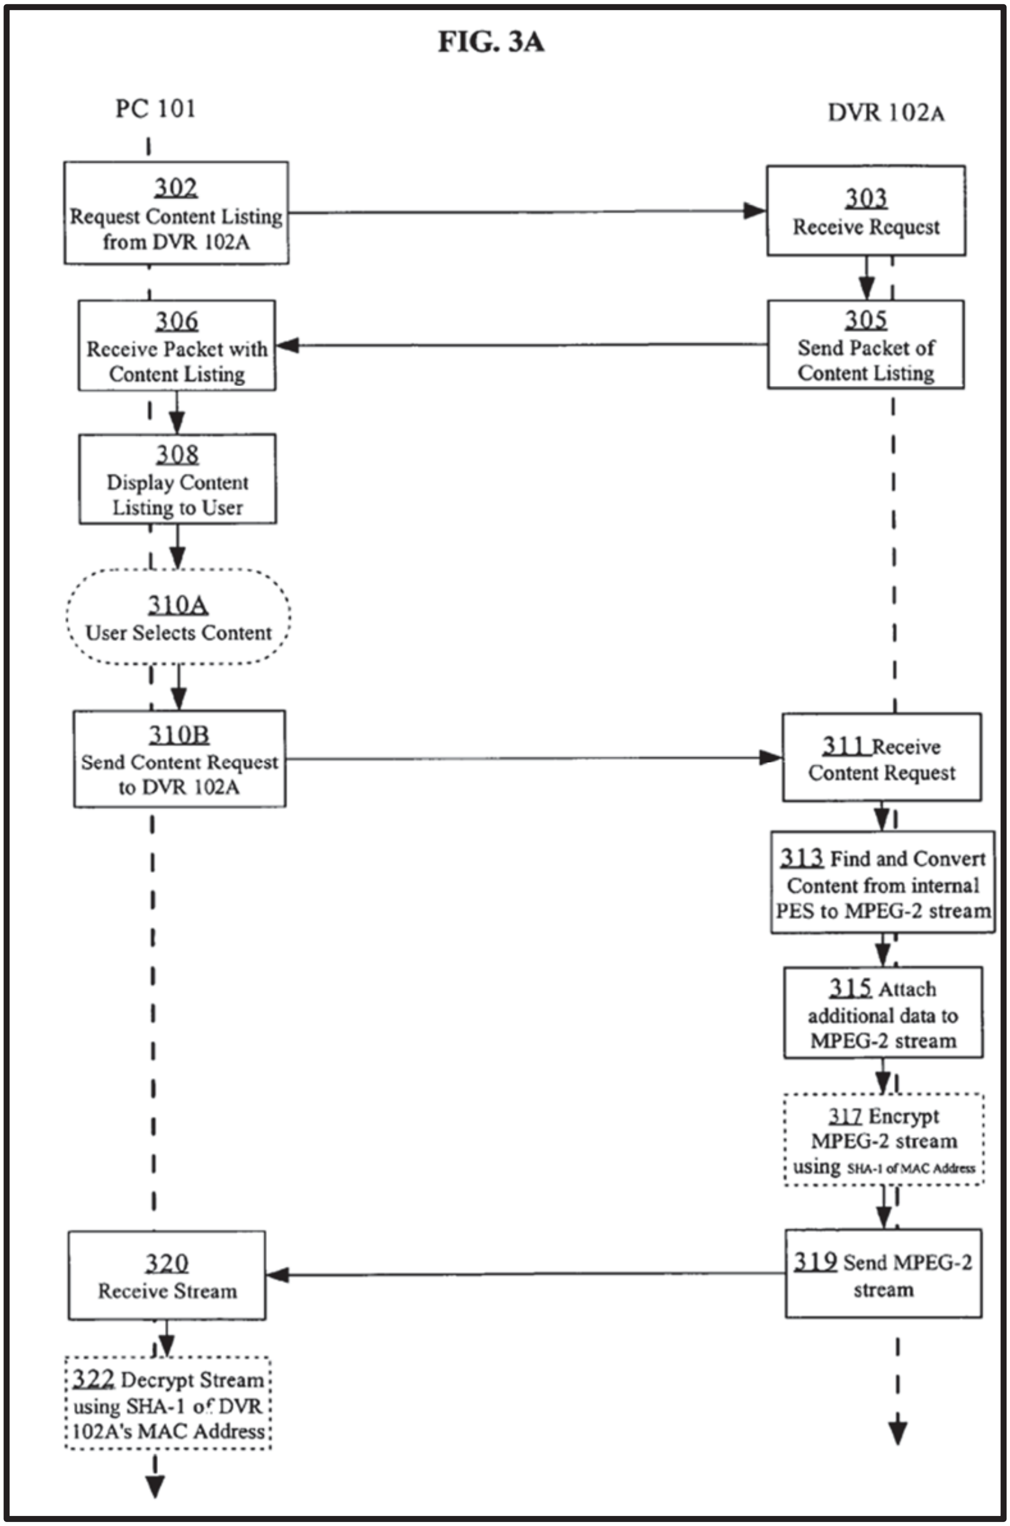

Ex.1007 [Van Hoff], FIG. 3A

D. Summary of Rathbone (Ex.1009)

177. "TiVo for Dummies" is a book written by Andy Rathbone and published in 2004 by Wiley Publishing. The "Dummies" series of books is well-known in at least the computing industry as providing resources and education on potentially complex technical topics to a wide audience, and POSITAs at the time of the '101 patent would have been familiar with books like Rathbone and other books in the series, and would have referred to such books and recommended such books to provide information on various technical topics. I understand that Rathbone qualifies as prior art under § 102(b).

178. Rathbone is a reference about the popular TiVo devices and describes many different features offered by TiVo devices at the time of its publication. As relevant to my analysis below, Rathbone teaches a tree structure (described as a hierarchical folder structure) for displaying and managing digital contents that was available in certain TiVo devices. For example, Rathbone explains that Series 2 TiVos "allow users to group shows into 'folders'" which can then be opened "to see all the recorded episodes of that particular show." Ex.1009 [Rathbone], 280-281.

E. Motivation to Combine

179. For the reasons set forth below, a POSITA would have been motivated to combine the teachings of Ito, Van Hoff, and Rathbone with the teachings of Roden.

180. A POSITA, when considering the teachings of Roden, would have also considered the teachings of Ito, Van Hoff, and Rathbone. All of these references relate to the storage and transfer of media content from various networked storage devices. As such, they are all analogous art to the '101 patent.

181. **First**, it would have been obvious, beneficial, and predictable to combine Van Hoff with Roden because the combination allows for the secure transfer of media content from a TiVo DVR (which is a commercial embodiment of the recording device in Roden) to a personal computer. Ex.1007 [Van Hoff], Abstract. A POSITA would know and would have found it obvious to combine the references because Roden specifically provides that an example of its recording device 112 is "a commercial product such as TiVoTM produced by TiVo Inc. of Alviso, Calif." Ex.1006 [Roden], ¶16. Van Hoff is also assigned to TiVo, Inc., and as such, a POSITA would have recognized its teachings to be applicable to Roden, given Roden's explicit example of a TiVo.

182. It would have also been obvious, beneficial, and predictable to apply Van Hoff's disclosure of stream-delivering content from a DVR to a PC with

Roden's recording device 112 because a POSITA would have understood the combination to result in pulling data from the storage device 118 for streaming to the PC 101, which addresses the problem acknowledged by Van Hoff that "DVR users are forced to view recorded TV programs at a TV set connected to the particular DVR which recorded the program"; Van Hoff's teachings address "DVR users desire to watch recorded TV programs away from the DVR that recorded the TV programs" and thus a POSITA would have recognized an explicit teaching, suggestion, or motivation to combine Van Hoff's teachings. Ex.1007 [Van Hoff], 1:47-57. Although Roden does not explicitly disclose that its connection to a television is a network connection, a POSITA would have known that recording devices, such as Roden's recording device 112, delivered streaming content over a network to other network-connected devices, and in particular, that TiVo devices (a commercial example of recording device 112 in Roden) streamed content over a network to other network-connected devices. Therefore, a POSITA would have known that Roden's recording device 112 could be adapted to stream deliver content from a TiVo to a PC.

183. A POSITA would have also had a reasonable expectation of success in the combination. The combination does not change the intended functionality of Roden, and indeed, would have been present in Roden's example TiVo devices, and explicitly teaches the concept of being allowed to stream content from a DVR

(as disclosed in Roden) to a PC and further to a handheld device. A POSITA would have found it obvious to apply Van Hoff's teachings to Roden because the combination merely amounts to applying a known technique to yield predictable results.

184. **Second**, it would have been obvious, beneficial, and predictable to combine Ito with Roden, where Ito teaches that if a backup request cannot be received due to a network failure, the content is not backed up. Based on these teachings of Ito, when combined with Roden, a POSITA would have understood that, because recording device 112 and storage device 118 in Roden are connected via a network, when a network failure occurs, Roden's recording device 112 would also not be able to save content to the storage device 118. Ito thus merely provides additional details of a content backup process that are analogous to Roden, which would have been well known to a POSITA. Thus, combining the teachings of Ito would have been nothing more than the combining prior art elements according to known methods (e.g., software programming) to yield predictable results, as well as the use of a known technique (not backing up data during a network failure as Ito teaches) applicable to a known device (Roden's recording device) to yield predicable results. A POSITA would have also had a reasonable expectation of success in the combination, as it does not change the intended functionality of

Roden, it simply adds detail on what would occur in Roden if a network failure occurs.

185. **Third**, it would have also been obvious, beneficial, and predictable to combine the teachings in Rathbone with Roden as doing so would have been nothing more than combining prior art elements according to known methods to yield predictable results. Again, Roden specifically identifies as an example of its recording device 112, "a commercial product such as TiVoTM produced by TiVo Inc. of Alviso, Calif." Ex.1006 [Roden], ¶16. Rathbone explains features of TiVo devices, and discloses a hierarchical folder organizational structure of such TiVo devices, which a POSITA would find obvious to combine with Roden's listing feature. Indeed, a POSITA would have expected Roden's devices to possess the "tree" organizational structure described by Rathbone because of Roden's explicit reference to TiVo devices, the subject of Rathbone. Further, tree structures were a well-known feature at the time of the '101 patent. There are only a finite number of ways to keep track of and display where media is stored in a networked media distribution system, and this combination would have been obvious because it represents known potential options with a reasonable expectation of success. A POSITA would have also had a reasonable expectation of success in the combination, as it does not change the intended functionality of Roden, it simply

adds provides additional detail on a commercial implementation of Roden's example device.

F. Independent Claim 1

[1.0] ***A server device for media, the server device for media comprising:***

186. In my opinion, Roden discloses or renders obvious the preamble of claim 1.

187. Roden discloses, as one example, content device 112, referred to as a "recording device 112" which includes "previously recorded program[s] saved on recording device 112." Ex.1006 [Roden], ¶31. Roden also discloses that "storage may apply to **any type of content such as documents, text files, email, images, video, sound, music, programming, and multi-media**. Ex.1006 [Roden], ¶12. Thus, Roden's recording device 112 discloses or renders obvious a "*server device for media.*"

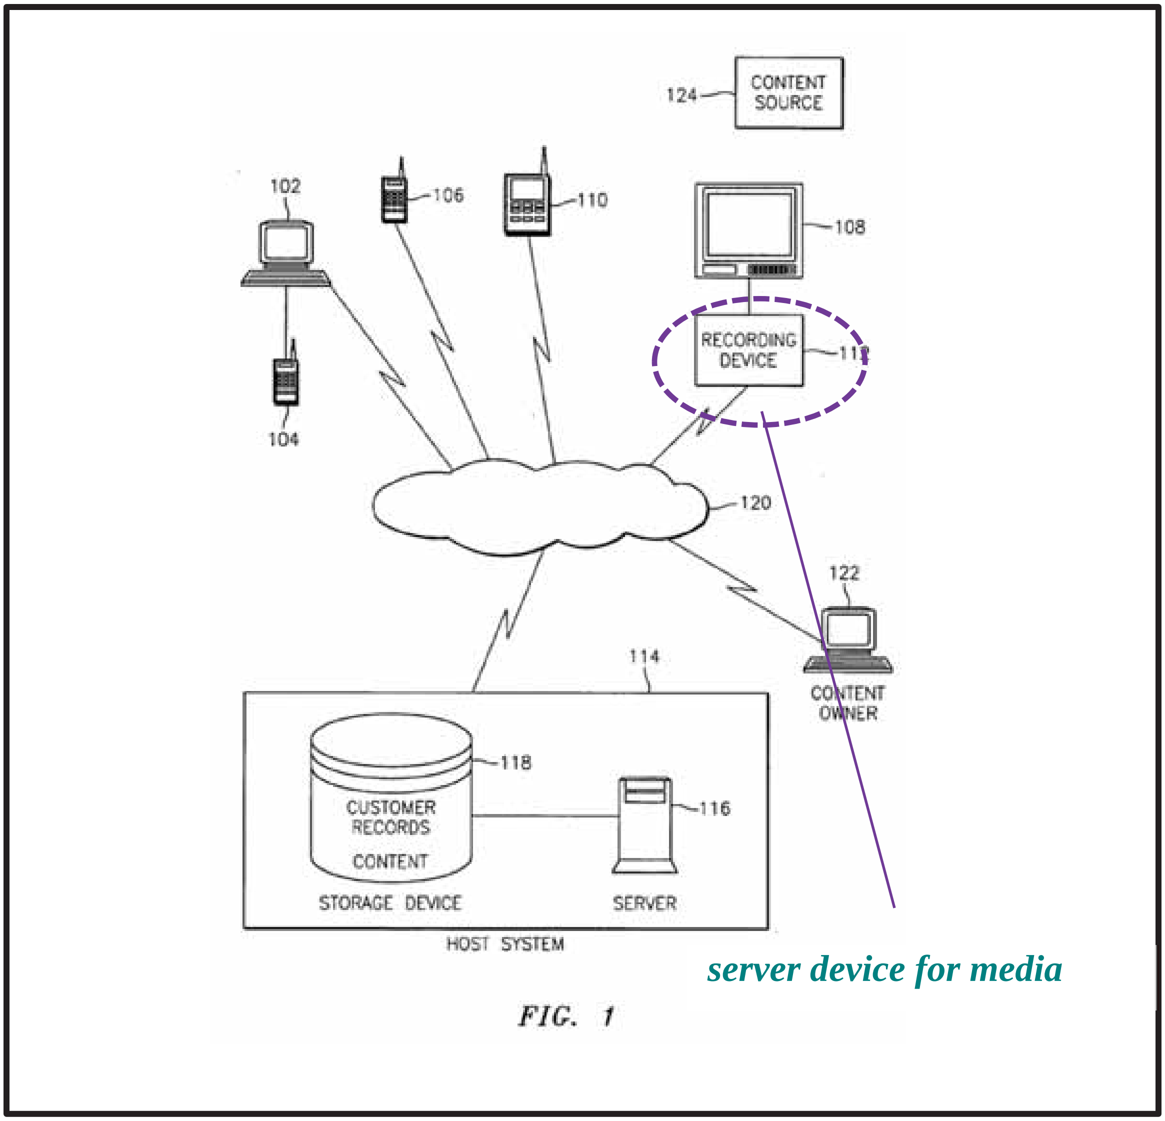

Ex.1006 [Roden], FIG. 1 (annotated)

[1.1] ***an internal storage device for storing digital contents, wherein the server device for media responds to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;***

188. In my opinion, the combined teachings of Roden and Van Hoff render obvious this limitation.

189. First, Roden teaches that its recording device 112 includes "*an internal storage device for storing digital contents*." Specifically, "[r]ecorder device 112 may comprise memory (e.g., an **internal hard drive** or DVD drive) whereby digital programming signals are received over the communications network 120 and **stored in the memory**. Recording device 112 may comprise a commercial product such as TiVo™ produced by TiVo Inc. of Alviso, Calif." Ex.1006 [Roden], ¶16.

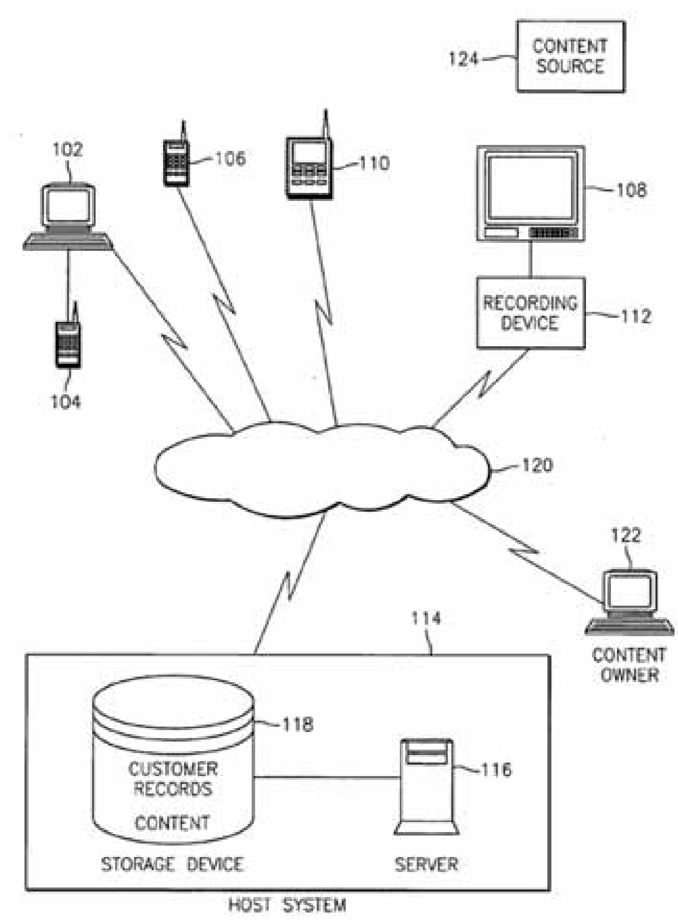

FIG. 1

Ex.1006 [Roden], FIG. 1

190. In Roden, the recording device 112 is coupled to a television 108, but Roden does not explicitly disclose that this connection is a network connection. However, prior to the ’101 patent, it was known that recording devices, such as recording device 112, delivered streaming content over a network to other network-connected devices, and in particular, that TiVo devices (a commercial example of recording device 112 in Roden) streamed content over a network to other network-connected devices. A POSITA would have known that a TiVo had an internal storage that could hold media that could then be streamed over the network to other network-connected devices.

191. Consistent with this knowledge in the art, Van Hoff (assigned to TiVo, Inc.) teaches a “method and apparatus for secure transfer and playback of multimedia content” which “enables the secure transfer of multimedia content from a digital video recorder (DVR)” (analogous to Roden’s recording device 112) “to a personal computer (PC) and further to a handheld device.” Ex.1007 [Van Hoff], Abstract. Van Hoff teaches “securely transferring multimedia content between devices in a computer network.” Ex.1007 [Van Hoff], 1:15-17.

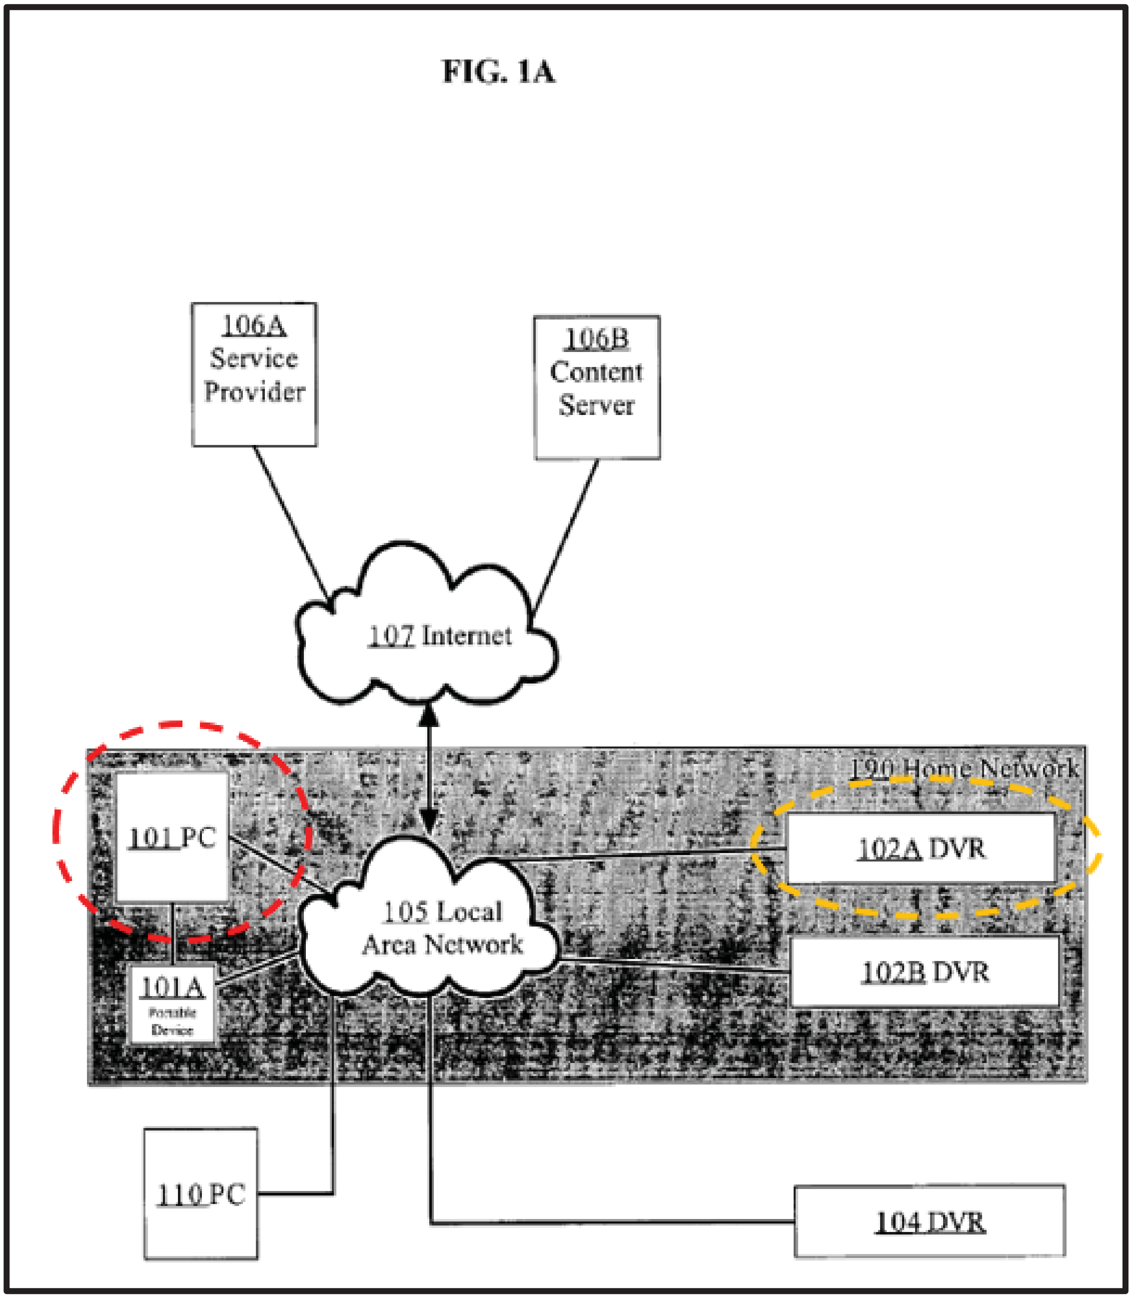

Ex.1007 [Van Hoff], FIG. 1A (annotated)

192. Van Hoff teaches the PC 101 sends a content request to the DVR 102A at step 310B, and that the "DVR 102A sends the **program stream** to PC 101." Ex.1007 [Van Hoff], 11:3-4; *see also* 4:15-31 (devices communicate over the LAN), 5:47-54 (PCs and DVR communicate over LAN), Fig. 3A, 8:57-11:20.

Thus, Van Hoff teaches "*the server device for media responds to a data transmission request from a network player*" (the content request from PC 101) "*by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network*" (the DVR 102A sends a program stream to the PC 101 over LAN 105).

193. A POSITA would know to combine Roden's disclosure of media streamed from a recording device with internal storage to a network-connected device over the network with Van Hoff's disclosure of a DVR sending media to stream to a PC. Thus, the combined teachings of Roden and Van Hoff render obvious this limitation.

[1.2.1] ***a transfer control unit adapted to transfer and store part of held digital contents in the internal storage device to a network storage device,***

194. In my opinion, Roden discloses or renders obvious this limitation.

195. Specifically, in Roden, "[t]he content storage and selection application may include business rules for monitoring and **removing selected content from a content device, managing the off-site network storage system** . . . The content storage and selection application includes logic that performs functions such **as determining what content will be stored locally on a user's content device and what content should be stored on the network** (i.e., storage device 118)." Ex.1006 [Roden], ¶24. The "storage device 118" corresponds to "*a network storage device*" because it is connected to a network and stores data.

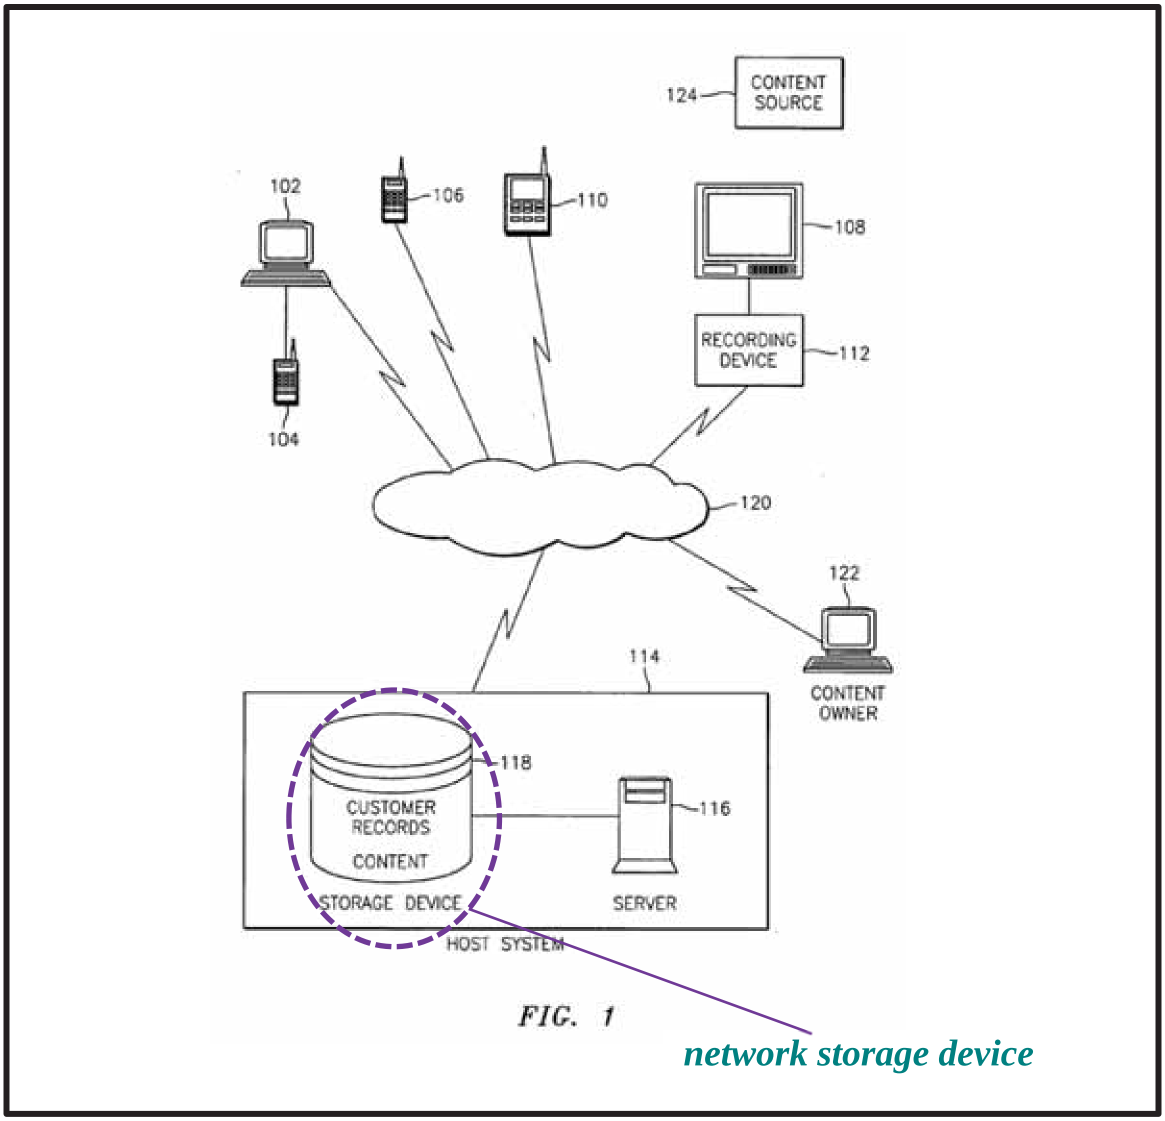

Ex.1006 [Roden], FIG. 1 (annotated)

196. Roden further explains that, in one example, "a content device (e.g., recording device 112) may be purged … when it reaches a pre-determined threshold (e.g., 80%)" and "if the threshold (or other criteria adopted above) has been met at step 204, the content storage and selection application removes content from the content device." Roden continues and states "the content may be saved to the storage device 118." Ex.1006 [Roden], ¶¶32-33. Roden's "content device" or

"recording device 112" corresponds to and renders obvious the recited "*transfer control unit*" at least because a POSITA would have recognized Roden's example of a TiVo as having a processor, memory, and software, consistent with other computing devices. *See, e.g.*, Ex.1009 [Rathbone], p. 25 (noting that TiVo devices had CPUs, RAM, and software). Thus, Roden's description of recording device 112 saving content to storage device 118 renders obvious "*a transfer control unit adapted to transfer and store part of held digital contents in the internal storage device to a network storage device*."

[1.2.2] ***wherein the network storage device is connected to the network and is capable of storing data,***

197. In my opinion, Roden discloses or renders obvious this limitation.

198. In Roden, "[i]n exemplary embodiments, host system 114 includes a **storage device 118 that is in communication with the server 116 over a network** such as a local area network, wide area network, virtual private network, etc. In accordance with exemplary embodiments, **storage device 118 refers to an off-site network storage system that stores content** on behalf of customers." Ex.1006 [Roden], ¶26. Thus, Roden's storage device 118 ("*the network storage device*") "*is connected to the network and is capable of storing data*" as recited.

[1.2.3] ***and wherein said transfer control unit does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device;***

199. In my opinion, the combined teachings of Roden and Ito render obvious this limitation.

200. A POSITA would have understood that, because recording device 112 and storage device 118 in Roden are connected via a network, when a network failure occurs, Roden's recording device 112 would not be able to save content to the storage device 118. And, while Roden discloses saving content from the recording device 112 to the storage device 118, it does not provide many details of this process. However, Ito provides additional details of a content backup process that are analogous to Roden and, as detailed above, a POSITA would have been motivated to combine these teachings with Roden.

201. Ito describes a technology "for generating a backup of digital content" similar to that taught by Roden. Ex.1008 [Ito], ¶1. In Ito's backup process, an HD recorder 100, which is analogous to Roden's recording device 112, includes a main control unit 118 that performs various processes, including a backup process described with reference to Figures 18 and 19.

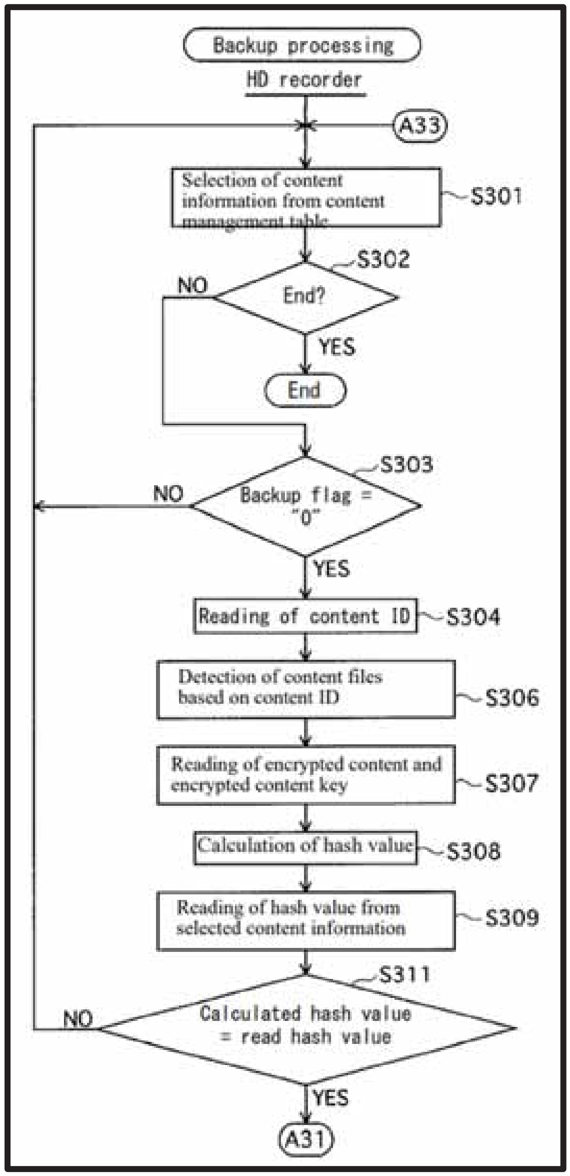

Ex.1008 [Ito], FIG. 18

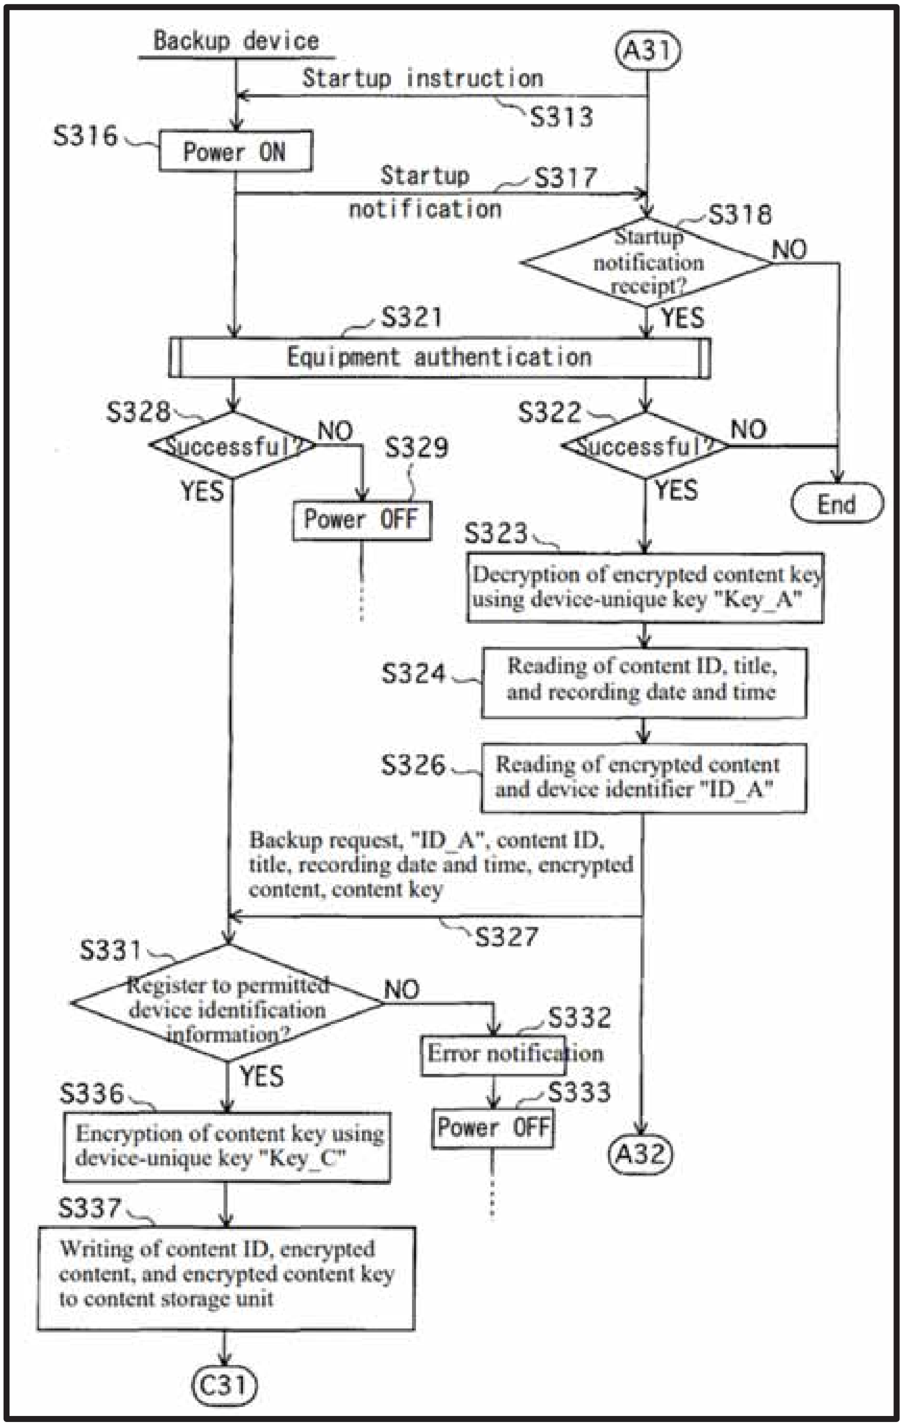

Ex.1008 [Ito], FIG. 19

202. In describing Figure 19, Ito explains a backup request sent from the HD recorder to a backup device 500 (analogous to Roden's storage device 118)

connected over a network. Ito explains that main control unit 118 transmits a backup request instructing a backup, the read device identifier 115 'ID_A', the content ID, the title, the recording date and time, the content key, and the encrypted content to the backup device 500." Ex.1008 [Ito], ¶¶120. However, Ito explains that, in some instances, "main control unit 118 receives an error notification…indicating that the backup request cannot be accepted." Ex.1008 [Ito], ¶¶121. Subsequently, the backup process ends, and the content is not sent from the HD recorder to the backup device. Ito teaches that the HD recorder 100 may fail to communicate with the backup device "due to a malfunction of the LAN 30" (i.e., a network failure occurs). Ex.1008 [Ito], ¶60.

203. That is, the Roden-Ito combination teaches that "*said transfer control unit*" (the recording device 112 of Roden as augmented by Ito's teachings) "*does not transfer, from the internal storage device to the network storage device, the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device*" (if the backup request taught by Ito cannot be received due to a network failure, the content is not backed up on Roden's storage device 118).

204. This is consistent with my understanding of the Patent Owner's interpretation of this claim language in its district court filings. In its allegations of infringement, I understand that the Patent Owner has stated: "Nest Aware allows

Nest Cam's video history to be stored in the cloud automatically. Users can then access and watch the events recorded by Nest Cam on their Google devices, such as tablets, smartphones, or smartwatches. The recorded events are not transferred when there's a power cut or when the Wi-Fi connectivity goes down." *See* Ex.1014 [E-3 chart], 9. In my view, this means that, when data is not transferred to a networked storage device because of a lack of network connectivity, then such a failure to transfer data meets the language of this claim limitation. Because the Roden-Ito combination similarly teaches that data is not transferred because of a network communication failure, then Roden-Ito renders obvious this claim language, when applying Patent Owner's view of the claims.[4]

205. Accordingly, it is my opinion that the combination renders obvious this limitation.

[1.3.1] ***a list information transmission unit adapted to respond to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,***

206. In my opinion, the combined teachings of Roden and Van Hoff render obvious this limitation.

---

[4] Although I set forth my understanding of Patent Owner's allegations in this paragraph, I am not expressing an agreement with Patent Owner's position that this concept is encompassed by the claim language, or that Patent Owner's allegations of infringement are correct in any sense.

207. Roden teaches the recording device 112 "may include a user interface that enables a user to build and browse one or more libraries of content." Ex.1006 [Roden], ¶22. Roden explains that the "user interface" may be accessed in one example by "selections via key depressions" or "web page user interface screen for a computer user" which a POSITA would have understood as requiring a request for the user interface, and thus Roden teaches a "*list information transmission unit adapted to respond to a list presentation request*" as recited. *See also* Ex.1009 [Rathbone], 116 ("Press the TiVo button twice to jump to TiVo's Now Playing list and see your list of recorded shows …."), 25 (CPU, RAM, and software in TiVos).

Ex.1009 [Rathbone], FIG. 6-12

208. Roden further explains that the "user interface may present a list of all content in the user's library on a display screen of the content device" and the "listing of content stored locally on the device and the content stored on the network may be integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user." Ex.1006 [Roden], ¶22. Thus, the user interface provided by Roden ("*transmitting list information to the network player*") includes "*the held digital contents of the server device for media*" (Roden's "content stored locally on the device" as well as the "content stored on the network") as recited, and the request is "*for the held digital contents of the server device for media*" from "*the network player*" (e.g., a PC, as taught by Van Hoff) . For example, Van Hoff teaches that PC 101 ("*the network player*") can "Request Content Listing from DVR 102A" ("*a list presentation request for the held digital contents of the server device for media from the network player*"). Subsequently, the DVR102 sends a "Packet of Content Listing" which is received by PC 101 and displayed. *See also* Ex.1007 [Van Hoff], 8:59-9:9.

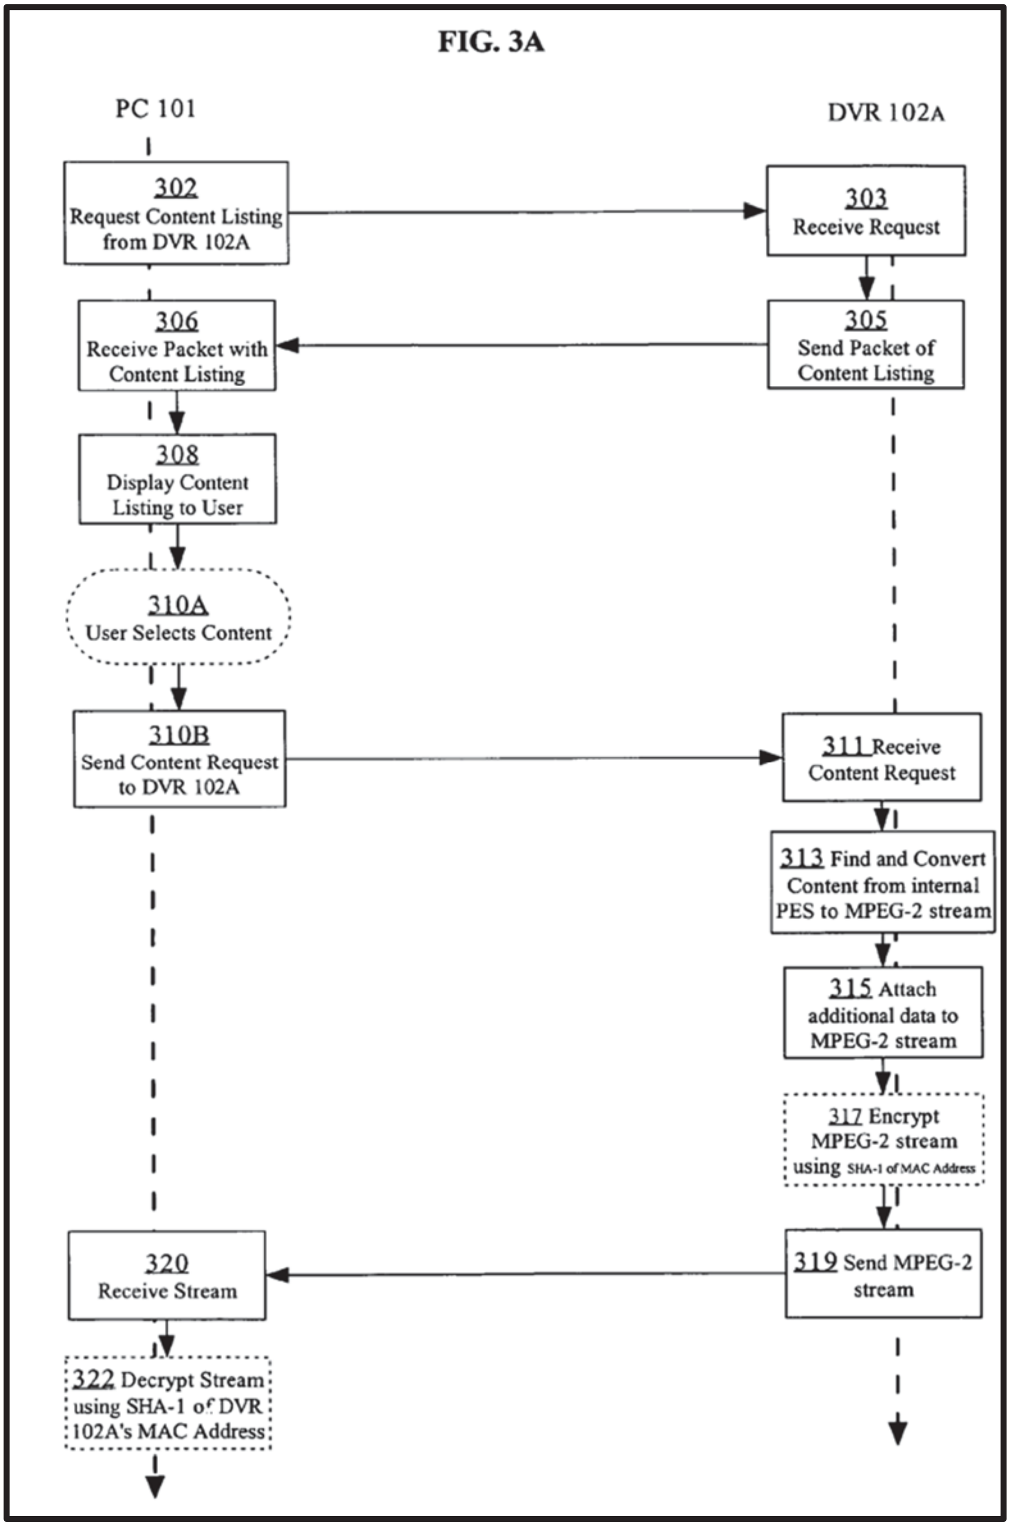

Ex.1007 [Van Hoff], FIG. 3A

209. Additionally, a POSITA would have found it obvious to have a list presentation of a user's media library that can be accessed from the display screen. This feature was already well-known in the art by the time of the '101 patent, as Rathbone shows. Thus, Roden's user interface discloses or renders obvious this limitation.

[1.3.2] ***wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,***

210. In my opinion, Roden discloses or renders obvious this limitation.

211. As detailed in the analysis of limitation [1.3.1], Roden teaches that "[t]he user interface may present a list of all content in the user's library on a display screen of the content device." Ex.1006 [Roden], ¶22. Roden explains that the listing ("*the list information*") may include "content stored locally on the device" ("*the digital contents left in the internal storage device*") "and the content stored on the network" ("*the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device*") "may be integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user." Ex.1006 [Roden], ¶22.

[1.3.3] ***and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;***

212. In my opinion, the combined teachings of Roden and Rathbone renders obvious this limitation.

213. Roden teaches, as detailed above, that the listing of content includes the "*digital contents in the internal storage device*" but does not explicitly describe that the list information maintains "*a tree structure*" of the digital contents. However, this was a well-known feature at the time of the '101 patent. Specifically, as detailed in Roden, an example of its recording device 112 is "a commercial product such as TiVoTM produced by TiVo Inc. of Alviso, Calif." Ex.1006 [Roden], ¶16.

214. A POSITA would have recognized that TiVo devices available at the time of the '101 patent presented recordings on the device in "*a tree structure*." For example, Rathbone explains that "Series 2 TiVos … allow users to group shows into 'folders.'" The book further explains that a user may "[o]pen a folder to see all the recorded episodes of that particular show …" Ex.1009 [Rathbone], 280-281 (Figure 20-5 caption).

Figure 20-4: Series 2 TiVos allow users to group shows into folders.

Ex.1009 [Rathbone], FIG. 20-4

Figure 20-5: Open a folder to see all the recorded episodes of that particular show or Auto-recording WishList.

Ex.1009 [Rathbone], FIG. 20-5

Thus, Rathbone teaches that it was known, in a recording device such as recording device 112, to "*maintain[] a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device*" as recited.[5] This is consistent with other literature about TiVo devices; for example, the 2005 Service Update Guide for TiVo Series2 devices states, "When Groups are on, programs are organized into **folders**" which a POSITA would have recognized as a characteristic of a tree-structured hierarchical file system. Ex.1019 [2005 Service Update Guide], p. 7.

215. Indeed, keeping track of files, or other digital contents, in a "*tree structure*" is a fundamental concept in computing known long-before the '101 patent. As one example, a 1992 book titled "File Structures" defines the term "Filesystem" as "The name used in UNIX to describe a collection of files and directories[6] organized into a tree-structured hierarchy." Ex.1018 [File Structures],

---

[5] As detailed in limitation [1.3.2], Roden's listing may include "content stored locally on the device" as well as content "stored on the network." Ex.1006 [Roden], ¶22. Thus, at all times, Roden's listing includes "*the digital contents in the internal storage device*," whether "*before transferring the digital contents to the network storage device*" or after, and thus Roden and Rathbone renders obvious this limitation irrespective of any temporal requirement.

[6] A POSITA would have recognized that "directories" and "folders" are two terms referring to largely equivalent concepts, and thus, when a document details organization of content into "folders" (e.g., Rathbone, or the TiVo Service Update Guide) such teachings indicate use of a tree-structured file system. *See, e.g.*, Ex.1017 [Microsoft Computer Dictionary], pp. 162 ("directory *n.* … In the

29. The book states: "No matter what computer system you have, even if it is a small PC, chances are there are hundreds or even thousands of files you have access to. To provide convenient access to such large numbers of files, your computer has some method for organizing its files. In UNIX this is called the *filesystem*. The UNIX filesystem is a **tree-structured organization** …." Ex.1018 [File Structures], 22. The book provides a figure of such a tree-structured organization:

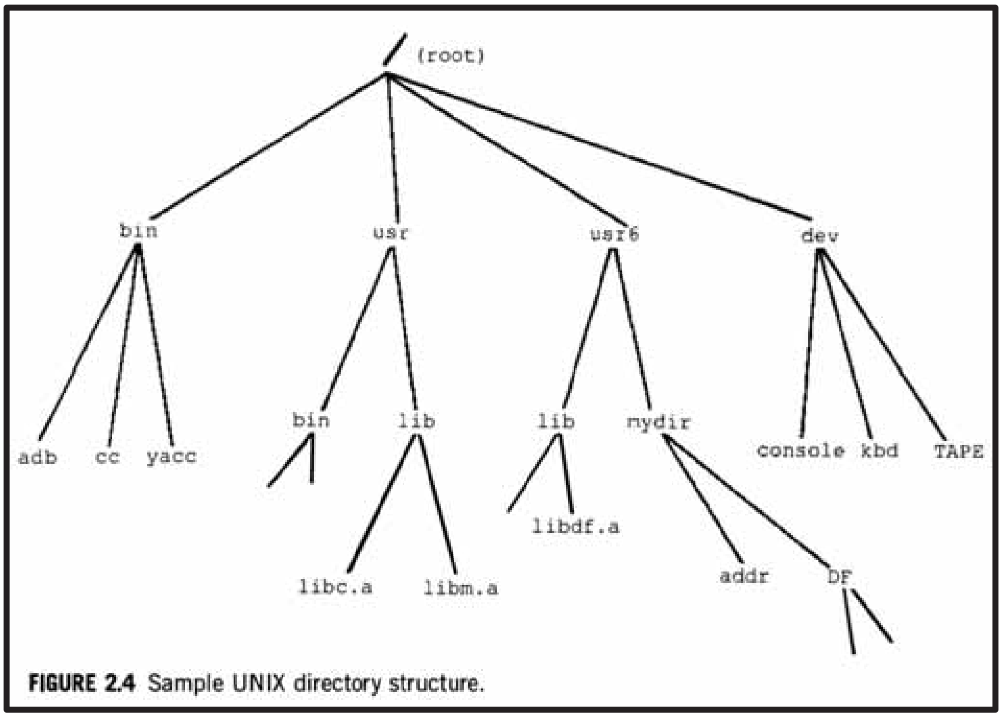

Ex.1018 [File Structures], FIG. 2.4

216. Notably, it was known that the Linux kernel was used for TiVo devices' operating systems (*see* Ex.1020 [Torvalds on TiVo] at 1: "TiVo uses the

---

Macintosh and Windows 9x operating systems, directories are called *folders*."), 220 ("folder *n.* … This container is called a directory in other systems.").

open source Linux operating system in its digital video recorders….”), and Linux was known as a Unix-like operating system. Thus, given Rathbone’s reference to a folder structure, in light of the knowledge of tree structures for file systems in Unix environments, a POSITA would have recognized that Roden’s devices “*maintain[] a tree structure of the digital contents in the internal storage device*” as recited.

217. Accordingly, the combination renders obvious this limitation.

[1.4] ***a search unit adapted to respond to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and***

218. In my opinion, Roden discloses or renders obvious this limitation.

219. Specifically, Roden teaches that “the content storage and selection application associates a link to the removed content within the user library so that the customer may continue to access the content (e.g., via a download over the network 120) at step 208. The demarcation of storage facility sources for content associated with a customer library (e.g., on-site content device storage, off-site network storage) may be imperceptible to the customer.” Ex.1006 [Roden], ¶34. Roden further teaches that “the user interface may present a list of all content in the user's library on a display screen of the content device. The listing of content stored locally on the device and the content stored on the network may be integrated and displayed to the user in such a way that the actual storage location

of the content (e.g., content device 102-112, storage device 118) is transparent to the user." Ex.1006 [Roden], ¶22. Roden also teaches that in response to "a request to access content stored in storage device 118," "the content storage and selection application grants access to the content …" Ex.1006 [Roden], ¶¶35-36; *see also* ¶31 ("searching and access of the content").

220. Thus, Roden teaches to a POSITA that it performs "*searching for a location where the held digital contents are currently stored*" because users can request and access content stored both locally on the recording device 112 as well as on the storage device 118, and the content storage and selection application implemented within the recording device 112 (corresponding to a "*search unit*") grants access to content upon a request (i.e., the Roden device searches for the content and provides the content) and renders obvious "*a search unit adapted to respond to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored*." *See also* Ex.1009 [Rathbone], 25 (CPU, RAM, and software in TiVos).

[1.5] ***a digital contents data transmission processing unit adapted to allow the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,***

221. In my opinion, Roden and Van Hoff render obvious this limitation.

222. As Roden teaches, "the content storage and selection application associates a link to the **removed content within the user library so that the customer may continue to access the content (e.g., via a download over the network 120) at step 208**." Ex.1006 [Roden], ¶34. That is, if content has been moved to the storage device 118, Roden teaches the content is still accessible via download over the network.

223. In combination with the teachings of Van Hoff of stream-delivering content from the DVR to the PC, a POSITA would have recognized that the recording device 112 of Roden, as augmented by the DVR teachings of Van Hoff, would pull data from the storage device 118 for streaming to the PC 101. *See also* Ex.1009 [Rathbone], 25 (CPU, RAM, and software in TiVos).

224. Thus, the combined teachings of Roden and Van Hoff render obvious "*a digital contents data transmission processing unit*" (e.g., Roden's recording device 112) "*adapted to allow the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device*" as recited.

[1.6] ***wherein the server device for media is a media player.***

225. In my opinion, Roden discloses or renders obvious this limitation.

226. Specifically, Roden's recording device 112 "*is a media player*" because content can be played from recording device 112 on television 108:

"recording device 112 refers to a system that communicates with television 108 and/or the associated set top box via a wireline or wireless technology, as well as with external systems such as host system 114 via a service provider network as shown in the system of FIG. 1. **Recording device 112 may receive instructions from a computer user on personal computer 102, from the set top box and/or television 108, and/or from host system 114 regarding which programs will be presented on television 108 or recorded by recorder device 112.**" Ex.1006 [Roden], ¶16. Thus, Roden discloses or renders obvious that the "*server device for media is a media player*" as recited.

G. Dependent Claim 2

[2.1] ***The server device for media according to claim 1, wherein said digital contents data transmission processing unit causes the network storage device to transmit the corresponding data to the server device for media, and then transmits the corresponding data received from the network storage device from the server device for media to the network player.***

227. In my opinion, Ito discloses or renders obvious this limitation.

228. Ito discloses that content offloaded to a backup device can be restored from the backup device: "[i]n the present specification, 'backup' refers to storing a copy of content in a backup device . . . A **recording and playback device** normally uses the content stored in itself, and when the **content stored in the recording and playback device** is lost, the recording and playback device obtains a copy from a backup device and restore the lost content." Ex.1008 [Ito], ¶8.

229. Ito teaches the process to restore data from a backup device. "The recording and playback device of the present invention is further provided with restore instruction acquisition means for acquiring a restore instruction indicating acquisition of the content stored by the backup device, restore request means for transmitting a transmission request for the content stored by the backup device to the backup device, and restoring means for receiving the content from the backup device and writing the received content to the storage means." Ex.1008 [Ito], ¶23, Figs. 16-17. Thus, Ito teaches causing "*the network storage device*" (the backup device) "*to transmit the corresponding data to the server device for media*" (transmitting the backup copy to the recording and playback device, e.g., the Roden recording device 112).

230. Roden and Van Hoff then teach that content is transmitted "*to the network player*" as detailed in limitation [1.1].

231. A POSITA would have incorporated the restoration teaching of Ito because a POSITA would have recognized that in certain circumstances, returning the content from the network storage (e.g., the backup device) to a DVR would have been preferable (e.g., in order to have content available when the network connection is intermittent or busy) rather than to rely on retrieving content from the backup device for playback. This would be beneficial if the user anticipates an interruption in network connectivity, e.g., the user is going offline or out of connection range to a network. The user would be motivated to move the content back to the local device so that it can be played in the absence of connectivity to the network storage device. Accordingly, the prior art renders obvious this claim.

H. Dependent Claim 3

[3.1] ***The server device for media according to claim 1, wherein said digital contents data transmission processing unit transmits the corresponding data and information for identifying the network storage device to the network player, and causes the network storage device to directly transmit the corresponding data to the network player.***

232. In my opinion, Roden in view of Van Hoff discloses or renders obvious this claim.

233. Roden teaches that "the content storage and selection application associates a link to the removed content within the user library so that the customer may continue to access the content." Ex.1006 [Roden], ¶34. A "link to the removed content," as Roden describes, would typically include information like a network location of the storage device (e.g., URL or IP address) along with the specific location within the storage device that the removed content exists (e.g., the sub-folders and file name for the removed content). This would have been well-known at the time of the '101 patent, for example, links on the Internet included similar information for users to retrieve web pages and other Internet-accessible content.

234. Roden's disclosure therefore teaches, to a POSITA, that the DVR sends, to the PC of Van Hoff (e.g., the network player), "*information for identifying*

the network storage device" (e.g., the link). Then, the network storage device directly sends the content to the PC (e.g., network player).

I. Dependent Claim 4

[4.1] ***The server device for media according to claim 1, further comprising a return control unit adapted to cause the digital contents corresponding to a predetermined condition among the digital contents which have been transferred to the network storage device to be returned from the network storage device to the internal storage device.***

235. In my opinion, Ito discloses or renders obvious this claim.

236. Ito discloses returning digital contents "*from the network storage device to internal storage device,*" through its teachings of the process to restore data from a backup device. "The recording and playback device of the present invention is further provided with restore instruction acquisition means for acquiring a restore instruction indicating acquisition of the content stored by the backup device, restore request means for transmitting a transmission request for the content stored by the backup device to the backup device, and restoring means for receiving the content from the backup device and writing the received content to the storage means." Ex.1008 [Ito], ¶23, Figs. 16-17.

237. As I stated above, a POSITA would have been motivated to return content from the network storage (e.g., the backup device) to a DVR in certain circumstances (e.g., in order to have content available when the network

connection is intermittent or busy) rather than to rely on retrieving content from the backup device for playback.

238. It would also have been obvious to a POSITA that a "*predetermined condition*" that would cause the return of digital contents from the network storage device to the internal storage device could be triggered by the user selecting a content item to restore. For example, Ito teaches that "the user selects the restore button 183" and then makes a "selection of the content to be restored." Ex.1008 [Ito], ¶¶175-189.

239. A POSITA would have incorporated the restore teachings of Ito for the same reasons as set forth in the analysis of claim 2.

J. Independent Claim 7

[7.0] ***A method for controlling a server device for media which is equipped with an internal storage device for storing digital contents, the method comprising the steps of:***

240. *See* [1.0].

[7.1] ***responding to a data transmission request from a network player by stream-delivering corresponding data in corresponding digital contents from the internal storage device to the network player during connection to a network;***

241. *See* [1.1].

[7.2.1] ***transferring and storing part of held digital contents in the internal storage device to a network storage device,***

242. *See* [1.2.1].

[7.2.2] ***wherein the network storage device is connected to the network and is capable of storing data,***

243. *See* [1.2.2].

[7.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents are not transferred from the internal storage device to the network storage device;***

244. *See* [1.2.3].

[7.3.1] ***responding to a list presentation request for the held digital contents of the server device for media from the network player by transmitting list information to the network player,***

245. *See* [1.3.1].

[7.3.2] ***wherein the list information lists the digital contents left in the internal storage device and the digital contents transferred from the internal storage device to the network storage device and stored in the network storage device,***

246. *See* [1.3.2].

[7.3.3] ***and wherein the list information maintains a tree structure of the digital contents in the internal storage device before transferring the digital contents to the network storage device;***

247. *See* [1.3.3].

[7.4] ***responding to a data transmission request for the held digital contents from the network player by searching for a location where the held digital contents are currently stored; and***

248. *See* [1.4].

[7.5] ***allowing the corresponding data in held digital contents to be stream-delivered from the network storage device to the network player, if the result of search shows the network storage device,***

249. *See* [1.5].

[7.6] ***wherein the server device for media is a media player.***

250. *See* [1.6].

XI. Ground 6: Claim 5 is obvious under § 103 over Roden in view of Van Hoff, Ito, and Rathbone, and in further view of Lamkin

A. Motivation to Combine

251. For the reasons set forth below, a POSITA would have been motivated to combine the teachings of Roden with those of Lamkin. It would have been obvious, beneficial, and predictable to apply Lamkin's teaching of a list showing where digital contents are stored, including whether digital contents are stored in internal storage or network storage, to Roden's disclosures of network data distribution and provision of a list of content stored in the user's library. This combination would achieve the benefits disclosed by Lamkin, of a "simplified … user interface" in its methods of "managing content," over a "distribution network," as disclosed by Roden and Lamkin. Ex.1005 [Lamkin], Abstract.

252. A POSITA, when considering the teachings of Roden, would have also considered the teachings of Lamkin. Both Roden and Lamkin relate to managing digital content storage over a network. Ex.1006 [Roden], Abstract; Ex.1005 [Lamkin], Abstract. As such, both Roden and Lamkin are analogous art to the '101 patent. *See, e.g.,* Ex.1001, 1:7-12.

253. While Roden also explains that the listing may include "content stored locally on the device and the content stored on the network," which "may be

integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user," Roden does not go into much detail about *how* or whether it identifies if each digital content is currently stored in the internal storage device or the network storage device. Ex.1006 [Roden], ¶22. Further, because Roden states that content listing "**may be** integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user," it contemplates the opposite, i.e., it contemplates providing the user with information indicating the "actual storage location of the content" but does not provide further detail on how this would appear or be accomplished.

254. A POSITA would have recognized at the time of the '101 patent that lists were a routine way to keep track of available digital contents that could be provided to end users and would have been a compatible method with Roden to provide the "actual storage location of the content" as contemplated by Roden. For example, Lamkin explicitly teaches "a simplified example of a user interface 520 according to some embodiments that identifies content added and/or altered since a previous network user access, within a time period, content not previously distributed, and other criteria or combinations of criteria. The user interface 520 can include a listing of content 522." Ex.1005 [Lamkin], ¶87. Lamkin also

discloses that the listing 522 includes information for identifying whether each digital content is currently stored in a source/client device or content that is distributed to and available from another device) in the list. Lamkin, Figure 5. Thus, Lamkin explicitly discloses a user interface that includes a listing of content that is available on various network-connected devices and information on the storage location of such content. Lamkin also provides an explicit motivation for doing so, as it provides a "simplified example of a user interface" and characterizes the invention as "methods, systems and apparatuses for use in managing content" on a network, which would have motivated a POSITA to use its list teachings to manage the digital contents available to a user to provide an easy method of managing such content. Ex.1005 [Lamkin], Abstract.

255. A POSITA would have recognized additional benefits of utilizing a list that specifies the location of content when managing digital contents, as Lamkin teaches, and a networked content storage system in the context of Roden's teachings. For example, in Lamkin's background section, it states that "[t]here has been a drastic increase in the number of consumer electronic devices capable of communicating with one or more computers or other consumer electronic devices." Ex.1005 [Lamkin], ¶4. As such, a consumer will need a way to "manage [that] media content over [said] network." Ex.1005 [Lamkin], ¶2. Thus, Lamkin's teachings would have provided well-known benefits by the time of the '101 patent.

256. A POSITA would have also had a reasonable expectation of success in combining these teachings. Specifically, given the close overlap in subject matter between Roden and Lamkin, a POSITA would have expected success implementing Lamkin's list feature of managing digital contents in Roden, which likewise deals with distributing digital contents across a networked media system ("determining what content will be stored locally on a user's content device and what content should be stored on the network.").

257. Further, a POSITA would have found it obvious to apply Lamkin's teachings to Roden's disclosure of managing media over a network because there was only a finite number of ways to keep track of and display where media is stored in a networked media distribution system. For example, while Roden provides a list where "the actual storage location…is transparent to the user" a POSITA would have recognized the alternative possibility, i.e., showing the user the actual storage location, to be an obvious option, especially given Lamkin, which explicitly teaches such an option. It would have been obvious to a POSITA to use a list to identify and inform the user where each piece of media was stored, in internal storage or network storage. Finally, a POSITA would have reasonably expected success in combining the teachings.

B. Dependent Claim 5

[5.1] ***The server device for media according to claim 1, wherein said list information transmission unit makes the list information to be transmitted to the network player***

258. In my opinion, Roden renders obvious this limitation.

259. Roden teaches that the recording device 112 "may include a user interface that enables a user to build and browse one or more libraries of content." Ex.1006 [Roden], ¶22. Roden further explains that the "user interface may present a list of all content in the user's library on a display screen of the content device" and the "listing of content stored locally on the device and the content stored on the network may be integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user." Ex.1006 [Roden], ¶22.

260. Thus, Roden's user interface discloses or renders obvious a list that is transmitted to the network player.

[5.2] ***include information for identifying whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player.***

261. In my opinion, the combined teachings of Roden and Lamkin render obvious this limitation.

262. Roden teaches that "[t]he user interface may present a list of all content in the user's library on a display screen of the content device." Ex.1006 [Roden], ¶22. While Roden also explains that the listing may include "content

stored locally on the device and the content stored on the network," which "may be integrated and displayed to the user in such a way that the actual storage location of the content (e.g., content device 102-112, storage device 118) is transparent to the user," Roden does not specifically identify whether each digital content is currently stored in the internal storage device or the network storage device, but a POSITA would have recognized it contemplates providing information that indicates the actual storage location of the content. Ex.1006 [Roden], ¶22.

263. Lamkin's user interface includes a listing of content that is available on various network-connected devices. Specifically, Lamkin's "listing 522 includes a listing of one or more content that can potentially be distributed over the local network 121 and/or remote network 140." Ex.1005 [Lamkin], ¶87; *see also* Ground 1, [1.3.2]. As Figure 5 shows, the listing 522 includes "*information for identifying whether each digital content is currently stored in the internal storage device*" (e.g., content that is stored in a **source/client device**) "*or the network storage device*" (e.g., content that is **distributed to and available from another device**) "*in the display list of the network player*":

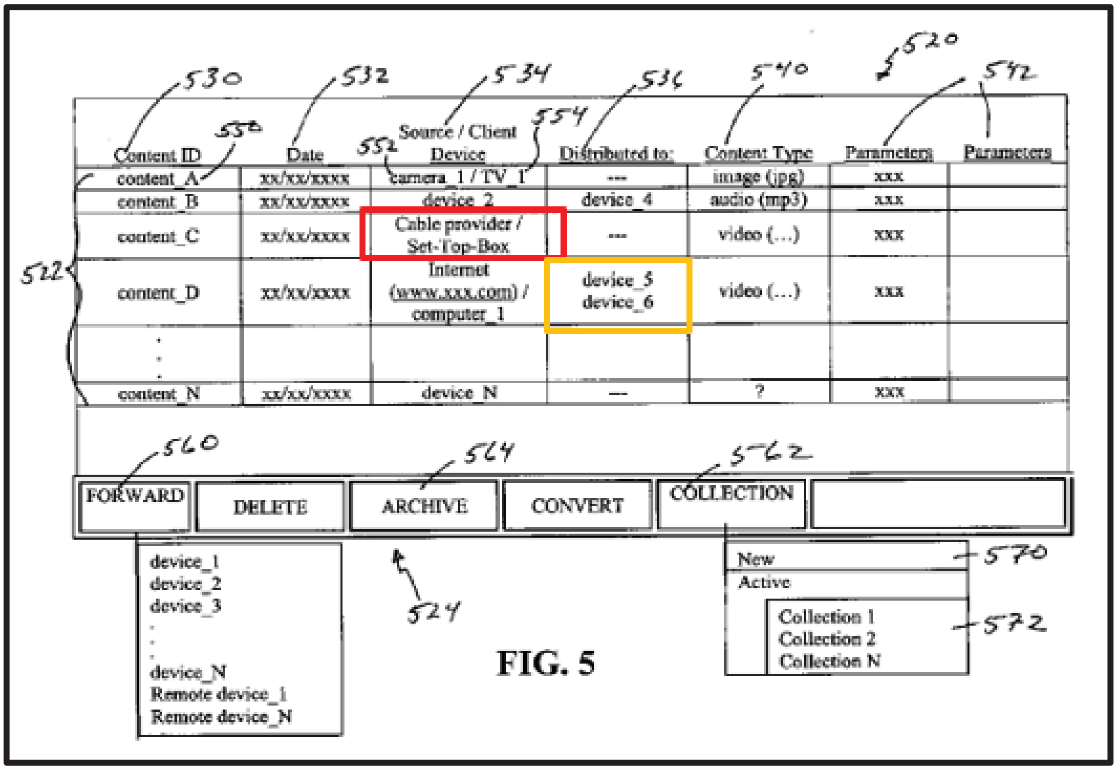

Ex.1005 [Lamkin], FIG. 5 (annotated)

264. Thus, Lamkin's user interface in Figure 5 renders obvious "*whether each digital content is currently stored in the internal storage device or the network storage device in the display list of the network player*" (again, Lamkin's providing of a user interface on the network player with a listing of digital contents, noting which network-connected device each content is located on).

265. Again, it would have been well-known to a POSITA that a list feature displaying where digital contents are stored in a networked media player was routine as of the time of the '101 patent. A POSITA would find it obvious to combine the information listed in Lamkin's Figure 5 with Roden's user interface list rendering this claim obvious.

XII. Ground 7: Claims 6, 8-10, and 12 are obvious under § 103 over Roden in view of Van Hoff, Ito, Rathbone, and in further view of Harris

A. Summary of Harris (Ex.1011)

266. The reference that I refer to as "Harris," is U.S. Patent 5,835,698. It was filed on September 20, 1996, and issued on November 10, 1998. Ex.1011 (Harris) Cover. Jack C. Harris is the first-named inventor. Ex.1011, Cover. On its face, it is assigned to Novell, Inc. Harris Cover.

267. I understand that Harris is § 102(b) prior art.

268. Harris teaches an invention that relates to "systems and methods for preserving data and continuing operations in the presence of network disconnections, and more specifically, to methods and apparatus for unilaterally controlling data and operations where preservation of data and continuation of operations is more important than some time delay of some significant extent." Ex.1011 [Harris], Abstract.

269. Harris also discloses "[o]ptionally, opportunity for user intervention may be provided, at a user's request." Ex.1011 [Harris], Abstract. Harris also discloses that "a user may frequently be presented with certain options. A selection is typically mandatory. Options may be to abort an operation, such as a transfer of data, or the like, to **retry to establish a communication link** or to ignore a failure of communication, attempting to proceed with other operations, if possible." Ex.1011 [Harris], 1:33-38. Harris also discloses "a user may also provide a signal

89 indicating a desire to select either an ignore signal 83 or **a retry signal** 85." Ex.1011 [Harris], 10:21-22.

B. Motivation to Combine

270. For the reasons set forth below, a POSITA would have been motivated to combine the teachings of Harris with the teachings of Roden, Van Hoff, Ito, and Rathbone. It would have been obvious, beneficial, and predictable to apply Harris's teachings regarding a manual override feature of an automatic system to the combination—for example, to provide a more customizable system, when addressing network "time outs" and other network errors.

271. A POSITA, when considering the '101 patent, would have also considered the teachings of Harris. Harris relates to managing content over a network or the internet. As such, Harris is analogous art.

272. A POSITA would have been specifically motivated to incorporate Harris's teachings because doing so would allow the user more control and specific tailoring of the system when a failed backup attempt occurred, meaning a user could retry the backup process irrespective of the earlier failure. As disclosed in Harris, "a user may frequently be presented with certain options. A selection is typically mandatory. Options may be to abort an operation, such as a transfer of data, or the like, to **retry to establish a communication link** or to ignore a failure of communication, attempting to proceed with other operations, if possible"

meaning that a user could override the prior failure and command the system to attempt the backup process again. Ex.1011 [Harris], 1:33-38. Roden and Ito do not disclose such retry procedures, but a POSITA would have been well aware of retrying failed computer operations, as explicitly taught by Harris. Generally speaking, it is standard protocol for an operating system to retry a failed order. This could occur in the context of a failure to establish a communication link, a failure to copy data, a failure to transfer data from one location to another, or a failed backup attempt. In the context of an operation that timed out, it is also well-known in the art that the operating system would retry that command. A POSITA would know that this was standard procedure for operating systems.

273. A POSITA would have also had a reasonable expectation of success in the combination. The teachings of Harris do not change the intended functionality of the combination, but rather adds an additional feature of user tailoring (retrying an operation) of automatic functionality. Any modifications to the combination's methods and software to accommodate Harris's additional information would have been relatively straightforward to a POSITA, as doing so would have been well-within the level of skill in the art. Accordingly, a POSITA would have found it obvious to apply Harris's teachings to the combination's methods because the combination merely amounts to applying a known technique to a method ready for improvement.

C. Independent Claim 6

Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are identical to Claims [1.0]-[1.2.2], [1.3.1]-[1.6]

274. *See* Ground 5 above. Claims [6.0]-[6.2.2], [6.3.1]-[6.6] are obvious for the same reasons discussed above for identical Claims [1.0]-[1.2.2], [1.3.1]-[1.6].

[6.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents from the internal storage device to the network storage device is transferred after obtaining permission from a user;***

275. This limitation is the only limitation in Claim 6 that differs from Claim 1. The combined teachings of Ito and Harris render obvious this limitation.

276. Ito does not explicitly teach that data that cannot be recovered if a network failure occurs during the transfer from internal storage to network storage, ***can*** be transferred after obtaining permission from a user, but this would have been well-known to a POSITA in the failed-backup-attempt context of Ito and prior to the invention. For example, a POSITA would have known that if a backup process fails, a user could retry the operation (e.g., obtaining permission from a user to transfer the data irrespective of the earlier failure).

277. Harris discloses "opportunity for user intervention may be provided, at a user's request." Ex.1011 [Harris], Abstract. Harris also discloses that "a user may frequently be presented with certain options. A selection is typically

mandatory. Options may be to abort an operation, such as a transfer of data, or the like, to **retry to establish a communication link** or to ignore a failure of communication, attempting to proceed with other operations, if possible." Ex.1011 [Harris], 1:33-38. Harris also discloses "a user may also provide a signal 89 indicating a desire to select either an ignore signal 83 or **a retry signal** 85." Ex.1011 [Harris], 10:21-22.

278. Accordingly, Harris renders obvious "*after obtaining permission from a user*" (Harris's providing of a user manually overriding and retrying a pervious failed backup attempt), and resulting in "*the digital contents . . . [being] transferred.*"

D. Dependent Claims 8-10

Claims 8, 9, 10 are identical to Claims 2, 3, 4, respectively

279. *See* Ground 5 above. Claims 8-10 are obvious for the same reasons discussed above for identical Claims 2-4.

E. Independent Claim 12

Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are identical to Claims [7.0]-[7.2.2], [7.3.1]-[7.6]

280. *See* Ground 5 above. Claims [12.0]-[12.2.2], [12.3.1]-[12.6] are obvious for the same reasons discussed above for identical Claims [7.0]-[7.2.2], [7.3.1]-[7.6].

[12.2.3] ***and wherein the digital contents that cannot be recovered if a network failure occurs during the transferring of the digital contents <u>from the internal</u>***

storage device to the network storage device is transferred after obtaining permission from a user;

281. The additional underlined portion is the only limitation in Claim 12 that differs from Claim 7. *See* [6.2.3].

XIII. Ground 8: Claim 11 is obvious under 35 U.S.C. § 103 over Roden, Van Hoff, Ito, and Rathbone in further view of Harris and Lamkin

282. *See* Ground 6 above. Claim 11 is obvious for the same reasons discussed above for identical Claim 5.

XIV. Conclusion

All statements made herein of my own knowledge are true, and all statements made on information and belief are believed to be true. Further, I am aware that these statements are made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001. I declare under penalty of perjury that the foregoing is true and correct.

Executed on July 9, 2025 (date), in Boston, MA (city, state).

Mark Crovella

EXHIBIT A

Mark E. Crovella

Duan Family Fellow
Professor, Computer Science
Professor, Faculty of Computing and Data Sciences

Boston University
665 Commonwealth Ave.
Boston, MA 02215

Current as of June 13, 2025

Education

Ph.D. Computer Science, University of Rochester, October 1994.

M.S. University of Rochester, May 1991.

M.S. State University of New York at Buffalo, May 1990.

B.S. Cornell University, May 1982.

Professional Experience

Professor Department of Computer Science, Boston University, 2006 to present. **Department Chair** 2013 – 2018; **Associate Professor** 2000 – 2006; **Assistant Professor** 1994 – 2000. On leave 2000 – 2001.

Professor Faculty of Computing & Data Sciences, Boston University, 2022 to present. **Duan Family Fellow** 2024 – present; **Chair of Academic Affairs** 2022 – present; **Founding Faculty Member**, 2020 – 2022; **Data Science Faculty Fellow** 2019 – 2022.

Affiliated Faculty Department of Electrical and Computer Engineering, Graduate Program in Bioinformatics, Center for Information Systems and Engineering, and Division of Systems Engineering, Boston University.

Chief Scientist Guavus, Inc., February 2012 to August 2014.

Visiting Faculty Laboratoire d'Informatique de Paris 6 (LIP6); 2003 – 2004 and 2018 – 2019; INRIA Paris 2018 – 2019; Laboratory of Information, Networking and Communication Sciences (LINCS) Paris 2018 – 2019.

Technical Director Network Appliance, Tewksbury, MA. November 2000 to April 2002.

Chief Architect Webmanage Technologies, Chelmsford, MA. May 2000 to acquisition by Network Appliance in November 2000.

Lead Research Scientist epicRealm, Inc., Richardson, TX. January to May 2001.

Senior Computer Scientist Calspan Corporation, Buffalo, NY. August 1984 to August 1994.

Honors

Duan Family Faculty Fellow, 2024, Boston University.

Distinguished Alumni Award, 2017, University at Buffalo, Department of Computer Science and Engineering.

Innovator of the Year, 2014, Boston University.

IETF/IRTF Applied Networking Research Prize for "Routing State Distance: A Path-based Metric for Network Analysis," ACM Internet Measurement Conference (IMC), 2012.

Best Student Paper Award for "Intrusion as (Anti)social Communication: Characterization and Detection," ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2012.

IEEE Fellow, 2011, "For contributions to the measurement and analysis of networks and distributed systems."

ACM Fellow, 2010, "For contributions to the measurement and analysis of networks and distributed systems."

ACM SIGMETRICS Test of Time Award for "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," (with Azer Bestavros). One of the three inaugural Test of Time Awards given by SIGMETRICS in 2010. The award is given for a paper published between 1973 and 1999 that has had the greatest impact over the subsequent period.

CNRS Fellowship, Centre National de la Recherche Scientifique (Department Sciences et Technologies de l'Information et de la Communication), for 2003-2004 academic year as Visiting Associate Professor at Laboratoire d'Informatique de Paris 6 (LIP6).

NSF CAREER Award, 1995.

ARPA Fellowship in High Performance Computing, 1993-1994.

DARPA Parallel Processing Fellowship, 1992-1993.

EDITORIAL BOARDS

Steering Committee, IEEE Transactions on Network Science and Engineering, 2016–2019.

Editor, Computer Communication Review, 2005 – 2008.

Editor, Computer Networks, 2004 – 2005.

Editor, IEEE/ACM Transactions on Networking, 2000 – 2006.

Editor, IEEE Transactions on Computers, 2000 – 2004.

FUNDING

"Large Scale Analysis of Configurations and Management Practices in the Domain Name System," NSF CISE/CNS NeTS Medium, 10/1/23 – 9/30/26, $525,000.

"Methods for Active Measurement of the Domain Name System," NSF CISE/CNS IMR, 10/1/23 - 9/30/26, $256,463.

Google Research Award, 10/1/20, $30,000.

"Traveling Back to the Future: Using Network Science to Unravel Evolutionary Conserved RNA Binding Proteins and Complexes Associated with Disease," BU Hariri Institute Research Incubation Award, 2/9/18, $35,280.

"Analytic Tools for Evolving Path-based Networks," NSF CISE/CNS NeTS Small, 8/15/16 – 7/31/19, $499,951.

"Structural Matrix Completion for Data Mining Applications," NSF CISE/III Small, 9/1/14 – 8/31/17, $499,702 with 1 other PI.

"Understanding Communication Strategies for Ad Hoc Networks," NSF CISE/CNS NeTS Small, 8/1/11 – 7/31/14, $390,853.

"Information and Software Assurance: A Coordinated Approach to Cyber-Situation Awareness Based on Traffic Anomaly Detection,", DoD ARO, 9/1/2011 - 8/31/2015, $800,000 with 3 other PIs.

"Securing the Open Softphone," NSF CISE/CNS TC Large, 7/1/10 – 6/30/15, $2,992,896, with 8 other PIs. REU Supplement, 4/13/11 – $73,125.

"New Directions in Network Dimensionality Reduction for Routing and Beyond," NSF CISE/CNS NeTS Small, 7/1/10 – 6/30/13, $450,000.

"Wide-Aperture Traffic Analysis for Internet Security," NSF CISE/CNS TC Medium, 9/1/09 – 8/31/13, $723,053, with 2 other PIs.

Sponsoring Advisor for Postdoctoral Fellow under the *Computing Innovations Fellowship Program,* NSF CIF, 9/1/09 – 8/31/11, $267,500.

"Passive Methods for Internet Topology Discovery," National Geospatial Intelligence Agency, 8/1/09 – 7/31/12, $360,000.

"Instrumentation and Measurement for GENI," NSF/BBN, 8/1/08 – 8/31/11, $80,932

"Whole-Network Anomaly Diagnosis," Intel, 1/1/05 – 12/31/07, $158,000.

"Coordinate Systems for the Internet," NSF CISE/ANI, 8/15/03 – 7/31/06, $222,035.

"Collaborative Research: Modular Strategies for Global internetwork Monitoring," with 6 other PIs, NSF CISE/CCR ITR, 9/1/03 – 8/31/08, $1,994,513.

"Research Support for IP Network Characterization," Sprint Advanced Technology Laboratories, 9/1/03 – 8/31/04, $29,592.

"Research Support for IP Network Characterization," with 3 other PIs, Sprint Advanced Technology Laboratories, 9/1/02 – 8/31/03, $80,000.

"A Control Theoretic Approach to the Design of Internet Traffic Managers," with 3 other PIs, NSF CISE/ANIR, 9/1/01 – 8/31/05, $719,994.

"Diagnosis and Control of Network Variability by Massively Accessed Servers," with 2 other Principal Investigators, NSF CISE/ANI, 9/1/00 – 8/31/05, $1,210,768.

"Wide Area Web Measurement," GTE Laboratories, 9/1/00 – 8/31/01, $20,000.

"Wide Area Web Measurement," Xerox PARC, 9/1/99 – 8/31/00, $15,000.

"COMMONWEALTH: Architecture and Protocols for Scalable WWW Service," with 2 other Principal Investigators, NSF CISE/CCR Experimental Software Systems, 7/15/97 – 6/30/00, $590,087.

"Research Support for Network Performance Characterization, Analysis, and Evaluation," Principal Investigator, Hewlett-Packard Company, 9/1/96 – 9/1/97: $69,404. Renewed for 9/1/97 – 9/1/98: $88,154. Renewed for 9/1/98 – 9/1/99: $29,873.

"Practical Performance Prediction for Parallel Programmers," Principal Investigator, NSF CISE/CCR CAREER, 9/1/95 – 5/31/98, $91,066

"Research Infrastructure for Parallel and Distributed Systems: Real-Time, Multimedia, and High-Performance," with 4 other Principal Investigators, NSF CISE Institional Infrastructure, $874,019

"Research Instrumentation for Real-Time, Multimedia and High-Performance Computing in Distributed Systems," with 4 other Principal Investigators, NSF CISE IGR, $142,744

"Graduate Assistantships in High Performance Computing," with 5 other Principal Investigators, US Dept. of Education, 9/1/95 – 8/31/98, $354,645

Professional Affiliations

Association for Computing Machinery (Fellow)
IEEE (Fellow)
ACM SIGCOMM and SIGMETRICS
IEEE Computer Society and Communication Society

Invited Conference Talks

Invited Speaker, *TMA Experts Summit,* Paris, June 18, 2019.

Invited Panelist, *The 3rd Future Network Development Conference,* Nanjing, China, May 22, 2019.

Keynote Speaker, *IEEE Workshop on Data Science for Networking (DS4N),* Boston, MA, December 14, 2017.

Invited Panelist, *CSE 50th Anniversary Celebration,* University at Buffalo, Department of Computer Science and Engineering, September 30, 2017.

Keynote Speaker, *DATANET 16: Workshop on Data-Driven Networking,* Arizona State University, May 13, 2016.

Invited Speaker, *Workshop on Mathematics of Data Analysis in Cybersecurity,* ICERM, Brown University, October 24, 2014.

Invited Speaker, *51st Annual Allerton Conference on Communication, Control, and Computing,* University of Illinois at Urbana-Champaign, October 3, 2013.

Invited Speaker, *Network Links: Connecting Social, Communication and Biological Network Analysis,* Institute for Mathematics and its Applications (IMA), Minneapolis, MN, February 29, 2012.

Invited Speaker, *Statistics of Networks Workshop,* Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, June 24-25, 2010.

Keynote Speaker, *NetEval 2009 Workshop, The First Workshop on Performance Evaluation of Next-Generation Networks,* Boston, MA, April 26, 2009.

Keynote Speaker, *2006 Euro-NGI Workshop on Visions of Future Generation Networks (EuroView2006),* Würzburg, Germany, August 1, 2006.

Invited Speaker, *Maryland Hybrid Networks Center,* Advanced Networks Colloquium Series, University of Maryland, December 16, 2005.

Keynote Speaker, *Workshop on End-to-End Monitoring Techniques and Services (E2EMON),* Nice, France, May 15, 2005.

Invited Speaker, *Workshop on Internet Measurement, Modeling, and Analysis,* Statistical Research Center for Complex Systems (SRCCS), Seoul National University, Seoul, South Korea, January 11, 2005.

Invited Speaker, *Workshop on Measurement, Modeling and Analysis of the Internet,* Institute for Mathematics and its Applications (IMA), Minneapolis, MN, January 12, 2004.

Keynote Speaker, *First Workshop on Performance of Computational Systems and Communication,* part of the XXII Congress of the Brazilian Computer Society, Florianopolis, Brazil, July 18, 2002.

Invited Speaker, Institute for Pure and Applied Mathematics (IPAM) Program on Large Scale Communication Networks, Los Angeles, CA, March 18-22, 2002.

Invited Speaker, *DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling,* New Brunswick, NJ, February 13-15, 2002.

Invited Speaker, *Workshop on Job Scheduling Strategies for Parallel Processing,* Cambridge, MA, June 16, 2001.

Invited Speaker, *COST 257 Final Conference,* Würzburg, Germany.

Invited Plenary Speaker, *Performance Tools 2000:* The 11th International Conference on Modelling Techniques and Tools for Computer Performance Evaluation, Schaumburg, Illinois, March 27 - 31, 2000.

Invited Speaker, *Institute for Mathematics and Its Applications Workshop: "Scaling Phenomena in Communication Networks,"* Minneapolis, MN, October 22-24, 1999.

Invited Speaker, Boston University Minority Engineers' Society, 2nd Annual Technical Exposition, Boston, MA, April 10, 1999.

Invited Speaker, Workshop on Intelligent Internet Infrastructure, IBM T.J. Watson Laboratory, Yorktown Heights, NY, March 22, 1999.

Invited Speaker, Workshop on Architectures, Performance, and Applications in Globally Distributed Web Services, Federal University of Minas Gerais (UFMG), Belo Horizonte, Brazil, March 11 - 12, 1998.

Invited Speaker, Long-Range Dependence and Self-Similar Network Traffic Session, Winter Simulation Conference, Atlanta, GA, Dec. 7-10, 1997

Invited Speaker, DIMACS/NSF Workshop on "End-to-End Network Modeling and Simulations," Princeton University, Princeton, NJ, October 23-25, 1997.

Invited Panelist, Hot Topic Session On Traffic Characterization, ITC-15, Washington DC, June 22-27, 1997

Invited Speaker, IFIP Working Group 7.3 Workshop, June 19-20, Orcas Island, WA.

Invited Panelist, "Engineering Implications of Long-Range Dependence in Teletraffic Data," MASCOTS '97, January 1997, Haifa, Israel

Invited Speaker, Applied Probability Session of Informs, the Operations Research Society, May 1997, San Diego, CA.

Invited Speaker, Joint Research Conference on Statistics in Quality, Industry and Technology, May 29-31 1996, NIST, Gaithersburg, MD. Sponsored by the *American Statistical Association*, the *Institute of Mathematical Statistics*, E. I. du Pont, and the *National Institute of Standards and Technology.*

Professional Activities

Chair, Scientific Advisory Board of the Max Plank Institute for Software Systems, 2024.

Scientific Advisory Board of the Max Plank Institute for Software Systems, 2016 – 2025.

Member, IETF/IRTF Applied Networking Research Prize Selection Committee, 2014 – 2016.

Executive Steering Committee of the Hariri Institute for Computing and Computational Science & Engineering, 2011 – 2014, 2024 – present.

Member of the External Review Visiting Committee for the Department of Computer Science at Amherst College, February 2013.

Organized Workshop: "GPU Programming with CUDA" at Boston University, July 24, 2012.

Member, SIGCOMM Rising Star Award Selection Committee, 2011.

Chair, SIGCOMM Test of Time Award Selection Committee, 2011.

Team Mentor for Conrad Awards, Cyber Security Challenge 2011.

External Judge for FCC's Open Internet Apps Challenge, 2011.

General Co-Chair, CoNEXT 2011.

Internet2 perfSONAR Workshop, Executive Committee, 2010.

Chair, ACM SIGCOMM 2007–2009.

Program Co-Chair, IFIP Networking 2010.

Internet2 Network Research Strategic Planning Team, 2009-present.

Internet2 Network Research Review Committee, `http://www.internet2.edu/networkresearch/nrrc.html`, 2009-present.

Advisory Board of CRAWDAD (Community Resource for Archiving Wireless Data At Dartmouth).

Steering Committee, Passive and Active Measurement Workshop (PAM), 2004-2006.

General Chair, Passive and Active Measurement Workshop (PAM) 2005.

Steering Committee, Internet Measurement Conference, 2003-2007.

ACM SIGCOMM Executive Committee, Award Committee Chair, 2003-2006. b

Program Comittee Chair, Internet Measurement Conference 2003.

Co-organizer (with C. Lindemann and M. Reiser), Workshop on Internet Performance Modeling, Dagstuhl Castle, Germany, September 26-29, 1999.

Co-organizer (with five others), BU/NSF Workshop on Internet Measurement, Instrumentation, and Characterization; Boston University, August 30, 1999.

Tutorials Chair, SIGMETRICS '99, Atlanta GA.

Program Committees:

- SIGCOMM: 2002, 2006, 2008, 2012, 2017
- SIGMETRICS: 1999, 2000, 2002, 2007, 2009, 2010, 2011, 2012, 2014, 2015, 2016
- CoNEXT: 2005, 2007, 2010, 2011

- Internet Measurement Conference: 2002, 2004, 2007, 2010, 2012, 2015, 2017, 2018, 2020, 2022
- ICDM: 2018
- The Web Conference: 2000, 2001, 2002, 2019, 2022 (Senior PC Member)
- KDD: 2017, 2023
- Intelligent Systems for Molecular Biology (ISMB): 2015, 2017
- First IMC Workshop on Internet Visualization: 2012
- Performance: 2011
- Passive and Active Measurement Workshop (PAM): 2009, 2010
- Infocom: 2009 (TPC Area Chair), 2010 (TPC Area Chair)
- IFIP Networking: 2009
- NetEval: 2009
- HotMetrics: 2008
- E2EMON: 2007
- First Workshop on Experimental Computer Science: 2007
- Workshop on Mining Network Data (MineNet-06): 2006
- Workshop on Large Scale Network Inference: 2005
- 6th Global Internet Symposium: 2001
- 6th International Workshop on Web Caching and Content Distribution: 2001
- Workshop on Caching, Coherence and Consistency: 2001
- Workshop on Network-Related Data Management: 2001
- ITC Specialist Seminar on IP Traffic Management and Modeling: 2000
- SPIE Performance and Control of Network Systems Conference: 1998, 1999
- 18th International Conference on Distributed Computing Systems (ICDCS-98): 1998
- Third IEEE Workshop on Services in Distributed and Networked Environments (SDNE '96): 1996

COURSE TEXTS AND MATERIALS

"Linear Algebra, Geometry, and Computation," Electronic Text, one-semester course on Linear Algebra from a Computer Science standpoint. 2015-2025. Available at `http://mcrovella.github.io/CS132-Geometric-Algorithms`. Includes contributions from Wayne Snyder and students.

"Computational Tools for Data Science," Electronic Text, one-semester course on applied methods in machine learning using Python. 2021-2022. Available at `http://mcrovella.github.io/CS506-Computational-Tools-for-Data-Science`. Includes contributions from Evimaria Terzi and George Kollios.

"Foundations of Data Science III," Electronic Text, one-semester course on probability and statistics for data science, 2023. Available at `https://mcrovella.github.io/DS122-Foundations-of-Data-Science-III/`. Co-developed with Pawel Przytycki and with contributions by Lisa Wobbes.

TUTORIALS AND SHORT COURSES

"Data Mining for Measurement Data," presented at the 5th PhD School on Traffic Monitoring and Analysis, associated with the 7th International Workshop on Traffic Monitoring and Analysis, Barcelona, April 21-22, 2015. Available at `https://github.com/mcrovella/mining-low-dim-network-data`.

"Traffic Modeling: Methods and Results for Single Links and Whole Networks," presented at SIGCOMM 2004, Portland OR.

"Traffic Modeling," short course presented at Aalborg University, February 2, 2004.

"Performance Characteristics of World Wide Web Information Systems," presented at SIGMETRICS '97, Seattle WA, June 1997; SIGMETRICS '98, Madison WI, June 1998; and Performance Tools '98, Palma de Mallorca, Spain, September 1998.

Consulting

Member of the Technical Advisory Board, Guavus, Inc. 2006 to 2017.

Consultant, Fujitsu Laboratories of America, Sunnyvale, CA. 2002 to 2004.

Consultant and Expert Witness, in litigation cases involving software patents, software trade secrets, and intellectual property. Ongoing.

Member of the Technical Advisory Board, epicRealm, Inc., Richardson, TX. March 2000 to May 2001.

Consultant and Member of the Technical Advisory Board, WebManage Technologies, Inc., Chelmsford MA. July 1999 to November 2000.

Founder and Principal Technical Officer, Commonwealth Network Technologies, Inc., Boston MA. Incorporated April 27, 1998; acquired by WebManage Technologies, October 12, 1999.

Consultant and Member of the Technical Advisory Board, VitalSigns Software Inc., Santa Clara, CA; acquired by International Network Services, Inc.; now part of Lucent. November 1996 to October 1999.

Patents

Methods for Storing and Reading Digital Data on a Set of DNA Strands, US Patent # 10,027,347, issued July 17, 2018.

Method and Apparatus for Whole-Network Anomaly Diagnosis and Method to Detect and Classify Network Anomalies Using Traffic Feature Distributions, US Patent # 8,869,276, issued October 21, 2014.

Compressed Set Representation for Sets as Measures in OLAP Cubes, US Patent # 8,533,167, issued September 10, 2013.

Method and Apparatus for Dynamic Resource Discovery and Information Distribution in a Data Network, US Patent # 7,747,741, issued June 29, 2010.

System for Creating and Distributing Prioritized List of Computer Nodes Selected as Participants in a Distribution Job, US Patent # 7,346,682, issued March 18, 2008.

Method and Apparatus for Election of Group Leaders in a Distributed Network, US Patent # 6,993,587, issued January 31, 2006.

Method and Apparatus for Scalable Distribution of Information in a Distributed Network, US Patent # 6,748,447, issued June 8, 2004.

Method and Apparatus for Reliable and Scalable Distribution of Data Files in Distributed Networks, US Patent # 6,718,361, issued April 6, 2004. European Patent # EP1305924 issued March 31, 2010.

Distributed Routing, US Patent # 6,370,584, issued April 9, 2002.

Method and Apparatus for Assigning Tasks in a Distributed Server System, US Patent # 6,223,205, issued April 24, 2001.

Doctoral Students

Ahmed Youssef, Ph.D. June 2023. "Computational Methods to Uncover Cell State Proteomes and Profile Protein Interaction Dynamics." Co-advised with Andrew Emili.

Bashir Rastegarpanah, Ph.D. August 2021. "Tools for Responsible Decision-Marking in Machine Learning."

Larissa Spinelli, Ph.D. June 2019. "Empirical Studies of Factors Affecting Opinion Dynamics."

Giovanni Comarela, Ph.D. July 2017. "On the Dynamics of Interdomain Routing in the Internet."

Natali Ruchansky, Ph.D. August 2016. "Matrix Completion with Structure." Co-advised with Evimaria Terzi.

Andrej Cvetkovski, Ph.D. May 2013. "Graph Embeddings for Low-Stretch Greedy Routing."

Gonca Gürsun, Ph.D. March 2013. "Inferring Hidden Features in the Internet."

Vijay Erramilli, Ph.D. December 2008. "Forwarding in Mobile Opportunistic Networks."

Nahur Fonseca, Ph.D. December 2008. "Stochastic Modeling Applied to Detection Problems in Network Protocols and Traffic."

Anukool Lakhina, Ph.D. August 2006. "Network-Wide Traffic Analysis: Methods and Applications."

Jun Liu, Ph.D. August 2002. "Characterizing Network Elements and Paths Using Packet Loss Behavior."

Paul Barford, Ph.D. June 2000. "Modeling, Measurement and Performance of World Wide Web Transactions."

Robert L. Carter, Ph.D. January 1998. "Performance Measurement and Prediction in Packet-Switched Networks: Techniques and Applications."

Other Doctoral Students Advised

Mengfei Cao, Ph.D, August 2016, "New Graph Metrics Improve Network-based Protein Function Prediction" Tufts University, Member of the Committee.

Bimal Viswanath, Ph.D, April 2016, "Towards Trustworthy Social Computing Systems." Universität des Saarlandes, Member of the Committee.

Yonjun Liao, Ph.D., January 2013, "Learning to Predict End-to-End Network Performance." Université de Liège, Belgium, Member of the Jury.

Diana Zeaiter Joumblatt, Ph.D., December 2012, "Prediction of User Dissatisfaction with Internet Application Performance at End-Hosts." Université Pierre et Marie Curie, Rapporteur.

Chrisil Arackaparambil, Ph.D., July 2011, "Anomaly Detection in Network Streams Through a Distributional Lens." Dartmouth College, Member of the Examining Committee.

Italo Cunha, Ph.D., July 2011, "Tracking Internet Routes and Path Reachability for Network Tomography." Université Pierre et Marie Curie (U. Paris VI), Rapporteur.

Hung Nguyen, Ph.D., February 2008, "User-level Internet Tomography and Overlay Routing." Ecole Polytechnique Fédérale de Lausanne (EPFL), Rapporteur.

Benoit Donnet, Ph.D., September 2006, "Algorithms for Large-Scale Topology Discovery." Université Pierre et Marie Curie (U. Paris VI), Président du Jury.

David Chua, Ph.D. (Mathematics), May 2006, "Statistical Analysis for Whole Networks." Second Reader.

Stilian Stoev, Ph.D. (Mathematics) May 2005.

Abhishek Kumar, Ph.D., October 2005, "Network Data Streaming: Algorithms for Network Monitoring and Coordination." Georgia Tech, Committee Member.

Marjan Bozinovski, Ph.D., May 2004, "Fault-Tolerant Platforms for IP-Based Session Control Systems." Aalborg University, Committee Member.

Artur Ziviani, Ph.D. December 2003, “Qualité de Service et Localisation d’Hotes.” University of Paris VI, Examinateur.

Liang Guo, Ph.D. April 2003, “Size-Aware Scheduling of TCP Flows.” Second Reader.

Alberto Medina, Ph.D. March 2003, “Practical Estimation of Internet Traffic Demands.” Third Reader.

Vladas Pipiras, Ph.D. (Mathematics) April 2002, “Symmetric Stable Self-Similar Processes.” Third Reader.

Khaled A. Harfoush, Ph.D. November 2001, “A Framework for End-to-End Characterization of Metric-Induced Network Topologies.” Committee Member.

Santiago M. Pericas-Geertsen, Ph.D. November 2001, “XML-Fluent Mobile Agents.” Committee Member.

Alia K. Atlas, Ph.D. August 1998, “Statistical Rate Monotonic Scheduling.” Second Reader.

Ghulam Bhatti, Ph.D. January 1998, “Parallel and Distributed Discrete Event Simulation: An Event-Reservation Approach.” Second Reader.

Sulaiman Mirdad, Ph.D. June 1997, “WebWave: Globally Load-Balanced, Fully-Distributed Caching”. Second Reader.

Gi Tae Kim, Ph.D. June 1997, “Framework for Adaptive Forward Erasure Recovery (AFER) for Fault-Tolerant Real-Time Communications”. Second Reader.

Carlos Cunha, Ph.D. June 1997, “Trace Analysis and its Application to Performance Enhancements of Distributed Information Systems”. Second Reader.

Susan Nagy, Ph.D. April 1997, “Admission Control and Scheduling Strategies for Real-Time Database Systems”. Second Reader.

Vadim Teverovksy, Ph.D. February 1997 (Mathematics), “Detection and Estimation of Long-Range Dependence.” Third Reader.

Euthimios Panagos, Ph.D. May 1995, “Client-Based Logging: A New Paradigm for Distributed Transaction Management”. Third Reader.

MASTERS STUDENTS

Megan Hacker, M.S. May 2012, “Methods for Pairing Mobile Devices Through Shared Sensing”

Lauren Walters, MSc. May 2004, “A Web Browsing Workload Model for Simulation,” University of Cape Town, South Africa, external examiner.

Alan Charbonneau, M.A. May 2003, “Simulation of an SJF Proxy”

Iana Vitkova, M.A. May 2003, “Simulation of an SJF Proxy”

Robert Frangioso, M.A. Sept 2002, “Connection Scheduling in Web Servers”

Xiaohui Zhang, M.S. May 1999, “Cachability of Web Objects”

Scott Harrison, M.S. May 1998, “DPR Monitor: A Web-based DPR Network Monitoring Tool”

Stanley Gambarin, M.S. May 1997, “Implementing a Single-Process Non-Multithreaded World Wide Web Server”

Kostas Magoutis, M.S. May 1996, “Economies of Scale In Partitioned Disk Caches with Self-Similar References”

Fred Bierhaus, M.S. May 1996, “Improving Interactive Traffic Through the EQL Driver”

BOOKS

M. Crovella and B. Krishnamurthy, “Internet Measurement: Infrastructure, Traffic, and Applications,” John Wiley and Sons, 2006, 512 pp.

Books Edited

M. Crovella and L. M. Feeney and D. Rubenstein (editors), "Proceedings of the 9th International IFIP TC 6 Networking Conference," Springer-Verlag, 410pp, 2010.

Journal Articles

Hui Peng, Sergei Kotelnikov, Megan E. Egbert, Shany Ofaim, Grant C. Stevens, Sadhna Phanse, Tatiana Saccon, Mikhail Ignatov, Shubham Dutta, Zoe Istace, Mohamed Taha Moutaoufik, Hiroyuki Aoki, Neal Kewalramani, Jianxian Sun, Yufeng Gong, Dzmitry Padhorny, Gennady Poda, Andrey Alekseenko, Kathryn A. Porter, George Jones, Irina Rodionova, Hongbo Guo, Oxana Pogoutse, Suprama Datta, Milton Saier, Mark Crovella, Sandor Vajda, Gabriel Moreno-Hagelsieb, John Parkinson, Daniel Segre, Mohan Babu, Dima Kozakov, and Andrew Emili, "Ligand interaction landscape of transcription factors and essential enzymes in *E. coli*," *Cell,* 2025.

A. Youssef, I. Paul, M. Crovella, and A. Emili, "DESP: Demixing Cell State Profiles from Dynamic Bulk Molecular Measurements," *Cell Reports Methods,* 2024.

L. Salamatian, S. Anderson, J. Matthews, P. Barford, W. Willinger, and M. Crovella, "A Manifold View of Connectivity in the Private Backbone Networks of Hyperscalers," *Communications of the ACM*, August 2023.

N. Kewalramani, A. Emili, M. Crovella, "State-of-the-Art Computational Methods to Predict Protein-Protein Interactions with High Accuracy and Coverage," *Proteomics,* 2023.

A. Youssef, F. Bian, P. Havugimana, M. Crovella, and A. Emili, "Dynamic Remodeling of *Escherichia coli* Interactome in Response to Environmental Perturbations," *Proteomics,* 2023.

L. Salamatian, S. Anderson, J. Mathews, P. Barford, W. Willinger, and M. Crovella, "A Manifold View of Connectivity in the Private Backbone Networks of Hyperscalers," *Communications of the ACM*, Research Highlights, 6(7), July 2023.

J. Jung, M. Guo, M. Crovella, J.G. McDaniel, and K.M. Warkentin. "Frog Embryos Use Multiple Levels of Temporal Pattern in Risk Assessment for Vibration-Cued Hatching," *Animal Cognition,* 2022.

J. N. Law, K. Akers, N. Tasnina, C. M. Della-Santina, S. Deutsch, M. Kshirsagar, J. Klein-Seetharaman, M. Crovella, P. Rajagopalan, S. Kasif, and T. M. Murali, "Interpretable Network Propagation with Application to Expanding the Repertoire of Human Proteins that Interact with SARS-CoV-2," *GigaScience,* 10(12), December 2021.

D. Lancour, J. Dupuis, R. Mayeux, J. Haines, M. Pericak-Vance, G. Schellenberg, M. Crovella, L. Farrer, and S. Kasif, "Analysis of Brain Region-Specific Co-Expression Networks Reveals Clustering of Established and Novel Genes Associated with Alzheimer Disease," *Alzheimer's Research & Therapy,* 2020.

J. Fan, X. C. Li, M. Crovella, and M. D. M. Leiserson, "Matrix (Factorization) Reloaded: Flexible Methods for Imputing Genetic Interactions with Cross-Species and Side Information," *OUP Bioinformatics,* 2020.

A. Ramirez, S. Dankel, B. Rastegarpanah, W. Cai, R. Xue, M. Crovella, Y.-H. Tseng, R. Kahn, and S. Kasif, "Single-Cell Transcriptional Networks in Differentiating Preadipocytes Suggest Drivers Associated with Tissue Heterogeneity," *Nature Communications,* 2020.

J. Fan, A. Cannistra, I. Fried, T. Lim, T. Schaffner, M. Crovella, B. Hescott, and M.D.M. Leiserson, "Functional Protein Representations from Biological Networks Enable Diverse Cross-Species Inference," *Nucleic Acids Research,* 2019.

D. Lancour, A. Naj, R. Mayeux, J. L. Haines, M. A. Pericak-Vance, G. C. Schellenberg, M. Crovella, L. A. Farrer, and S. Kasif, "One for All and All for One: Improving Replication of Genetic Studies through Network Diffusion," *PLoS Genetics*, Vol 14, April 2018.

A. Cvetkovski and M. Crovella, "Low-stress Data Embedding in the Hyperbolic Plane Using Multidimensional Scaling," *Applied Mathematics and Information Sciences,* Vol 11, Number 1, January 2017.

A. Cvetkovski and M. Crovella, "Multidimensional Scaling in the Poincaré Disk," *Applied Mathematics and Information Sciences,* Vol 10, Number 1, January 2016.

M. Cao, H. Zhang, J. Park, N. Daniels, M. Crovella, L. Cowen, and B. Hescott, "Going the Distance for Protein Function Prediction: A New Distance Metric for Protein Interaction Networks," *PLoS One,* Vol 8, Number 10, October 2013.

A. Cvetkovski and M. Crovella, "On the Choice of a Spanning Tree for Greedy Embedding of Network Graphs," *Networking Science,* February 2013.

B. Donnet, P. Raoult, T. Friedman, and M. Crovella, "Deployment of an Algorithm for Large-Scale Topology Discovery," *IEEE Journal on Selected Areas in Communications, Special Issue on Sampling the Internet*, Vol 24, Number 12, December 2006.

D. Chua, E. Kolaczyk, and M. Crovella, "Network Kriging," *IEEE Journal on Selected Areas in Communications, Special Issue on Sampling the Internet*, 2006.

B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint-Based Geolocation of Internet Hosts," *IEEE/ACM Transactions on Networking,* Vol 14, Number 6, December 2006.

R. Fonseca, V. Almeida, and M. E. Crovella, "Locality in a Web of Streams," *Communications of the ACM*, 48(1), pp. 82–88, January 2005.

M. Fayed, P. Krapivsky, J. Byers, M. Crovella, D. Finkel, and S. Redner, "On the Emergence of Highly Variable Distributions in the Autonomous System Topology," *Computer Communcation Review*, 33(2), pp. 41–49, July 2003.

A. Lakhina, J. W. Byers, M. E. Crovella, and I. Matta, "On the Geographic Location of Internet Resources," *IEEE Journal on Selected Areas in Communications, Special Issue on Internet and WWW Measurement, Mapping, and Modeling*, 2003.

P. Barford and M. E. Crovella, "Critical Path Analysis of TCP Transactions," *IEEE/ACM Transactions on Networking*, 9 (3), pp. 238–248.

M. E. Crovella, C. Lindemann, and M. Reiser, "Internet Performance Modeling: the state of the art at the turn of the century," *Performance Evaluation,* 42 (2–3), pp. 91-108, September 2000.

R. L. Carter and M. E. Crovella, "On the Network Impact of Dynamic Server Selection," *Computer Networks,* 31 (23–24), pp. 2529–2558.

M. Harchol-Balter, M. E. Crovella and C. D. Murta, "On Choosing a Task Assignment Policy for a Distributed Server System," *Journal of Parallel and Distributed Computing,* Special Issue on Software Support for Distributed Computing, September 1999.

P. Barford and M. E. Crovella, "Measuring Web Performance in the Wide Area," *Performance Evaluation Review,* Special Issue on Network Traffic Measurement and Workload Characterization, August 1999.

P. Barford, A. Bestavros, A. Bradley, and M. E. Crovella, "Changes in Web Client Access Patterns: Characteristics and Caching Implications," *World Wide Web,* Special Issue on World Wide Web Characterization and Performance Evaluation, 1999.

M. E. Crovella and M. S. Taqqu, "Estimating the Heavy Tail Index from Scaling Properties," *Methodology and Computing In Applied Probability,* 1(1), 1999.

M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," *IEEE/ACM Transactions on Networking,* 5(6):835–846, December 1997.

R. L. Carter and M. E. Crovella, "Measuring Bottleneck Link Speed in Packet-Switched Networks," *Performance Evaluation,* 27&28:297–318, October 1996.

L. A. Crowl, M. E. Crovella, T. J. LeBlanc, and M. L. Scott, "The Advantages of Multiple Parallelizations in Combinatorial Search," *Journal of Parallel and Distributed Computing,* special issue on Data-Parallel Algorithms and Programming, 21(1):110–123, April 1994.

Book Chapters

M. E. Crovella, M. S. Taqqu, and A. Bestavros, "Heavy-Tailed Probability Distributions in the World Wide Web," Chapter 1 of *A Practical Guide To Heavy Tails,* pages 3–25, Chapman & Hall, New York, 1998.

M. E. Crovella and L. Lipsky, "Simulations with Heavy-Tailed Workloads," Chapter 3 of *Self-Similar Network Traffic and Performance Evaluation,* Wiley/Wiley-Interscience, 1999.

K. Park, G. Kim and M. Crovella, "The Protocol Stack and its Modulation Effect on Self-Similar Traffic," Chapter 16 of *Self-Similar Network Traffic and Performance Evaluation,* Wiley/Wiley-Interscience, 1999.

M. E. Crovella, "Performance Characteristics of the World Wide Web," in *White Book on Performance Evaluation,* Günter Haring, Christoph Lindemann, and Martin Reiser, editors, Springer-Verlag, 1999.

Invited Papers

M. E. Crovella, "Performance Evaluation with Heavy Tailed Distributions," in *Lecture Notes in Computer Science 1786*, Proceedings of the 11th International Conference on Modelling Tools and Techniques for Computer and Communication System Performance Evaluation (TOOLS '2000), pp. 1-9, March 2000. Revised version appeared in *Job Scheduling Strategies for Parallel Processing : 7th International Workshop,* pp. 1-10, June 16, 2001.

Refereed Conference Papers

P. Calais, G. Franco, Z. Tang, T. Nikas, W. Meira Jr., E. Terzi, M. Crovella, "Disentangling Text and Math in Word Problems: Evidence for the Bidimensional Structure of Large Language Models' Reasoning," Findings of the Association for Computational Linguistics, June, 2025.

C. Kranig, E. Pauley, W.-S. Wang, P. Barford, M. Crovella, and J. Sommers, "Toward a Representative DNS Data Corpus: A Longitudinal Comparison of Collection Methods, *Proceedings of the Network Traffic Measurement and Analysis Conference,* Copenhagen, Denmark, June 2025. Recipient of the TMA Community Contribution Award.

W.-S. Wung, C. Kranig, E. Pauley, P. Barford, M. Crovella, and J. Sommers, "Squatspotting: Towards the Systematic Measurement of Typosquatting Techniques," *Proceedings of IFIP Networking,* Limassol, Cyprus, May 2025.

A. Anderson, A. S. Mondal, P. Barford, M. Crovella, and J. Sommers, "An Elemental Decomposition of DNS Name-to-IP Graphs," *Proceedings of INFOCOM*, Vancouver, Canada, May 2024.

G. Franco, M. Crovella, and G. Comarela, "Dependence and Model Selection in LLP: The Problem of Variants," *Proceedings of KDD*, Los Angeles, CA, August 2023.

K. Quinn, E. Terzi, and M. Crovella, "Characterizing Covid Waves via Spatio-Temporal Decomposition," *Proceedings of KDD*, Washington DC, August 2022.

L. Salamatian, S. Anderson, J. Matthews, P. Barford, W. Willinger, and M. Crovella, "Curvature-based Analysis of Network Connectivity in Private Backbone Infrastructures," *Proceedings of ACM SIGMETRICS / IFIP Performance,* Mumbai, India, June 2022.

B. Rastegarpanah, K. P. Gummadi, and M. Crovella, "Auditing Black-Box Prediction Models for Data Minimization Compliance," *Proceedings of NeurIPS,* Online, December 2021. **Selected as a Spotlight Presentation (3% acceptance rate).**

L. Amichi, A. Carneiro Viana, M. Crovella, and A. A. F. Loureiro, "From Motion Purpose to Perceptive Spatial Mobility Prediction," *Proceedings of ACM SIGSPATIAL*, Beijing, China, November 2021.

G. Grassi, C. Barakat, M. Crovella, and R. Teixeira, "Leveraging Website Popularity Differences to Identify Performance Anomalies," *Proceedings of Infocom,* Online, May 2021.

L. Amichi, A. Carneiro Viana, M. Crovella, and A. A. F. Loureiro, "Understanding Individuals' Proclivity for Novelty Seeking," *Proceedings of ACM SIGSPATIAL,* Online, November 2020.

J. Fan, X. C. Li, M. Crovella, and M. D. M. Leiserson, "Matrix (Factorization) Reloaded: Flexible Methods for Imputing Genetic Interactions with Cross-Species and Side Information," *Proceedings of the 19th European Conference on Computational Biology,* Online, August 2020.

B. Rastegarpanah, M. Crovella, and K. P. Gummadi, "Fair Inputs and Fair Outputs: the Incompatibility of Fairness in Privacy and Accuracy," *Proceedings of the Workshop on Fairness in User Modeling, Adaptatation, and Personalization (FairUMAP),* Online, July 2020.

L. Spinelli and M. Crovella, "How YouTube Leads Privacy-Seeking Users Away from Reliable Information," *Proceedings of the Workshop on Fairness in User Modeling, Adaptatation, and Personalization (FairUMAP),* Online, July 2020.

J. Law, M. Kshirsagar, M. Crovella, P. Rajagopalan, J. Klein-Seetharaman, S. Kasif and T. M. Murali, "A Network-Based Label Propagation Framework to Reposition Drugs for COVID-19 (Poster)," Proceedings of RECOMB, Online, June 2020.

L. Amichi, A. Viana, M. Crovella, and A. A. F. Loureiro, "Mobility profiling: Identifying Scouters in the Crowd (Poster)," *Proceedings of the CoNEXT Student Workshop,* Orlando, FL, December 2019.

B. Rastegarpanah, K. P. Gummadi, and M. Crovella, "Fighting Fire with Fire: Using Antidote Data to Improve Polarization and Fairness of Recommender Systems," *Proceedings of WSDM 2019,* Melbourne, Australia, February 2019.

G. Comarela, R. Durairajan, P. Barford, D. Christenson, and M. Crovella, "Assessing Candidate Preference through Web Browsing History," *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD),* London, UK, August 2018. Press coverage: BU Today: `http://www.bu.edu/today/2018/better-election-predictions`.

L. Spinelli and M. Crovella, "Unravelling the Dynamics of Online Ratings," Proceedings of the 4th International Symposium on Social Media Mining and Analysis (SMMA 2018), Exeter, UK, June 2018.

M.D.M. Leiserson, J. Fan, A. Cannistra, I. Fried, T. Lim, T. Schaffner, M. Crovella, and B. Hescott, "A Multi-Species Functional Embedding Integrating Sequence and Network Structure," Proceedings of RECOMB, Paris, France, April 2018.

B. Rastegarpanah, K. Gummadi, and M. Crovella, "Exploring Explanations for Matrix Factorization Recommender Systems" (Position Paper), FATREC Workshop on Responsible Recommendation, Como Italy, August 2017.

L. Spinelli and M. Crovella, "Closed-Loop Opinion Formation", *Proceedings of the 9th International ACM Web Science Conference (WebSci),* Troy, NY, June 2017.

N. Ruchansky, E. Terzi, and M. Crovella, "Targeted Matrix Completion," *Proceedings of the SIAM International Conference on Data Mining (SDM),* Houston, TX, April 2017.

C. Mavroforakis, D. Erdös, M. Crovella, and E. Terzi, "Active Positive-Definite Matrix Completion," *Proceedings of the SIAM International Conference on Data Mining (SDM),* Houston, TX, April 2017.

G. Comarela, E. Terzi, and M. Crovella, "Detecting Unusually-Routed ASes: Methods and Applications," *Proceedings of the ACM Internet Measurement Conference (IMC),* Santa Monica, CA, November 2016.

N. Ruchansky, M. Crovella, and E. Terzi, "Matrix Completion with Queries," *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD),* Sydney, Australia, August 2015.

G. Comarela and M. Crovella, "Identifying and Analyzing High Impact Routing Events with PathMiner," *Proceedings of the ACM Internet Measurement Conference (IMC),* Vancouver, BC, November 2014.

B. Viswanath, M. A. Bashir, M. Crovella, S. Guha, K. Gummadi, B. Krishnamurthy, and A. Mislove, "Towards Detecting Anomalous User Behavior in Online Social Networks," *Proceedings of the 23rd USENIX Security Symposium*, San Diego, CA, August 2014.

Y. Zhang, T. Lappas, M. Crovella, and E. Kolaczyk, "Online ratings: convergence towards a positive perspective?," *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP),* Florence, Italy, May 2014.

G. Comarela, G. Gürsun, and M. Crovella, "Studying Interdomain Routing over Long Timescales," *Proceedings of the ACM Internet Measurement Conference (IMC),* Barcelona, Spain, October 2013.

L. Spinelli, M. Crovella and B. Eriksson, "AliasCluster: A Lightweight Approach to Interface Disambiguation," *Proceedings of the Global Internet Symposium,* Turin, Italy, April 2013.

B. Eriksson and M. Crovella, "Understanding Geolocation Accuracy using Network Geometry," in *Proceedings of the Infocom 2013 Miniconference,* Turin, Italy, April 2013.

J. M. Agosta, J. Chandrasekar, M. Crovella, N. Taft, and D. Ting, "Mixture Models of Endhost Network Traffic," in *Proceedings of the Infocom 2013 Miniconference,* Turin, Italy, April 2013.

G. Gürsun, N. Ruchansky, E. Terzi, and M. Crovella, "Routing State Distance: A Path-based Metric for Network Analysis," in *Proceedings of the ACM Internet Measurement Conference (IMC),* Boston, MA, November 2012. **Winner of a 2013 IETF/IRTF Applied Networking Research Prize.**

G. Gürsun and M. Crovella, "On Traffic Matrix Completion in the Internet," in *Proceedings of the ACM Internet Measurement Conference (IMC),* Boston, MA, November 2012.

T. Lappas, M. Crovella, and E. Terzi, "Selecting a Characteristic Set of Reviews," in *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2012,* Beijing, China, August 2012.

Q. Ding, N. Katenka, P. Barford, E. Kolaczyk, and M. Crovella, "Intrusion as (Anti)social Communication: Characterization and Detection," in *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD) 2012,* Beijing, China, August 2012. **Best Student Paper Award.**

G. Gürsun, N. Ruchansky, E. Terzi, and M. Crovella, "Inferring Visibility: Who's (not) Talking to Whom?," in *Proceedings of SIGCOMM 2012,* Helsinki, Finland, August 2012.

B. Eriksson and M. Crovella, "Estimating Intrinsic Dimension Via Clustering," in *IEEE Statistical Signal Processing Workshop (SSP),* Ann Arbor, MI, August 2012.

C. Thomas, J. Sommers, P. Barford, D. Kim, A. Das, R. Segebre, M. Crovella, "A Passive Measurement System for Network Testbeds," in *Tridentcom 2012,* Thessaloniki, Greece, June 2012.

G. Comarela, M. Crovella, V. Almeida and F. Benevenuto, "Understanding Factors that Affect Response Rates in Twitter," in *ACM Hypertext 2012,* Milwaukee, WI, June 2012.

A. Cvetkovski and M. Crovella, "Low-Stretch Greedy Embedding Heuristics," in *Fourth International Workshop on Network Science for Communication Networks (NetSciCom 2012),* Orlando, FL, March 2012.

V. Manfredi, M. Crovella, and J. Kurose, "Understanding Stateful vs Stateless Communication Strategies for Ad hoc Networks," in *Proceedings of Mobicom 2011,* Las Vegas, NV, September 2011.

G. Gürsun, M. Crovella, and I. Matta, "Describing and Forecasting Video Access Patterns," in *Proceedings of the Infocom 2011 Miniconference,* Shanghai, China, April 2011.

V. Bharti, P. Kankar, L. Setia, G. Gürsun, and M. Crovella, "Inferring Invisible Traffic," in *Proceedings of CoNEXT 2010,* Philadelphia, PA, December 2010.

J. Sommers, P. Barford and M. Crovella, "Router Primitives for Programmable Active Measurement," in *Workshop on Programmable Routers for Extensible Services of Tomorrow (PRESTO '09)*, Barcelona, Spain, August, 2009.

A. Cvetkovski and M. Crovella, "Hyperbolic Embedding and Routing for Dynamic Graphs," in *Proceedings of Infocom 2009*, Rio de Janeiro, Brazil, April 2009.

V. Erramilli and M. Crovella, "Forwarding in Opportunistic Networks with Resource Constraints," in *Proceedings of the Fourth ACM Workshop on Challenged Networks (CHANTS 08)*, San Francisco, CA, September, 2008.

N. Fonseca, M. Crovella, and K. Salamatian, "Long Range Mutual Information," in *Proceedings of the First Workshop on Hot Topics in Measurement and Modeling of Computer Systems (Hotmetrics '08),* Annapolis, MD, June, 2008.

V. Erramilli, A. Chaintreau, M. Crovella, and C. Diot, "Delegation Forwarding," in *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (ACM MobiHoc),* Hong Kong SAR, China, May 2008.

P. Chhabra, C. Scott, E. D. Kolaczyk, and M. Crovella, "Distributed Spatial Anomaly Detection," in *Proceedings of Infocom 2008*, Phoenix, AZ, April 2008.

V. Erramilli, A. Chaintreau, M. Crovella, and C. Diot, "Diversity of Forwarding Paths in Pocket Switched Networks," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, San Diego, October, 2007.

B. Eriksson, P. Barford, R. Nowak, and M. Crovella, "Learning Network Structure from Passive Measurements," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, San Diego, October, 2007.

V. Erramilli, M. Crovella, and N. Taft, "An Independent-Connection Model for Traffic Matrices," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Rio de Janeiro, October, 2006.

X. Li, F. Bian, M. Crovella, C. Diot, R. Govindan, G. Iannaccone, and A. Lakhina, "Detection and Identification of Network Anomalies Using Sketch Subspaces," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Rio de Janeiro, October, 2006.

A. Lakhina, M. Crovella, and C. Diot, "Detecting Distributed Attacks using Network-Wide Flow Traffic," in *Proceedings of FloCon 2005 Analysis Workshop*, New Orleans, LA, September, 2005.

A. Lakhina, M. Crovella, and C. Diot, "Mining Anomalies Using Traffic Feature Distributions," in *Proceedings of ACM SIGCOMM*, Philadelphia, PA, August 2005.

D. B. Chua, E. D. Kolaczyk, and M. Crovella, "A Statistical Framework for Efficient Monitoring of End-to-End Network Properties," (Poster) in *Proceedings of ACM SIGMETRICS*, Banff, Canada, June 2005.

B. Donnet, P. Raoult, T. Friedman, and M. Crovella, "Efficient Algorithms for Large-Scale Topology Discovery," in *Proceedings of ACM SIGMETRICS*, Banff, Canada, June 2005.

A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot, "Traffic Matrices: Balancing Measurements, Inference and Modeling," in *Proceedings of ACM SIGMETRICS*, Banff, Canada, June 2005.

B. Donnet, T. Friedman, and M. Crovella, "Improved Algorithms for Network Topology Discovery," in *Lecture Notes in Computer Science, Proceedings of the 2005 Passive and Active Measurement Workshop*, Boston, MA, March, 2005.

D. B. Chua, E. D. Kolaczyk, and M. Crovella, "Efficient Monitoring of End-to-End Network Properties," in *Proceedings of Infocom 2005*, Miami, FL, March 2005.

N. Fonseca and M. Crovella, "Bayesian Packet Loss Detection for TCP," in *Proceedings of Infocom 2005*, Miami, FL, March 2005.

B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint-Based Geolocation of Internet Hosts," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Taormina, Italy, October 2004.

A. Lakhina, M. Crovella, and C. Diot, "Characterization of Network-Wide Anomalies in Traffic Flows," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference*, Taormina, Italy, October 2004.

A. Lakhina, M. Crovella, and C. Diot, "Diagnosing Network-Wide Traffic Anomalies," in *Proceedings of ACM SIGCOMM*, Portland, OR, August 2004.

A. Lakhina, M. Crovella, and C. Diot, "Exploring the Subspace Method for Network-Wide Anomaly Diagnosis," (Poster) in *Network Troubleshooting Workshop: Research, Theory and Operations*, Portland, OR, August 2004.

A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft, "Structural Analysis of Network Traffic Flows," in *Proceedings of ACM SIGMETRICS / Performance 2004*, New York, NY, June 2004.

L. Tang and M. E. Crovella, "A Geometric Exploration of the Landmark Selection Problem," in *Lecture Notes in Computer Science 3015, Proceedings of the 2004 Passive and Active Measurement Workshop (PAM2004)*, pp. 63–72, Juan-les-Pins, France, April 2004.

L. Tang and M. E. Crovella, "Virtual Landmarks for the Internet," in *Proceedings of the 2003 Internet Measurement Conference*, Miami, Florida, October 2003.

R. Fonseca, V. Almeida, M. E. Crovella, and B. Abrahão, "On the Intrinsic Locality Properties of Web Reference Streams," in *Proceedings of IEEE Infocom 2003,* San Francisco, California, April 2003.

A. Lakhina, J. Byers, M. E. Crovella, and P. Xie, "Sampling Biases in IP Topology Measurements," in *Proceedings of IEEE Infocom 2003,* San Francisco, California, April 2003.

M. E. Crovella and E. Kolaczyk, "Graph Wavelets for Spatial Traffic Analysis," in *Proceedings of IEEE Infocom 2003,* San Francisco, California, April 2003.

J. Liu, I. Matta, and M. E. Crovella, "End-to-End Inference of Loss Nature in a Hybrid Wired/Wireless Environment," in *Proceedings of Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'03)*, Sophia-Antipolis, France, March 2003.

A. Lakhina, J. W. Byers, M. E. Crovella, and Ibrahim Matta, "On the Geographic Location of Internet Resources," (Extended Abstract) in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2002,* Marseilles, France, November 2002.

J. Liu and M. Crovella, "Using Loss Pairs to Discover Network Properties," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2001,* pp. 127–138, San Francisco, CA, November 2001.

P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the Marginal Utility of Network Topology Measurements," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop 2001,* pp. 5–17, San Francisco, CA, November 2001.

L. Guo, I. Matta, and M. E. Crovella, "How Does TCP Generate Pseudo-Self-Similarity?," in *Proceedings of MASCOTS 2001*, Cincinnati, Ohio, August 2001.

P. Barford and M. E. Crovella, "Critical Path Analysis of TCP Transactions," in *Proceedings of the 2000 ACM SIGCOMM Conference*, Stockholm, Sweden, pp. 127–138, September 2000. Also appeared by invitation as "Análisis de secuencia crítica de transacciones TCP" in *Proceedings of SIGCOMM Latin America,* San Antonio de Belén, Costa Rica, April 2001.

M. E. Crovella, R. Frangioso, and M. Harchol-Balter "Connection Scheduling in Web Servers," in *1999 USENIX Symposium on Internet Technologies and Systems (USITS '99).*

P. Barford and M. E. Crovella, "A Performance Evaluation of Hyper Text Transfer Protocols," in *Proceedings of ACM SIGMETRICS '99,* pp. 188–197, Atlanta, Georgia, May 1999.

A. Bestavros, M. E. Crovella, J. Liu, and D. Martin, "Distributed Packet Rewriting and its Application to Scalable Server Architectures," in *Proceedings of the Sixth International Conference on Network Protocols (ICNP '98),* pp. 290–297, October, 1998.

M. Harchol-Balter, M. E. Crovella, and C. D. Murta, "On Choosing a Task Assignment Policy for a Distributed Server System," in *Proceedings of Performance Tools '98, Lecture Notes in Computer Science,* No. 1469, pp. 231–242, September 1998.

P. Barford and M. E. Crovella, "Generating Representative Web Workloads for Network and Server Performance Evaluation,"in *Proceedings of ACM SIGMETRICS/Performance '98,* pp. 151–160, Madison, Wisconsin, July 1998.

M. E. Crovella, M. Harchol-Balter, and C. D. Murta, "Task Assignment in a Distributed System: Improving Performance by Unbalancing Load," in *Proceedings of SIGMETRICS '98 (Poster paper),* Madison, Wisconsin, July 1998.

M. E. Crovella and P. Barford, "The Network Effects of Prefetching," in *Proceedings of IEEE Infocom '98, the Seventeenth Annual Joint Conference of the IEEE Computer and Communication Societies,* San Francisco, California, pp. 1232–1240, April 1998.

M. E. Crovella and L. Lipsky, "Long-Lasting Transient Conditions in Simulations with Heavy-Tailed Workloads," in *Proceedings of the 1997 Winter Simulation Conference*, December 1997.

K. Park, G. Kim, and M. E. Crovella, "On the Effect of Traffic Self-Similarity on Network Performance," in *Proceedings of the SPIE International Conference on Performance and Control of Network Systems*, November, 1997.

P. Fiorini, L. Lipsky, and M. E. Crovella, "Consequences of Ignoring Self-Similar Data Traffic In Communications Modeling," In *Tenth International Conference on Parallel and Distributed Computing Systems (PDCS-97),* New Orleans, LA, pp. 322-327, October 1997.

R. L. Carter and M. E. Crovella, "Server Selection using Dynamic Path Characterization in Wide-Area Networks," in *Proceedings of IEEE Infocom '97, the Sixteenth Annual Joint Conference of the IEEE Computer and Communication Societies,* Kobe, Japan, April 1997.

V. Almeida, A. Bestavros, M. Crovella, and A. de Oliveira, "Characterizing Reference Locality in the WWW," in *Proceedings of the Fourth International Conference on Parallel and Distributed Information Systems (PDIS '96)*, December 1996.

K. Park, G. Kim, and M. E. Crovella, "On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic," in *Proceedings of the Fourth International Conference on Network Protocols (ICNP '96),* pp. 171–180, October, 1996.

R. L. Carter and M. E. Crovella, "Measuring Bottleneck Link Speed in Packet-Switched Networks," in *PERFORMANCE '96, the International Conference on Performance Theory, Measurement and Evaluation of Computer and Communication Systems,* October 1996.

M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes," in *Proceedings of the 1996 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems,* May 1996. **Winner of a 2010 ACM SIGMETRICS Test of Time Award.**

M. E. Crovella and R. L. Carter, "Dynamic Server Selection in the Internet," In *Proceedings of the Third IEEE Workshop on the Architecture and Implementation of High Performance Communication Subsystems (HPCS '95),* August 1995.

A. Bestavros, R. L. Carter, M. E. Crovella, C. R. Cunha, A. Heddaya, and S. A. Mirdad, "Application-Level Document Caching in the Internet," In *Proceeedings of the Second Workshop on Services in Distributed and Networked Environments (SDNE '95),* June 1995.

M. E. Crovella and T. J. LeBlanc, "Parallel Performance Prediction Using Lost Cycles Analysis," In *Proceedings of Supercomputing '94,* November, 1994.

R. Bianchini, M. E. Crovella, L. Kontothanassis, and T. J. LeBlanc, "Software Interleaving," In *Proceedings of the 6th IEEE Symposium on Parallel and Distributed Processing*, October, 1994.

M. E. Crovella and T. J. LeBlanc, "Performance Debugging using Parallel Performance Predicates," In *3rd ACM/ONR Workshop on Parallel and Distributed Debugging*, May, 1993, pp. 140-150.

M. Crovella, R. Bianchini, T. LeBlanc, E. Markatos, and R. Wisniewski, "Using Communication-to-Computation Ratio in Parallel Performance Prediction and Program Design," *Proceeedings of the Fourth IEEE Symposium on Parallel and Distributed Processing*, December 1992, pp. 238-245.

M. Crovella, P. Das, C. Dubnicki, T. LeBlanc, and E. Markatos, "Multiprogramming on Multiprocessors," *Proceeedings of the Third IEEE Symposium on Parallel and Distributed Processing*, December 1991, pp. 590–597, Also published as TR 385, Department of Computer Science, University of Rochester, February 1991 (revised May 1991).

E. Markatos, M. Crovella, P. Das, C. Dubnicki, and T. LeBlanc, "The Effects of Multiprogramming on Barrier Synchronization," *Proceeedings of the Third IEEE Symposium on Parallel and Distributed Processing*, December 1991, pp. 662-669. Also TR 380, University of Rochester Computer Science Department, May 1991.

N. Bashir, M. Crovella, G. DeTitta, F. Han, H. Hauptman, J. Horvath, H. King, D. Langs, R. Miller, T. Sabin, P. Thuman and D. Velmurugan, "Parallel Solutions to the Phase Problem in X-Ray Crystallography," *Proceedings of The Fifth Distributed Memory Computing Conference*, IEEE Computer Society Press, 1990, pp. 513–521.

M. Crovella, M. S. Summers, E. M. Camardo, and S. S. Weissman, "Exploiting a Restricted Application Domain in SDE Design," *Advance Working Papers of the Second International Workshop on Computer Aided Software Engineering,* IEEE Computer Society Press, July 1988, pp. 2-3–2-4.

J. G. Neal and M. E. Crovella, "An Expert System for Tactical Support Jammer Configuration," *Proceedings of the 1987 IEEE National Aerospace and Electronics Conference (NAECON '87),* May 1987, pp. 1308–1314.

J. G. Neal, M. E. Crovella, and C. Y. Thielman, "A Knowledge-Based Approach to an EW Mission Planning Task," *Proceedings of the Knowledge Based Information Systems Symposium (KBIS-87),* SHAPE Technical Centre, The Hague, Netherlands, 1987.

J. G. Neal, D. F. Funke, and M. E. Crovella, "JAMEX: An Expert Advisor on Tactical Support Jammer Configuration," *Proceedings of the Second Annual Conference on Expert Systems in Government,* IEEE Computer Society Press, November 1986, pp. 383–394.

Refereed Presentations

I. Fried, A. Cannistra, C. Casey, A. Piel, T. Lim, M. Crovella, B. Hescott, "CANDL: Coarsely Aligning Networks with Diffusion and Landmarks," Theme Presentation, ISMB/ECCB, Dublin, July 2015.

Unrefereed Publications

R. Canetti, A. Cohen, C. Conley, M. Crovella, S. L. Dogan, M. Gaboardi, W. Hartzog, R. Van Loo, C. T. Robertson, K. Baird Silbaugh, "National Telecommunications and Information Administration: Comments from Researchers at Boston University and the University of Chicago," Boston Univ. School of Law Research Paper No. 23-24, June 2023.

M. Crovella, "The Skillful Interrogation of the Internet," in *Computer Communication Review,* vol. 49, no. 3, October 2019.

kc Claffy, M. Crovella, T. Friedman, C. Shannon and N. Spring, "Community-oriented network measurement infrastructure (CONMI) workshop report," in *Computer Communication Review,* vol. 36, no. 2, pp. 41–48, April, 2006.

M. E. Crovella, "How Big is the Average Web Page?" in *Inclusions, the Arts of Mathematics,* Journal of the Mathematical Association of America Student Chapter at Boston University, May 1999.

T. J. LeBlanc, M. L. Scott, B. Marsh, E. Markatos, C. Dubnicki, M. Crovella, and T. Becker, "The Psyche Parallel Operating System," *IEEE TC on Operating Systems Newsletter,* vol. 6, no. 1, 1992, pp. 11-13.

TECHNICAL REPORTS

G. Grassi, C. Barakat, M. Crovella, and R. Teixeira, "Leveraging Website Popularity Differences to Identify Performance Anomalies," (Extended Version), Available at `https://hal.inria.fr/hal-03109717`, January 2021.

J. N. Law, N. Tasnina, M. Kshirsagar, J. Klein-Seetharaman, M. Crovella, P. Rajagopalan, S. Kasif, and T. M. Murali, "Identifying Human Interactors of SARS-CoV-2 Proteins and Drug Targets for COVID-19 using Network-Based Label Propagation," arXiv:2006.01968, June 2, 2020.

B. Rastegarpanah, M. Crovella, and K. P. Gummadi, "Fair Inputs and Fair Outputs: the Incompatibility of Fairness in Privacy and Accuracy," arXiv:2005.09209, May 19, 2020.

M.D.M. Leiserson, J. Fan, A. Cannistra, I. Fried, T. Lim, T. Schaffner, M. Crovella, and B. Hescott, "A Multi-Species Functional Embedding Integrating Sequence and Network Structure," bioRxiv, December 4, 2017.

J.-M. Agosta, J. Chandrasekar, M. Crovella, N. Taft, and D. Ting, "Mixture Models of Endhost Network Traffic," arxiv:1212.2744, December 13, 2012.

M. Hacker, M. Crovella, and L. Reyzin, "Secure Pairing of Mobile Devices," BUCS-TR-2012-011, May 16, 2012.

A. Cvetkovski and M. Crovella, "Multidimensional Scaling in the Poincare Disk," arxiv:1105.5332, May 26, 2011.

A. Lakhina, M. Crovella, and C. Diot, "Diagnosing Network-Wide Traffic Anomalies," BUCS-TR-2004-008, February 24, 2004. Also published as a Sprint Research Report.

A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft, "Structural Analysis of Network Traffic Flows," BUCS-TR-2003-021, November 10, 2003. Also published as Sprint Research Report RR03-ATL-110708.

M. Fayed, P. Krapivsky, J. Byers, M. Crovella, and D. Finkel, "On the Size Distribution of Autonomous Systems," BUCS-TR-2003-001, January 17, 2003.

R. Fonseca, V. Almeida, M. E. Crovella, and B. Abrahão, "On the Intrinsic Locality Properties of Web Reference Streams," BUCS-TR-2002-022, July 15, 2002.

A. Lakhina, J. Byers, M. E. Crovella, and P. Xie, "Sampling Biases in IP Topology Measurements," BUCS-TR-2002-021, July 15, 2002.

M. E. Crovella and E. Kolaczyk, "Graph Wavelets for Spatial Traffic Analysis," BUCS-TR-2002-020, July 15, 2002.

A. Lakhina, J. W. Byers, M. E. Crovella, and Ibrahim Matta, "On the Geographic Location of Internet Resources," BUCS-TR-2002-015, 2002.

J. Liu, I. Matta, and M. E. Crovella, "End-to-End Inference of Loss Nature in a Hybrid Wired/Wireless Environment," BUCS-TR-2002-008, March 14, 2002.

M. E. Crovella, R. Frangioso, and M. Harchol-Balter "Connection Scheduling in Web Servers," BUCS-TR-1999-003, March 31, 1999.

M. Harchol-Balter, M. E. Crovella, and C. D. Murta, "On Choosing a Task Assignment Policy for a Distributed Server System," MIT Laboratory for Computer Science, MIT-LCS-TR-757, 1998.

P. Barford and M. E. Crovella, "Generating Representative Web Workloads for Network and Server Performance Evaluation," BUCS-TR-1997-006, November 4, 1997.

M. E. Crovella, M. Harchol-Balter, and C. D. Murta, "Task Assignment in a Distributed System: Improving Performance by Unbalancing Load," BUCS-TR-1997-018, October 31, 1997.

M. Harchol-Balter, M. E. Crovella, and C. D. Murta, "To Queue or Not to Queue: When Queueing is Better Than Timesharing in a Distributed System," BUCS-TR-1997-017, October 31, 1997.

K. Park, G. Kim, and M. E. Crovella, "On the Effect of Traffic Self-Similarity on Network Performance," CSD-TR 97-024, Dept. of Computer Sciences, Purdue University, July 1997.

M. Crovella and P. Barford, "The Network Effects of Prefetching," BUCS-TR-97-002, February 7, 1997.

K. Park, G. Kim, and M. E. Crovella, "On the Relationship Between File Sizes, Transport Protocols, and Self-Similar Network Traffic," BUCS-TR-96-016, August 7, 1996.

V. Almeida, A. Bestavros, M. Crovella, and A. deOliveira, "Characterizing Reference Locality in the WWW," BUCS-TR-96-011, June 21, 1996.

A. Bestavros, M. Chen, M. Crovella, A. Heddaya, S. Sclaroff, and J. Cowie, "Responsive Web Computing: Resource Management, Protocol Techniques, and Applications (A research statement)," BUCS-TR-96-008, March 21, 1996.

R. L. Carter and M. E. Crovella, "Dynamic Server Selection Using Bandwidth Probing in Wide-Area Networks," BUCS-TR-96-007, Boston University Computer Science Department, March 18, 1996.

R. L. Carter and M. E. Crovella, "Measuring Bottleneck Link Speed in Packet-Switched Networks," Technical Report BUCS-TR-96-006, Boston University Computer Science Department, March 15, 1996.

M. E. Crovella and A. Bestavros, "Explaining World Wide Web Traffic Self-Similarity," Technical Report BUCS-TR-95-015 (Revised), Boston University Department of Computer Science, October 1995.

M. E. Crovella, T. J. LeBlanc, and W. Meira Jr., "Performance Measurement and Modeling with the Lost Cycles Toolkit," Technical Report 580, University of Rochester Computer Science Department, June 1995.

C. Cunha, A. Bestavros, and M. Crovella, "Characteristics of WWW Client-based Traces," Technical Report BUCS-TR-95-010, Boston University Department of Computer Science, April 1995.

A. Bestavros, R. L. Carter, M. E. Crovella, C. R. Cunha, A. Heddaya, and S. A. Mirdad, "Application-Level Document Caching in the Internet," Technical Report BUCS-TR-95-002, Boston University Department of Computer Science, February 1995.

M. E. Crovella, "Performance Prediction and Tuning of Parallel Programs," Technical Report 573 and Ph.D. Thesis, Computer Science Dept., U. Rochester, February 1995.

M. E. Crovella and T. J. LeBlanc, "The Search for Lost Cycles: a New Approach to Parallel Program Performance Evaluation," Technical Report 479, University of Rochester Computer Science Department, December 1993.

L. A. Crowl, M. E. Crovella, T. J. LeBlanc, and M. L. Scott, "Beyond Data Parallelism: The Advantages of Multiple Parallelizations in Combinatorial Search," Technical Report 92-80-06, December 1992, Department of Computer Science, Oregon State University. Revised version appears as Technical Report 451, U. Rochester Department of Computer Science, April 1993.

R. Bianchini, M. E. Crovella, L. Kontothanassis, and T. J. LeBlanc, "Alleviating Memory Contention in Matrix Computations on Large-Scale Shared-Memory Multiprocessors," Technical Report 449, U. Rochester Computer Science Department, April 1993.

R. Bianchini, M. E. Crovella, L. Kontothanassis, and T. J. LeBlanc, "Memory Contention in Scalable Cache-Coherent Multiprocessors," Technical Report 448, U. Rochester Computer Science Department, April 1993.

D. Bergmark and M. Crovella, "Case Study in KSR Programming: Finding Outliers by the Minimum Volume Ellipsoid Method," Technical Report CTC92TR107, October 1992, Center for Theory and Simulation in Science and Engineering, Cornell University, Ithaca, NY 14853.

Recent Talks

"Beyond accuracy: understanding the performance of LLMs on exams designed for humans,"

- UFMG, Department of Computer Science, March 8, 2024.

"Estimating Traffic Flows that Can't Be Measured"

- Google Telemetry Seminar Series, May 17, 2021.

"Assessing Candidate Preference Through Web Browsing History"

- Sorbonne Université (UPMC) LIP6, April 15, 2019.
- INRIA Sophia-Antipolis, March 20, 2019.
- INRIA Saclay, December 13, 2018.
- INRIA Paris, November 8, 2018.
- LINCS Laboratory, Paris, October 10, 2018.

"Human Behavior is Low Dimensional"

- TMA Experts Summit 2019, Paris, June 18, 2019.
- Chinese Academy of Sciences, Institute of Computing Technology, Beijing, May 20, 2019.
- Huawei Network AI Workshop, Paris, April 9, 2019.
- LINCS Laboratory, Paris, February 13, 2019.
- Wesleyan University Department of Mathematics and Computer Science, April 13, 2017.
- IEEE Workshop on Data Science for Networking (DS4N), December 14, 2017.

"The Statistics of Human Activity and Implications for Security"

- BU Statistics Club, Professor Speaker Series, November 10, 2016.

"Analyzing the Evolution of Internet Routing through Path Analysis"

- Cisco FAST Seminar Series, July 17, 2015.
- Raytheon BBN Technologies, SDP Seminar, September 28, 2015.

"Reversing the Curse of Dimensionality for Security Analysis"

- NYU Poly Department of Computer Science, December 5, 2014.

"Anomaly Detection and Dimensionality Reduction"

- Workshop on Mathematics of Data Analysis in Cybersecurity, ICERM, Brown University, October 24, 2014.

"Targeted Matrix Completion"

- 51st Annual Allerton Conference on Communication, Control, and Computing, University of Illinois at Urbana-Champaign, October 3, 2013.

"Better Routing with Hyperbolic Geometry"

- Laboratory of Information, Networking and Communication Sciences (LINCS) Paris, December 19, 2012.

"A Fine Grained Distance Metric for Small Worlds"

- Université de Liège, Liège Belgium, January 11, 2013.
- Laboratory of Information, Networking and Communication Sciences (LINCS) Paris, Mar 20, 2012.

"Inferring Invisible Traffic"

- Williams College, Dept. of Computer Science, Williamstown, MA, November 22, 2013.
- Tufts University, Dept. of Computer Science, Medford, MA, November 7, 2013.
- University of Miami, Dept. of Computer Science, Miami, FL, March 2, 2011.
- International Computer Science Institute (ICSI), Berkeley, CA, July 20, 2010.
- Statistics of Networks Workshop, Isaac Newton Institute for Mathematical Sciences, Cambridge, UK, June 24, 2010.

"Network Monitoring: Challenges and Opportunities"

- Indian Institute of Technology (IIT), Delhi, India, June 9, 2008.

"Exploiting Feature Distributions in Anomaly Diagnosis"

- Dartmouth, November 28, 2005.
- MIT CSAIL-LIDS Networking and Systems Seminar, Cambridge MA, April 11, 2005.

"Network Anomaly Diagnosis"

- Boston Area Networking and Networkers Annual Summit (BANANAS-2005), BBN Technologies, Cambridge MA, January 20, 2005.

"Applying the Subspace Method to Network Traffic Analysis"

- Statistical Research Center for Complex Systems (SRCCS), Workshop on "Internet Measurement, Modeling, and Analysis," Seoul National University, Seoul, South Korea, January 11, 2005.
- Electronics and Telecommunications Research Institute (ETRI), Daejeon, South Korea, January 13, 2005.

"Diagnosing Network-Wide Traffic Anomalies"

- Columbia University, Nov. 17, 2004.
- Workshop on Internet Signal Processing (WISP), CAIDA, Nov. 11, 2004.
- Ecole Normale Supérieure, Paris, Apr. 13, 2004.
- France Telecom R&D, Paris, Mar. 18, 2004.

"Structural Analysis of Network Traffic Flows"

- Project Metropolis Meeting, École Nationale Supérieure des Télécommunications de Bretagne (ENST), Rennes, Feb. 10, 2004.
- Institute for Mathematics and its Applications, Workshop on "Measurement, Modeling and Analysis of the Internet," Minneapolis, MN, January 12, 2004.

"Virtual Landmarks for the Internet"

- Intel Research, Cambridge UK, October 1, 2003.
- Ecole Polytechnique Fédérale de Lausanne (EPFL), November 10, 2003.
- Laboratoire d'Informatique de Paris VI (LIP6), November 14, 2003.

"Sampling Biases in IP Topology Measurements"

- Séminaire Graphes, Réseaux et Modélisation, Laboratoire de Recherche en Informatique, Université d'Orsay, January 8, 2004.
- IBM T.J. Watson Labs, April 11, 2003.
- BU Statistics Colloquium, April 10, 2003.
- National Institute of Statistical Sciences, March 28, 2003.