

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

META PLATFORMS, INC.,
Petitioner

v.

WEPLE IP HOLDINGS LLC,
Patent Owner.

Case IPR2026-00081
U.S. Patent No. 12,131,357 B2
Issue Date: October 29, 2024

Title: MOBILE DEVICE STREAMING MEDIA APPLICATION

DECLARATION OF MARK CROVELLA, PH.D.

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

TABLE OF CONTENTS

	Page
I. INTRODUCTION AND QUALIFICATIONS	6
A. Qualifications and Experience	6
B. Materials Considered.....	8
II. PERSON OF ORDINARY SKILL IN THE ART	9
III. STATEMENT OF LEGAL PRINCIPLES.....	12
A. Claim Construction	12
B. Obviousness (§ 103).....	14
IV. THE '1357 PATENT	20
A. Overview of the Specification.....	20
B. The Challenged Claims	24
V. APPLICATION OF THE PRIOR ART TO CHALLENGED CLAIMS	25
A. Brief Summary and Overview of Prior Art.....	26
1. Fosnacht (EX1003)	26
2. Kirkpatrick (EX1004)	33
3. Sharma (EX1005)	40
B. Ground 1: Claims 1-26 Are Obvious Over Fosnacht in View of Kirkpatrick.....	43
1. Independent Claim 1	43
(a) “receiving, at one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;” (Claim 1[a])	44
(b) “storing the plurality of media messages on the one or more servers;” (Claim 1[b])	63

TABLE OF CONTENTS
(continued)

Page

(c)	“providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application;” (Claim 1[c]).....	64
(d)	“providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.” (Claim 1[d])	84

TABLE OF CONTENTS
(continued)

	Page
2. Claim 2: “The method of claim 1, wherein respective ones of the plurality of media messages received at the one or more servers include an image or video captured by the one or more mobile applications via a corresponding one or more cameras of the corresponding one or more mobile devices.”	88
3. Claim 3: “The method of claim 2, wherein the media message included in the first feed includes one or more of an attached image or attached audio that is not native to the image or video captured by the one or more mobile applications via the corresponding one or more cameras of the corresponding one or more mobile devices.”	91
4. Claim 4: “The method of claim 1, wherein the link is a first link, wherein the media message included in the first feed is encoded with a second link, wherein interaction with the second link via the mobile application enables access to a webpage.”	95
5. Claim 5: “The method of claim 1, wherein presentation of the first feed by the mobile application includes presentation of an advertisement prompt with the media message included in the first feed, wherein interaction with the advertisement prompt via the mobile application causes presentation of an advertisement.”	96
6. Claim 6: “The method of claim 1, wherein presentation of the first feed by the mobile application includes presentation of an item purchase prompt with the media message included in the first feed, wherein interaction with the item purchase prompt via the mobile application enables presentation of information to initiate purchase of an item.”	105

TABLE OF CONTENTS
(continued)

	Page
7. Claim 7: “The method of claim 1, wherein one or more stored media messages are associated with expiration information to determine when the one or more stored media messages can no longer be provided in the first feed.”	107
8. Independent Claim 8	110
9. Dependent Claims 9-13.....	112
10. Independent Claim 14	112
11. Dependent Claims 15-20.....	115
12. Independent Claim 21	115
13. Dependent Claims 22-26.....	118
C. Ground 2: Claims 5, 10, 18, and 23 Are Obvious Over Fosnacht in View of Kirkpatrick and Sharma	118
V. NO SECONDARY CONSIDERATIONS OF NONOBVIOUSNESS.....	124
VI. CONCLUSION.....	126

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

I, Mark Crovella, Ph.D., declare as follows:

I. INTRODUCTION AND QUALIFICATIONS

A. Qualifications and Experience

1. I received a Ph.D. in Computer Science from the University of Rochester in October 1994. I received a Master of Science degree in Computer Science, also from the University of Rochester, in 1991, as well as a Master of Science degree from State University of New York at Buffalo in 1990. I received a Bachelor of Science Degree from Cornell University in 1982.

2. For the last 18 years, I have been Professor in the Department of Computer Science at Boston University. I served as Department Chair from 2013-2018. I was an Associate Professor between 2000 and 2006 and an Assistant Professor between 1994 and 2000. Starting in 2022, I have been a Professor on the Faculty of Computing & Data Sciences at Boston University, of which I was one of the founding members. I presently serve as Chair of Academic Affairs for the Faculty of Computing & Data Sciences. I was also a Data Science Faculty Fellow between 2019 and 2022. I regularly teach a course in Computer Networks, which is a senior/graduate-level introduction to networking with a focus on the Internet.

3. My career in the area of computing and data analysis began in 1984 when I started work at Calspan Corporation in Buffalo, NY where I was eventually

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

a Senior Computer Scientist. Between 2000 and April 2002, I was the Chief Architect for WebManage Technologies of Chelmsford, MA, a company that provided solutions for managing and analyzing Internet sites such as websites and monitoring website traffic. WebManage was later acquired by Network Appliance (NetApp), where I served as Technical Director. Between February 2012 and August 2014, I also served as Chief Scientist at Guavus Inc., a telecom analytics company.

4. I have been awarded several United States patents, including the following examples:

U.S. Pat. No.	Title
10,027,347	Methods for storing and reading digital data on a set of DNA strands
8,869,276	Method and Apparatus for Whole-Network Anomaly Diagnosis and Method to Detect and Classify Network Anomalies Using Traffic Feature Distributions
8,533,167	Compressed Set Representation for Sets as Measures in OLAP Cubes
7,747,741	Method and Apparatus for Dynamic Resource Discovery and Information Distribution in a Data Network
7,346,682	System for Creating and Distributing Prioritized List of Computer Nodes Selected as Participants in a Distribution Job
6,993,587	Method and Apparatus for Election of Group Leaders in a Distributed Network
6,748,447	Method and Apparatus for Scalable Distribution of Information in a Distributed Network

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

U.S. Pat. No.	Title
6,718,361	Method and Apparatus for Reliable and Scalable Distribution of Data Files in Distributed Networks
6,370,584	Distributed Routing
6,223,205	Method and Apparatus for Assigning Tasks in a Distributed Server System

5. Based on my above-described over three decades of experience in Computing and Data Sciences and the acceptance of my publications (200+) and professional recognition by societies in my field, I believe that I am qualified to be an expert in the field of the challenged patent.

6. My Curriculum Vitae is attached as **Appendix A**.

7. I have been retained by counsel for Petitioner to provide my expert opinion in connection with the above-captioned proceeding as set forth herein. I am being compensated for my time working on this matter at my standard hourly rate plus expenses. I do not have any personal or financial stake or interest in the outcome of the present proceeding, and my compensation is not dependent on the outcome of this IPR and in no way affects the substance of my statements in this Declaration.

B. Materials Considered

8. The analysis that I provide in this Declaration is based on my education, research, and experience, as well as the documents I have considered. In forming

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

my opinions, I have read and considered U.S. Patent No. 12,131,357 B2 (“1357 patent”) (**EX1001**) and its prosecution history. I have cited to the following documents in my analysis below:

Exhibit No.	Description of Document
1001	U.S. Patent No. 12,131,357 B2 to Mary Anne Fletcher (filed January 25, 2024, issued October 29, 2024) (“ 1357 ” or “ 1357 patent ”)
1003	U.S. Patent Application Publication No. 2009/0271283 A1 to Fred M. Fosnacht et al. (filed February 13, 2009, published October 29, 2009) (“ Fosnacht ”)
1004	U.S. Patent No. 8,826,145 B1 to Ficus Kirkpatrick et al. (filed November 11, 2008, issued September 2, 2014) (“ Kirkpatrick ”)
1005	U.S. Patent Application Publication No. 2007/0156520 A1 to Jayesh Sharma (filed December 30, 2005, published July 5, 2007) (“ Sharma ”)
1006	Excerpts from <i>Microsoft Computer Dictionary</i> (5th ed. 2002)
1007	Excerpts from Erin Jansen, <i>NetLingo: The Internet Dictionary</i> (2002)
1008	Excerpts from Adi Rome et al., <i>Multimedia on Symbian OS: Inside the Convergence Device</i> (2008) (“ Rome ”)
1009	Excerpts from David Fox et al., <i>HTML Web Publisher’s Construction Kit</i> (1995) (“ Fox ”)
1010	Anthony Bruno, <i>A Bigger Hand? Apple’s iPad Offers Larger Mobile Platform for Music and Media</i> , <i>Billboard Magazine</i> , Feb. 6, 2010
1011	U.S. Patent No. 8,218,073 to Bill Cheng et al. (filed January 28, 2008, issued July 10, 2012)

II. PERSON OF ORDINARY SKILL IN THE ART

9. I understand that, under the patent laws in effect before the America Invents Act (“AIA”) of 2011, an assessment of claims of a patent filed before the

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

AIA took effect should be undertaken from the perspective of a person of ordinary skill in the art as of the claimed priority date (*i.e.*, the “time the invention was made”). I have preliminarily, for purposes of my Declaration only, assumed that priority date to be February 12, 2010, the filing date of the earliest provisional application to which the ’1357 patent claims priority (provisional application no. 61/304,331).¹

10. I have also been advised that to determine the appropriate level of a person having ordinary skill in the art, the following factors may be considered: (1) the types of problems encountered by those working in the field and prior art solutions thereto; (2) the sophistication of the technology in question, and the rapidity with which innovations occur in the field; (3) the educational level of active workers in the field; and (4) the educational level of the inventor.

¹ I express no opinion on whether the ’1357 patent is actually entitled to the benefit of the 2010 provisional application to which it claims priority (61/304,331). It is unnecessary for me to opine on this issue as the prior art references relied upon here predate that date. In the event entitlement to the provisional filing date may later become relevant, for example based on arguments later presented by Patent Owner, I reserve my right to provide opinions on that issue at that time.

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

11. The '1357 patent states that it provides “a system and process for coordinating a programmed media stream.” ('1357, 2:31-33.) The content can “relate[] to consciousness messages such as prayers, requests for charity, thoughts or similar messages” as well as “breaking news, music videos, comedy or similar content.” ('1357, 2:33-37.) “A user can view these feeds or streams of these streams by downloading a mobile application or browsing to a website.” ('1357, 2:40-42.) “The application or website can also be used to create, schedule and pay for media content airtime for a message or program.” ('1357, 2:42-44.)

12. In my opinion, a person of ordinary skill in the art would have possessed a bachelor's degree in electrical engineering, computer science, or similar field, with two years combined experience in designing and implementing applications for mobile devices, as well as server-based systems for interacting with those mobile devices over a computer network such as the Internet. A person could also have qualified as a person of ordinary skill in the art with some combination of (1) more formal education (such as a master's of science degree) and less technical experience, or (2) less formal education and more technical or professional experience.

13. My opinions regarding the level of ordinary skill in the art are based on, among other things, my experience in the fields of computer science and

engineering, my understanding of the basic qualifications that would be relevant to an engineer or scientist tasked with investigating methods and systems in the relevant area, and my familiarity with the backgrounds of colleagues and students, both past and present.

14. Although my qualifications and experience exceed those of the hypothetical person having ordinary skill in the art defined above, my analysis and opinions regarding the '1357 patent have been based on the perspective of a person of ordinary skill in the art as of February 2010.

III. STATEMENT OF LEGAL PRINCIPLES

A. Claim Construction

15. I understand from counsel that a purpose of claim construction is to determine how a person of ordinary skill in the art would have understood the claim terms to mean. I am informed that claim terms are generally given their ordinary and customary meaning, which is the meaning that the term would have to a person of ordinary skill in the art in question as of the effective filing date.

16. I understand that the person of ordinary skill in the art is deemed to read the claim term not only in the context of the particular claim in which the disputed term appears, but in the context of the entire patent, including the specification. I understand that the patent specification, under the legal principles, has been

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

described as the single best guide to the meaning of a claim term, and is thus highly relevant to the interpretation of claim terms. And I understand for terms that do not have a customary meaning within the art, the specification usually supplies the best context of understanding the meaning of those terms.

17. I further understand that other claims of the patent in question, both asserted and unasserted, can be valuable sources of information as to the meaning of a claim term. Because the claim terms are normally used consistently throughout the patent, the usage of a term in one claim can often illuminate the meaning of the same term in other claims. Differences among claims can also be a useful guide in understanding the meaning of particular claim terms.

18. I understand that the prosecution history can further inform the meaning of the claim language by demonstrating how the inventors understood the invention and whether the inventors limited the invention in the course of prosecution, making the claim scope narrower than it otherwise would be. Extrinsic evidence, such as dictionaries, may also be consulted in construing the claim terms.

19. I understand that, in *Inter Partes* Review (IPR) proceedings, a claim of a patent shall be construed using the same claim construction standard that would be used to construe the claim in a civil action filed in a U.S. district court (which I understand is called the “*Phillips*” claim construction standard), including

construing the claim in accordance with the ordinary and customary meaning of such claim as understood by one of ordinary skill in the art and the prosecution history pertaining to the patent.

20. I have been instructed by counsel to apply the “*Phillips*” claim construction standard for purposes of interpreting the claims in this proceeding, to the extent they require an explicit construction. The description of the legal principles set forth above thus provides my understanding of the “*Phillips*” standard as provided to me by counsel.

21. For purposes of my analysis here, I do not believe express claim constructions are necessary because the prior art renders the claims obvious under the plain and ordinary meaning of the claim language, and I am not aware of any claim construction disputes in the pending litigation that would require that I take such a position at this time. For the claim term “**mobile application**,” my analysis below notes the possibility that this term could be construed broadly or narrowly, but my analysis accounts for both possibilities thus making explicit construction unnecessary. I also reserve the right to respond to any claim construction positions that may be advanced by Patent Owner in this proceeding.

B. Obviousness (§ 103)

22. I understand from counsel that a patent claim is obvious if, as of the

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

effective filing date, it would have been obvious to a person having ordinary skill in the field of the technology (the “art”) to which the claimed subject matter belongs.

23. I understand that the following factors should be considered in analyzing obviousness: (1) the scope and content of the prior art; (2) the differences between the prior art and the claims; and (3) the level of ordinary skill in the pertinent art. I also understand that certain other facts known as “secondary considerations” such as commercial success, unexplained results, long felt but unsolved need, industry acclaim, simultaneous invention, copying by others, skepticism by experts in the field, and failure of others may be utilized as indicia of nonobviousness. I understand, however, that secondary considerations should be connected, or have a “nexus,” with the invention claimed in the patent at issue.

24. I understand that a reference qualifies as prior art for obviousness purposes when it is analogous to the claimed invention. The test for determining what art is analogous is: (1) whether the art is from the same field of endeavor, regardless of the problem addressed, and (2) if the reference is not within the field of the inventor’s endeavor, whether the reference still is reasonably pertinent to the particular problem with which the inventor is involved.

25. I understand that a person of ordinary skill in the art is assumed to have knowledge of all prior art. I understand that one skilled in the art can combine

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

various prior art references based on the teachings of those prior art references, the general knowledge present in the art, or common sense. I understand that a motivation to combine references may be implicit in the prior art, and there is no requirement that there be an actual or explicit teaching to combine two references. Thus, an obviousness analysis need not seek out precise teachings directed to the specific subject matter of the challenged claim because one may take into account the inferences and creative steps that a person of ordinary skill in the art would employ to combine the known elements in the prior art in the manner claimed by the patent at issue. I understand that one should avoid “hindsight bias” and *ex post* reasoning in performing an obviousness analysis. But this does not mean that a person of ordinary skill in the art for purposes of the obviousness inquiry does not have recourse to common sense.

26. I understand that when determining whether a patent claim is obvious in light of the prior art, neither the particular motivation for the patent nor the stated purpose of the patentee is controlling. The primary inquiry has to do with the objective reach of the claims, and that if those claims extend to something that is obvious, then the entire patent claim is invalid.

27. I understand one way that a patent can be found obvious is if there existed at the time of the invention a known problem for which there was an obvious

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

solution encompassed by the patent's claims. I understand that a motivation to combine various prior art references to solve a particular problem may come from a variety of sources, including market demand or scientific literature. I understand that a need or problem known in the field at the time of the invention can also provide a reason to combine prior art references and render a patent claim invalid for obviousness. I understand that familiar items may have obvious uses beyond their primary purpose, and that a person of ordinary skill in the art will be able to fit the teachings of multiple prior art references together like the pieces of a puzzle. I understand that a person of ordinary skill is also a person of at least ordinary creativity. I understand when there is a design need or market pressure to solve a problem and there are a finite number of identified, predictable solutions, a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this finite number of predictable solutions leads to the anticipated success, I understand that the invention is likely the product of ordinary skill and common sense, and not of any sort of innovation. I understand that the fact that a combination was obvious to try might also show that it was obvious, and hence invalid, under the patent laws. I understand that if a patent claims a combination of familiar elements according to known methods, the combination is likely to be obvious when it does no more than yield predictable results. Thus, if a person of

ordinary skill in the art can implement a predictable variation, an invention is likely obvious. I understand that combining embodiments disclosed near each other in a prior art reference would not ordinarily require a leap of inventiveness.

28. I understand that obviousness may be shown by demonstrating that it would have been obvious to modify what is taught in a single piece of prior art to create the patented invention. Obviousness may also be shown by demonstrating that it would have been obvious to combine the teachings of more than one item of prior art. I understand that a claimed invention may be obvious if some teaching, suggestion, or motivation exists that would have led a person of ordinary skill in the art to combine the invalidating references. I also understand that this suggestion or motivation may come from the knowledge of a person having ordinary skill in the art, or from sources such as explicit statements in the prior art. I understand that when there is a design need or market pressure, and there are a finite number of predictable solutions, a person of ordinary skill may be motivated to apply common sense and his skill to combine the known options in order to solve the problem.

29. I understand the following are examples of approaches and rationales that may be considered in determining whether a piece of prior art could have been combined with other prior art or with other information within the knowledge of a person having ordinary skill in the art:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

- (1) Some teaching, motivation, or suggestion in the prior art that would have led a person of ordinary skill to modify the prior art reference or to combine prior art reference teachings to arrive at the claimed invention;
- (2) Known work in one field of endeavor may prompt variations of it for use in the same field or a different field based on design incentives or other market forces if the variations would have been predictable to a person of ordinary skill in the art;
- (3) Combining prior art elements according to known methods to yield predictable results;
- (4) Applying a known technique to a known device, method, or product ready for improvement to yield predictable results;
- (5) Applying a technique or approach that would have been “obvious to try” (choosing from a finite number of identified, predictable solutions, with a reasonable expectation of success);
- (6) Simple substitution of one known element for another to obtain predictable results; or
- (7) Use of a known technique to improve similar products, devices, or methods in the same way.

30. I understand that, when determining whether a claimed combination is

obvious, the correct analysis is not whether one of ordinary skill in the art, writing on a blank slate, would have chosen the particular combination of elements described in the claim. Instead, I understand the correct analysis considers whether one of ordinary skill, facing the wide range of needs created by developments in the field of endeavor, would have seen a benefit to selecting the combination claimed.

31. I understand that the test for obviousness is not whether the features of a secondary reference may be bodily incorporated into the structure of the primary reference. The test for obviousness, in other words, is not whether the references could be physically combined but whether the claimed inventions are rendered obvious by the teachings of the prior art as a whole.

IV. THE '1357 PATENT

A. Overview of the Specification

32. The '1357 patent states that it provides “a system and process for coordinating a programmed media stream.” ('1357, 2:31-33.) The content can “relate[] to consciousness messages such as prayers, requests for charity, thoughts or similar messages” as well as “breaking news, music videos, comedy or similar content.” ('1357, 1:33-37.) “A user can view these feeds or streams of these streams by downloading a mobile application or browsing to a website.” ('1357, 2:40-42.) “The application or website can also be used to create, schedule and pay for media

content airtime for a message or program.” (’1357, 2:42-44.)

33. Figure 1 of the ’1357, reproduced below, “illustrates an example media coordination system **100** according to the present invention which allows network connected devices, such as mobile devices **106** (cell phones, portable music players, tables, laptops, etc.), to coordinate times for specific media messages/programs to be streamed to other application users (e.g., a feed of message posts or a stream of real-time media data)”:




Figure 1

(’1357, Fig. 10; 3:50-56.) The system includes a “computer server **102** (e.g., a processor, RAM, a hard drive, an operating system, web serving software and database software)” with which “mobile devices **106**” can communicate via “a wide area network **104** such as the internet or wireless cell phone data network.” (’1357, 3:56-62.)

34. In terms of general system operation, Figure 10 of the ’1357 patent, reproduced below, “illustrates a process of coordinating media content times according to an embodiment of the present invention”:




Figure 10

(’1357, Fig. 10; 3:5-7.) As shown in the figure above, the process includes selecting

and uploading a desired media message and associated media message time, storing the media message data at a server in a database, and then providing the media message data at the scheduled message time by, for example, making the message data “available for download to each mobile device **106** running the application.” (See ’1357, 4:43-61.)

35. Figure 2 of the ’1357 patent, reproduced below, “illustrates a current media content interface according to an embodiment of the present invention”:




Figure 2

(’1357, Fig. 2; 2:55-56.) “[C]urrent media message interface **110** [] displays the media message currently scheduled by the server **102**.” (’1357, 5:8-11.) Among

other things, the interface also “includes a comments button **111** that displays a comment interface for leaving messages related to the current media message.” (’1357, 5:24-26.) And “selecting the switch stream button **115** allows the user to switch between” “multiple feeds or streams are available to the mobile application” such as “specific breaking news, religion, music, or pets.” (’1357, 5:29-33.)

36. I discuss additional aspects of the ’1357 patent in my analysis of the claims below.

B. The Challenged Claims

37. This Declaration addresses claims 1-26 of the ’1357 patent. Independent claim 1 is representative and recites:

1. A method, comprising:
 - [a] receiving, at one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;
 - [b] storing the plurality of media messages on the one or more servers;
 - [c] providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first

feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and

- [d] providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.

(’1357, 7:49-8:16 (Claim 1; bracketed notation (e.g., [a]) added).)

38. I address the claims further in my detailed analysis in **Part V** below.

V. APPLICATION OF THE PRIOR ART TO CHALLENGED CLAIMS

39. I have reviewed and analyzed the prior art references and materials listed in **Part I.B** above. In my opinion the claims of the ’1357 patent would have

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

been obvious to a person of ordinary skill in the art based on the following combinations of the prior art.

Ground	Claims	Basis for Challenge Under § 103
1	1-26	Fosnacht (EX1003) in view of Kirkpatrick (EX1004)
2	5, 10, 18, 23	Fosnacht and Kirkpatrick, in further view of Sharma (EX1005)

40. I am informed by counsel that each of the references cited in the grounds above qualifies as prior art to the challenged claims because each reference was filed and/or published before the earliest filing date to which the '1357 patent can claim priority (February 12, 2010).

A. Brief Summary and Overview of Prior Art

1. Fosnacht (EX1003)

41. Fosnacht is a published patent application entitled “Network Media Distribution.” Fosnacht discloses a media distribution system in which a web server allows distribution of “media items” to users through a web site. (Fosnacht, Abstract.) I am informed that Fosnacht qualifies as prior art because it was filed on February 13, 2009 and published on October 29, 2009, both dates being before the priority date of the challenged claims.

42. Fosnacht discloses a web site and service similar in many respects to well-known Internet services, such as YouTube, that allow users to upload videos

and other types of media items that can be shared with other users. (Fosnacht, ¶¶0003, 0006.) Figure 1 discloses the primary components of media distribution system 100:




FIG. 1

(Fosnacht, Fig. 1.) “FIG. 1 is a block diagram of an exemplary media distribution system **100**. Media distribution system **100** includes web server **102**, network **104**, publisher computing system **106**, user computing systems **108**, and third-party service computing system **110**.” (Fosnacht, ¶0066.) “Web server **102** is a computing system that is in data communication with network **104**.” (Fosnacht, ¶0067.) It is thus capable of communicating with the other computing systems shown in Figure 1. (*Id.*) “Some embodiments include multiple Web servers **102**.” (*Id.*) Fosnacht explains that “[i]n some embodiments, Web server **102** stores and transmits web site data. The web site data is requested by a computing system, such as user computing system **108**.” (Fosnacht, ¶0075.)

43. As shown in Figure 1, web server **102** includes media database **103**. “Media database **103** stores digital data including a plurality of media items, such as audio recordings, video recordings, or other digital media.” (Fosnacht, ¶0068.) Although Figure 1 shows media database **103** as part of web server **102**, in another embodiment “media database **103** is a database server that is in data communication with web server **102**.” (*Id.*)

44. Figure 1 also shows multiple user computing systems **108**. “User computing system **108** is a computing system through which a user communicates with network **104**. In some embodiments user computing system **108** is a personal

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

computer, a laptop computer, a tablet computer, a handheld computer, a personal digital assistant, a cell phone, or other digital devices.” (Fosnacht, ¶0077.)

45. One of the key features of Fosnacht is the ability of users to upload a media item (such as a video) to web server **102** through user computing system **108**. (Fosnacht, ¶¶0349-0355.) Figure 41 shows the upload user interface:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

The screenshot shows the 'mycatholicvoice' website interface. At the top, a navigation bar includes links for Media, Groups, Blogs, People, and AboutUs, along with a search bar and user options like 'Upload johndoe (0)' and 'Logout'. A secondary banner reads: 'MyCatholicVoice Beta site (that's Web-speak for "we're working out the kinks") - We welcome your comments and feedback'. The main content area is titled 'Upload a File' (4100) and contains several sections: 'Upload a File' (1302) with instructions on upload times and review; 'Select File Type' (4104) with icons for Music, Video, Photo, Document, PodCast, and Spoken Word; 'Attach to a group' (4106) with a dropdown menu currently set to 'No Group'; 'Describe' (4108) with fields for Title, Description, and a link to 'locate this media item'; 'Select Category' (4110) with a grid of checkboxes for various categories including Age / Gender, Catholic Life, Events, Lifestyle, Other, Pop Culture, School / Universities, and Spirituality; 'Allow other users to download this item?' (4112) with a checkbox for 'Yes, allow others to download this item.'; and 'Select File' (4114) with 'Browse' (4120) and 'Upload' (4122) buttons. On the right side, there is a quote by Saint Jane Frances de Chantal: 'It is a great consolation to surrender ourselves wholly to God and to know that he sees and penetrates to the very secret of our hearts.' The footer contains links for 'NewEvents', 'Advisory Board', 'resources to help you get started.', 'ACTA Publications', and 'Medic Media / Contemplative-Life'.

FIG. 41

(Fosnacht, Fig. 41.) “FIG. 41 is a screen shot of an example upload page of the example web site of FIG. 11.” (Fosnacht, ¶0056.) Fosnacht explains that

“[i]nstructions for uploading a file are provided in general instruction section **4102**. After reviewing the instructions, the user selects a file type from file type selection section **4104**. Examples include music, video, photo, document, podcast, and spoken word.” (Fosnacht, ¶0350.) Other sections in Figure 41 provide additional options for the uploaded media item, which are mostly self-explanatory. Group attachment section **4106** allows the user to associate the media item with a group. (Fosnacht, ¶0351.) Description section **4108** allows the user to provide a title and description for the media item. (Fosnacht, ¶0352.) Category selection section **4110** allows the user to associate the media item with one or more categories. (Fosnacht, ¶0353.) Publication settings section **4112** allows the user to specify whether other users should be able to download the media item (in addition to merely being able to view the item through a media viewer). (Fosnacht, ¶0354.) File section **4114** allows the user to select a file associated with the media item. (Fosnacht, Fig. 41.)

46. Once a media item has been uploaded, it can be made available to other users through a number of different features of the web site. For example, Figure 4 shows a welcome page **402** that includes media section **408** that allows the user to access featured, most viewed, or newest media items:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 4

(Fosnacht, Fig. 4.) Media section 408, as shown, “includes a featured tab, most viewed tab, and a newest tab.” (Fosnacht, ¶0095.) “If the featured tab is selected, welcome page 402 displays a set of featured content in media section 408. This is also the default tab. If the most viewed tab is selected, welcome page 402 displays a set of the most viewed content, arranged in an order from the content that has been viewed the most to the content that has been viewed the least. If the newest tab is selected, then welcome page 402 displays a set of the content that was most recently

added, arranged in an order from the newest content to the oldest content.” (Fosnacht, ¶0095.) Fosnacht discloses several other ways in which media items can be presented to users, which I will discuss in more detail in my discussion of the claim limitations below. As demonstrated below, Fosnacht discloses a media distribution system markedly similar to the one later described in the ’1357 patent, but Fosnacht provides far more detail about the implementation of its disclosed system. The similarities between the two are further highlighted by the fact that, like the ’1357 patent, Fosnacht focuses on embodiments for distributing prayers and other religious-themed messages. (*Compare* Fosnacht, *e.g.*, Figs. 4-11, ¶¶0118, 0158, 0274 *with* ’1357, 1:58-2:35, 7:14-16.)

2. Kirkpatrick (EX1004)

47. Kirkpatrick is a patent entitled, “Unified Web and Application Framework.” Kirkpatrick discloses techniques for “integrating software models defining native applications and access to web content.” (Kirkpatrick, 1:46-48.) I am informed that Kirkpatrick qualifies as prior art to the ’1357 patent because it issued from an application filed on November 11, 2008, before the priority date of the challenged claims.

48. My Declaration cites Kirkpatrick for limitations that recite use of one or more “**mobile applications**” executing on one or more mobile devices. As I will

explain below, these limitations are disclosed and rendered obvious by Fosnacht alone under the plain and ordinary meaning of “mobile application,” which would have been understood to include a software application such as a web browser executing on a mobile device. I nevertheless have cited Kirkpatrick in the event a narrower interpretation limited to so-called “native” mobile applications is applied to the challenged claims.

49. Kirkpatrick describes techniques for creating what are commonly called “native applications” for accessing content and features available through a web site. (Kirkpatrick, 3:26-29.) The term “native application” typically refers to a software application (usually for a mobile device) that was written specifically for the particular type of device or operating system on which it executes. A key feature of “native” applications is that they can typically execute and function without requiring use of a web browser application on the device. I note that native applications can incorporate web browser functionality, for example by allowing a user to view certain web pages or web page content from within the native application. This functionality is typically provided to native applications through application programming interfaces (APIs) offered by mobile device operating systems, which allow native applications to render and display HTML and web pages without requiring that the user activate the web browser application on the

mobile device.

50. Today, native applications are used by hundreds of millions of smartphone users each day; for example, mobile applications such as YouTube, TikTok, Instagram, eBay, among others, are native applications that can provide functionality to users without requiring use of a web browser or access to a web site. Most of the “apps” available from the App Store (for Apple iOS devices) or the Google Play Store (for Android devices) are considered native applications.

51. One of the key reasons some developers prefer native applications over applications provided via a web browser is that native applications can be more tailored to the particular characteristics of the device. As Kirkpatrick explains, a native application can, among other things, “format the content in a way that is more conducive to display on a small screen of the mobile device (e.g., a screen of a cell phone).” (Kirkpatrick, 3:35-37.) Web-based applications, in contrast, typically rely on standard technologies used by a web browser such as HTML, CSS, and JavaScript, which typically do not provide the same level of access to low-level operating system features that is available to native applications.

52. Mobile device adoption exploded after 2007 due to widespread availability of modern smartphones such as Apple’s iPhone and devices running Google’s Android operating system, and the opening of app “stores” allowing

developers to create and widely distribute applications for these types of devices. These developments caused many developers to begin transitioning from web sites to native applications for some of their mobile device users. Kirkpatrick describes one technique for facilitating that transition.

53. Kirkpatrick, which was assigned to Google (the makers of the popular Android mobile operating system) describes a unified framework for providing web-based features through “native” applications (which Kirkpatrick also refers to as “non-browser” applications). Figure 1A shows the primary components of the disclosed system in one embodiment:




FIG. 1A

(Kirkpatrick, Fig. 1A.)

54. Figure 1A above shows “cell phone **102** that has a unified framework system **104** that can integrate native applications and web content. The system **100** also includes a web server **106** that stores one or more web pages **108** that have content **110** (e.g., text, images, video, etc.).” (Kirkpatrick, 4:10-15.) In the example shown above, the user has selected a Uniform Resource Identifier (URI) for the website, www.cnn.com. (Kirkpatrick, 4:15-19.) Kirkpatrick explains that “[t]he unified framework system **104**,” which is contained in cell phone **102**, “can have

several components that are activated when the user selects the URI. The components may include an activity monitor **114** that detects the selection of the URI, an activity manager **116** having an activity register **118**, and an application **120** (shown here as an application titled ‘Headlines’).” (Kirkpatrick, 4:19-25.)

55. “Once the activity monitor **114** detects the selection of the URI, the monitor **114** can forward the URI to the activity manager **116**. The activity manager can access the activity register **118** to determine which application is associated with the URI.” (Kirkpatrick, 4:26-30.) In Figure 1A, the exemplary native/non-browser application **120** is called “Headlines.” The URI is then forwarded to the application **120**, which requests web content from the web server **106** as shown in arrow **122**, which returns web content **110** to the application **120** as shown in arrow **124**. (Kirkpatrick, 4:30-41.) “The Headlines application, in some implementations, can reformat or otherwise process the web content and display the results to the user interface **112** as indicated by the arrow **126**.” (Kirkpatrick, 4:46-48.)

56. Figure 1B shows another embodiment in which the native/non-browser application **160** is be used to upload a photo to web server **156**:




FIG. 1B

(Kirkpatrick, Fig. 1B.)

57. Figure 1B is similar in some respects to Figure 1A except that Figure 1B describes a non-browser application **160** called “FlickU” that accesses the well-known photo sharing website, www.flickr.com. (Kirkpatrick, 4:52-60.) As shown, photo **164** is taken using camera **162** of mobile device **152**, which is then stored in the local data store **166** of the device. (Kirkpatrick, 4:66-5:16.) The FlickU application **160** then requests information from web server **156** about uploading photos, as shown by arrow **178**. (Kirkpatrick, 5:29-39.) The FlickU application **160**

then retrieves the photo from the local data store **166**, and uploads photo **184** to web server **156** as shown by arrow **190**. (Kirkpatrick, 5:40-63.)

58. Kirkpatrick discloses yet another embodiment in which a non-browser application called “YouTube Enhanced” is provided for accessing the popular YouTube video sharing website. In this embodiment, all URIs that point to the YouTube website (*i.e.* all URIs that include “www.youtube.com”) can be forwarded to YouTube Enhanced for processing. (Kirkpatrick, 6:21-35.)

59. I will provide more information about Kirkpatrick in my discussion of the claim limitations below.

3. Sharma (EX1005)

60. Sharma is a published patent application entitled “Using Search Query Information to Determine Relevant Ads for a Landing Page of an Ad.” Sharma discloses an Internet advertising technique for delivering relevant ads to users. (Sharma, Abstract, ¶0009.) I am informed that Sharma qualifies as prior art to the ’1357 patent because it was filed on December 30, 2005 and published on July 7, 2007, both dates being several years before the priority date of the challenged claims.

61. My Declaration cites Sharma in connection with **Ground 2**, which addresses dependent claims 5, 10, 18, and 23. These claims depend from claims 1, 8, 14, and 21, respectively, and each recite that “**presentation of the first feed by**

the mobile application includes presentation of an advertisement prompt with the media message included in the first feed, wherein interaction with the advertisement prompt via the mobile application causes presentation of an advertisement.” As shown below, in my opinion, these claims are disclosed and rendered obvious by Fosnacht in view of Kirkpatrick, but I have cited Sharma in the event it is argued that Fosnacht does not provide sufficient detail about the operation of its disclosed banner advertisements. Sharma discloses a technique for allowing an advertiser to deliver relevant ads to users based on search query information provided by a user. (Sharma, Abstract, ¶0009.)

62. As was well-known to skilled artisans, and widely known to laypersons who had used the web before the priority date, websites frequently presented one or more Internet advertisements such as banner ads that, when clicked (for example through hypertext links), would cause a user’s web browser to present a web page associated with the advertiser. (Sharma, ¶0049.) This advertiser page is frequently called a “landing” page. (Sharma, ¶¶0049, 0059; *see also* Erin Jansen, *NetLingo: The Internet Dictionary* (2002) (EX1007), p.237 (“**landing page** Another name for a page on a Web site to which people are referred from other Web sites, especially from ad banners... Usually, a company will design one landing page per banner so the follow-up content matches the sell on the particular banner ad.”) (boldface in

original); *see also id.*, p.29 (definition of “ad banner”).)

63. My Declaration cites Sharma for this basic and well-known sequence. For example, Figure 6 shows an exemplary web page **600** that, in one embodiment, shows search results **610** and one or more ads **620**. (Sharma, ¶0067.)





FIGURE 6

(Sharma, Fig. 6 (highlighting added).) Sharma explains that ads may be, for example, “so-called banner ads,” among other types. (Sharma, ¶0033.)

64. As shown in Figure 6 in highlighting, “[a] user has selected, as indicated by click **630**, ad **1**.” (Sharma, ¶0067.) “Ad **1** includes a reference (e.g., a hypertext link) to a landing Web page.” (Sharma, ¶0068.) Accordingly, clicking on ad **630** brings up the landing page shown in Figure 7:



(Sharma, Fig. 7.) “FIG. 7 illustrates an exemplary landing Web page **700**. It **700** includes content **710** and additional ads **720**.” (Sharma, ¶0068.)

65. I will provide more information about Sharma in my discussion of the claim limitations below.

B. Ground 1: Claims 1-26 Are Obvious Over Fosnacht in View of Kirkpatrick

1. Independent Claim 1

66. The preamble of claim 1 recites “[a] **method, comprising....**” The prior art discloses and renders obvious a method of distributing media messages as described below.

- (a) **“receiving, at one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;” (Claim 1[a])**

67. Fosnacht discloses a method that includes the step of receiving, at one or more servers, video messages from applications executable on one or more computing devices. Figure 1 of Fosnacht discloses the primary components of media distribution system **100**:




FIG. 1

(Fosnacht, Fig. 1.) “FIG. 1 is a block diagram of an exemplary media distribution system **100**. Media distribution system **100** includes web server **102**, network **104**, publisher computing system **106**, user computing systems **108**, and third-party

service computing system **110**.” (Fosnacht, ¶0066.) “Web server **102** is a computing system that is in data communication with network **104**.” (Fosnacht, ¶0067.) It is thus capable of communicating with the other computing systems shown in Figure 1. (*Id.*) “Some embodiments include multiple Web servers **102**.” (*Id.*) “In some embodiments, Web server **102** stores and transmits web site data. The web site data is requested by a computing system, such as user computing system **108**.” (Fosnacht, ¶0075.)

68. Figure 1 shows multiple user computing systems **108**. “User computing system **108** is a computing system through which a user communicates with network **104**. In some embodiments user computing system **108** is a personal computer, a laptop computer, a tablet computer, a handheld computer, a personal digital assistant, a cell phone, or other digital devices.” (Fosnacht, ¶0077.) Figure 1 also shows publisher computing system **106**, which is “a computing system through which a commercial publisher communicates with network **104**,” and like user computer system **108**, can include “a personal computer or a handheld computer.” (Fosnacht, ¶0072.)

69. Fosnacht thus discloses “**one or more servers**” in the form of web server **102**, shown in Figure 1 above as also including media database **103**, which in turn stores “**a plurality of media messages**,” which take the form of media items

such as video recordings. (Fosnacht, ¶¶0067, 0068.) “Media database **103** stores digital data including a plurality of media items, such as audio recordings, video recordings, or other digital media.” (Fosnacht, ¶0068.) Although Figure 1 shows media database **103** as part of web server **102**, in another embodiment “media database **103** is a database server that is in data communication with web server **102**.” (*Id.*)

70. The step of “**receiving, at one or more servers, a plurality of media messages,**” occurs when one or more users or publishers upload media items to web server **102**, which results in the media items being stored in media database **103**. (Fosnacht, ¶¶0349-0355.) Figure 41 shows an exemplary user interface allowing upload of a media item:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

The screenshot shows the 'mycatholicvoice' website interface. At the top, there is a navigation bar with links: Media, Groups, Blogs, People, AboutUs, and a search bar. Below the navigation bar, a banner reads: 'MyCatholicVoice Beta site (that's Web-speak for "we're working out the kinks") - We welcome your comments and feedback'. The main content area is divided into several sections:

- Upload a File (1302):** A text box explaining the upload process: 'Uploads typically take 1 to 5 minutes per MB on a high-speed internet connection. Once uploaded, your file will be converted into one of our standard file formats. The item will go through our review process to ensure it meets our Terms of Use. This process can sometimes take up to 24 hours. For more details about uploading media see the Help Center.'
- Select File Type (4104):** A section titled 'What type of file are you uploading?' with radio buttons for Music, Video, Photo, Document, PodCast, and Spoken Word.
- Attach to a group (4106):** A section titled 'By attaching this media to a group that you have permission to, this item will now show under that group's media. You can only attach to groups that you are a member of and have media access too.' It includes a dropdown menu labeled 'Assign this post to:' with 'No Group' selected.
- Describe (4108):** A section titled 'what are these?' with input fields for Title, Description, and a link to 'locate this media item.'
- Select Category (4110):** A section titled 'what are these?' with multiple columns of checkboxes for categories: Age / Gender, Catholic Life, Events, Lifestyle, Other, Pop Culture, School / Universities, and Spirituality.
- Allow other users to download this item? (4112):** A section with a checkbox labeled 'Yes, allow others to download this item.'
- Select File (4114):** A section with 'Browse' and 'Upload' buttons, and a text input field.

On the right side of the page, there is a sidebar with the 'mycatholicvoice' logo, a landscape image, and a quote from Saint Jane Frances de Chantal: 'It is a great consolation to surrender ourselves wholly to God and to know that he sees and penetrates to the very secret of our hearts.'

At the bottom of the page, there is a footer with links: NewEvents, Advisory Board, resources to help you get started, ACTA Publications, and Media Media / Contemplative-Life.

FIG. 41

(Fosnacht, Fig. 41.) "FIG. 41 is a screen shot of an example upload page of the example web site of FIG. 11." (Fosnacht, ¶0056.)

71. The upload user interface of Figure 41 is not only utilized by users but also by publishers. (Fosnacht, ¶0356 (“... the publisher uploads a media item through user upload page **4100** (shown in FIG. 41) or a similar page in some embodiments.”).) Fosnacht explains that “[i]nstructions for uploading a file are provided in general instruction section **4102**. After reviewing the instructions, the user selects a file type from file type selection section **4104**. Examples include music, video, photo, document, podcast, and spoken word.” (Fosnacht, ¶0350.) Other sections in Figure 41 provide additional options for the uploaded media item, which are mostly self-explanatory. Group attachment section **4106** allows the uploader to associate the media item with a group. (Fosnacht, ¶0351.) Description section **4108** allows the uploader to provide a title and description for the media item. (Fosnacht, ¶0352.) Category selection section **4110** allows the uploader to associate the media item with one or more categories. (Fosnacht, ¶0353.) Publication settings section **4112** allows the uploader to specify whether other users should be able to download the media item (in addition to merely being able to view the item through a media viewer). (Fosnacht, ¶0354.) File section **4114** allows the uploader to select a file associated with the media item. (Fosnacht, Fig. 41.)

72. The claimed step of “**receiving**” a “**media message**” occurs in Fosnacht when the media item (such as a video) is uploaded to media database **103** which, as

noted, is part of web server **102** (“**one or more servers**”). Fosnacht explains that, after the user selects a file for upload, “[t]he upload button **4122** is then selected to upload the selected file to the server and to save the media item and associated information in the media database (e.g., **103** shown in FIG. 1).” (Fosnacht, ¶0355; *see also id.*, ¶0068 (“Web server **102** includes media database **103**.”).) It would also have been apparent and obvious that Fosnacht further discloses receiving and storing in a server “**a plurality of media messages**,” such as multiple video media items uploaded by one or more users. (Fosnacht, ¶¶0068 (“Media database **103** stores digital data including a plurality of media items, such as audio recordings, video recordings, and other digital media.”), 0311 (“Media content tabs **3506** display media content associated with the user. Examples include those media items tagged by the user as being favorite, media items uploaded by the user, and media items purchased by the user.”).) Fosnacht therefore discloses “**receiving, at one or more servers, a plurality of media messages**.”

73. Fosnacht also discloses and renders obvious that the plurality of media messages can be received from “**one or more mobile applications executable on a corresponding one or more mobile devices**.” As noted, Figure 1 shows multiple user computing systems **108**, each of which “is a computing system through which a user communicates with network **104**.” (Fosnacht, ¶0073.) Fosnacht further

discloses that user computing system **108** includes “a browser software application,” such as the well-known Microsoft Internet Explorer web browser. (Fosnacht, ¶0075.) “The browser operates to display a user interface at user computing system **108** to display the web site to the user.” (*Id.*; *see also id.*, ¶0089.) Fosnacht further explains user computing system **108** can be a mobile device, such as “a laptop computer, a tablet computer, a handheld computer, a personal digital assistant, a cell phone,” or “other digital devices configured to communicate digital data across network **104**.” (Fosnacht, ¶0073.) Figure 1 as noted also shows publisher computing system **106**, which is “a computing system through which a commercial publisher communicates with network **104**,” and like user computer system **108**, can include “a handheld computer.” (Fosnacht, ¶0072.) The claimed “**one or more mobile applications executable on a corresponding one or more mobile devices application**,” therefore, take the form of at least one web browser application executable on at least one mobile computing system, such as a laptop computer, tablet, handheld computer, etc. And as explained above, the web browser on a mobile device can be used to upload media items to web server **102**, for example, by utilizing the upload interface in Figure 41 discussed above. Fosnacht therefore discloses “**receiving, at one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more**

mobile devices,” as claimed.

74. I observe that the '1357 patent does not provide any express definition of “**mobile application.**” The term “**application**” in computer science would have been generally understood to refer to a software program designed to perform a specific task. (See *Microsoft Computer Dictionary* (5th ed. 2002) (EX1006), p.31 (“**application** *n.* A program designed to assist in the performance of a specific task, such as word processing, accounting, or inventory management.”); see also Erin Jansen, *NetLingo: The Internet Dictionary* (2002) (EX1007), p.42 (“**application aka ‘app’ ‘program’ ‘software’** Used interchangeably with program and software, this is a general term for a program that performs specific tasks, such as word processing, database management, e-mail sending or retrieval, or Web browsing.”) (boldface in original).) A web browser executing on a mobile device, such as a laptop computer, handheld computer, tablet, cell phone, etc., qualifies as a mobile “application” based on its plain and ordinary meaning. I note that this conclusion is consistent with claims of the '1357 patent that expressly recite that interaction with a second link “**via the mobile application enables access to a webpage.**” ('1357, Claims 4, 9, 17, 22.) Enabling access to a webpage is precisely what a mobile web browser application does.

75. Nevertheless, I am informed by counsel that Patent Owner may attempt

to argue that “**mobile application**” should be more narrowly construed as a software program (other than a general-purpose web browser application) written specifically for a particular mobile device, which was commonly referred to as a “native” mobile application by persons of ordinary skill in the art. I note that nothing in the patent requires that a “**mobile application**” be a non-browser or native application, but in the event the Patent Owner makes this argument, this claim limitation would still have been obvious over Fosnacht in view of Kirkpatrick.

76. Kirkpatrick describes a series of techniques for creating what are commonly called “native applications” for accessing content and features available through a web site. (Kirkpatrick, 3:26-29.) The term “native application,” as noted, refers to a software application (typically for a mobile device) written specifically for the particular type of device on which it executes. A key feature of “native” applications is that they can typically execute and function without requiring use of a web browser application (although many native applications can also utilize the device’s general purpose web browser to present certain pages). Native applications are used by millions of Apple iPhone and Google Android smartphone users each day; for example, mobile applications such as YouTube, TikTok, eBay, among others, are examples of native applications that can provide functionality to users without requiring a web browser or access to a web site. Most of the “apps” available

from the App Store (for Apple iOS devices) or the Google Play Store (for Android devices) are considered native applications.

77. One of the key reasons some developers prefer native applications over applications provided via a web browser is that native applications can be more tailored to the particular characteristics of the device. As Kirkpatrick explains, a native application can, among other things, “format the content in a way that is more conducive to display on a small screen of the mobile device (e.g., a screen of a cell phone).” (Kirkpatrick, 3:35-37.) Web-based applications, in contrast, typically rely on standard technologies used by a web browser such as HTML, CSS, and JavaScript, which typically do not provide the same level of access to low-level operating system features that is available to native applications.

78. Mobile device adoption exploded after 2007 due to widespread availability of modern smartphones such as Apple’s iPhone and devices running Google’s Android operating system, and the opening of app “stores” allowing developers to create and widely distribute applications for these types of devices. These developments caused many developers to begin transitioning from web sites to native applications for some of their mobile device users. Kirkpatrick describes one technique for facilitating that transition.

79. Kirkpatrick, which was assigned to Google (the makers of the popular

Android mobile operating system), describes a unified framework for providing web-based features through “native” applications (which Kirkpatrick also refers to as “non-browser” applications). Figure 1A shows the primary components of the disclosed system in one embodiment:




FIG. 1A

(Kirkpatrick, Fig. 1A.)

80. Figure 1A above shows “cell phone 102 that has a unified framework system 104 that can integrate native applications and web content. The system 100 also includes a web server 106 that stores one or more web pages 108 that have

content **110** (e.g., text, images, video, etc.).” (Kirkpatrick, 4:10-15.) In the example shown above, the user has selected a Uniform Resource Identifier (URI) for the website, www.cnn.com. (Kirkpatrick, 4:15-19.) Kirkpatrick explains that “[t]he unified framework system **104**,” which is contained in cell phone **102**, “can have several components that are activated when the user selects the URI. The components may include an activity monitor **114** that detects the selection of the URI, an activity manager **116** having an activity register **118**, and an application **120** (shown here as an application titled ‘Headlines’).” (Kirkpatrick, 4:19-25.)

81. “Once the activity monitor **114** detects the selection of the URI, the monitor **114** can forward the URI to the activity manager **116**. The activity manager can access the activity register **118** to determine which application is associated with the URI.” (Kirkpatrick, 4:26-30.) In Figure 1A, the exemplary native/non-browser application **120** is called “Headlines.” The URI is then forwarded to the application **120**, which requests web content from the web server **106** as shown in arrow **122**, which returns web content **110** to the application **120** as shown in arrow **124**. (Kirkpatrick, 4:30-41.) “The Headlines application, in some implementations, can reformat or otherwise process the web content and display the results to the user interface **112** as indicated by the arrow **126**.” (Kirkpatrick, 4:46-48.)

82. Figure 1B shows another embodiment in which the native/non-browser

application **160** is be used to upload a photo to web server **156**:




FIG. 1B

(Kirkpatrick, Fig. 1B.)

83. Figure 1B is similar in some respects to Figure 1A except that Figure 1B describes a different exemplary non-browser application **160** called “FlickU” that accesses the well-known photo sharing website, www.flickr.com. (Kirkpatrick, 4:52-60.) As shown, photo **164** is taken using camera **162** of mobile device **152**, which is then stored in the local data store **166** of the device. (Kirkpatrick, 4:66-5:16.) The FlickU application **160** then requests information from web server **156**

about uploading photos, as shown by arrow **178**. (Kirkpatrick, 5:29-39.) The FlickrU application **160** then retrieves the photo from the local data store **166**, and uploads photo **184** to web server **156** as shown by arrow **190**. (Kirkpatrick, 5:40-63.)

84. Kirkpatrick discloses a further embodiment in which a non-browser application called “YouTube Enhanced” is provided for accessing the popular YouTube video sharing website. In this embodiment, all URIs that point to the YouTube website (*i.e.* all URIs that include “www.youtube.com”) can be forwarded to YouTube Enhanced for processing. (Kirkpatrick, 6:21-35.)

85. ***Rationale and Motivation to Combine (Fosnacht and Kirkpatrick):*** It would have been obvious to combine Fosnacht with Kirkpatrick. The combination would have predictably resulted in the media distribution system of Fosnacht, adapted to allow a mobile device to access the system through a native/non-browser application instead of through a web browser application. As applied to limitation 1[a], therefore, the proposed combination would have resulted in web server **102** of Fosnacht receiving media messages “**from one or more mobile applications executable on a corresponding one or more mobile devices,**” *i.e.*, through one or more native/non-browser applications executing on a mobile user computing system **108**. Accordingly, the features of the web site in Fosnacht and relied upon for claim 1, including the upload user interface in Figure 41 and the upload of a media message

to web server **102** described above, as well as the other user interface features of Fosnacht relied upon for other limitations described below, would have been performed by a native/non-browser mobile application executing on the mobile device using the techniques of Kirkpatrick.

86. For example, in one exemplary implementation of the proposed combination utilizing the unified framework in Kirkpatrick, all of the URIs for the media distribution web site in Fosnacht (*i.e.*, all URIs containing the string `www.mycatholicvoice.com`) would have been processed by a native/non-browser application on user computing system **108** (or publisher computing system **106**) instead of the web browser. I note that this is consistent with an exemplary embodiment in Kirkpatrick describing the use of a native/non-browser application for accessing features of the popular YouTube video sharing website, which provides similar functionality to the media distribution website described in Fosnacht. (Kirkpatrick, 6:21-28 (“For example, if the application **206** is entitled ‘YouTube Enhanced,’ the application **206** can request that it handle all URIs that include the pattern `www.youtube.com/*`, where the ‘*’ indicates a wildcard that can stand for any number of characters. If an activity includes a reference to `www.youtube.com`, the activity will be forwarded to the YouTube Enhanced application for processing.”).) I note that YouTube was (and remains today) a well-

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

known video sharing website that provides features similar to Fosnacht, which itself specifically refers to YouTube as providing similar features. (Fosnacht, ¶0003 (“Some web sites are focused on the distribution of user-generated and uploaded content. An example is YouTube (www.youtube.com), which distributes videos that are typically generated and uploaded by users.”).)

87. Fosnacht and Kirkpatrick are analogous references in the same field of endeavor of the '1357 patent, which discloses systems and processes for distributing media messages to mobile devices over the Internet. Fosnacht, for example, discloses embodiments for “distribution of media across a network,” such as “audio, video, visual, audio-visual, multimedia, interactive media, documents, spreadsheets, databases, and any other type of media capable of being represented in digital form.” (Fosnacht, ¶0006.) The analogousness of Fosnacht is further highlighted by the fact that, like the '1357 patent, it discloses embodiments for distributing religious or prayer-oriented media messages. (Fosnacht, *e.g.*, Figs. 4-11.) Kirkpatrick similarly discloses techniques for creating native/non-browser mobile applications for distributing media such as photos and videos over the Internet. (Kirkpatrick, 5:17-63 (FlickU example), 6:21-35 (YouTube Enhanced example).) Both references would also have been reasonably pertinent to problems facing the inventor. Fosnacht, for example, discloses techniques for receiving, storing and distributing

media items such as videos to users through one or more feeds – key features of the alleged invention recited in the challenged claims. Fosnacht also discloses painstaking implementation details about its media distribution system, including screenshots of key user interfaces, that would have been pertinent to the inventor in carrying out and implementing her alleged invention. Kirkpatrick, for its part, would have been reasonably pertinent to designing a native/non-browser mobile application (to the extent required by the claims), including how to design a mobile application that interacts with a web server over the Internet.

88. A person of ordinary skill in the art would have had ample motivation to combine. As explained above, it was widely known to skilled artisans that native/non-browser mobile applications provided benefits over web-based applications that rely on a web browser, including the ability to have fuller access to the features and capabilities of the mobile device. Kirkpatrick itself echoes these motivations to combine by describing its advantages as including, among others, “seamless integration between viewing locally stored content and remote content” and the ability to intercept URIs “so that content from a web page is displayed within an application optimized for the displaying device.” (Kirkpatrick, 2:30-36.) For example, it was well-known that native/non-browser applications were often more capable of displaying information more effectively on smaller screen mobile devices

as compared to a web site viewed through a mobile web browser. Kirkpatrick echoes this advantage by explaining, with reference to the YouTube Enhanced example, that the non-browser YouTube Enhanced application “can be optimized for display on small screen devices, while the YouTube web page may not be easily viewed on a small screen.” (Kirkpatrick, 8:43-46.) Figure 7B of Kirkpatrick, as another example, shows an exemplary user interface in which a user accessing YouTube from a web browser is encouraged to download “an Enhanced YouTube Application for optimum viewing.” (Kirkpatrick, Fig. 7B, 11:4-10.)

89. Market forces would have also provided a powerful motivation to combine, as native mobile applications were becoming extraordinarily popular. By January 2010, for example, the Apple iPhone had been purchased by tens of millions of people and the Apple App Store had more than 140,000 apps. (Anthony Bruno, *A Bigger Hand? Apple’s iPad Offers Larger Mobile Platform for Music and Media*, Billboard Magazine, Feb. 6, 2010 (**EX1010**), p.12.) The commercial desirability of native/non-browser mobile applications, coupled with the other benefits described above, would have motivated a skilled artisan to adapt the media distribution web site of Fosnacht to be accessible via a native/non-browser application.

90. Finally, a person of ordinary skill in the art would have had at least a reasonable expectation of success. As noted, native/non-browser applications were

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

commonplace before the priority date, with more than 140,000 iPhone apps available in the App Store before the priority date. Techniques for developing mobile applications were also well-known, with operating system providers having released software development kits (SDKs) that provided application programming interfaces (APIs) and other tools to allow developers to create native/non-browser applications utilizing the features of mobile devices. In fact, I note that the '1357 patent specification provides no details whatsoever about how to implement the claimed mobile application, further suggesting that techniques for developing mobile applications were already well-known. A person of ordinary skill in the art could have implemented the proposed combination using standard and conventional techniques, without undue experimentation.

(b) “storing the plurality of media messages on the one or more servers;” (Claim 1[b])

91. As noted previously, Fosnacht discloses **“storing the plurality of media messages on the one or more servers,”** *i.e.*, when uploaded media items are stored on media database **103**, which is part of web server **102**. (Fosnacht, ¶¶0355 (“The upload button **4122** [of Figure 41] is then selected to upload the selected file to the server and to save the media item and associated information in the media database (e.g., **103** shown in FIG. 1).”), 0068 (“Web server **102** includes media database **103**.”).)

- (c) **“providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application;” (Claim 1[c])**

92. The '1357 patent does not provide a definition of **“feed”** but describes it generally as a series or group of media messages posted by others that can be made available to a user to view. ('1357, 3:50-56 (“FIG. 1 illustrates an example media coordination system **100** according to the present invention which allows network connected devices... to coordinate times for specific media messages/programs to be streamed to other application users (e.g., a feed of message posts or a stream of real-time media data).”).) Fosnacht discloses a number of different feeds for displaying media items to users of the system, which disclose the claimed first and second feeds recited in this claim limitation.

93. For example, welcome page **402** in Figure 4 shows at least three different feeds for showing “Featured,” “Most Viewed,” and “Newest” media items:




FIG. 4

(Fosnacht, Fig. 4.) “Welcome page **402** includes welcome section **404** and media section **408**.” (Fosnacht, ¶0093.) The key feature shown in Figure 4 is media section **408** which, as shown, “includes a featured tab, most viewed tab, and a newest tab.” (Fosnacht, ¶0095.) “If the featured tab is selected, welcome page **402** displays a set of featured content in media section **408**. This is also the default tab. If the most

viewed tab is selected, welcome page **402** displays a set of the most viewed content, arranged in an order from the content that has been viewed the most to the content that has been viewed the least. If the newest tab is selected, then welcome page **402** displays a set of the content that was most recently added, arranged in an order from the newest content to the oldest content.” (Fosnacht, ¶0095.) Each of these tabs discloses a distinct “**feed**” that includes a subset of media items from the plurality of media items (such as video recordings).

94. Fosnacht discloses several other examples of feeds that include a subset of media items that can be presented to users. In one embodiment, users can choose to become members of one or more content centered communities (or groups) and can access associated community pages. (Fosnacht, ¶0114.) A user can access a particular community or group page by activating its link from the user’s home page. (Fosnacht, ¶0114 (“If the user was a member of one or more communities, a link would be displayed. Selection of the link would cause web site **300** to display the associated community page.”).) “Examples of groups that would benefit from a content centered community page include a parish, diocese, church group, youth organization, Bible study group, educational group, religious order, or other groups.” (Fosnacht, ¶0158.) A community or group page provides information about the community and lists media items available to community members. (Fosnacht, Fig.

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

10, ¶¶0157-0158.) Figure 10 below shows an exemplary community page 1002:




FIG. 10

(Fosnacht, Fig. 10.) An exemplary “feed” from Figure 10 is shown on the right side of the page, which presents a list of media items associated with the community (e.g. “Video Title 1” through “Video Title 5”).

95. As yet another example of a “feed,” Fosnacht discloses a list of favorite

media items, to which users can add media items using a “Save to Favorites” link. (Fosnacht, ¶0122.) “Once saved as a favorite, the user can quickly return to this media item by selecting it from the favorites section (e.g., **508**) of the MyHome page **502** (shown in FIG. 5), or by navigating to a favorites page that displays the media items that the user has identified as favorites.” (*Id.*) A user’s favorites can also be accessed from the user’s profile page. (Fosnacht, ¶¶306, 310-311, Fig. 15 (media content tab **3506**)).) In either case, the list of favorites provides another exemplary “**feed**” that includes a subset of media items.

96. As yet another example of a feed, Fosnacht discloses a feature known as a “ticker zone” **1304** that can be displayed on the home page of the website, as shown in Figures 13 and 14. As shown in Figure 14, ticker zone **1304** can present a number of media items in sequence:




FIG. 14

(Fosnacht, Fig. 14.) “Ticker zone **1304** is a section of home page **1300** where certain content within web site **1100** can be featured. Any content can be featured including a particular media item, a collection of media items from a particular user or publisher, a selection of media items from multiple users or publishers... or any other content available through or associated with web site **1100**. Ticker zone **1304** is preferably updated often so that the information contained therein is always fresh, new, or different, even if a user visits frequently.” (Fosnacht, ¶0174; *see also id.*, ¶0183 (“Ticker zone **1304** can be used to feature both publisher content and user generated content.”).) Ticker zone **1304** includes “display region **1404** [that]

typically provides a graphical image and brief description of the associated media content. In other possible embodiments an audio or video clip is provided.” (Fosnacht, ¶0177.) “In some embodiments display region **1404** scrolls through a plurality of featured content. For example, the first content is first featured in display region **1404** for a period of time. After the period of time has elapsed, the second content is next displayed. The process continues until all featured content has been displayed.” (Fosnacht, ¶0178.) Ticker zone **1304** therefore provides another example of a “**feed**” that includes a subset of media items from the plurality of media items.

97. As yet another example, Fosnacht discloses a media browsing interface in Figure 9 that allows the user to browse media items and apply a series of filters to further specify the media items displayed. (Fosnacht, Fig. 9, ¶¶0145-0153.)

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 9

(Fosnacht, Fig. 9.) "Browsing section 906 [of Figure 9] displays a list of media items that match the criteria specified in filter section 904. In this example, media is displayed as snapshots, including media snapshots 922, 924, and 926. Media

snapshots display information about an associated media item.” (Fosnacht, ¶0151.)

These media snapshots include, among other things, a thumbnail and title, as well as the name of the person or publisher who uploaded it. (*Id.*) A list of media items provided to the user in response to the user’s search criteria therefore provides yet another example of a “**feed**,” as claimed.

98. Fosnacht therefore discloses “**providing the first feed from the one or more servers**,” which takes place when web server **102** provides one or more web pages to a user computing system **108**, the web pages including a “**feed**” with a subset of previously-uploaded and previously-stored media items to display to the user in a first feed. (Fosnacht, ¶0075 (“In some embodiments, Web server **102** stores and transmits web site data. The web site data is requested by a computing system, such as user computing system **108**. In response to the request, web server **102** transmits the web site data to user computing system **108** across network **104**.”).) For example, each of the examples listed above including the featured, most viewed and newest tabs in media section **408** of Figure 4, the list of video media items in the community page in Figure 10, the media items available through favorites section **508** of Figure 5, ticker zone **1304** in Figures 13-14, and browsing section **906** of Figure 9, provides a “**feed**” that lists a first subset of media items. The claimed “**first feed**” could be mapped to any of the examples described above; to take just one

example, the “**first feed**” could be mapped to the subset of media items available through the community page in Figure 10, or the subset of items in the browsing section of **906** in Figure 9, but many other combinations are also possible for mapping the first feed. As explained above for limitation 1[a], Fosnacht further discloses that user computing system **108**, in one embodiment, may be “**a mobile device.**” (Fosnacht, ¶0073.) And as fully explained for limitation 1[a], it would have been obvious for the web site of Fosnacht, which includes the claimed “**first feed,**” to have been provided to and presented through a “**mobile application executing on a mobile device,**” under Fosnacht alone or under the further combination with Kirkpatrick.

99. The claimed step of “**providing the first feed**” occurs under the proposed combination when the server in Fosnacht (a) sends a page to the user showing available media items in the first feed, and then (b) allows the user to view one or more selected media items from the first feed through media pages associated with that media item. With respect to step (a), this occurs in Fosnacht when the server sends a page listing available media items in the “**first feed,**” such as the listing of items in the featured, most viewed and newest tabs in media section **408** of Figure 4, the community page in Figure 10, favorites section **508** of Figure 5, ticker zone **1304** in Figures 13-14, and browsing section **906** of Figure 9. With

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

respect to step (b), Fosnacht discloses a media page that the user can access by selecting a particular media item, which presents a media item to the user. (Fosnacht, ¶¶0118 (“If the user wants to get more information about the media item, the user selects the media item. Web site **300** then displays the media page associated with the selected media item.”), 0119, 0181 (for Ticker Zone: “When a user wants more information about particular featured content, the user selects more info button **1412**... If the featured content is a media item, then the associated media item page is displayed.”), 0194 (“Each displayed search is also linked to the associated media page....”), 0311 (“The media item is displayed as a thumbnail image **3514**, which is linked to the associated media page.”).) .) It would have been obvious, with respect to each of the feeds identified for this limitation that list media items, for the user to have accessed the media page for a media item by clicking on its associated link. A skilled artisan would have been motivated to provide this functionality across all of the feeds in Fosnacht to provide a consistent user interface for allowing the user to access the media page for a media item. Figure 7 shows one example of media page **702** for an audio media item:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 7

(Fosnacht, Fig. 7, ¶¶0119-0120.) “Media page 702 is a page where additional information about a selected media item is displayed. The media item is typically either user-generated content (such as previously uploaded by a user) or content provided by a commercial publisher.” (Fosnacht, ¶0120.) Figure 8 shows a similar media page but for a video media item with substantially the same features. (Fosnacht, Fig. 8, ¶0138 (“FIG. 8 is a screen shot of an exemplary media page 802. Media page 802 is similar to media page 702, except that media page 802 relates to

a video media item, rather than an audio media item.”).)

100. In the example of Figure 7 above, the media item is an audio entitled “A Catholic Prayer Companion Complete Collection.” (Fosnacht, ¶0120.) “Media player section **704** [in Figure 7] includes an audio player **720** that plays at least part of the audio recording to the user, such as through a speaker or headphones associated with the user's computing system.” (*Id.*) Figure 8 also includes a media player for video. (Fosnacht, ¶0138 (“In this embodiment, media player section **804** includes a section for a video player to display a video.”).)

101. Figure 21 provides another example of a media page, also for a video:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 21

(Fosnacht, Fig. 21.) “FIG. 21 is a media page associated with a video provided by a user (user-generated content).” (Fosnacht, ¶0224.) Figure 21 includes, similar to Figures 7 and 8 discussed above, media player section **2104** for playing the video. (Fosnacht, ¶0239 (“User generated content is freely available through web site **1100**.

As a result, media player section **2104** of media page **2100** includes the full content of this video uploaded by a user. To play the media, the user either clicks on the video display **2120** or clicks on the play button in video controls **2122**. The video then plays in media player section **2104**.”.)

102. Fosnacht (alone or under the combination with Kirkpatrick) therefore discloses “**providing the first feed from the one or more servers to a mobile application executable on a mobile device, the mobile application configured to present the first feed,**” by providing a page with a list of media items in the first feed, and by allowing the user to display one or more selected media items on a media item page accessible from the first feed. I note that, under this mapping, the user viewing a particular feed may take a further action to view a particular video media item (such as clicking on its associated link), but nothing in claim 1 precludes the user taking additional action to cause the media item to be presented. Although this mapping includes two web pages, this is entirely consistent with the ’1357 patent, which shows a first screen (Figure 2) for displaying a current media item, and a second screen (Figure 3) for displaying information about the creator of the media item and providing a share button **128** for sharing the media item. (’1357, Figs. 2-3, 5:8-50.) I note that the language of claim 1 does not specify any particular form or manner in which the first feed must be provided to, and presented by, a

mobile application.

103. Fosnacht further discloses that **“presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed.”** I will cover each of these claimed prompts below.

“Comment Prompt”

104. When a media item is uploaded, the uploader can choose to enable comments for the item. (Fosnacht, Fig. 42, ¶0359 (“Prompts section **4202** prompts the publisher to identify whether comments should be allowed to be posted on a media page for this media item, and whether users should be allowed to rate the media item.”).) The media page shown in Figure 21 above, for example, includes a comments section **2110**. “Comments section **2110** [in Figure 21] is similar to comment section **1910** shown in FIG. 19. In this example, several users have provided comments. The comments are displayed in comment windows **2130**.” (Fosnacht, ¶0241.) “Comments section **1910** allows users to leave a comment regarding the media item. Comments are displayed publicly below the comment field (as shown in FIG. 21).” (Fosnacht, ¶0229.) Fosnacht therefore discloses that **“presentation of the first feed by the mobile application includes presentation of a comment prompt with the media message included in the first feed,”**

because the media message presented from the first feed is accompanied by comments section **2110** that allows the user to enter a comment.

105. Fosnacht further discloses and renders obvious that **“interaction with the comment prompt via the mobile application enables a comment to be associated with the media message.”** This is because, when a user enters a comment through the comment prompt, the comment is associated with the media message. (Fosnacht, ¶¶0229 (“Comments section **1910** allows users to leave a comment regarding the media item. Comments are displayed publicly below the comment field (as shown in FIG. 21).”), 0241 (“Comments section **2110** is similar to comment section **1910** shown in FIG. 19. In this example, several users have provided comments. The comments are displayed in comment windows **2130**.”).)

“Sharing Prompt”

106. Fosnacht also discloses and renders obvious that presentation of the first feed **“includes presentation of a sharing prompt with a media message included in the first feed,”** which is provided in the media page itself, more specifically through sharing options **730** shown in Figure 7, sharing options **830** in Figure 8, and the sharing links feature available through tool bar **2106** in Figure 21. In reference to Figure 7, Fosnacht explains, “[m]edia player section **704** also includes a sharing options link **730**. Clicking on this link causes web site **300** to

display a sharing options page.” (Fosnacht, ¶0125; *see also id.*, Fig. 8 (showing sharing options feature for video media page).) Fosnacht explains, with reference to sharing options link **730**, that “[e]xamples of sharing options include the ability to link the media item to a third-party service (e.g., **110** shown in FIG. 1), to generate an e-mail to a friend regarding the media item, or other sharing options.” (Fosnacht, ¶0125.) Similarly, Figure 21 shows tool bar **2106**, which provides various options including sharing features including “share:”, “embed:”, and “permalink:” (Fosnacht, Fig. 21.) Fosnacht explains these linking features result in generation of a link (e.g. a URL) to enable access to the media item. (Fosnacht, ¶0237 (“Sharing links **2034** allow users to interconnect the media page **1900** with external systems or third-party services (e.g., **110**, shown in FIG. 1). In this example, links are provided to allow sharing with third party services such as facebook, digg, and stumble upon. In addition, URLs are provided to allow a user to easily embed the media item in their own or other web sites (embed), or to link directly to the media page (permalink) from their own or other web sites.”).) The term “URL” refers to Uniform Resource Locator, the standard address used to uniquely identify and locate a resource on the Internet (such as a web page). It would have been apparent that the media page in Fosnacht would have been identified by at least its own URL. (See David Fox et al., *HTML Web Publisher’s Construction Kit* (1995) (**EX1009**),

p.13 (1995) (“The World Wide Web gives each article, newsgroup, database, computer, file, or other Internet fragment its very own URL.”).) Fosnacht further discloses that the first feed **“includes presentation of a sharing prompt with a media message,”** because as explained, the media page containing the sharing prompt is accessed by the user when the user selects a media item to display from the list of media items in the first feed.

107. Fosnacht further discloses and renders obvious that interaction with the sharing prompt **“facilitates generation of a link configured to enable access to the media message,”** as claimed. This is because, as explained, the sharing prompt in Fosnacht available through Figure 7, 8, or 21, all result in generation of a link to enable access to the media item. The disclosures above, for example, confirm that Fosnacht generates a link (*e.g.* permanent link or permalink) for the media page of a particular media item, allowing the media page to be accessed through that link, for example, on another site. It thus would have been obvious that the link generated in Fosnacht enables access to the media item because it allows the user to link to the media page for the media item, or to embed the media item in their own sites. In either case, the sharing prompt in Fosnacht facilitates generation of **“a link configured to enable access to the media message.”**

“Bookmark Prompt”

108. Like the sharing prompt just discussed, the claimed “**bookmark prompt**” appears on the media page for a selected media item, and more specifically, as a “Save to Favorites” link available on the media page. For example, with reference to Figure 7, Fosnacht explains that “[t]he Save to Favorites link is selectable to save this media item to the user’s list of favorite media items. Once saved as a favorite, the user can quickly return to this media item by selecting it from the favorites section (e.g., **508**) of the MyHome page **502** (shown in FIG. 5), or by navigating to a favorites page that displays the media items that the user has identified as favorites.” (Fosnacht, ¶0122.)

109. I note that the ’1357 patent provides only a single sentence describing the claimed bookmark feature, which only discloses displaying bookmarked media messages and does not describe the claimed feature of bookmarking a media message using a bookmark prompt. (’1357, 5:26-28 (“A past media messages button **113** [in Figure 2] is also included for displaying a list of past media messages or bookmarked media messages.”).) Nevertheless, a person of ordinary skill in the art would have understood “**bookmark**” in this context to refer generally to a feature common in web browsers in which a user could store a reference or shortcut to a piece of content, allowing the user to easily return to it at a later time. Bookmarking was (and remains today) a well-known feature of web browsers for allowing users

to create a list of their favorite and/or often-visited web pages.

110. The “Save to Favorites” feature of Fosnacht thus discloses “**a bookmark prompt**” with which interaction “**enables the media message to be listed with another bookmarked media message through the first mobile application.**” This is because the Save to Favorites feature adds the current media item “to the user’s list of favorite media items.” (Fosnacht, ¶0122.) “Once saved as a favorite, the user can quickly return to this media item by selecting it from the favorites section (e.g., **508**) of the MyHome page **502** (shown in FIG. 5), or by navigating to a favorites page that displays the media items that the user has identified as favorites.” (*Id.*; (underlining added); *see also id.*, Fig. 5 (showing “Favorites” section **508**).)

- (d) “**providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.**” (Claim 1[d])

111. The analysis for this limitation is similar in many respects to limitation 1[c] above. A second feed, like the first feed discussed for limitation 1[c], is provided when web server **102** provides the web page to a user computing system

108 that includes the second feed including a second subset of the previously-uploaded and previously-stored media messages. (Fosnacht, ¶0075.) As explained for limitation 1[a] and 1[c], it would have been obvious that the second feed (just like the first feed) would have been provided “**to the mobile application,**” which would have been configured to present the feed to the user through the user interface by displaying a page showing a list of media items available in the second feed, and allowing the user to select one or more of those media items to display on associated media item pages.

112. As explained for limitation 1[c], Fosnacht discloses several examples of feeds including the featured, most viewed and newest tabs in media section **408** of Figure 4, the list of video media items in the community page in Figure 10, the media items available through favorites section **508** of Figure 5, ticker zone **1304** in Figures 13-14, and browsing section **906** of Figure 9, as each provides a “**feed**” that lists a subset of media items. The claimed “**first**” and “**second**” feeds could be mapped to any two of the examples described above; to take just one example, the “**first feed**” could be mapped to the feed of media items available through the community page in Figure 10, and the “**second feed**” to the subset of items in the browsing section of **906** in Figure 9, but many other combinations are also possible for mapping the two feeds.

113. It would also have been obvious that the “**first feed**” and the “**second feeds**” differ from each other, as each shows a different subset of media items. I note that the claim does not appear to require that there be no overlap whatsoever in terms of the media items between the two feeds. The feeds in Fosnacht differ from one another because each feed serves a different purpose (*e.g.*, to list media items based on the ones that are new, featured, popular, favorite, based on a particular community, meeting the user’s search filters, and so forth), and thus, it would have been obvious that each of these feeds would have presented a different subset of media items. And even if the claim required that the two feeds provide a disjoint listing of media items with no overlap, it would have been obvious that the first and second feeds in Fosnacht would have met this limitation. For example, Fosnacht confirms, in one embodiment, that media items uploaded to a community may be designated as private to that community. (Fosnacht, ¶0157 (“Typically, each community has at least one administrator that is allowed to customize the community for the participants. Some embodiments of a content centered community page(s) include one or more of the following... an ability to add media items and other content and designate each individual item as private to that community....”).) Under this embodiment, if the first feed was a community- or group-based page with uploaded media items designated as private to that community, then those media

items would be unavailable on the second feed, thus preventing any overlap between media items available in the two feeds.

114. It would also have been obvious that **“the mobile application [is] configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application,”** as claimed. This is because Fosnacht discloses an interactive user interface that allows the user to navigate from one feed to another. For example, the “Featured,” “Most Popular” and “Newest” feeds in Figure 4 described in connection with limitation 1[c], are activated by user-selectable tabs through the user interface. (Fosnacht, Fig. 4, ¶¶0095 (“If the featured tab is selected, welcome page **402** displays a set of featured content.... If the most viewed tab is selected, welcome page **402** displays a set of the most viewed content... If the newest tab is selected, then welcome page **402** displays a set of the content that was most recently added[.]”).) Similarly, the community page and the media browsing pages described above in connection with limitation 1[b] are activated by user-selectable links. (Fosnacht, ¶¶0114 (“The Communities section **514** (also sometimes referred to herein as a group) includes links to the community pages of any communities that the user is a member.”), 0145 (“FIG. 9 is a screen shot of an exemplary media browsing page **902**. Media browsing page **902** is displayed, for example, when a user selects Explore from the quick links

bar (e.g., **504** described herein with reference to FIG. 5).”) Fosnacht therefore discloses and renders obvious that the mobile application is configured to **“switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application,”** for example, when the user switches from one feed to another by making the appropriate selections through the user interface. And as explained for limitation 1 [a], it would have been obvious that navigation through the user interface of Fosnacht would have taken place in response to user interaction with **“the mobile application,”** as claimed, based on Fosnacht alone or in view of Kirkpatrick.

2. Claim 2: “The method of claim 1, wherein respective ones of the plurality of media messages received at the one or more servers include an image or video captured by the one or more mobile applications via a corresponding one or more cameras of the corresponding one or more mobile devices.”

115. As explained for claim 1, the media items in Fosnacht that users can upload to web server **102** (**“the plurality of media messages”**) can include images or videos. (Fosnacht, ¶0083 (“In some embodiments, web server **102** includes a database of media items stored in computer storage media. The database includes a plurality of media items, such as videos, audio recordings, images, and possibly other media items.”).) Fosnacht further discloses that media items can include “web camera recordings.” (Fosnacht, ¶0143 (“Although media pages **702** and **802** have

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

been described with reference to audio media and video media, respectively, it is recognized that other embodiments include media pages that include other media times [sic; items], such as... web camera recordings... or any other type of digital media capable of being conveyed across a communication network.”.) It would have been obvious to a person of ordinary skill in the art that the term “web camera” (also commonly known as a “webcam”) includes a camera connected or otherwise built-into a computer. For example, external web cameras that could connect to computers (such as the Logitech QuickCam) were popular, and before 2010, were commonly used with desktop and laptop computers to capture still images and video. (See U.S. Patent No. 8,218,073 B1 (filed January 25, 2008) (**EX1011**), 1:9-11 (“Web cameras (or webcams) are cameras that are configured to collect video or still images for transmission over a network, such as the Internet.”), 2:24-34 (“FIG. 1 is a simplified diagram of a webcam **100** according to one embodiment of the present invention... The webcam may also include... a mount **130** for standing the webcam on a desk or for attaching the webcam to a monitor, laptop computer or the like.”).) It therefore would have been obvious based on Fosnacht alone that the video item uploaded by a user would have included **“video captured by the one or more mobile applications via a corresponding one or more cameras of the corresponding one or more mobile devices,”** *i.e.*, video captured on a webcam of

the device as disclosed in Fosnacht, executing the mobile application.

116. Claim 2 would also have been obvious in further view of Kirkpatrick. Kirkpatrick also discloses, in connection with the FlickU embodiment, the ability of the mobile application to use a mobile device's camera to capture and photo and upload it to a web server. (Kirkpatrick, 4:57-5:63.) Although this embodiment involves uploading of photos, it would have been obvious to also adapt this technique to videos. For example, Kirkpatrick discloses a mobile device **900** (Figure 9) that "implements the unified framework feature," which includes a built-in camera **907**. (Kirkpatrick, 11:63-65, 14:18-23, Fig. 9.) "The camera **907** allows the device **900** to capture digital images, and may be a scanner, a digital still camera, a digital video camera, or other digital input device." (Kirkpatrick, 14:18-20.) It therefore would have been obvious, under the proposed combination, for the video of a video media item uploaded to web server **102** in Fosnacht to have been captured by a mobile application via a camera on the mobile device.


117. The ability to use a built-in camera to capture video was well-known to persons of ordinary skill in the art and a standard feature of well-known mobile device operating systems. (See, e.g., Adi Rome, *Multimedia on Symbian OS: Inside the Convergence Device* (2008) (**EX1008**), p.56, § 4.1.2 ("On Symbian smartphones with a built-in camera and microphone, it is possible to record a video. The video

and audio can be recorded to a file for subsequent playback or could be streamed off the device on a server elsewhere.”.) A person of ordinary skill in the art would have been motivated to implement Fosnacht such that the video of a video media item uploaded to web server **102** was captured by a mobile application via a camera on the mobile device itself, as this would have allowed the user to more spontaneously and immediately share videos with other users using the mobile device. This would have been viewed as preferable to requiring that the user record a video using another device and then transfer the video to the mobile device before uploading, which would have added complexity and technical hurdles that may have discouraged the user from recording and sharing videos.

3. **Claim 3: “The method of claim 2, wherein the media message included in the first feed includes one or more of an attached image or attached audio that is not native to the image or video captured by the one or more mobile applications via the corresponding one or more cameras of the corresponding one or more mobile devices.”**

118. The limitations of claim 3 would have been obvious based on the disclosures in Fosnacht. As explained in claim 1, Fosnacht makes clear that a media item **“included in the first feed,”** could be a video. For example, Fosnacht discloses the ability of a user to upload a video media item (**“media message”**). (Fosnacht, ¶0239 (“To play the media, the user either clicks on the video display **2120** or clicks on the play button in video controls **2122**. The video then plays in media player

section 2104.”.) Figure 41 of Fosnacht illustrates:



(Fosnacht, Fig. 41 (excerpt, highlighting added).)

119. Fosnacht further discloses and renders obvious that the video media message may include “**one or more of an attached image or attached audio that is not native to the image or video,**” captured by the mobile application via the camera on the mobile device. The claim thus only requires either an attached image or an attached audio, and Fosnacht discloses at least the former.

120. More specifically, Fosnacht discloses the ability to attach a “thumbnail” image to a previously-uploaded media item, and the thumbnail can be any arbitrary image chosen by the uploader. (Fosnacht, ¶¶0117, 0369.) For example, Figure 45 shows a thumbnail definition page 4500 to associate any thumbnail image with a previously uploaded media item by simply uploading an image file to the server. (Fosnacht, Fig. 45, ¶0369.) “Thumbnail definition page 4500 [of Figure 45] is used by a publisher to define the image that should be used for the thumbnail display. To do so, the browse button is selected from thumbnail definition page and the desired file selected. The upload button 4504 is then selected to upload the selected file and

to save the changes to the server. A preview of the currently selected thumbnail, if any, is shown in preview window **4506**.” (Fosnacht, ¶0369.) Under this mapping, therefore, the “**media message included in the first feed**” would have included not just the previously-uploaded video media item, but also the thumbnail image received by the server. I note that, although Figure 45 shows an option available to publishers, and although not required by the claim, it would have been obvious to provide the same thumbnail selection options for non-publisher users in order to give those users more customization over presentation of media items they upload.

121. It would also have been obvious that the uploaded thumbnail image in Fosnacht would have been “**not native to the... video**,” as recited in the claim. This is because nothing in Fosnacht requires that the uploaded image have been captured by the camera on the mobile device – or even by a camera at all. Fosnacht confirms that, in one embodiment, “[t]he thumbnail image is, for example, a single frame from a video clip, a cover image from a media case (such as a CD or DVD cover image), an image from the publisher, or any other image associated with the media.” (Fosnacht, ¶0117.) Fosnacht further explains that, in one embodiment, “[t]he thumbnail image 740 [in Figure 7] is an image associated with the full content media item, such as a CD cover image, a graphic, a logo, or any other image.” (Fosnacht, ¶0131.) It would have been obvious that a CD or DVD cover image, a graphic, or logo, etc., as

described in Fosnacht would have included images that were not captured by a camera on the mobile device. It would have been obvious, for example, to create the thumbnail image created artificially such as through drawing, rendering, or other generation of the image on the computer, thus resulting in image data attached to the video that is “**not native to the... video**” that was captured using the camera through a mobile application executing on the mobile device.

122. Alternatively, even in those cases where a thumbnail associated with an uploaded video message *was* a real-world image originally captured by a digital camera, it would have been obvious for the thumbnail to have been captured by a camera other than the one on the mobile device, such as the case of a CD or DVD cover that incorporates a photo professionally taken using an entirely different digital camera device. And finally, it would also have been obvious that, even if a thumbnail image associated with an uploaded video message *was* obtained from the digital camera on the mobile device, that thumbnail image may not even have been part of the part of the video media item previously uploaded by the user (“**the media message**”), and thus, “**not native**” to the video previously captured by the camera. All of these examples provide independent ways in which it would have been obvious for “**the media message included in the first feed,**” to have included “**an attached image**” that is “**not native to the image or video captured by the one or**

more mobile applications via the corresponding one or more cameras of the corresponding one or more mobile devices,” as claimed.

4. **Claim 4: “The method of claim 1, wherein the link is a first link, wherein the media message included in the first feed is encoded with a second link, wherein interaction with the second link via the mobile application enables access to a webpage.”**

123. As explained for claim 1[c], the **“first link”** is the link generated by the sharing prompt described in Fosnacht. For purposes of claim 4, the **“second link”** refers to a different link on the media page for the media item that allows the user to access a webpage. More specifically, Fosnacht discloses that **“the media message included in the first feed is encoded with a second link,”** because the media message included in the first feed is presented along with a link associated with the creator of the media message, such as a link to its web site. For example, Figure 7 shows an “Also From Publisher” section **710** that “displays a recommendation for one or more other media items from the same publisher that may be of interest to the user.” (Fosnacht, ¶0135; *see also id.*, Fig. 7.) “The information displayed in section **710** includes a title of the media item, a thumbnail image of the media item, and a description of the media item. If desired, the user can click on the title of the media item to be brought to a separate media page where additional information, a snack version of the media item, and purchasing information is presented. Alternatively,

the Also From Publisher section **710** includes a link to the publisher's web site." (Fosnacht, ¶0135.) The same feature also appears in the Figure 8 media page for videos. (Fosnacht, ¶¶0138, 0142.) Fosnacht therefore also discloses that the media message is encoded with a second link, wherein **"interaction with the second link via the mobile application enables access to a webpage."**

5. **Claim 5: "The method of claim 1, wherein presentation of the first feed by the mobile application includes presentation of an advertisement prompt with the media message included in the first feed, wherein interaction with the advertisement prompt via the mobile application causes presentation of an advertisement."**

124. Fosnacht discloses several examples in which presentation of the first feed includes **"presentation of an advertisement prompt"** that **"causes presentation of an advertisement."** I note that the term **"advertising prompt"** does not appear in the written description of the '1357 patent; the embodiment that relates to this claim limitation instead describes a banner advertisement that can be displayed on the user interface for the current media message. ('1357, 5:16-19 ("Optionally, the current media message interface includes an advertisement banner **112** [Figure 2] for displaying advertisements from a downloaded source, such as from Google AdSense or Apple in-app advertising."), Fig 2 (showing banner **112**).) This would have indicated, to a person of ordinary skill in the art, that an **"advertising prompt"** at least includes a conventional advertisement banner, which

was a common fixture on countless websites and webpages years before the priority date of the challenged claims. Fosnacht discloses several examples in presentation of the first feed includes an advertisement banner or prompt.

125. For example, as explained fully for claim 1, Fosnacht displays a media page (Figure 21) that can be reached by clicking on a media item presented on each “**first feed**,” as fully described for limitation 1[c]. This page shows at least two advertisement banners:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2





FIG. 21

(Fosnacht, Fig. 21 (highlighting added).) Media page 2100 in Figure 21 displays advertisements 2116 and 2118. (Fosnacht, Fig. 21, ¶0238.) These qualify as advertising prompts because they display and provide access to advertisements through the media page associated with the displayed media item (“an

advertisement prompt with the media message included in the first feed”).

These advertisements are also clickable; advertisement **2116** for example displays a Google Ad for a fundraising service and is clickable to provide “More Info” about the service. (Fosnacht, Fig. 21 (**2116**); *see also id.*, Fig. 19 (“Click Now >” button of advertisement **1918**).) For example:



(Fosnacht, Figs. 21 & 19 (excerpts; highlighting added).)

126. Fosnacht discloses several other examples in which presentation of the first feed includes “**presentation of an advertisement prompt**,” as claimed. For example, as noted for claim 1, Figure 10 shows a community interest page that provides an example of the claimed “**first feed**” for presenting media items, which includes at least two advertisement prompts:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 10

(Fosnacht, Fig. 10 (highlighting added).) Figure 10, as shown above, shows sections for presenting for graphic and text advertisements. As another example, as noted for claim 1, Figure 9 shows a page with media items provided in response to a search criteria, which provides another example of the claimed “**first feed**”:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2




FIG. 9

(Fosnacht, Fig. 9 (highlighting added).) As shown, Figure 9 includes a section for presenting a Google advertisement. Figures 10 and 9 discussed above thus provide further examples of advertising prompts that display and provide access to advertisements through the first feed (“an advertisement prompt with the media message included in the first feed”). Although I have identified certain exemplary

pages in Fosnacht that expressly show advertising banners, it would have been obvious to provide an advertising banner on other pages in Fosnacht that present further examples of the claimed “**first feed**,” such as the pages in Figures 4-5 and 13-14 discussed above. A skilled artisan would have been motivated to provide advertising banners across these other pages in Fosnacht to increase opportunities to present advertising (and thus increase the opportunities for raising revenue) from the media distribution system. Fosnacht acknowledges, in fact, that advertisers may pay each time an advertisement is displayed (Fosnacht, ¶0110), thus providing a further incentive to increase the number of opportunities for presenting advertisements through the media distribution system. In all cases, the advertisement prompts in Fosnacht are presented with the media message because the banner advertisement would have been displayed on the same page and alongside the media item in the first feed. This is consistent with the embodiment in the ’1357 patent, which shows the banner advertisement **112** displayed on the same page as a current media item **114**. (’1357, Fig. 2.)

127. For each of these advertisement prompts, moreover, Fosnacht discloses that “**interaction with the advertisement prompt via the mobile application causes presentation of an advertisement**,” in at least two ways. Under a first mapping, the claimed “**advertisement prompt**” corresponds to any of the banner

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

advertisements shown on Figures 9, 10, and 21 above, and the claimed interaction takes place when the mobile application retrieves and then displays the correct advertisement from the server in those regions. This mapping is entirely consistent with the embodiment in the '1357 patent, which describes presenting banner advertisements from a third party source such as Google AdSense. ('1357, 5:15-18 (“Optionally, the current media message interface includes an advertisement banner **112** for displaying advertisements from a downloaded source, such as from Google AdSense or Apple in-app advertising.”).) I further note that the term “**advertising prompt**” does not appear in the written description, and the patent does not suggest that any further interaction with the user (such as clicking on the ad) is required to cause display of the claimed advertisement. The interaction, in other words, occurs when the user navigates to the media item page that presents the advertising banner.

128. Second, even if an explicit and further user action were required to cause presentation of an advertisement, the claimed “**interaction with the advertisement prompt**” could be mapped to an explicit click on the advertisement (such as clicking on “MORE INFO” for advertisement **2116** in Figure 21), causing *further* advertising material to be presented. Under this mapping, therefore, clicking on any one of the advertisement banners in Fosnacht would have caused presentation of a *further* advertisement, such as more information about the fundraising service

offered in advertisement **2118** of Figure 21, more information about the downloadable music offered in advertisement **2116** of Figure 21, more information about the weight loss service offered in Figure 9, and so forth.

129. It would also have been obvious that the user clicking on any of these advertisements would have caused presentation of a further advertisement because the mere act of clicking on an advertisement on a web page rarely, if ever, would have by itself resulted in consummation of a commercial transaction in which the advertised product or service was purchased. As would have been obvious to a skilled artisan, or anyone who had used the Internet and clicked on an advertisement before the priority date, clicking on an advertisement would have taken the user to another page (such as the advertiser's page) in which additional advertising and promotional material would have been presented encouraging the user to consummate a commercial transaction by purchasing the advertised product or service (or some other product or service offered by the advertiser). I have also provided a separate **Ground 2**, below, which relies on Sharma to further confirm that it would have been obvious to have presented additional advertisements to a user in response to clicking on an advertisement in Fosnacht.

6. **Claim 6: “The method of claim 1, wherein presentation of the first feed by the mobile application includes presentation of an item purchase prompt with the media message included in the first feed, wherein interaction with the item purchase prompt via the mobile application enables presentation of information to initiate purchase of an item.”**

130. Fosnacht discloses several examples of “an item purchase prompt,” available through the media page provided during presentation of the first feed, that enables presentation of information to initiate purchase of an item. For example, the media page of Figure 7 provides an item purchase prompt to purchase an item, *e.g.*, a full content item. More specifically, Figure 7 shows full content section **708** that “includes thumbnail image **740**, price display **742**, purchase button **744**, information display **746**, and related links section **748**. Thumbnail image **740** is an image associated with the full content media item, Such as a CD cover image, a graphic, a logo, or any other image.” (Fosnacht, ¶0131.) “Full content section **708** also includes price display **742** that displays the cost to purchase the full content media item. If the user desires to purchase the media item at that price (or less, such as if the user has a coupon), the user selects purchase button **744**. Web site **300** then guides the user through the purchase of the media item, or allows the user to save the media item in the shopping cart for later purchase.” (Fosnacht, ¶0132.) The same feature is also available through the media page of Figure 8. (Fosnacht, *e.g.*, Fig. 8, ¶¶0138, 0142 (full content section **808**).) Fosnacht therefore discloses

“presentation of an item purchase prompt with the media message included in the first feed, wherein interaction with the item purchase prompt via the mobile application enables presentation of information to initiate purchase of an item.”

As fully explained for limitation 1[a], it would have been obvious that the user interface features of Fosnacht, including display of the media page with the item purchase prompt, would have been presented via the mobile application.

131. Additionally, the media page of Figure 21 includes a Related Media section **2114** that shows media items related to the media item being displayed through the media page (**“the media message included in the first feed”**). (Fosnacht, Fig. 21, ¶0243.) “A related media item can be related for various reasons, such as by having been uploaded by the same user or publisher, by sharing a common category or tag(s), or by sharing common keywords or other metadata.” (Fosnacht, ¶0231.) Fosnacht further confirms that some of the related media items “may be publisher content that is available for purchase.” (Fosnacht, ¶0243.) A media item in media section **2114**, therefore, provides **“an item purchase prompt,”** that the user can interact with to **“enable[] presentation of information to initiate purchase of an item.”** For example, when the user selects the media item in media section **2114**, they can be brought to the media page of that item where purchasing information is presented. (Fosnacht, ¶0135 (“If desired, the user can click on the

title of the media item to be brought to a separate media page where additional information, a snack version of the media item, and purchasing information is presented.”).)

7. **Claim 7: “The method of claim 1, wherein one or more stored media messages are associated with expiration information to determine when the one or more stored media messages can no longer be provided in the first feed.”**

132. Fosnacht also discloses that stored media messages may be “**associated with expiration information,**” as described below. Fosnacht discloses a number of options for customizing aspects of the media item after upload. (Fosnacht, ¶00356 (“FIGS. 42-45 illustrate additional example pages that are presented to a publisher after the publisher has uploaded a media item to the server.”).) Figure 42 includes several of these options, including a starting and ending date the media item will be available on the website:

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

4200 →

4202

mycatholicvoice

My Media :: Edit Media Settings Publisher Settings

General Snack Dessert Thumbnail

Publisher Information

Would you like to enable comment posting for this media item?
☐ Yes ☒ No

Would you like to enable User ratings for this media item?
☐ Yes ☒ No

Start Date: May 21, 2006 End Date: May 22, 2006

Media Status?
Public

Save

“If He who was without sin prayed, how much more ought sinners to pray?”
St. Cyprian

Traveling Dog.net

4204

FIG. 42

(Fosnacht, Fig. 42.) “FIG. 42 is a media editing page 4200 associated with a previously uploaded media item. The media editing page 4200 allows the publisher to define certain aspects of how the media item will be used by the web site.” (Fosnacht, ¶0357.) The key feature of Figure 42 here is publication date selection

section **4204**, which “allows the publisher to identify a starting date and an ending date for the media item to be available through the web site.” (Fosnacht, ¶0360.) Figure 42 discloses options available for publishers, and although not required by the claim, it would have been obvious to provide the same date selection options for non-publisher users in order to give those users more timing control over presentation of media items they upload to the media distribution system.

133. Fosnacht therefore discloses that **“one or more stored media messages are associated with expiration information to determine when the one or more stored media messages can no longer be provided in the first feed,”** as claimed, with **“expiration information”** including at least the ending date specifying when the media item will no longer be available. (*See also* Fosnacht, p.58 (showing “start_date” and “end_date” fields for media items).) With respect to the claimed **“feed,”** as I will explain below in my discussion of limitation 1[c], Fosnacht discloses several different feeds available through the website, including feeds for displaying featured, most viewed, and newest media items, among many other examples as described below. (Fosnacht, ¶0095, Fig. 4.) Because expiration information in Fosnacht allows the publisher to specify the dates “the media item [will] be available through the web site” (Fosnacht, ¶0360), it specifies when the media item can no longer be provided in the first feed.

8. Independent Claim 8

134. Claim 8 is substantially similar to claims 1 and 2, addressed above, with the differences shown below (in underlining).

Claim 1	Claim 8
1. A method, comprising:	8. A method, comprising:
	<u>generating a plurality of media messages at one or more mobile applications executable on a corresponding one or more mobile devices, wherein respective ones of the plurality of media messages include an image or video captured by the one or more mobile applications via a corresponding one or more cameras of the corresponding one or more mobile devices;</u>
receiving, at one or more servers, <u>a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;</u>	receiving, at one or more servers, <u>the plurality of media messages from the one or more mobile applications;</u>
storing the plurality of media messages on the one or more servers;	storing the plurality of media messages on the one or more servers;
providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the	providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

Claim 1	Claim 8
first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and	first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and
providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.	providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.

135. As shown, claim 8 is substantially similar to claim 1 except that it adds

an additional limitation that from dependent claim 2: **“generating a plurality of media messages at one or more mobile applications executable on a corresponding one or more mobile devices, wherein respective ones of the plurality of media messages include an image or video captured by the one or more mobile applications via a corresponding one or more cameras of the corresponding one or more mobile devices.”** Claim 8 is thus obvious for the same reasons as claim 1 and 2 above. As explained for claim 2, the step of **“generating a plurality of media messages”** occurs when the mobile applications on the mobile devices capture the image or video data for the media messages, as explained for claim 3 above.

9. Dependent Claims 9-13

136. Dependent claims 9-11 are substantially similar to dependent claims 4-6, respectively. Dependent claim 12 is substantially similar to dependent claim 3. Dependent claim 13 is substantially similar to dependent claim 7. Claims 9-13 are therefore obvious for the same reasons provided above.

10. Independent Claim 14

137. Claim 14 is substantially similar to claim 1 but written as a system claim, as shown by the table below (with differences shown in underlining):

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

Claim 1	Claim 14
1. A <u>method</u> , comprising:	14. A <u>system</u> , comprising:
	<u>one or more servers including memory, machine-readable instructions, and one or more processors, the one or more processors configured to execute the machine-readable instructions to perform operations comprising:</u>
receiving, at one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;	receiving, at the one or more servers, a plurality of media messages from one or more mobile applications executable on a corresponding one or more mobile devices;
storing the plurality of media messages on the one or more servers;	storing the plurality of media messages on the one or more servers;
providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a	providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

Claim 1	Claim 14
link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and	link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and
providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.	providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.

138. As shown, the only difference between claim 1 and claim 16 is that the latter recites “**one or more servers including memory, machine-readable instructions, and one or more processors, the one or more processors configured to execute the machine-readable instructions,**” which in my opinion, adds nothing of significance. Fosnacht discloses that web server 102 (“**one or more servers**”) includes “a processor and memory, the memory containing computer executable instructions, which when executed by the processor cause the processor

to generate web page data and to communicate the web page data across the data communication network” (Fosnacht, ¶0007), and as explained for claim 1, the web page and web page data in Fosnacht render obvious the limitations of claim 1.

11. Dependent Claims 15-20

139. These claims are substantially similar to claims 2-7, respectively, and this would have been obvious for the same reasons.

12. Independent Claim 21

140. Claim 21 is essentially an amalgam of limitations from claims 1, 8, and 14, with the differences shown in the table below (in underlining):

Claim 1	Claim 21
1. A <u>method</u> , comprising:	24. A <u>system</u> , comprising:
	<u>one or more mobile applications executable on a corresponding one or more mobile devices, the one or more mobile applications configured to create a plurality of media messages, wherein respective ones of the plurality of media messages include an image or video captured via a corresponding one or more cameras of the corresponding one or more mobile devices; and</u>
	<u>one or more servers including memory, machine-readable instructions, and one or more processors, the one or more processors configured to execute the machine-readable instructions to perform operations comprising:</u>

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

Claim 1	Claim 21
receiving, at one or more servers, a plurality of media messages from one or more mobile applications <u>executable on a corresponding one or more mobile devices</u> ;	receiving, at the one or more servers, <u>the</u> plurality of media messages from <u>the</u> one or more mobile applications;
storing the plurality of media messages on the one or more servers;	storing the plurality of media messages on the one or more servers;
providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and	providing a first feed from the one or more servers to a mobile application executable on a mobile device, the first feed including a first subset of media messages selected from among the plurality of media messages stored on the one or more servers, the mobile application configured to present the first feed, wherein presentation of the first feed by the mobile application includes presentation of a comment prompt, a sharing prompt, and a bookmark prompt with a media message included in the first feed, wherein interaction with the comment prompt via the mobile application enables a comment to be associated with the media message, wherein interaction with the sharing prompt via the mobile application facilitates generation of a link configured to enable access to the media message, wherein interaction with the bookmark prompt via the mobile application enables the media message to be associated with another bookmarked media message through the mobile application; and

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

Claim 1	Claim 21
providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.	providing a second feed from the one or more servers to the mobile application, the second feed including a second subset of media messages selected from among the plurality of media messages stored on the one or more servers, the second feed differing from the first feed, the mobile application configured to present the second feed, the mobile application configured to switch a presentation of the first feed to a presentation of the second feed in response to a user interaction with the mobile application.

141. As shown, the sole difference between claim 21 and claim 1 is the addition of two limitations. The first limitation, **“one or more mobile applications executable on a corresponding one or more mobile devices, the one or more mobile applications configured to create a plurality of media messages, wherein respective ones of the plurality of media messages include an image or video captured via a corresponding one or more cameras of the corresponding one or more mobile devices,”** is substantially similar to the first limitation of claim 8, and is obvious for the same reasons. The second limitation, **“one or more servers including memory, machine-readable instructions, and one or more processors, the one or more processors configured to execute the machine-readable instructions,”** is substantially similar to the first limitation of claim 14, and is

obvious for the same reasons.

13. Dependent Claims 22-26

142. Dependent claims 22-24 are substantially similar to dependent claims 4-6, respectively, and are obvious for the same reasons provided above. Dependent claim 25 is substantially similar to dependent claim 3, and is obvious for the same reasons. Dependent claim 26 is substantially similar to dependent claim 7, and is obvious for the same reasons.

C. Ground 2: Claims 5, 10, 18, and 23 Are Obvious Over Fosnacht in View of Kirkpatrick and Sharma

143. As I explained for **Ground 1**, Fosnacht discloses multiple examples of displaying banner advertisements to users of Fosnacht's media distribution system, and I further explained that it would have been obvious that clicking on any of these ads would have resulted in presentation of an advertisement. Nevertheless, I note that Fosnacht does not provide specific details relating to user interaction with these ads – presumably because this functionality would have been plainly apparent and obvious to a skilled artisan and even laypersons who had even a basic understanding of the Internet and how web advertisements operated. Nevertheless, in the event it is argued that the prior art cited for **Ground 1** is insufficiently detailed to cover the limitations of claims 5, 10, 18, and 23, **Ground 2** confirms that these claims are obvious in further view of Sharma.

144. Sharma discloses a technique for allowing an advertiser to deliver relevant ads to users based on search query information provided by a user. (Sharma, Abstract, ¶0009.) But for purposes of my Declaration, I have cited Sharma to confirm exceptionally well-known behavior of Internet ads.

145. Sharma confirms that, as was well-known, websites could present one or more Internet advertisements that, when clicked-on (for example through hypertext links), would cause a user's web browser to present a web page associated with the advertiser. (Sharma, ¶0049.) This advertiser page is frequently called a "landing" page. (Sharma, ¶¶0049, 0059; *see also* Erin Jansen, *NetLingo: The Internet Dictionary* (2002) (EX1007), p.237 ("landing page Another name for a page on a Web site to which people are referred from other Web sites, especially from ad banners... Usually, a company will design one landing page per banner so the follow-up content matches the sell on the particular banner ad.") (boldface in original); *see also id.*, p.29 (definition of "ad banner").) Sharma confirms this basic and well-known sequence.

146. For example, Figure 6 shows an exemplary web page 600 that, in one embodiment, shows search results 610 and one or more ads 620. (Sharma, ¶0067.)





FIGURE 6

(Sharma, Fig. 6 (highlighting added).) Sharma explains that ads may be, for example, “so-called banner ads,” among other types. (Sharma, ¶0033.)

147. As shown in Figure 6 in highlighting, “[a] user has selected, as indicated by click **630**, ad **1**.” (Sharma, ¶0067.) “Ad **1** includes a reference (e.g., a hypertext link) to a landing Web page.” (Sharma, ¶0068.) Accordingly, clicking on ad **630** brings up the landing shown in Figure 7:



700
FIGURE 7

(Sharma, Fig. 7.) “FIG. 7 illustrates an exemplary landing Web page **700**. It **700** includes content **710** and additional ads **720**.” (Sharma, ¶0068.) Sharma therefore confirms that a landing page, *i.e.*, the page presented when a user clicks on a banner ad, can present one or more additional advertisements. Sharma therefore discloses **“presentation of an advertisement prompt... wherein interaction with the advertisement prompt... causes presentation of an advertisement,”** as claimed.

148. *Rationale and Motivation to Combine (Fosnacht and Kirkpatrick, with Sharma)*: It would have been obvious to combine the prior art for **Ground 1** (Fosnacht and Kirkpatrick), with Sharma. The combination would have predictably resulted in the media distribution system according to the Fosnacht-Kirkpatrick

combination, further adapted such that when a user clicked-on any of the banner advertisements shown in Figures, 9, 10, and 21 of Fosnacht (as described for **Ground 1**), the user would have been taken to a landing page in which additional advertisements would have been presented, as described in Sharma. Under the proposed combination, therefore, “**presentation of the first feed... includes presentation of an advertisement prompt with the media message included in the first feed,**” *i.e.*, a banner advertisement displayed by Fosnacht as described for **Ground 1**, “**wherein interaction with the advertisement prompt... causes presentation of an advertisement,**” *i.e.*, clicking on the banner advertisement would have caused presentation of an advertisement, for example through a landing page as described in Sharma. Under the combination with Kirkpatrick, as explained for **Ground 1**, the interaction with the advertising prompt and the presentation of an advertisement would have occurred “**via the mobile application.**”

149. As explained in my discussion of **Ground 1**, Fosnacht and Kirkpatrick are analogous references. Sharma is also an analogous reference in the field of systems and processes for distributing content to devices over the Internet, and would have been reasonably pertinent to problems facing the inventor. For example, Sharma provides a description of the basic functioning of banner advertisements, which is a feature of the alleged invention and reflected in the challenged claims for

this ground. ('1357, 5:16-19 (“Optionally, the current media message interface includes an advertisement banner **112** [Figure 2] for displaying advertisements from a downloaded source, such as from Google AdSense or Apple in-app advertising.”), Fig 2 (showing banner **112**).) The disclosures in Sharma would have been reasonably pertinent in understanding how banner advertisements are displayed and the various ways they can behave when they are clicked on or otherwise activated by users.

150. A person of ordinary skill in the art would have been motivated to combine. As noted for **Ground 1**, Fosnacht already discloses banner advertisements and specifically discloses they can be clicked on to obtain more information (such as clicking “MORE INFO” for advertisement **2116** in Figure 21 of Fosnacht). Fosnacht does not disclose the mechanics of what occurs when a user clicks on a banner advertisement, but the basic functionality of Internet banner advertisements was well-understood as described in Sharma. Sharma thus picks up where Fosnacht leaves off by confirming that it would have been obvious that activating any of the advertisement banners in Fosnacht would have caused presentation of an advertisement, such as the additional advertisements shown in Figure 7 of Sharma. A person of ordinary skill in the art would have been amply motivated to adapt the system of Fosnacht to provide this functionality in order to provide additional

opportunities to provide further promotional material and thus raise additional revenue from user interactions with the banner ads presented through the Fosnacht media distribution system.

151. Finally, a person of ordinary skill in the art would have had at least a reasonable expectation of success. The mechanics of displaying banner advertisements and linking to landing pages was a well-known fixture of webpages since at least the late 1990s, more than a decade before the earliest priority date of the '1357 patent. I note that the '1357 patent itself provides no detailed description about the operation of an advertisement banner and specifically notes that it can come from known “Google AdSense or Apple in-app advertising.” ('1357, 5:16-19.) A person of ordinary skill in the art would have found the disclosures in Sharma more than sufficient to implement the features of the proposed combination described above, without undue experimentation.

V. NO SECONDARY CONSIDERATIONS OF NONOBVIOUSNESS

152. As explained in **Part III.B**, I understand that “secondary considerations” may be used as indicia of nonobviousness, provided that any such consideration has a nexus to the invention claimed. I understand from Petitioner’s counsel that Patent Owner in the related district court litigation has not yet identified any evidence related to secondary considerations with a nexus to the claimed

Declaration of Mark Crovella, Ph.D. in Support of
Petition for *Inter Partes* Review of
U.S. Patent No. 12,131,357 B2

invention. Nor am I aware of information, such as commercial success, unexplained results, long felt but unsolved need, industry acclaim, simultaneous invention, copying by others, skepticism by experts in the field, and failure of others, suggesting that the claims addressed in this Declaration are not obvious. To the extent Patent Owner later provides information it claims relates to secondary considerations, I reserve the right to supplement my analysis and opinions to comment on it.

VI. CONCLUSION

153. In signing this Declaration, I recognize that the Declaration will be filed as evidence in a contested case before the Patent Trial and Appeal Board of the United States Patent and Trademark Office. I also recognize that I may be subject to cross-examination in this proceeding. If required, I will appear for cross-examination at the appropriate time. I reserve the right to offer opinions relevant to the invalidity of the challenged claims at issue and/or offer testimony in support of this Declaration.

154. I hereby declare that all statements made herein of my own knowledge are true and that all statements are made on information and belief are believed to be true, and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under 18 U.S.C. § 1001.

Dated: October 30, 2025

Respectfully submitted,



Mark Crovella, Ph.D.
Boston, Massachusetts

APPENDIX 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 September 18, 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

“Mechanistic Interpretability of Large Language Models,” Open Philanthropy, 2025 – 2026, \$52,646.

“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 Institutional 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 Committee 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 Jomblatt**, 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, Examineur.
- 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-Geertsens**, 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 Communication 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. Crawl, 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

- N. Tasnina, M. Crovella, and T. M. Murali, "Provenance Tracing in Network Diffusion Algorithms," *Proceedings of the Pacific Symposium on Biocomputing (PSB)*, Hawaii, January, 2026.
- G. Franco and M. Crovella, "Pinpointing Attention-Causal Communication in Language Models," *Proceedings of NeurIPS*, San Diego, CA, December 2025
- S. Jasina, L. Salamatian, P. Barford, M. Crovella, and W. Willinger, "A Breath of Fresh Air: Visualizing How Networks 'Breathe'", *Proceedings of the 24th ACM Workshop on Hot Topics in Networking (HotNets 2025)*, College Park, MD, November, 2025.
- S. Buxbaum, L. Tassis, L. Boschelli, D. Christenson, G. Comarela, M. Crovella, and M. Varia, "Privacy-Preserving Machine Learning on Web Browsing for Public Opinion," *Proceedings of The International Symposium on Cyber Security, Cryptology and Machine Learning (CSCML)*, December, 2025.
- P. Calais, G. Franco, Z. Tang, T. Nikas, W. Meira Jr., E. Terzi, and 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, Adaptation, 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, Adaptation, 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. Erdős, 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 MAS-COTS 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 *Proceedings 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," *Proceedings 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," *Proceedings 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," *Proceedings 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. Gummedi, "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.