

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re Patent of: Roy Shkedi Attorney Docket No.: 58717-0006IP1
U.S. Patent No.: 8,677,398
Issue Date: Mar. 18, 2014
Appl. Serial No.: 13/167,605
Filing Date: Jun. 23, 2011
Title: SYSTEMS AND METHODS FOR TAKING ACTION
WITH RESPECT TO ONE NETWORK-CONNECTED
DEVICE BASED ON ACTIVITY ON ANOTHER DEVICE
CONNECTED TO THE SAME NETWORK

DECLARATION OF DR. MICHAEL D. MITZENMACHER

I declare that all statements made herein on my own knowledge are true and that all statements made on information and belief are believed to be true, and further, that these statements were made with the knowledge that willful false statements and the like so made are punishable under Section 1001 of Title 18 of the United States Code.

Date: 7/17/25

By: Michael Mitzenmacher
Dr. Michael Mitzenmacher

TABLE OF CONTENTS

I.	Introduction.....	3
II.	Qualifications and Background Information	3
III.	Overview of the '398 Patent.....	6
IV.	Teachings of Eldering.....	8
V.	A POSITA's Knowledge of IPv6	19
VI.	Combination of Eldering and A POSITA's Knowledge of IPv6	25
	A. Explanation of At Least One Way a POSITA Would Have Known to Combine Eldering and Their Knowledge of IPv6	26
	B. A POSITA's Motivation to Combine Eldering with their Knowledge of IPv6	29
VII.	Teachings of Banga	32
VIII.	Combination of Eldering, A POSITA's Knowledge of IPv6, and Banga.....	34
	A. Explanation of At Least One Way a POSITA Would Have Known to Combine Eldering, Their Knowledge of IPv6, and Banga.....	35
	B. A POSITA's Motivation to Combine Eldering with their Knowledge of IPv6 and Banga.....	41
IX.	Manner in Which The Prior Art References Render the '398 Patent Claims Obvious.....	43
	A. [GROUND 1] Eldering-IPv6-Banga Combination Renders Obvious Claims 1-4, 6-18, 20-44.	43
X.	Level of Ordinary Skill in the Art	64
XI.	Legal Principles	66
	A. Terminology.....	66
	B. Legal Standards.....	66
	C. Anticipation.....	67
	D. Obviousness	67
XII.	Conclusion	71

I. INTRODUCTION

1. I have been retained on behalf of LiveIntent, Inc (“Petitioner”) to offer technical opinions related to 8,677,398 (“the ’398 Patent”) (EX1001). I understand that Petitioners are requesting that the Patent Trial and Appeal Board (“PTAB” or “Board”) to institute an *inter partes* review (“IPR”) proceeding of the ’398 Patent. I have been asked to provide my independent analysis of the ’398 Patent based on the prior art publications cited in this declaration, and my own knowledge, education, experience, and expertise in the relevant fields.

II. QUALIFICATIONS AND BACKGROUND INFORMATION

2. My qualifications are summarized here and explained in more detail in my curriculum vitae, which is attached as Appendix A to this report.

3. I am currently employed as a Professor of Computer Science at Harvard University. Specifically, I am the Thomas J. Watson, Sr. Professor of Computer Science in the School of Engineering and Applied Sciences. I joined the faculty of Harvard as an Assistant Professor in January 1999. I was promoted to Associate Professor in 2002 and to Professor in 2005. In 2010, I began a three-year term as Area Dean, which is essentially equivalent to what other schools call Department Chair of Computer Science, and held that position through June 2013. I also served as the Area Co-Chair for Computer Science from July 1, 2018, to June 30, 2019. My work address is 150 Western Ave, Allston, MA 02134.

4. I received my undergraduate degree in Mathematics and Computer Science from Harvard College in 1991. I received a Certificate of Advanced Study in Mathematics from Cambridge University in 1992. I received a Ph.D. in Computer Science from the University of California at Berkeley in 1996. From August 1996 to January 1999, I was employed as a Research Scientist at Digital Systems Research Center.

5. I have published over 250 research papers in computer science conferences and journals. My general area of research includes algorithms and data structures, as well as the analysis of probabilistic processes. I am listed as an inventor or co-inventor on 19 issued patents and am the co-author of a textbook, entitled *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, published by Cambridge University Press. I am a Fellow of the Institute of Electrical and Electronics Engineers (IEEE) and a Fellow of the Association for Computing Machinery (ACM). In addition, I have consulted for several companies while I have been a professor, such as Microsoft, Mitsubishi, Akamai, ITA software, Fluent Mobile, Adverplex, and e-Harmony.

6. I teach courses on various computer science topics to undergraduate and graduate students at Harvard. Among other topics, I have regularly taught a class on algorithms and data structures for information, networking, and communication to Harvard undergraduate and graduate students.

7. I have both done research in the area of electronic commerce, as well as worked for companies related to electronic commerce. For example, I have written papers regarding how the use of Groupon can affect ratings (such as Yelp ratings) for businesses, on the issue of information asymmetry in pay-per-bid auction sites, and on algorithms for finding statistically significant frequent item sets. Also, as examples, I have served on the scientific advisory board for e-harmony (a dating service with millions of users), and worked at companies such as Google, ITA software (which was bought by Google and serves as the basis for Google flights) and Adverplex/Cogolabs (which works in the space of web and Internet advertising).

8. Much of my work involves issues relating to networking and communications. I regularly serve on program committees for conferences in networking, algorithms, and communication. For example, I have served on the program committee multiple times for the SIGCOMM conference, which is described on the conference homepage as follows: “SIGCOMM is the flagship annual conference of the ACM Special Interest Group on Data Communication (SIGCOMM) on the applications, technologies, architectures, and protocols for computer communication.” <https://sigcomm.org/events/sigcomm-conference>. Similarly, I have served several times on the Program Committee for NSDI, the USENIX Symposium on Networked Systems Design and Implementation. I have

written papers on networking that have been published in the IEEE/ACM Transactions on Computer Networking, the SIGCOMM conference, the INFOCOM conference, and other major venues for networking research. My graduate course entitled “Algorithms at the end of the wire” covers many subjects at the intersection of networking and algorithms.

III. OVERVIEW OF THE '398 PATENT

9. The '398 Patent discloses a method for delivering targeted television advertisements based on online behavior. *See* EX1001, Abstract.

10. Figure 7 of the '398 Patent is reproduced below:




FIG. 7

11. Figure 7 depicts central ad server (CAS) 40 and online site 50 connected to Internet 10; modem 32 connected to Internet service provider (ISP) 22, computer 34 and set-top box (STB) 36; and STB 36 connected to TV 38 and television provider (TVP) 24. *See Ex. 1001, 12:47–13:21.* An identifier can be assigned to STB 36, such as an equipment serial number, a device MAC address, a username, a pseudonym, a tag, or other identifying code or data element, and may include the IP address STB 36 is using when in contact with CAS 40. *See id.* at 15:1–14.

Once a STB is confirmed to have been using a certain IP address at a certain date and time (for example by having a uniquely tagged STB consecutively communicate with the CAS using the same IP address . . .), the CAS database can be searched to find all computers, mobile devices, and other STBs that used the same IP address at the same date and time, and those will be assumed to be devices that are part of the same household and same LAN. Then, an association or link is created among the devices, by linking the device identifiers or tags with the STB's identifier in the CAS database

Once the association is created, profile information (both behavioral and demographic) collected about the computers and mobile devices before and after the association was created could be used to target ads to the STB, regardless of the IP addresses used by the computers and mobile devices (such as cell phones) to access the web at the time the profiles were collected.

EX1001, 18:24–50.

12. I have not analyzed the propriety of any of the priority claims made by the '398 Patent. However, I have applied the earliest possible priority date of the '398 Patent—April 17, 2007—as the invention date (Critical Date) in the following analysis.

IV. TEACHINGS OF ELDERING

13. Eldering describes a “a system, method and apparatus for targeting advertisements (ads) to subscribers.” EX1004, [0024]. More specifically, Eldering teaches “a system and method that is capable of correlating ads with

subscribers based on a plurality of criteria and also a need for a system and method for correlating ads with subscriber data that may be distributed over a plurality of locations.” EX1004, [0019]. “The ads are targeted to subscribers by correlating subscriber profiles with ad profiles.” EX1004, [0024]. “The subscriber profiles identify characteristics and/or traits associated with the subscriber and the ad profiles identify characteristics and/or traits about an intended target market for the ad.” EX1004, [0024].

14. As shown in FIG. 5 (reproduced below), the system described by Eldering comprises two major parts: the Secure Profiling Server (SPS) and the Secure Correlation Server (SCS). EX1004, [0085].

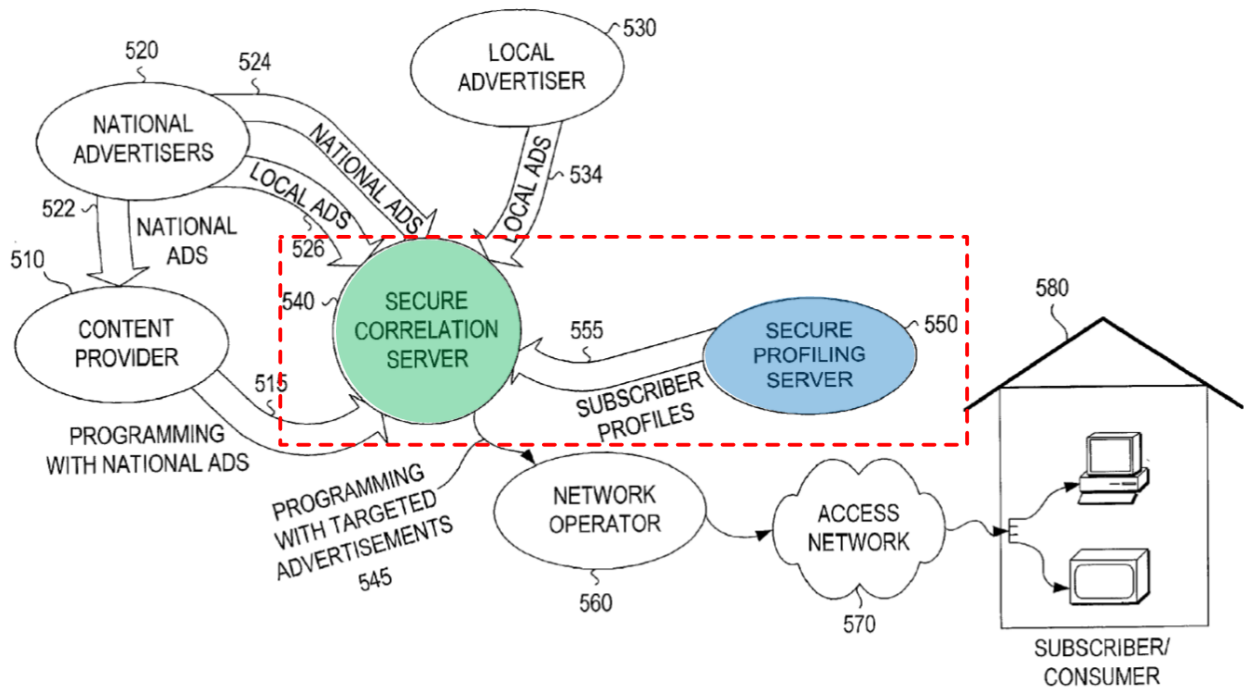


FIG. 5

EX1004, FIG. 5 (Annotated)

15. “The exemplary system includes content providers 510, national advertisers 520, local advertisers 530, a Secure Correlation Server™ (SCS) 540, a Secure Profiling System (SPS) 550, a network operator, an access network and subscribers 580.” EX1004, [0082]. “[S]ubscriber profiles are generated by a Secure Profiling Server (SPS). The characteristics and/or traits associated with the subscriber profile can be retrieved from a plurality of sources.” EX1004, [0022]. “Ad profiles and subscriber profiles are received by a Secure Correlation Server™ (SCS). The SCS correlates the ad profiles with one or more subscriber profiles or one or more group of subscribers.” EX1004, [0025].

16. In the context of Eldering, a POSITA would have understood or at least found obvious that the SPS and/or SCS are programmed hardware computer systems. For example, Eldering describes that the SPS and SCS are together connected to subscribers via an access network 570, which may include the Internet. *See* EX1004, [0085]-[0086]. A POSITA would have recognized that systems connected to the Internet and configured to perform operations like those described with respect to the SPS and SCS would have been programmed hardware computer systems. More to the point, the SCS and SPS are “servers,” which a POSITA would have understood or at least found obvious would have been implemented by programmed hardware computer systems.

17. Furthermore, when describing the operation of the SPS, Eldering teaches that “Applicant’s co-pending U.S. application Ser. No. 09/268,519, filed on Mar. 12, 1999 entitled ‘Consumer Profiling System’ (Atty. Docket No. T706-00), describes in further detail, the collection and aggregation, summation and characterization of subscriber purchases,” and that this application is “incorporated by reference in its entirety.” EX1004, [0108]. The ’519 application issued as U.S. Patent No. 6,298,348. *See* EX1015, 1. With respect to its FIG. 4, the ’348 Patent describes “a computer system for a realization of the consumer profiling system.” EX1015, 8:66-67. It goes on to describe that this computer system includes a “system bus 422 [a] transports data amongst the CPU 203, the RAM 204, Read Only Memory—Basic Input Output System (ROM-BIOS) 406 and other components.” EX1015, 8:67-9:3. In other words, Eldering explicitly describes a computer system that carries out the functions of the SPS, which relies upon standard computer system hardware. *See* EX1015, 8:67-9:16. This computer system would include a non-transitory tangible medium encoding computer-readable instructions that, when applied to one or more computers, instruct the one or more computers to perform a method. *See id.*

18. Further, a POSITA would have understood or at least found obvious that the SPS and SCS could be distinct computer systems or could be components of the same computer system. Eldering illustrates the SPS and SCS as distinct

components within FIG. 5, and describes that the SCS “receives subscriber profiles 555 from the SPS 550.” *See* EX1004, FIG. 5, [0085]-[0087]. Based on these teachings, a POSITA would have understood that the SPS and SCS can be distinct computer systems. However, Eldering also describes that more details about the “collection of subscriber selection data and the generation of subscriber viewing characteristics” (*i.e.*, the role of the SPS) can be found in “Applicant's co-pending U.S. application . . . Ser. No. 09/205,653 filed on Dec. 3, 1998 entitled ‘Client-Server Based Subscriber Characterization System’ (Atty. Docket No. T703-00),” which Eldering “incorporate[s] in [its] entirety.” EX1004, [0106]. The ’653 application issued as U.S. Patent No. 6,457,010. EX1016, 1. The ’010 Patent describes that profile creation is accomplished by the Secure Correlation Server (SCS), without the existence of a separate and distinct SPS. *See* EX1016, FIG. 1, 5:14-19 (“In generating a subscriber profile, the SCS 100 receives from a user 120 commands in the form of a volume control signal 124 or program selection data 122 which can be in the form of a channel change but may also be an address request which requests the delivery of programming from a network address.”). Based on these teachings, a POSITA would have understood that the SPS and SCS could have been implemented as a single computer system.

19. Whether the SPS is part of the SCS or distinct, Eldering describes that “[t]he subscriber profiles are generated by a Secure Profiling Server (SPS).”

EX1004, [0025]. “[T]he SPS 550 generates profiles of the [] subscribers 580 that are used to form groups and thus correlate ads.” EX1004, [0087]. “According to one embodiment, the profiles are formed in advance and forwarded to the SCS 540 where they are matched with ads.” EX1004, [0088]. Eldering describes several types of profiles that might be generated by the SPS and utilized by the SCS to correlate ads. *See* EX1004, [0101]-[0105].

20. For example, Eldering describes a “session profile,” which is “a profile, such as a viewing characteristics profile, that is associated with a single viewing session, wherein the initiation and completion of a viewing session can be determined in various manners.” EX1004, [0104]. The “viewing characteristics” for a “session” may include TV viewing characteristics (*e.g.*, characteristics associated with a user interacting with a set-top box while watching cable television) or Internet viewing characteristics (*e.g.*, characteristics associated with a user interacting with a computer while viewing content on the Internet, such as “sites visited, click throughs, bookmarks and other commands applicable to Internet surfing”). *See* EX1004, [0090]-[0092]. Eldering also describes a “signature profile,” which is “associated with a compilation of viewing sessions.” EX1004, [0105]; *see also id.* [0101], [0109].

21. The signature profile and the session profiles that make it up are for a “subscriber,” which Eldering teaches may be an individual (*i.e.*, “a single

subscriber”) or a household with multiple individuals (*i.e.*, “a household of subscribers”). *See* EX1004, [0101]-[0102] (“subscriber’ -a single subscriber, a household of subscribers, or some combination of subscribers”), [0156] (“According to one embodiment, it is possible for an advertiser to determine the applicability of an ad to a subscriber (individual/household) or group . . . ”); EX1016, 4:18-24, 4:52-61 (describing household profiles that “average characteristics of a household of multiple subscribers” over a number of sessions). And Eldering teaches that a household may include multiple types of devices, all of which may be monitored to collect session data. *See* EX1004, [0084] (“As should be obvious to one of ordinary skill in the art, there are numerous characteristics by which subscribers can be grouped, including but not limited to geographic, . . . viewing habits, . . . [and] Internet surfing habits. . . . These characteristics can be gathered from a multitude of different sources, may be generated within, or a combination thereof”) [0090] (teaching that a subscriber’s viewing characteristics may include “data from a TV viewing characteristics database 612 and an Internet viewing characteristics database 614”); [0092] (describing multiple types of devices that may be monitored); FIG. 5 (showing multiple devices within a household). Taken together, these teachings would have taught a POSITA that a signature profile may be a collection of session profiles for all of the devices and people within a household.

22. Furthermore, Eldering describes that additional information about “collection of subscriber selection data and the generation of subscriber viewing characteristics” is provided in U.S. Appl. No. 09/205,653, which issued as U.S. Pat. No. 6,457,010 and which Eldering incorporates by reference in its entirety. EX1004, [0109]; EX1016, 1. Eldering’s ’010 Patent describes an example embodiment in which the SCS system “can be realized as part of a client-server architecture.” EX1016, 14:52-55. For example, “residence 1800 contains a personal computer (PC) 1820 as well as the combination of a television 1810 and a set-top 1808, which can request and receive programming.” EX1016, 14:55-59. This equipment forms a “client side of the network as defined herein.” EX1016, 14:59-61. “The server side receives the requested programming which is displayed on PC 1820 or television 1810 according to which device made the request.” EX1016, 15:7-10. All of these requests, regardless of the device from which they originate, are recorded by the SCS system for use in creating a household profile. *See* EX1016, 14:52-15:19. Indeed, as of the Critical Date, it was known to create household profiles that track all of the devices associated within a given household. *See, e.g.*, EX1006, [0054] (describing “a system and associated methods to provide a more robust ability to target advertising to consumers across all access devices and mediums and associated services that run over them”), [0058] (describing an example of a household profile that encompasses a mother and her daughters’

browsing and viewing history), [0080] (describing that the system includes “an Internet profiling database 204 which supplies the surfing and content consumption history for the subscriber across multiple access devices”), [0086] (describing that “profile subscriber history data information may be obtained from disparate access networks, even those operated by independent network service providers, which may be combined to make a ‘complete view’ (stored in a user profile) of subscriber usage history across multiple access devices and networks plausible).

Accordingly, a POSITA would have understood or at least found it obvious that Eldering’s household signature profile would have included session data for all of the members of the household across all of the household’s devices.

23. Having described that a “subscriber” may be a “household” (*see* EX1004, [0101]-[0102]), Eldering teaches that, in order to compile a signature profile, the viewing sessions must be “determined to be associated with one another.” EX1004, [0105]. Based on this teaching, a POSITA would have understood and found obvious that, in order to create a signature profile for a household, Eldering’s system would need to associate the session profiles created with respect to each of the devices within the household. Where the SPS is located locally (*e.g.*, within a household’s set-top box), this association would be straightforward, as all data would be collected within the household LAN. *See* EX1004, [0125]. However Eldering describes that session data collection and

profile creation may be performed at “a third party location connected to the access network,” in which case session data would be received from multiple households and would need to be differentiated. *See* EX1004, [0125].

24. For example, Eldering’s FIG. 5 illustrates that the SCS and SPS are connected to a household’s local network via an “access network.” EX1004, [0073], [0085]-[0086]. Eldering describes that “[t]he access network 150 may be a cable TV (CTV) network, a Switched Digital Video (SDV) network or other networks now known or later discovered and may have a hybrid fiber-coax (BFC) architecture, a satellite-based architecture, an Internet-based architecture, digital subscriber line (xDSL) architecture, fiber to the curb (FTTC) or fiber to the home (FTTH), or other architectures now known or later discovered.” EX1004, [0073]. In other words, a POSITA would have understood and found obvious that the devices in household 580 could have been connected to the SCS and SPS across the Internet (*e.g.*, via the subscriber’s ISP). In a case where the SCS and SPS are connected via the Internet, the SCS and SPS would not be in the subscriber’s local-area network.

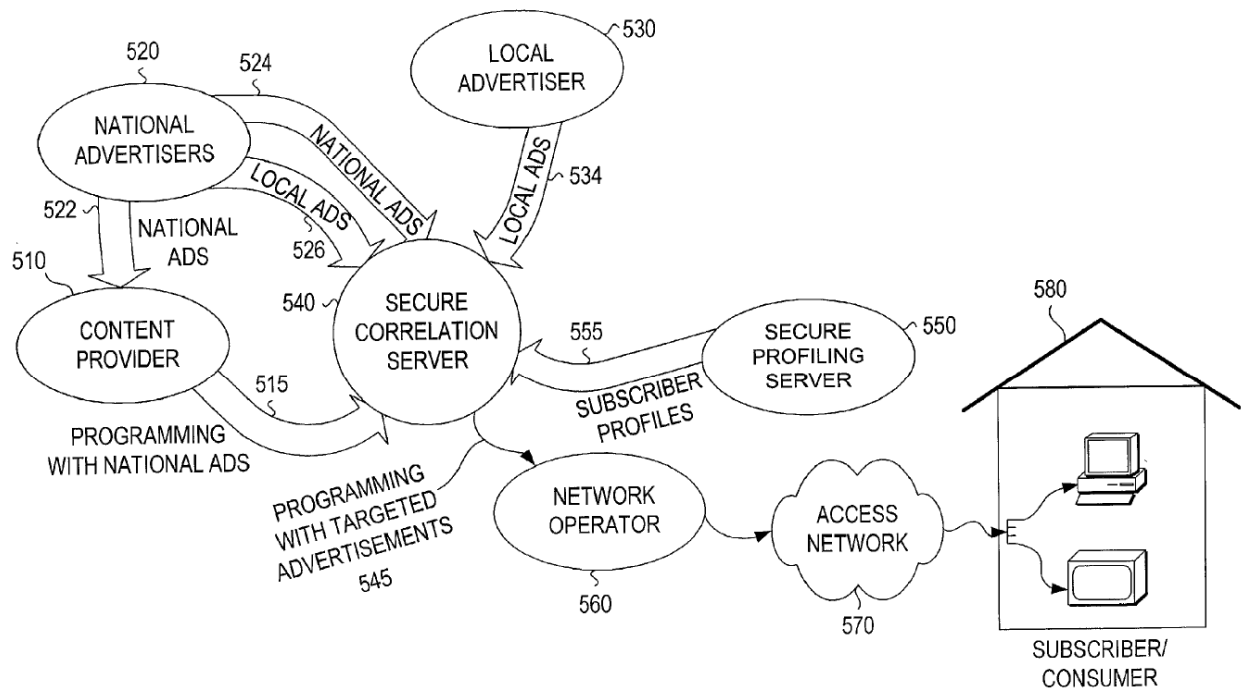


FIG. 5

EX1004, FIG. 5 (Highlighted)

25. For those embodiments where profiles are generated and managed at a third party location, Eldering’s SPS would need to determine what household—and thus which signature profile—received session profiles for a given device should be associated. However, Eldering does not require or otherwise limit a methodology for associating each device with its respective household.

26. Furthermore, Eldering describes that “each subscriber will not be identified by personal information, such as name, but instead will be identified by some unique identification, which may include but is not limited to customer number, media access control (MAC) ID, and Internet protocol (IP) address.”

EX1004, [0091]¹. Thus, Eldering describes that, in order to maintain privacy, the subscriber profiles are identified, for example, by an IP address, which a POSITA would have understood as identifiers for a subscriber's device from which session data is obtained.

27. Accordingly, Eldering teaches a need to associate devices with a household in order to create a household signature profile, and describes that the session data obtained from the household's devices may be identified by an IP address of the device. *See* EX1004, [0091], [0101]-[0105].

V. A POSITA'S KNOWLEDGE OF IPV6

28. As of the Critical Date, a POSITA would have known that at least one way of associating a device with a household was through the use of an Internet Protocol version 6 (IPv6) address, which was a type of IP address that provided a

¹ A POSITA would have understood that the clause "which may include but is not limited to" means the "and" here isn't a conjunction — these are all things that can but don't need to be included (individually). This is further evidenced by the overall teachings of Eldering that subscriber data may be identified with any one or more of these unique identifiers. For example, Eldering describes a case in which "each subscriber is only identified by MAC ID." EX1004, [0116].

globally unique identifier for a device that included information about the network on which the device was located and the MAC address for the device.

29. I am personally familiar with the Internet Engineering Task Force (IETF), and have personal experience with IETF and procedures used by IETF to gather industry members and to distribute information to the same. Specifications for IPv6—including the documents I reference herein (*see generally* EX1007, EX1008, EX1009, EX1010, EX1011)—were defined by working groups within the Internet Engineering Task Force (IETF), and were published in a set of Requests for Comments (RFCs). IETF is a standards setting organization that has been well known to POSITAs since the 1980s. It was and is “the principal body engaged in the development of new Internet standard specifications.” EX1017, 3. “Because of the unrestrictive membership policies, IETF members come from all over the world and from many different parts of the Internet industry.” EX1017, 4. At the time of the Critical Date, I was personally aware of in-person IETF meetings held at least three times a year, which were attended by many industry members who qualified as POSITAs (*see* Section X, *infra*), and mailing lists used by IETF to distribute information to even more POSITAs. *See* EX1017, 2-4. RFCs were and still are “the IETF’s main technical documentation series.” EX1017, 3.

The RFC Editor edits, formats, and publishes Internet Drafts as RFCs, working in conjunction with the IESG. An important secondary role is to provide one definitive repository for all RFCs (see <http://www.rfc-editor.org>). Once an RFC is published, it is never revised. If the standard it describes changes, the standard will be re-published in another RFC that “obsoletes” the first.


EX1017, 7. Thus, the date of publication of an RFC or one of its drafts—including EX1007, EX1008, EX1009, EX1010, EX1011—corresponds to when that RFC was published and made available to POSITAs, and leads me to believe that the RFCs and drafts I reference herein are authentic copies of the same RFCs and drafts available to POSITAs at the Critical Date.

30. Publication of the RFCs and drafts I reference herein by the IETF and their availability on the Internet would have been well known by POSITAs as of the Critical Date. As noted above, a POSITA would know that all IETF RFCs were available at the IETF website, <http://www.rfc-editor.org>. See EX1017, 7, 19. “That site also has links to other RFC collections, many with search capabilities.” EX1007, 19. In fact, a POSITA visiting this well-known site at or near the Critical Date would have had the option to search the RFC index “for a match on number, author, title, and/or keywords,” and located the RFCs and drafts I reference herein based on the same. EX1018, 1. Further, RFC announcements were mailed to any POSITA who requested to be added to the IETF announcement mailing list, [ietf-](#)

announce@ietf.org. EX1017, 8. In other words, the RFCs and drafts I reference herein were both keyword indexed for search on a widely known website, as well as contemporaneously announced via e-mail to thousands of persons skilled in the art of the '398 Patent. In my own experience as of the Critical Date, IETF RFCs were also indexed by Google, and I regularly searched for and found them through Google searches. Therefore, based on my personal experience, I believe that a POSITA would have been able to search Google for the RFCs I reference herein based on keywords contained in their text (*e.g.*, “RFC” and “IPv6”).

31. IPv6 was “a new version of the Internet Protocol, designed as the successor to IP version 4 (IPv4).” EX1007, 2. “IPv6 increases the IP address size from 32 bits to 128 bits, to support more levels of addressing hierarchy, a much greater number of addressable nodes, and simpler auto-configuration of addresses.” EX1007, 2. “The IPv6 address, as currently defined, consists of 64 bits of ‘network number’ and 64 bits of ‘host number’.” EX1008, 2. “The requirements for IPv6 agreed to in 1993 included a plan to be able to address approximately 2^{40} networks and 2^{50} hosts; the 64/64 split effectively accomplishes this,” by providing for up to 2^{64} networks and 2^{64} hosts. EX1008, 2.

32. In practice, a global unicast address for a node (*e.g.*, a user’s device on a network) using the 64/64 split is structured as follows:



EX1010, 2 (“The IPv6 address, as currently defined, consists of 64 bits of ‘network number’ and 64 bits of ‘host number.’”). The global “routing prefix is a value assigned to identify a site (a cluster of subnets/links), the subnet ID is an identifier of a subnet within the site, and the interface ID is a modified EUI-64 format.”

EX1010, 3.

33. For household internet subscribers, the global routing prefix is a unique identification for the subscriber’s local network in their home. The Internet Architecture Board (IAB) and the Internet Engineering Steering Group (IESG) of the IETF “recommended that the address block given to a single edge network . . . be recursively sub-netted be a 48-bit prefix.” EX1008, 3. In this case, each global prefix is 48-bits, 45 of which are variable in which to identify edge networks. EX1008, 3, 5. “That is, the number of available prefixes is 2 to the power 45 or about 35 trillion (35,184,372,088,832).” EX1008, 5. These edge networks range from “large enterprises” to “home network subscribers.” EX1008, 3. The IETF’s “expectation [was] that personal, home networks will become the norm.” EX1008, 3. Thus, when a household signs up for Internet connectivity with an ISP, the IETF initially recommended that ISP’s assign that “home network subscriber” and its local network (*e.g.*, devices sharing a connection to an ISP router) a globally

unique 48-bit routing prefix. EX1008, 3 (“In particular, we recommend: . . . Home network subscribers, connecting through on-demand or always-on connections should receive a /48”). As of the Critical Date, it was common for a household subscriber to connect all devices to a single router, but even in cases where a given household employed multiple subnets, a POSITA would have considered the collection of those subnets a LAN, particularly where all of the subnets located within the same residence and accessed the Internet through the same connection to the ISP (*e.g.*, ISP router).

34. The standard size of the globally unique routing prefix is set by the ISP and may differ from the IETF recommendation (*e.g.*, some ISPs use 56-bit prefixes for households). *See, e.g.*, EX1019, 4-5. However, the principal is the same regardless of the size of the prefix, because the “network number” was always 64 bits. EX1010, 2.

35. Any device that connects to that local household network (or one of its subnets, where a subscriber organizes their network this way) can be assigned a global unicast address for communication on the Internet that is, in part, a combination of (1) the global routing prefix for that home network and (2) a modified EUI-64 format interface ID for the device. EX1008, 2-3; EX1010, 2-3. “Modified EUI-64 format based Interface identifiers may have global scope when derived from a global token (*e.g.*, IEEE 802 48-bit MAC or IEEE EUI-64

identifiers [EUI64]).” EX1009, 8. In other words, a device’s interface ID is a transform of its unique 48-bit MAC address such that its IPv6 address includes the device’s MAC address. *See, e.g.*, EX1012, [0136] (“In IPv6, the possibility arises, that since the IP address includes the MAC address, that the IP address may become an easier and quicker way to derive a device-specific identification”); EX1009, 8-9.

36. Thus, when a device, such as a computer or set-top box, is connected to the Internet through an ISP subscriber’s household router, the device may be assigned an IPv6 address, which includes information to uniquely identify the local household network through which the device accesses the Internet (*i.e.*, the global routing prefix) and information indicating the device’s unique MAC address (*i.e.*, the interface ID).

VI. COMBINATION OF ELDERING AND A POSITA’S KNOWLEDGE OF IPV6

37. As I discuss above, Eldering teaches a need to associate devices within a household in order to create a household signature profile, and describes that the session data obtained from the household’s devices may be identified by a MAC ID and/or IP address of the device. *See* EX1004, [0091], [0101]-[0105]. However, Eldering does not require or otherwise limit a methodology for associating each device with its respective household. As of the Critical Date, a

POSITA would have known that at least one way of associating a device with a household was through the use of the global routing prefix portion of an Internet Protocol version 6 (IPv6) address. *See* EX1008, 3; EX1010, 3.

A. Explanation of At Least One Way a POSITA Would Have Known to Combine Eldering and Their Knowledge of IPv6

38. A POSITA would have found it obvious that at least one advantageous way to associate session profiles would have been for Eldering's system to identify those devices being monitored that were using IPv6 addresses, and to identify the session data as part of a particular household LAN with at least the IPv6 address (a subset of IP addresses used at the time of the '398 Patent), consistent with the teachings of Eldering. *See* EX1004, [0091] (teaching that viewing activity data "will be identified by some unique identification, which may include but is not limited to . . . Internet protocol (IP) address"). At the time of the '398 Patent, it was known to use IPv6 addresses to identify data and profiles in content targeting systems. *See, e.g.*, EX1012, [0134]-[0137]. Thus, a POSITA would have found it obvious that, in at least some cases, the session data of a customer device would have been identified with an IPv6 address of the device, and that Eldering's system would have checked for the use of IPv6 addresses for the devices being monitored for session data in order to associate them with a household signature profile. *See* EX1004, [0105] (describing a signature profile as

being “associated with a compilation of viewing sessions that are determined to be associated with one another”)

39. Specifically, when a customer device in Eldering’s system connected to home networks that utilized IPv6, the device would have been assigned an IPv6 address that included the global routing prefix assigned to the subscriber’s network by the ISP and an interface ID that was a transformation of the device’s unique MAC ID. *See* EX1010, 3; EX1008, 3, 5. When the SPS received session data from that device and the device was connected to the subscriber’s home network during the session, the SPS would have received an identifying IPv6 address that included the global routing prefix of the subscriber’s home network. *See* EX1004, [0091]; EX1008, 3. A POSITA would have found it obvious that the SPS would have stored the IPv6 address associated with each session when compiling sessions to create a household signature profile. *See, e.g.*, EX1005, [0054]; EX1006, [0054], [0086]. Each time the SPS received new session data, the SPS would have compared the global routing prefix received with the new session data to the global routing prefixes stored in the household signature profiles from prior sessions. *See, e.g.*, EX1004, [0072], [0105]; EX1005, [0030], [0054], [0072], [0074]. If there was no match, the SPS would have created a new household signature profile.

40. A POSITA would have understood and found obvious that Eldering's SPS would have been able to utilize the global routing prefix portion of the device IPv6 address associated with the session data to determine which devices are associated with a given household and store the session data in association with that household's signature profile. *See* EX1004, [0105] (describing a signature profile as being "associated with a compilation of viewing sessions that are determined to be associated with one another"). When a given subscriber device is connected to the subscriber's local household network, network traffic would have been routed to that device's IPv6 address, which will include the global routing prefix shared by all of the devices connected to the subscriber's local household network. And the global routing prefix would have been assigned to all of the devices comprising the subscriber's local household network, including any router (*e.g.*, an "ISP router") that may, for example, serve to connect the subscriber's devices to the Internet. *See* EX1008, 2-3, EX1010, 2-3, EX1011, 17-20.

41. Of course, Eldering's system could have utilized other identifying information (*e.g.*, customer number) to identify the household profile with which a given viewing session should be associated. However, a POSITA would have recognized that data received at profiling systems is not always received with a complete set of identifying information (*e.g.*, because the device being connected to the LAN was new to the profiling system and had not yet been fully configured

or because the device employed security software capable of blocking the sharing of identifying information).

42. Implemented in the way I describe above, the Eldering-IPv6 combination represents the predictable and expected result of applying, in the context of Eldering’s profile-based system, an obvious mechanism for determining the household with which a device is associated, as provided by IPv6. From a POSITA’s perspective in 2007, integrating the supplemental teachings of IPv6 with Eldering’s disclosure would have involved nothing more than applying known solutions to known problems in the field of network-based profiling for targeting advertisements to achieve predictable outcomes. The fundamental aspects of Eldering’s disclosure would remain intact, bolstered by the well-known and beneficial teachings of IPv6—not least because Eldering already teaches identifying subscriber data by IP address, and IPv6 was a well-known subtype of IP addresses. *See* EX1004, [0091] (“each subscriber . . . will be identified by some unique identification, which may include . . . Internet protocol (IP) address”).

B. A POSITA’s Motivation to Combine Eldering with their Knowledge of IPv6

43. A POSITA would have recognized that Eldering provides a reason *why* using IPv6 would be advantageous (*i.e.*, building a household signature profile based on information included in each message) and the POSITA’s knowledge of

IPv6 provides the *how* (*i.e.*, the global routing prefix portion of each device's IP address). To this end, there are several reasons that a POSITA would have been motivated to make the abovedescribed Eldering-IPv6 combination.

44. First, IPv6 provides one example of implementation details that Eldering does not otherwise limit to any specific approach. As discussed above, Eldering teaches that a "signature profile" is a "profile associated with a compilation of viewing sessions that are determined to be associated with one another." EX1004, [0105]. However, Eldering does not further limit how this association process occurs and leaves the implementation details to POSITAs. IPv6 provides one example of implementation details a POSITA would have found obvious to utilize when implementing the association of viewing sessions based on IP address. *See* EX1004, [0091], [0105]. And the fact that IP addresses were transitioning from IPv4 to IPv6 would have helped "future-proof" the Eldering-IPv6 system. EX1007, 2.

45. Second, a POSITA would have recognized that associating all devices used in a household in the manner described by Eldering would have allowed for better ad targeting. For example, it was known that "[a]uction-based advertising across multiple screens or access devices based on a correlation or comparison of the advertisement product category with the product interest correlation score for the product/service for a particular subscriber allows advertisers to tightly control

and select what degree of advertisement contact should be made with a subscriber (e.g., one to one TV advertisement, text message to phone) and what price contacting that subscriber with an advertisement is worth.” EX1006, [0077]. By tracking all of the devices for a household, the advertisement targeting capabilities were enhanced.

46. In one example scenario:

[D]uring the summer, Anna and her sister have watched a children’s television show on a particular television channel every afternoon. It is their favorite show. This is reflected in the product interest correlation scores for Anna and her sister and their mom is selected to receive an advertisement in a product category related to the show. When summer ends, as their mom goes shopping for school supplies, she receives a targeted video advertisement on her WiFi enabled Personal Digital Assistant (PDA) or data capable cellular telephone indicating the availability of the children’s television show signature back packs and lunch boxes, providing her with directions on how to get to ‘aisle 12’ in the store (physical merchant location) where the product is located.

EX1006, [0058]; *see also id.*, [0059]-[0060] (describing additional examples). By utilizing the household signature profiles, Eldering’s system would have been able to track the entire subscriber household and select suitable cross-device advertisements.

47. Third, a POSITA would have known that utilizing the global routing prefix portion of the device IPv6 address would have been one of a limited number of ways of identifying the household with which received session data should be associated. Eldering describes that at least one unique identifier that might be used would be a “customer number,” which might have been used to identify all subscribers within a household. *See* EX1004, [0091]. However, utilizing “customer numbers” may become unwieldy to manage across numerous distinct ISPs and may also present privacy concerns. Because every device, regardless of ISP, can have an IPv6 address with a distinct global routing prefix, and the IPv6 address of a device is information that is already transmitted with data requests (and thus does not present additional privacy concerns), the IPv6 address is an advantageous choice among the limited number of ways of identifying the household with which received session data should be associated.

VII. TEACHINGS OF BANGA

48. In addition to tracking devices based on their IP addresses to build profiles, it was well known to combine the IP address with a persistent cookie stored on the device that can be used to associate the device with one or more profiles. For example, like Eldering, Banga describes a system “utilizing a user profile containing user preferences” that “can identify media or media types which are most relevant to the user.” EX1005, [0021]. In Banga’s system, a user

interacts with content (*e.g.* a website) using a network access device (*e.g.*, a “desktop[], laptop[], personal digital assistant[], [or] cellular telephone[]”) that is connected to the Internet by a service provider (*e.g.*, an ISP). EX1005, [0022]-[0027]. “[T]he service provider 120 assigns an external IP address to each network access device 110 seeking access to the communications network 130.” EX1005, [0026].

49. Banga describes a client relationship (CR) server 150 that is configured to identify any number of network access devices 110 seeking access to a network, and generates one or more user profiles. EX1005, [0028]. “Each user profile identifies a user (indirectly via the network access device 110 the user is using) through the current IP address and device identifier (or other information) associated with the network access device 110.” EX1005, [0041]. The IP address of a network access device may then be used to retrieve the user profile associated with the network access device and utilize the profile to serve directed content (*e.g.*, an advertisement). EX1005, [0049], [0056]-[0060], [0072].

50. Though Banga’s system does not require the placement of cookies on a user’s network access device to identify the device, Banga recognized that, “[b]ecause most of the current advertisement industry is based on using cookies, the use of a cookie and/or a device identifier will assist in transitioning the advertisement industry away from the current cookie-only system.” EX1005,

[0031]. “As a result, each network access device 110 may be identified by a hybrid cookie/device identifier combination.” EX1005, [0031]. Indeed, Banga recognizes that different IP addresses may be assigned to a device by different access points through which a user may connect to the Internet (*e.g.*, a router in a coffee shop or a hotspot in the park). EX1005, [0034]-[0035]. A POSITA would have understood and found obvious that a cookie can be used to store these various IP addresses to better track and authenticate the user over time. *See, e.g.*, EX1013, [0094]-[0095], [0099]-[0104]; EX1014, [0120]-[0121], [0133].

VIII. COMBINATION OF ELDERING, A POSITA’S KNOWLEDGE OF IPV6, AND BANGA

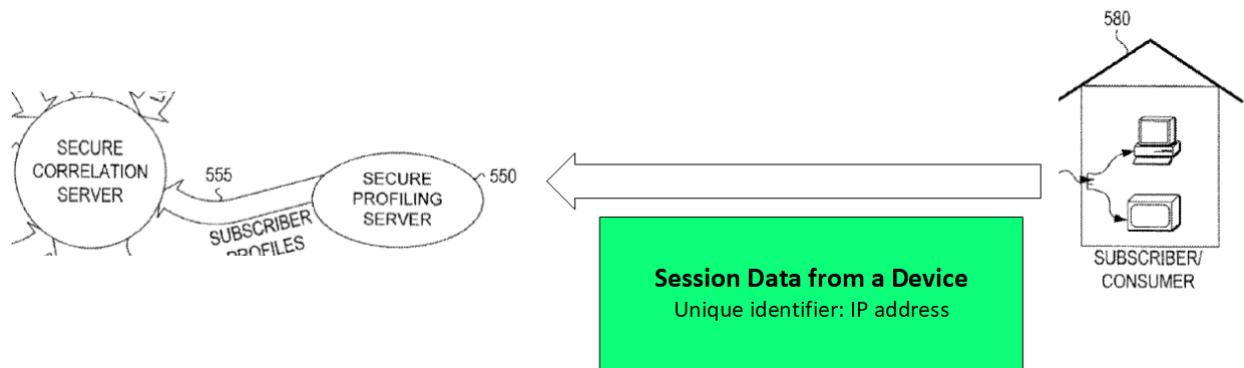
51. In addition to utilizing an IPv6 address to identify a device and associate its session data with a particular profile, a POSITA would have recognized the benefits of utilizing additional information to track the device and associated subscriber. As described in Section III.A.3, *supra*, Banga, like Eldering, teaches tracking a user’s device based on the device’s IP address. EX1005, [0041], [0049], [0056]-[0060], [0072]. However, Banga recognized that, “[b]ecause most of the current advertisement industry is based on using cookies, the use of a cookie and/or a device identifier will assist in transitioning the advertisement industry away from the current cookie-only system.” EX1005,

[0031]. “As a result, each network access device 110 may be identified by a hybrid cookie/device identifier combination.” EX1005, [0031].

A. Explanation of At Least One Way a POSITA Would Have Known to Combine Eldering, Their Knowledge of IPv6, and Banga

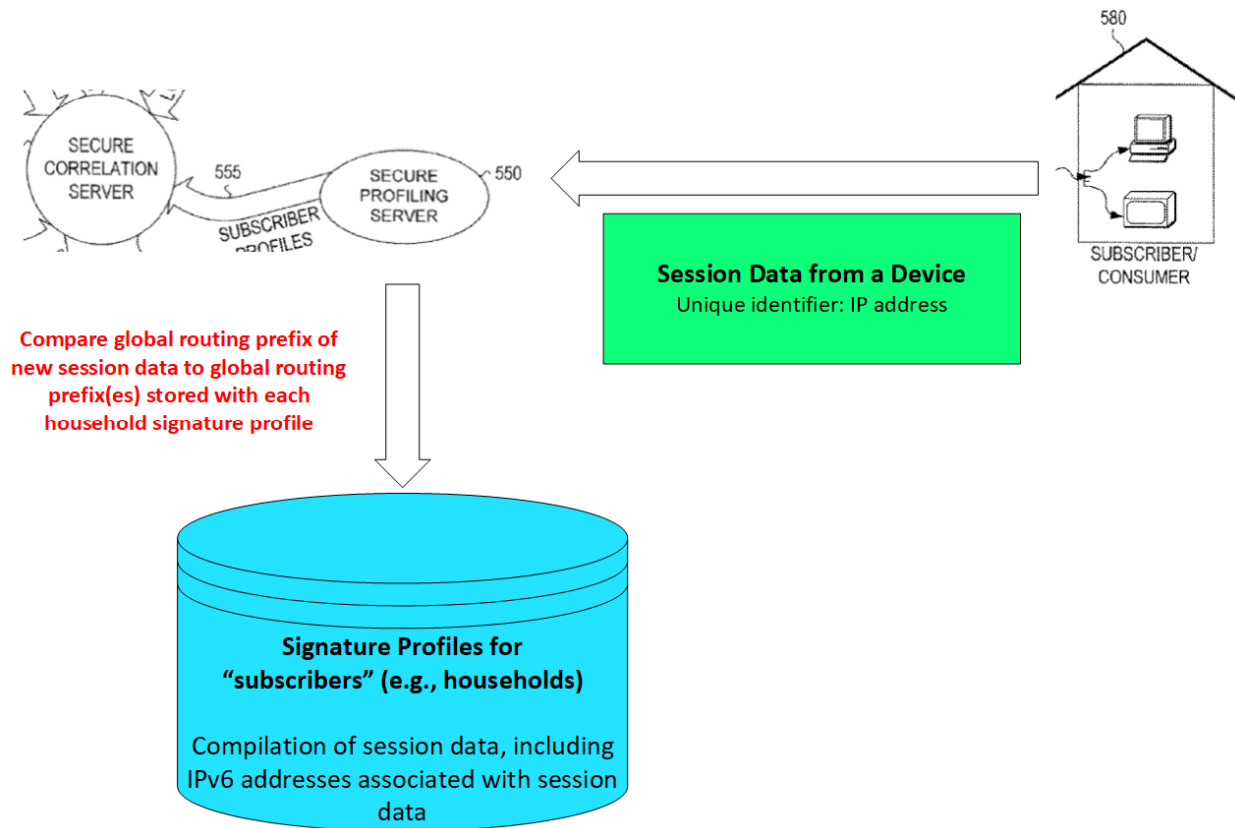
52. Accordingly, the Eldering-IPv6-Banga combination represents the predictable and expected result of applying, in the context of Eldering’s profile-based advertising system, an additional tracking mechanism for a subscriber’s devices, as provided by Banga. As described in greater detail below, when a subscriber’s device was directed to send session data to the SPS, a POSITA would have found it obvious for the SPS to check whether there was a cookie stored on the customer device that identified the customer device with respect to the SPS, and, if not, the SPS would have stored a cookie on the customer device for future identification, even where the customer device joined from another LAN. *See* EX1005, [0007], [0034]-[0035].

53. Such a combination is shown in the following diagram, where Eldering’s SPS (alone or in combination with the SCS) would receive session data (*e.g.*, in the form of a session profile) from a subscriber’s device, and the session data would be identified by an IPv6 address associated with the device.



**Diagram of Elderling-IPv6-Banga Combination
Based on EX1004, FIG. 5**

54. Elderling describes that a “signature profile” is a “compilation of viewing sessions.” EX1004, [0105]; *see also id.* [0101], [0109]. In other words, Elderling describes adding information from each session to a subscriber’s signature profile. *See id.* As I described above with respect to the combination of Elderling and a POSITA’s knowledge of IPv6, a POSITA would have found it obvious that, among other collected information, the SPS would have stored the IPv6 address associated with each session when compiling sessions to create a household signature profile. *See, e.g.,* EX1005, [0054]; EX1006, [0054], [0086]. Each time the SPS received new session data, the SPS would have compared the global routing prefix received with the new session data to the global routing prefixes stored in the household signature profiles from prior sessions. *See, e.g.,* EX1004, [0072], [0105]; EX1005, [0030], [0054], [0072], [0074]. If there was no match, the SPS would have created a new household signature profile. This process of comparison is illustrated in the following diagram:



**Diagram of Elderling-IPv6-Banga Combination
Based on EX1004, FIG. 5**

55. Accordingly, to facilitate the association of session data described by Elderling (*see* EX1004, [0105]; EX1016, 4:18-24, 4:52-61, 14:52-15:19), a POSITA would have understood that each time Elderling's SPS (either alone or in combination with the SCS) received new session data, it would add that session data to the household signature profile, along with the associated IPv6 address. Through its storage as part of the signature profile, the new session data portion of the stored profile indicates an association between the IPv6 addresses of the subscriber household devices, and thus the signature profile itself—or at least the

new session data portion thereof—is electronic indicia of an electronic association. As a result, the updated signature profile with the new session data will include the IPv6 address for each of the household devices from which session data has been collected (*e.g.*, each of the TV+STB and computer illustrated in FIG. 5).

56. In addition to storing the device’s IPv6 address, it would further have been obvious to utilize Banga’s hybrid cookie/device identifier combination when storing the session data. When a device provides session data to the SPS and it does not have a cookie associated with the SPS, the SPS can set a cookie that contains an identifier associated with the household signature profile to which the session data corresponds. *See* EX1005, [0031]. This would have allowed Eldering’s SPS to match a device more easily with the appropriate household signature profile, even if the device later moves to another network (*e.g.*, a customer taking her laptop to a coffee shop). *See* EX1005, [0042], [0053]-[0054].

57. For example, Eldering describes that “Internet transactions are not restricted to computers as one can connect to the Internet with wireless phones, personal digital assistance, and other devices now known to those skilled in the art or later discovered.” EX1004, [0087]. Thus, Eldering describes collecting session data (*e.g.*, Internet transaction data) from a mobile device (*e.g.*, a wireless phone or PDA). *Id.* A POSITA would have understood and found obvious that a wireless telephone and PDA are capable of moving between networks, and therefore be

connected to any one of them (*e.g.*, the subscriber's household network) intermittently. *See, e.g.*, EX1005, [0034]-[0035] (describing different types of access points at different locations to which a mobile device can connect in order to access the Internet).

58. Once the customer device was associated with a household signature profile, Eldering's SCS would have utilized the signature profile to serve all of the customer devices associated with the household signature profile. Specifically, Eldering describes that "the SCS 540 correlates ads and subscribers based on ad characteristics that are received from advertisers and subscriber profiles generated in the SPS 550." Eldering, [0136]. "The SCS correlates the ad profiles with one or more subscriber profiles or one or more group of subscribers. The correlation can be performed by applying an operator to the subscriber profiles in the form of market vectors to determine if a particular ad is applicable to the subscriber." Eldering, [0025]. Eldering describes either that: (1) the SCS can insert ads into the program stream; or (2) that the SCS can direct the subscriber's device to insert ads into the program stream. *See* EX1004, [0086] ("In one embodiment, the SCS 540 creates presentation streams 545 that have the same programming but targeted ads in place of the default ad"), [0030] ("Alternatively, the ad insertion may be done by the node or by the subscriber (via a PVR)"), [0172]-[0174].

59. Because Eldering describes that the SPS creates profiles based on data from multiple different devices, session data collected from one subscriber device (e.g., a computer) would have been used to build the signature household profile that in turn is used to identify targeted ads served to any of the subscriber's devices (e.g., the computer and a television). See EX1004, [0101]-[0105] (“The viewing characteristics may be maintained for viewing sessions, a compilation of viewing sessions, set time durations (i.e., 30 day window), for households, individual subscribers, different combinations of subscribers, other parameters obvious to those skilled in the art, or some combination thereof. The viewing characteristics profile 740 may be represented in vector, table or graphical form and can be the basis for targeting ads and creating subscriber groups.”), [152]-[154], [0156] (describing “determin[ing] the applicability of an ad to a subscriber (individual/household)”). In fact, Eldering's FIG. 5 shows a household with a computer and a TV and FIG. 6 shows the SPS receiving information from “a TV viewing characteristics database 612 and an Internet viewing characteristics database 614.” EX1004, [0090]; *see also id.*, [0073] (“The program stream 145 may be delivered to a personal computer, a TV or any other display means available at the subscriber end.”).

60. A POSITA would have known of similar systems that would have supported this configuration of the Eldering-IPv6-Banga combination. In one such

system, it was known to send targeted advertisements to one subscriber device based on profile information collected from another subscriber device. EX1006, Abstract, [0020], [0022]-[0024], [0043]. For example,

[W]hile doing his homework, Sean, a high school student has tuned in to a hip hop Internet radio channel from his laptop computer. Sean's product interest correlation score is thus increased in this product category, "rap music". Later on in the evening, when he turns on his cell phone to chat with a friend, he receives a targeted text message advertisement for a ring tone of a rap artist's latest track.

EX1006, [0059].

61. From a POSITA's perspective in 2007, integrating the supplemental teachings of Banga with Eldering's disclosure in the manner I describe above would have involved nothing more than applying known solutions to known problems in the field of network-based profiling for targeting advertisements to achieve predictable outcomes. The fundamental aspects of Eldering's disclosure would remain intact, bolstered by the well-known and beneficial teachings of both IPv6 and Banga to successfully arrive at the '398 patent's claimed invention, and the integration of both IPv6 and Banga would not have affected each other.

B. A POSITA's Motivation to Combine Eldering with their Knowledge of IPv6 and Banga

62. The following independent reasons support the Eldering-IPv6-Banga combination.

63. First, as described by Banga, utilizing a hybrid cookie/device identifier combination “the use of a cookie and/or a device identifier will assist in transitioning the advertisement industry away from the current cookie-only system.” EX1005, [0031]. By utilizing Banga’s hybrid cookie/device identifier, the Eldering-IPv6-Banga combination would be able to more readily integrate with existing advertising systems.

64. Second, Banga recognizes that different IP addresses may be assigned to a device by different access points through which a user may connect to the Internet (*e.g.*, a router in a coffee shop or a hotspot in the park). EX1005, [0034]-[0035]. Further, Banga recognized that “users can erase their cookies manually or purchase anti-spyware and anti-virus programs which either prevent cookies from downloading or delete cookies entirely (*e.g.*, cookie churn).” EX1005, [0007]. A POSITA would have understood and found obvious that, by utilizing a combination of both a device identifier (*e.g.*, an IPv6 address) with a cookie, the Eldering-IPv6-Banga combination is more robust, as it does not rely solely upon either piece of information. In fact, a cookie can be used to store the various IP addresses assigned to a device and thus better track and authenticate the user over time. *See, e.g.*, EX1013, [0094]-[0095], [0099]-[0104] (teaching that a cookie can store multiple IP addresses); EX1014, [0120]-[0121], [0133] (teaching that storing

a user's various IP addresses allows for better tracking and may be used to authenticate the user).

65. Third, a POSITA would have found obvious that utilizing a cookie with an identifier associated with the household signature profile would reduce the number of comparisons that need to be made when identifying which household profile session data corresponds to, because any given household signature profile may include all of the IPv6 addresses associated with all of the subscriber's devices. By including an identifier associated with the household signature profile, the cookie includes a reference indicia of the common global routing prefix, because the household signature profile would contain at least the global routing prefix associated with a household's network in order to make the associations of session data to a household signature profile, as I described in more detail above.

IX. MANNER IN WHICH THE PRIOR ART REFERENCES RENDER THE '398 PATENT CLAIMS OBVIOUS

A. [GROUND 1] Eldering-IPv6-Banga Combination Renders Obvious Claims 1-4, 6-18, 20-44.

66. A POSITA would have found claims 1-4, 6-18, 20-44 obvious in view of the Eldering-IPv6-Banga combination that I describe above, as indicated in the table below. Each of Eldering, IPv6, and Banga were analogous to and in the same field as the '398 Patent, as the '398 Patent, Eldering, and Banga both relate to internet profiling and targeting advertisements (*see, e.g.*, EX1001, abstract;

EX1004, abstract; EX1005, abstract), and the '398 Patent relies upon IP addresses of which IPv6 is a subset (*see, e.g.*, EX1001, abstract; EX1007, 2).

Element	Mapping of Relevant Portion of Eldering-IPv6-Banga Combination to Claim Element
[1.pre] A method implemented using a programmed hardware computer system, the method comprising:	A POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination provides <i>a method implemented using a programmed hardware computer system</i> [<i>e.g.</i> , the SPS, alone or in combination with the SCS].
[1.a] receiving, at the computer system, an electronic identifier of a first device;	A POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would <i>receiv[e], at the computer system</i> [<i>e.g.</i> , the SPS or a combination of the SCS and SPS], <i>an electronic identifier of a first device</i> [<i>e.g.</i> , an IPv6 address of a subscriber's device from which session data is received].
[1.b.i] with the computer system, automatically generating and storing electronic indicia of an association between the first device identifier and an electronic identifier of a second device based on automatically recognizing that each of the first and second devices was connected, independently of the other, to a common local area network,	A POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would have <i>automatically recogniz[ed] that each of the first</i> [<i>e.g.</i> , a subscriber's computer connected to the subscriber's home network, as shown in Eldering's FIG. 5] <i>and second devices</i> [<i>e.g.</i> , a subscriber's TV connected to the subscriber's home network via an STB or PVR, as shown in Eldering's FIG. 5] <i>was connected, independently of the other, to a common local area network</i> [<i>e.g.</i> , that the SPS recognized that the computer and set-top box are connected to the same local household network and therefore should be compiled into the same household signature profile based on comparing the global

routing prefix of the IPv6 address of a first device from which session data is received to the global routing prefix of the IPv6 address for a second device from which session data was previously received].

There are at least two ways in which the Eldering-IPv6-Banga combination ***automatically generate[s] and store[s] electronic indicia of an association between the first device identifier and an electronic identifier of a second device.***

First, by adding the session data and associated IPv6 to the household signature profile, the Eldering-IPv6-Banga combination would ***automatically generate and store electronic indicia of an association between the first device identifier and an electronic identifier of a second device*** [e.g., generating and storing a portion of the signature profile corresponding to session data, which through its storage as part of the signature profile indicates an association between the IPv6 addresses of the subscriber household devices].

Second, once the SPS identifies the household signature profile to which session data corresponds based on a comparison of the global routing prefix of the IPv6 address of the device from which the session data was received, a POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would, ***with the computer system*** [e.g., with the SPS or a combination of the SCS and SPS], ***automatically generate and store electronic***

	<p><i>indicia of an association between the first device identifier and an electronic identifier of a second device</i> [e.g., creating and setting a cookie in the device from which session data is received that includes an identifier associated with the corresponding household signature profile].</p>
<p>[1.b.ii] wherein the computer system is connected to the local area network through the Internet but is not in the local area network; and</p>	<p>A POSITA would have understood and found obvious that, in the Eldering-IPv6-Banga combination, <i>the computer system is connected to the local area network through the Internet but is not in the local area network</i> [e.g., the computer and TV+STB shown in FIG. 5 of Eldering are connected to the same ISP router in order to connect to the access network, and the SPS and SCS are located at a third-party location on the Internet].</p>
<p>[1.c] with the computer system, based on the electronic indicia of the association between the first and second device identifiers, automatically sending an electronic transmission that causes another programmed hardware computer system to take an action, based on first electronic profile data associated with the first device identifier, with respect to the second device, which is indicated at the time of the action by the second device identifier</p>	<p>In a case where the SPS and SCS are separate computer systems, a POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would, <i>with the computer system</i> [e.g., the SPS], <i>based on the electronic indicia of the association between the first and second device identifiers</i> [e.g., when the second subscriber device is later viewing content and the SCS determines that the second subscriber device needs a targeted ad, the SPS will retrieve the household signature profile that includes the unique identifier set in the cookies of each of the devices earlier determined to be associated with the household, as well as the device's IPv6 addresses and associated session data], <i>automatically sending an electronic transmission that causes another programmed hardware computer system to take an action</i> [e.g., the SPS sending</p>

subscriber profile information to the SCS—the other programmed hardware computer system—which then utilizes the profile information to select an advertisement for insertion into the user’s “program stream”], **based on first electronic profile data associated with the first device identifier, with respect to the second device** [e.g., the SCS uses prior session data collected from a first subscriber device, such as a computer, to select an ad to be displayed on a second subscriber device, such as a TV], **which is indicated at the time of the action by the second device identifier** [e.g., the SCS inserts a targeted ad into the program stream of the second subscriber device, and when the TV is connected to the subscriber’s household network, the targeted ad will be sent to the IPv6 address associated with the TV that included the global routing prefix of the subscriber’s household network].

In a case where the SPS and SCS are part of the same computer system, a POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would, **with the computer system** [e.g., the SPS in combination with the SCS], **based on the electronic indicia of the association between the first and second device identifiers** [e.g., when the second subscriber device is later viewing content and the SCS determines that the second subscriber device needs a targeted ad, the SPS will retrieve the household signature profile that includes the unique identifier set in the cookies of each of the devices earlier determined to be associated with the household, as well as the device’s IPv6

	<p>addresses and associated session data], automatically sending an electronic transmission that causes another programmed hardware computer system to take an action [e.g., the combination of the SPS and SCS sending instructions to the subscriber's STB/PVR—the other programmed hardware computer system—to insert ads at the subscriber location], based on first electronic profile data associated with the first device identifier, with respect to the second device [e.g., the combination of the SPS and SCS uses prior session data collected from a first subscriber device, such as a computer, to select an ad to be displayed on a second subscriber device, such as a TV], which is indicated at the time of the action by the second device identifier [e.g., the targeted ad are inserted into the program stream of the second subscriber device such as a TV, and when the TV is connected to the subscriber's household network, the targeted ad will be sent to the IPv6 address associated with the TV that included the global routing prefix of the subscriber's household network].</p>
<p>[2] The method of claim 1 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common</p>	<p>In the Eldering-IPv6-Banga combination, <i>the electronic association is recognized based on a common IP address, or portion thereof</i> [e.g., the subscriber's devices are determined to be associated with the same household signature profile based on the global routing prefix of their IPv6 address], <i>wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</i> [e.g., when a given subscriber device is connected to the</p>

IP address ² for each of the first and second devices when each is connected to the local area network.	subscriber's local household network, network traffic will be routed to that device's IPv6 address, which will include the global routing prefix shared by all of the devices connected to the subscriber's local household network].
[3] The method of claim 2 wherein the common IP address ³ comprises an IP address, or portion thereof, of a modem or router by which the first and second devices are connected to the local area network.	In the Eldering-IPv6-Banga combination, a POSITA would have at least found it obvious that <i>the common IP address comprises an IP address, or portion thereof</i> , [e.g., the common global routing prefix portion of the IPv6 address] <i>of a modem or router by which the first and second devices are connected to the local area network</i> [e.g., the ISP router that connects the subscriber's local household network to the Internet, which would have

² For the antecedent basis of this phrase to make sense, this phrase would need to refer back to the “common IP address, or portion thereof.” That is, when the association is made based on a common portion of the IP address, there may not be a full common IP address, so a POSITA would have understood this claim to mean that network traffic is routed to an IP address including the common portion.

³ As with claim [2], when the association is made based on a common portion of the IP address and there is not a full common IP address, a POSITA would have understood this phrase to mean an IP address including the common portion.

	<p>an IPv6 address with the same global routing prefix as all other devices connected to the subscriber's local household network].</p>
<p>[4] The method of claim 2 wherein the electronic indicia of the electronic association include or reference indicia of the common IP address or portion thereof.</p>	<p><u>Under the First Mapping of Element [1.b.i]</u></p> <p>In the Eldering-IPv6-Banga combination, <i>the electronic indicia of the electronic association include or reference indicia of the common IP address or portion thereof</i> [e.g., through its storage as part of the signature profile, session data portion of the stored profile indicates an association between the IPv6 addresses of the subscriber household devices, and thus the signature profile itself—or at least the new session data portion thereof—is an electronic indicia of the electronic association, and the signature profile includes the IPv6 addresses that share a common global routing prefix].</p> <p><u>Under the Second Mapping of Element [1.b.i]</u></p> <p>In the Eldering-IPv6-Banga combination, <i>the electronic indicia of the electronic association include or reference indicia of the common IP address or portion thereof</i> [e.g., by including an identifier associated with the household signature profile, the cookie includes a reference indicia of the common global routing prefix, because the household signature profile would contain at least the global routing prefix associated with a household's network in order to make the associations of session data to a household signature profile].</p>

<p>[6] The method of claim 1 wherein the first device is an online user interface device and the second user interface device is a television set-top box.</p>	<p>In the Eldering-IPv6-Banga combination, a POSITA would have at least found it obvious that <i>the first device is an online user interface device</i> [e.g., the computer in Eldering's FIG. 5] <i>and the second user interface device is a television set-top box</i> [e.g., the television in Eldering's FIG. 5].</p>
<p>[7] The method of claim 1 wherein the first device comprises a mobile device that is connected to the local area network only intermittently.</p>	<p>In the Eldering-IPv6-Banga combination, a POSITA would have at least found it obvious that <i>the first device comprises a mobile device that is connected to the local area network only intermittently</i> [e.g., a wireless telephone and PDA, which Eldering teaches may be the subscriber device, are capable of moving between networks, and therefore be connected to any one of them (e.g., the subscriber's household network) intermittently].</p>
<p>[8] The method of claim 1 wherein the action comprises automatically causing, with the computer system, a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.</p>	<p>In the Eldering-IPv6-Banga combination, a POSITA would have at least found it obvious that <i>the action comprises automatically causing, with the computer system</i> [e.g., either the SPS under the first mapping of Element 1.c or the combination of the SPS and SCS under the second mapping of Element 1.c], <i>a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data</i> [e.g., the SCS uses prior session data collected from a first subscriber device, such as a computer, to select an ad to be displayed on a second subscriber device, such as a TV, and either inserts the ad itself or instructs the subscriber's device to insert the ad]</p>

<p>[9] The method of claim 1 wherein the electronic indicia of the electronic association comprise one or more electronic records in a database accessible to the computer system.</p>	<p>By adding the session data and associated IPv6 to the household signature profile, the Eldering-IPv6-Banga combination would have been configured such that <i>the electronic indicia of the electronic association comprise one or more electronic records in a database accessible to the computer system</i> [e.g., the session data portions of the household signature profile, which include the IPv6 addresses of the household devices from with the session data is received].</p>
<p>[10] The method of claim 1 wherein the electronic indicia of the electronic association comprise one or more tags or cookies readable by the computer system that are stored on the first device or second device.</p>	<p>A POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would have been configured such that <i>the electronic indicia of the electronic association comprise one or more tags or cookies readable by the computer system that are stored on the first device or second device</i> [e.g., the cookie stored on each device matched to the household via its IPv6 address]</p>
<p>[11] The method of claim 10 further comprising, using the computer system, automatically storing, in a memory in the first device or second device, the one or more tags or cookies.</p>	<p>A POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would, <i>using the computer system, automatically stor[e], in a memory in the first device or second device, the one or more tags or cookies</i> [e.g., Eldering’s SPS storing a cookie on the user’s browser per the teachings of Banga].</p>
<p>[12] The method of claim 1 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>a POSITA would have understood and found obvious that the Eldering-IPv6-Banga combination would have been configured such that <i>the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device</i> [e.g., Eldering teaches that “each subscriber <u>will not be identified by</u></p>

	<p><u>personal information</u>, such as name, but instead will be identified by some unique identification, which may include but it not limited to customer number, media access control (MAC) ID, and Internet protocol (IP) address.”] EX1004, [0091].</p>
<p>[13.pre] A method implemented using a programmed hardware computer system, the method comprising:</p>	<p>Element [13.pre] is substantially the same as Element [1.pre].</p>
<p>[13.a] (a) based on first electronic profile data associated with an electronic identifier of a first device, automatically causing, with the computer system, an action to be taken with respect to a second device that is indicated at the time of the action by an electronic identifier electronically associated with the first device identifier;</p>	<p>Element [13.a] is substantially the same as at least portions of Element [1.c].</p>
<p>[13.b] (b) wherein the electronic association between the first and second device identifiers is based on connection, before the action, of each of the first and second devices, independently of the other, to a common local area network, wherein the computer system is connected to the local area network through the Internet but is not in the local area network.</p>	<p>Element [13.b] is substantially the same as at least portions of Elements [1.b.i] and [1.b.ii].</p>

<p>[14.a] The method of claim 13 further comprising: (i) receiving, at the computer system, the first device identifier; and</p>	<p>Element [13.a] is substantially the same as at least portions of Element [1.a].</p>
<p>[14.b] (ii) with the computer system, automatically generating and storing electronic indicia of the association between the first and second device identifiers based on automatically recognizing that, before the action, each of the first and second devices was connected, independently of the other, to the common local area network.</p>	<p>Element [14.b] is substantially the same as at least portions of Element [1.b.i].</p>
<p>[15] The method of claim 13 wherein the first device is connected to the local area network only intermittently.</p>	<p>Claim [15] is substantially the same as at least portions of Claim [7].</p>
<p>[16] The method of claim 13 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</p>	<p>Claim [16] is substantially the same as at least portions of Claim [2].</p>
<p>[17] The method of claim 16 wherein the common IP address comprises an IP address, or portion thereof, of a modem or router by which the first and</p>	<p>Claim [17] is substantially the same as at least portions of Claim [3].</p>

second devices are connected to the local area network.	
[18] The method of claim 16 wherein the electronic indicia of the electronic association include or reference indicia of the common IP address or portion thereof.	Claim [18] is substantially the same as at least portions of Claim [4].
[20] The method of claim 13 wherein the first device is an online user interface device and the second user interface device is a television set-top box.	Claim [20] is substantially the same as at least portions of Claim [6].
[21] The method of claim 13 wherein the first device comprises a mobile device that is connected to the local area network only intermittently.	Claim [21] is substantially the same as at least portions of Claim [7].
[22] The method of claim 13 wherein the action comprises automatically causing, with the computer system, a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.	Claim [22] is substantially the same as at least portions of Claim [8].
[23] The method of claim 13 wherein the electronic indicia of the electronic association comprise one or more electronic records in a database accessible to the computer system.	Claim [23] is substantially the same as at least portions of Claim [9].

<p>[24] The method of claim 13 wherein the electronic indicia of the electronic association comprise one or more tags or cookies readable by the computer system that are stored on the first device or second device.</p>	<p>Claim [24] is substantially the same as at least portions of Claim [10].</p>
<p>[25] The method of claim 24 further comprising, using the computer system, automatically storing, in a memory in the first device or second device, the one or more tags or cookies.</p>	<p>Claim [25] is substantially the same as at least portions of Claim [11].</p>
<p>[26] The method of claim 13 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>Claim [26] is substantially the same as at least portions of Claim [12].</p>
<p>[27.pre] A system comprising one or more computers, wherein the one or more computers are programmed and connected:</p>	<p><i>See</i> [1.pre]. In the context of Eldering, a POSITA would have understood or at least found obvious that the SPS and/or SCS are <i>one or more computer</i> that are <i>programmed and connected</i>.</p>
<p>[27.a] (a) to receive at one or more of the computers an electronic identifier of a first device;</p>	<p>Element [27.a] is substantially the same as at least portions of Element [1.a].</p>
<p>[27.b.i] (b) to generate and store automatically electronic indicia of an association between the first device identifier and an electronic identifier of a second device based on automatically</p>	<p>Element [27.b.i] is substantially the same as at least portions of Element [1.b.i].</p>

<p>recognizing that each of the first and second devices was connected, independently of the other, to a common local area network,</p>	
<p>[27.b.ii] wherein the computer system is connected to the local area network through the Internet but is not in the local area network; and</p>	<p>Element [27.b.ii] is substantially the same as at least portions of Element [1.b.ii].</p>
<p>[27.c] (c) based on the electronic indicia of the association between the first and second device identifiers, to send automatically an electronic transmission that causes another programmed hardware computer system to take an action, based on first electronic profile data associated with the first device identifier, with respect to the second device, which is indicated at the time of the action by the second device identifier.</p>	<p>Element [27.c] is substantially the same as at least portions of Element [1.c].</p>
<p>[28] The system of claim 27 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</p>	<p>Claim [28] is substantially the same as at least portions of Claim [2].</p>

<p>[29] The system of claim 27 wherein the action comprises automatically causing a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.</p>	<p>Claim [29] is substantially the same as at least portions of Claim [8].</p>
<p>[30] The system of claim 27 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>Claim [30] is substantially the same as at least portions of Claim [12].</p>
<p>[31.pre] A system comprising one or more computers, wherein the one or more computers are programmed and connected:</p>	<p>Element [31.pre] is substantially the same as Element [27.pre].</p>
<p>[31.a] (a) based on first electronic profile data associated with an electronic identifier of a first device, to cause automatically an action to be taken with respect to a second device that is indicated at the time of the action by an electronic identifier electronically associated with the first device identifier;</p>	<p>Element [31.a] is substantially the same as at least portions of Element [1.c].</p>
<p>[31.b] (b) wherein the electronic association between the first and second device identifiers is based on connection, before the action, of each of the first and</p>	<p>Element [31.b] is substantially the same as at least portions of Elements [1.b.i] and [1.b.ii].</p>

<p>second devices, independently of the other, to a common local area network, wherein the computer system is connected to the local area network through the Internet but is not in the local area network.</p>	
<p>[32.a] The system of claim 31 wherein the one or more computers are further programmed and connected: (i) to receive at one or more of the computers the first device identifier; and</p>	<p>Element [32.a] is substantially the same as at least portions of Element [1.a].</p>
<p>[32.b] (ii) to generate and store automatically electronic indicia of the association between the first and second device identifiers based on automatically recognizing that, before the action, each of the first and second devices was connected, independently of the other, to the common local area network.</p>	<p>Element [32.b] is substantially the same as at least portions of Element [1.b.i].</p>
<p>[33] The system of claim 31 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</p>	<p>Claim [33] is substantially the same as at least portions of Claim [2].</p>

<p>[34] The system of claim 31 wherein the action comprises automatically causing a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.</p>	<p>Claim [34] is substantially the same as at least portions of Claim [8].</p>
<p>[35] The system of claim 31 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>Claim [35] is substantially the same as at least portions of Claim [12].</p>
<p>[36.pre] An article comprising a non-transitory tangible medium encoding computer-readable instructions that, when applied to one or more computers, instruct the one or more computers to perform a method comprising:</p>	<p><i>See [1.pre]. In the context of Eldering, a POSITA would have understood or at least found obvious that the SPS and/or SCS include an article comprising a non-transitory tangible medium encoding computer-readable instructions that, when applied to one or more computers, instruct the one or more computers to perform a method.</i></p>
<p>[36.a] (a) receiving at one or more of the computers an electronic identifier of a first device;</p>	<p>Element [36.a] is substantially the same as at least portions of Element [1.a].</p>
<p>[36.b.i] (b) automatically generating and storing electronic indicia of an association between the first device identifier and an electronic identifier of a second device based on automatically recognizing that each of the first and second devices was</p>	<p>Element [36.b.i] is substantially the same as at least portions of Element [1.b.i].</p>

<p>connected, independently of the other, to a common local area network,</p>	
<p>[36.b.ii] wherein the computer system is connected to the local area network through the Internet but is not in the local area network; and</p>	<p>Element [36.b.ii] is substantially the same as at least portions of Element [1.b.ii].</p>
<p>[36.c] (c) based on the electronic indicia of the association between the first and second device identifiers, automatically sending an electronic transmission that causes another programmed hardware computer system to take an action, based on first electronic profile data associated with the first device identifier, with respect to the second device, which is indicated at the time of the action by the second device identifier.</p>	<p>Element [36.c] is substantially the same as at least portions of Element [1.c].</p>
<p>[37] The medium of claim 36 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</p>	<p>Claim [37] is substantially the same as at least portions of Claim [2].</p>
<p>[38] The medium of claim 36 wherein the action comprises automatically causing a selected</p>	<p>Claim [38] is substantially the same as at least portions of Claim [8].</p>

<p>advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.</p>	
<p>[39] The medium of claim 36 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>Claim [39] is substantially the same as at least portions of Claim [12].</p>
<p>[40.pre] An article comprising a non-transitory tangible medium encoding computer-readable instructions that, when applied to one or more computers, instruct the one or more computers to perform a method comprising:</p>	<p>Element [40.pre] is substantially the same as Element [36.pre].</p>
<p>[40.a] (a) based on first electronic profile data associated with an electronic identifier of a first device, automatically causing an action to be taken with respect to a second device that is indicated at the time of the action by an electronic identifier electronically associated with the first device identifier;</p>	<p>Element [40.a] is substantially the same as at least portions of Element [1.c].</p>
<p>[40.b] (b) wherein the electronic association between the first and second device identifiers is based on connection, before the action, of each of the first and</p>	<p>Element [40.b] is substantially the same as at least portions of Elements [1.b.i] and [1.b.ii].</p>

<p>second devices, independently of the other, to a common local area network, wherein the computer system is connected to the local area network through the Internet but is not in the local area network.</p>	
<p>[41.a] The medium of claim 40 wherein the method that the one or more computers are instructed to perform using the encoded instructions further comprises: (i) receiving at one or more of the computers the first device identifier; and</p>	<p>Element [41.a] is substantially the same as at least portions of Element [1.a].</p>
<p>[41.b] (ii) automatically generating and storing electronic indicia of the association between the first and second device identifiers based on automatically recognizing that, before the action, each of the first and second devices was connected, independently of the other, to the common local area network.</p>	<p>Element [41.b] is substantially the same as at least portions of Element [1.b.i].</p>
<p>[42] The medium of claim 40 wherein the electronic association is recognized based on a common IP address, or portion thereof, wherein network traffic is routed via the common IP address for each of the first and second devices when each is connected to the local area network.</p>	<p>Claim [42] is substantially the same as at least portions of Claim [2].</p>

<p>[43] The medium of claim 40 wherein the action comprises automatically causing, with the computer system, a selected advertisement to be directed to the second device, wherein the advertisement is selected based at least in part on at least a portion of the first electronic profile data.</p>	<p>Claim [43] is substantially the same as at least portions of Claim [8].</p>
<p>[44] The medium of claim 40 wherein the electronic indicia of the electronic association include no personally identifiable information pertaining to a user of the first device or second device.</p>	<p>Claim [44] is substantially the same as at least portions of Claim [12].</p>

X. LEVEL OF ORDINARY SKILL IN THE ART

67. My analysis and conclusions set forth in this Declaration are based on the perspective of a person of ordinary skill in the art having the level of knowledge and skill as of the invention date of the '398 Patent (POSITA).

68. I was asked to consider the patent claims and the prior art through the eyes of a POSITA at the Critical Date. I understand that the factors considered in determining the ordinary level of skill in a field of art include the level of education and experience of persons working in the field; the types of problems encountered in the field; the teachings of the prior art, and the sophistication of the

technology at the Critical Date. I understand that a POSITA is not a specific real individual, but rather is a hypothetical individual having the qualities reflected by the factors above. I understand that a POSITA would also have knowledge from the teachings of the prior art, including the art cited in this Declaration.

69. Taking these factors into consideration, based on my knowledge and experience in the field and my review of the '398 Patent and file history, I believe that a POSITA in this matter would have had at least a combination of (1) at least a bachelor's degree in computer science, computer engineering, electrical engineering, or a related field, and (2) at least two years of industry experience with e-commerce computing systems, including systems that support targeted advertising. More education could substitute for work experience, and *vice versa*.

70. Based on my experiences, I have a good understanding of the capabilities of a POSITA. Indeed, I have taught, participated in organizations with, and worked closely with many such persons over the course of my career. Based on my knowledge, skill, and experience, I have an understanding of the capabilities of a POSITA. For example, from my industry experience, I am familiar with what a POSITA would have known and found predictable in the art. From teaching and supervising my undergraduate, graduate, and post-graduate students, I also have an understanding of the knowledge that a person with this academic experience

possesses, and, more generally, people with various levels of academic experience possess. Furthermore, I possess those capabilities myself.

XI. LEGAL PRINCIPLES

A. Terminology

71. I understand that claim terms are generally given their plain and ordinary meaning in light of the patent's specification and file history as understood by a POSITA as of the invention date—here, the Critical Date. In that regard, I understand that the best indicator of claim meaning is its usage in the context of the patent specification as understood by a POSITA. I further understand that the words of the claims should be given their plain meaning unless that meaning is inconsistent with the patent specification or the patent's history of examination before the U.S. Patent and Trademark Office (Office). I also understand that the words of the claims should be interpreted as they would have been interpreted by a POSITA at the Critical Date.

B. Legal Standards

72. I have been informed by Counsel and understand that documents and materials that qualify as prior art can render a patent claim unpatentable as being anticipated or obvious.

73. I have been informed by Counsel and understand that all prior art references are to be looked at from the viewpoint of a person of ordinary skill in

the art at the time of the invention, and that this viewpoint prevents one from using his or her own insight or hindsight in deciding whether a claim is anticipated or rendered obvious.

C. Anticipation

74. I understand that patents or printed publications that qualify as prior art can be used to invalidate a patent claim as anticipated or as obvious.

75. I understand that, once the claims of a patent have been properly construed, the second step in determining anticipation of a patent claim requires a comparison of the properly construed claim language to the prior art on a limitation-by-limitation basis.

76. I understand that a prior art reference “anticipates” an asserted claim, and thus renders the claim invalid, if all limitations of the claim are disclosed in that prior art reference, either explicitly or inherently (i.e., necessarily present).

D. Obviousness

77. I understand that even if a patent is not anticipated, it is still invalid if the differences between the claimed subject matter and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a POSITA.

78. I have been informed by Counsel and understand that a claim is unpatentable for obviousness and that obviousness may be based upon a

combination of references. I have been informed by Counsel and understand that the combination of familiar elements according to known methods is likely to be obvious when it does no more than yield predictable results. However, I have been informed by Counsel and understand that a patent claim composed of several elements is not proved obvious merely by demonstrating that each of its elements was, independently, known in the prior art.

79. I have been informed by Counsel and understand that when a patented invention is a combination of known elements, a court determines whether there was an apparent reason to combine the known elements in the fashion claimed by the patent at issue by considering the teachings of prior art references, the effects of demands known to people working in the field or present in the marketplace, and the background knowledge possessed by a person having ordinary skill in the art.

80. I have been informed by Counsel and understand that a patent claim composed of several limitations is not proved obvious merely by demonstrating that each of its limitations was independently known in the prior art. I am informed by Counsel and understand that identifying a reason those elements would be combined can be important because inventions in many instances rely upon building blocks long since uncovered, and claimed discoveries almost of necessity will be combinations of what, in some sense, is already known. I am informed by

Counsel and understand that it is improper to use hindsight in an obviousness analysis, and that a patent's claims should not be used as a "roadmap."

81. I have been informed by Counsel and understand that an obviousness inquiry requires consideration of the following factors: (1) the scope and content of the prior art, (2) the differences between the prior art and the claims, (3) the level of ordinary skill in the art, and (4) any so called "secondary considerations" of non-obviousness, which include: (i) "long felt need" for the claimed invention, (ii) commercial success attributable to the claimed invention, (iii) unexpected results of the claimed invention, and (iv) "copying" of the claimed invention by others.

82. I have been informed by Counsel and understand that an obviousness evaluation can be based on a single reference or a combination of multiple prior art references. I understand that the prior art references themselves may provide a suggestion, motivation, or reason to combine, but that the nexus linking two or more prior art references is sometimes simple common sense. I have been informed by Counsel and understand that obviousness analysis recognizes that market demand, rather than scientific literature, often drives innovation, and that a motivation to combine references may be supplied by the direction of the marketplace.

83. I have been informed by Counsel and understand that if a technique has been used to improve one device, and a person of ordinary skill at the time of

invention would have recognized that it would improve similar devices in the same way, using the technique is obvious unless its actual application is beyond his or her skill.

84. I have been informed by Counsel and understand that practical and common sense considerations should guide a proper obviousness analysis, because familiar items may have obvious uses beyond their primary purposes. I have been informed by Counsel and understand that a person of ordinary skill looking to overcome a problem will often be able to fit together the teachings of multiple prior art references. I have been informed by Counsel and understand that obviousness analysis therefore takes into account the inferences and creative steps that a person of ordinary skill would have employed at the time of invention.

85. I have been informed by Counsel and understand that a proper obviousness analysis focuses on what was known or obvious to a person of ordinary skill at the time of invention, not just the patentee. Accordingly, I understand that any need or problem known in the field of endeavor at the time of invention and addressed by the patent can provide a reason for combining the elements in the manner claimed.

86. I have been informed by Counsel and understand that a claim can be obvious in light of a single reference, without the need to combine references, if

the elements of the claim that are not found explicitly or inherently in the reference can be supplied by the common sense of one of skill in the art.

87. I have been informed by Counsel and understand that there must be a relationship between any such secondary considerations and the invention, and that contemporaneous and independent invention by others is a secondary consideration supporting an obviousness determination.

88. In sum, my understanding is that prior art teachings are properly combined where one of ordinary skill having the understanding and knowledge reflected in the prior art and motivated by the general problem facing the inventor, would have been led to make the combination of elements recited in the claims. Under this analysis, the prior art references themselves, or any need or problem known in the field of endeavor at the time of the invention, can provide a reason for combining the elements of multiple prior art references in the claimed manner.

89. I have been informed by Counsel and understand that in an IPR, “the petitioner shall have the burden of proving a proposition of unpatentability,” including a proposition of obviousness, “by a preponderance of the evidence.”

XII. CONCLUSION

90. I am over the age of 18 and am competent to write this declaration. I have personal knowledge, or have developed knowledge of these technologies based upon education, training, or experience, of the matters set forth herein.

91. I am not and never have been, an employee of LiveIntent. I received no compensation for this declaration beyond my normal hourly compensation based on my time actually spent analyzing the '398 Patent, the prior art publications cited below, and issues related thereto, and I will not receive any added compensation based on the outcome of any IPR or other proceeding involving the '398 Patent.

92. In writing this Declaration, I have considered the following: my own knowledge and experience, including my work experience in the fields of networking, algorithms and data structures, electronic commerce, artificial intelligence, computer systems, and databases; my experience in teaching those subjects; and my experience in working with others involved in those fields.

93. My opinions, as explained above, are based on my education, experience, and expertise in the fields relating to the '398 Patent. Any figures that appear within this document have been prepared with the assistance of Counsel and reflect my understanding of the '398 Patent and the prior art discussed above. In addition, I have analyzed the following publications and materials, in addition to other materials I refer to or cite in my declaration:

Exhibit Number	Citation
EX1001	U.S. Patent No. 8,677,398 to Shkedi (“the '398 Patent”)

EX1002	Excerpts from the Prosecution History of the '398 Patent (“the Prosecution History”)
EX1004	U.S. Pub. No. 2002/0123928 to Charles A. Eldering <i>et al.</i> (“Eldering”)
EX1005	U.S. Pub. No. 2006/0265507 to Jaz Banga <i>et al.</i> (“Banga”)
EX1006	U.S. Pub. No. 2007/0244750 to Grannan <i>et al.</i> (“Grannan”)
EX1007	Internet Engineering Task Force (IETF) Request for Comment (RFC) 2460, <i>Internet Protocol, Version 6 (IPv6)</i> <i>Specification</i> (Dec. 1998)
EX1008	IETF RFC 3177, <i>IAB/IESG Recommendations on IPv6</i> <i>Address Allocations to Sites</i> (Sep. 2001)
EX1009	IETF RFC 3513, <i>Internet Protocol Version 6 (IPv6)</i> <i>Addressing Architecture</i> (Apr. 2003)
EX1010	IETF RFC 3587, <i>IPv6 Global Unicast Address Format</i> (Aug. 2003)
EX1011	Internet Draft v08 of IETF RFC 2462, IPv6 Stateless Address Autoconfiguration, draft-ietf-ipv6-rfc2462bis-08.txt (May 12, 2005)
EX1012	U.S. Pub. No. 2008/0133327 to Shah Ullah

EX1013	U.S. Pub. No. 2007/0038568 to Todd Greene <i>et al.</i>
EX1014	U.S. Pub. No. 2008/0201733 to Kent Ertugrul <i>et al.</i>
EX1015	U.S. Patent No. 6,298,348 to Charles A. Eldering
EX1016	U.S. Patent No. 6,457,010 to Charles A. Eldering and M. Lamine Sylla
EX1017	Wayback Capture of P. Hoffman and S. Harris, “The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force,” available at https://web.archive.org/web/20070214120950/http://www.ietf.org:80/tao.html (captured on Feb. 14, 2007)
EX1018	Wayback Capture of RFC Editor available at https://web.archive.org/web/20070217105548/http://www.rfc-editor.org:80/rfc.html (captured on Feb. 17, 2007)
EX1019	IETF RFC 6177, <i>IPv6 Address Assignment to End Sites</i> (Mar. 2011)

94. I currently hold the opinions set expressed in this declaration. But my analysis may continue, and I may acquire additional information and/or attain supplemental insights that may result in added observations.

APPENDIX A

Michael Mitzenmacher
michaelm@eecs.harvard.edu
33 Cary Avenue, Lexington MA 02421
617-496-7172

Research Interests Design and Analysis of Algorithms; Machine Learning Systems; Networks and Data Transmission; Databases.

Education UNIVERSITY OF CALIFORNIA AT BERKELEY, Berkeley, CA
Ph.D. in Computer Science awarded December, 1996.
Dissertation: **The Power of Two Choices in Randomized Load Balancing.**
Advisor: Alistair Sinclair. GPA: 4.0/4.0

CAMBRIDGE UNIVERSITY, Cambridge, England
Attended as one of ten recipients of the Churchill Fellowship.
Cambridge C.A.S. in Mathematics with highest distinction awarded June 1992.

HARVARD COLLEGE, Cambridge, MA
B.A. in Mathematics with Computer Science, summa cum laude, awarded June 1991.

Employment HARVARD UNIVERSITY, Cambridge, MA *Spring 1999-present*
Assistant professor (from Jan. 1999 -July 2002), **Associate professor** (from July 2002-January 2005), **Professor** (from Jan. 2005-present), **Area Dean for Computer Science** (from July 2010-June 2013), **Area Co-Chair for Computer Science** (from July 2018-June 2019).

Teach undergraduate courses including “Introduction to Algorithms and Data Structures” and “Discrete Mathematics for Computer Science”, and graduate courses covering topics in randomized algorithms, algorithms for networks, compression, coding, cryptography, and information retrieval.

DIGITAL SYSTEMS RESEARCH CENTER, Palo Alto, CA *Fall 1996-Winter 1998*
Research scientist. Projects included work on information retrieval on the Web, erasure codes, error-correcting codes, on-line algorithms, and load balancing. Co-inventor for twelve submitted patents.

SANTA CLARA UNIVERSITY, Santa Clara, CA *Spring 1997*
Guest professor for the undergraduate class “Introduction to Algorithms.”

Consultant: I consult on intellectual property issues as an expert witness and in other capacities. As an expert witness, I have testified in multiple trials. I have also consulted for several technology companies and research laboratories, including Adverplex (Cogolabs), Akamai, AT&T, Digital Fountain, eHarmony, Fluent Mobile (Fiksu), Google, Huawei, ITA Software, JobSync, Microsoft, Mitsubishi Research Laboratories, Weta Digital, and Yahoo.

Funding NSF CCF-2101140: AF: Small: Algorithms and Data Structures with Predictions. PI: Michael Mitzenmacher. Total grant: \$400,000. 7/21-6/24.

NSF CNS-2107078: CNS Core: Medium: Approximation and Randomization in the Programmable Data Plane. co-PIs: Minlan Yu. Total grant: \$1,200,000. 7/21-6/25.

NSF DMS-2023528: Foundations of Data Science Institute (Harvard University). 9/20-8/25.

NSF CCF-1563710: CIF: NeTS: Medium: Collaborative Research: Unifying Data Synchronization. co-PIs: David Starobinksi, Ari Trachtenberg. Total grant: \$400,000. 7/16-6/20.

NSF CCF-1535705: AitF: FULL: Collaborative Research: Better Hashing for Applications: From Nuts & Bolts to Asymptotics. co-PIs: David Andersen. Total grant: \$250,000. 9/15-8/19.

NSF CCF-1320231: AF: Small: Data Synchronization : Theory, Algorithms, and Practice PI: Michael Mitzenmacher. Total grant: \$399,370. 9/13-8/16.

NSF CNS-1228598: TWC: Medium: Collaborative: Privacy-Preserving Distributed Storage and Computation. co-PIs: Michael Goodrich, Roberto Tamassia. Total grant: \$400,000. 8/12-8/16.

NSF IIS-0964473: HCC: Medium: Collaborative Research: Data-Parallel Hash Tables: Theory, Practice and Applications. co-PIs: Nina Amenta. Total grant: \$171,095. 8/10-7/13.

NSF CCF-0915922: AF : Small : The Theory and Practice of Hash-Based Algorithms and Data Structures. PI: Michael Mitzenmacher. Total grant: \$441,956. 8/09-7/12.

NSF CNS-0721491: NeTS FIND: A Network-Wide Hashing Infrastructure for Monitoring and Measurement. PI: Michael Mitzenmacher. Total grant: \$330,000. 9/07-8/10.

NSF CCF-0634923: Towards a Basic Understanding of Channels with Synchronization Errors. PI: Michael Mitzenmacher. Total grant: \$200,000. 9/06-8/09.

NSF CCR-0121154: ITR/SY Algorithmic Issues in Large Scale Dynamic Networks. co-PIs: Eli Upfal, Brown. Subcontract to Harvard. Total grant: \$502,000. 9/01-8/06.

NSF CCR-0118701: Low Density Parity Check Codes for Channels with Memory. co-PIs: Alek Kavcic. Total grant: \$510,000. 9/01-8/04.

NSF CCR-9983832: Dynamic Processes and Network Algorithms (CAREER). \$200,000. 7/00-6/04.

Dean's Competitive Fund for Promising Scholarship. 2018. \$49,333.

Google University Research Program. 8/13-8/14. PIs: Eddie Kohler and Michael Mitzenmacher \$56,000.

Google University Research Program. 12/11-12/12. \$25,000.

Yahoo! University Research Program. 7/11-6/12. \$10,000.

Google University Research Program. 12/09-12/10. \$60,000.

Yahoo! University Research Program. 9/09-8/10. \$10,000.

Google University Research Program. 8/08-7/09. \$75,000.

Yahoo! University Research Program. 9/07-8/08. \$25,000.

Yahoo! University Research Program. 9/06-8/07. \$50,000.

Cisco University Research Program. 8/08-7/09. \$80,000.

Cisco University Research Program. 8/07-7/08. \$83,000.

Cisco University Research Program. 8/06-7/07. \$80,000.

Cisco University Research Program. 8/05-7/06. \$72,000.

Alfred P. Sloan Research Fellowship. \$40,000. Awarded in 2000.

IBM Faculty Research Grant. \$10,000 Awarded in 2005.

IBM Faculty Research Grant. \$10,000 Awarded in 2003.

Mitsubishi Electronic Research Laboratory. \$10,000 for undergraduate research projects. Awarded in 2002.

Honors

Applied Cryptography and Network Security Conference Test of Time Paper Award (2025)

ACM Paris Kanellakis Theory and Practice Award (2020)

IEEE Fellow (2020)

IEEE Senior Member (2018)

ACM Fellow (2014)

ACM Symposium on Parallelism in Algorithms and Architectures Best Paper Award (2014)

World Wide Web Conference Best Paper Award (2014)

Royal Academy of Engineering Distinguished Visiting Fellowship (2010)

ACM SIGCOMM Test of Time Paper Award (2009)

IEEE Information Theory Society Best Paper Award (2002)

Alfred P. Sloan Research Fellowship (2000)

Sakrisson Award (for Ph.D. thesis at Berkeley) (1997)

NDSEG Graduate Fellowship (1992-95)

NSF Graduate Fellowship Winner (1991)

Churchill Fellowship (1991-92)

Hoopess Prize (for senior thesis at Harvard) (1991)

Phi Beta Kappa (1990)

Harvard Distinction in Teaching Award (1990)

Goldwater Fellowship (1989-91)

Professional Activities

Editorships and Boards:

Advisory Board, Hariri Institute at Boston University (2018-2021)

Editorial Board, Leibniz International Proceedings in Informatics (2009-2019)

Editorial Board, Communications of the ACM (2013-2018)

Science Board, Santa Fe Institute (2013-2019)

Guest Editor, SIAM Journal of Computing special issue for STOC 2009

SIAM Journal on Computing, Editor (2006-2010)

Internet Mathematics, Managing Editor

Guest Editor, Theory of Computing System special issue for SPAA 2002

Journal of Interconnection Networks, Editor

Organizing Committees:

Sublinear Algorithms Workshop Organizing Committee, FODSI (2022)
Algorithms with Predictions Workshop Organizing Committee, STOC (2022)
Machine Learning for Algorithms Workshop Organizing Committee, FODSI (2021)
SIGACT Chair (2015-2018)
Dagstuhl Seminar: Theory and Applications of Hashing (2017)
EADS Summer School on Hashing: Theory and Applications (2014)
ICERM Workshop on Stochastic Graph Models (2014)
WAW 2013 Organizing Committee
SIGACT Vice-Chair (2009-2012)
General Chair, STOC (2010)
Workshop on Randomized Algorithms and Random Graphs (2010)
DARPA Information Science and Technology Study Groups (2007-2008)
SIGACT Committee for the Advancement of Theoretical Computer Science (2005-2008)
MSRI Workshop on Models of Real-World Random Networks (2005)
FOCS 2003 Local Arrangements Chair
Second Workshop on Randomized Algorithms and Random Graphs (2003)
Workshop on Algorithms and Models for the Web Graph (2002)
First Workshop on Randomized Algorithms and Random Graphs (2002)
DIMACS Workshop on Quality of Service Issues in the Internet (2001)
BU/NSF Workshop on Internet Measurement, Instrumentation, and Characterization (1999)
FOCS 1998 Local Arrangements Chair

Program Committees:

NSDI 2025 Program Committee
HotNets 2024 Program Committee
WADS 2023 Program Committee
ACM CoNEXT 2023 Program Committee
NeurIPS 2022 Program Committee
ACM CoNEXT 2022 Program Committee
ICML 2022 Program Committee
SIGCOMM 2022 Program Committee
NeurIPS 2021 Program Committee
ICML 2021 Program Committee
ACDA 2021 Program Committee
ICLR 2021 Program Committee
APoCS 2021 Program Committee
ACM CoNEXT 2020 Program Committee
NeurIPS 2020 Program Committee
ICML 2020 Program Committee
APoCS 2020 Program Committee
NeurIPS 2019 Program Committee
SIGCOMM 2019 Program Committee
SOSA 2019 Program Committee (Co-Chair)
ACM CoNEXT 2018 Program Committee
ANALCO 2018 Program Committee
SIGCOMM 2017 Program Committee
NSDI 2017 Program Committee
ACM CoNEXT 2016 Program Committee
SIGCOMM 2016 Program Committee
ALENEX 2016 Program Committee (Co-Chair)
ICALP 2016 Program Committee (Track C PC Chair)

ACM CoNEXT 2015 Program Committee
ICALP 2015 Program Committee
ITCS 2015 Program Committee
SIGCOMM 2014 Program Committee
ACM CoNEXT 2014 Program Committee
WSDM 2014 Senior Program Committee
WAW 2013 Program Committee (Co-Chair)
STOC 2013 Executive Committee
WSDM 2013 Senior Program Committee
SIGCOMM 2012 Program Committee
ISIT 2012 Program Committee
NSDI 2012 Program Committee
WAW 2011 Program Committee
ACM CoNEXT 2010 Program Committee
SIGCOMM 2010 Program Committee
NSDI 2010 Program Committee
FUN 2010 Program Committee
LATIN 2010 Program Committee
NetSciCom 2009 Program Committee
SIGCOMM 2009 Program Committee
STOC 2009 Program Committee (PC Chair)
NSDI 2009 Program Committee
ICALP 2008 Program Committee
SPAA 2008 Program Committee
WAW 2007 Program Committee
Analysis of Algorithms 2007 Program Committee
ANALCO 2007 Program Committee
ICALP 2007 Program Committee
RANDOM 2007 Program Committee
NCA 2006 Program Committee
STOC 2006 Program Committee
PODC 2005 Program Committee
ANALCO 2005 Program Committee
PODC 2004 Program Committee
SIGCOMM 2004 Program Committee
IPTPS 2004 Program Committee
Data Compression Conference 2004 Program Committee
ESA 2004 Program Committee
WAW 2003 Program Committee
FOCS 2003 Program Committee
Data Compression Conference 2003 Program Committee
ESA 2002 Program Committee
PODC 2002 Program Committee
SPAA 2002 Program Committee
STOC 2002 Program Committee
Data Compression Conference 2002 Program Committee
ALENEX 2002 Program Committee
HiPC 2001 Program Committee
ISAAC 2000 Program Committee
RANDOM 2000 Program Committee
STACS 2000 Program Committee

FOCS 99 Program Committee
RANDOM 98 Program Committee

Reviewer for several conferences, journals, and grant panels.

Books

Probability and Computing: Randomized Algorithms and Probabilistic Analysis, by Michael Mitzenmacher and Eli Upfal. Published by Cambridge University Press in 2005. (2nd edition available in 2017.) This is a textbook meant for an advanced undergraduate or beginning graduate class. The textbook has been used in courses at Brown, Harvard, U. C. Berkeley, Univ. of Victoria, Tufts, Univ. of Mass. at Amherst, Purdue, U. Penn., and several other universities.

Conference and Journal Publications

“Don’t Stop Me Now: Embedding Based Scheduling For LLMs,” with R. Shahout, E. Malach, C. Liu, W. Jiang, and M. Yu. To appear in *The Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

“Parallel Peeling of Invertible Bloom Lookup Tables in a Constant Number of Rounds,” with M. Goodrich and R. Kitagawa. To appear in *Proceedings of the 50th International Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM)*, 2025.

“Efficient d -ary Cuckoo Hashing at High Load Factors by Bubbling Up,” with W. Kuzmaul. To appear in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, 2025.

“Leveraging Parameterized Chernoff Bounds for Simplified Algorithm Analyses,” with M. Dillencourt and M. Goodrich. *Information Processing Letters*, 187:106516, 2025.

“Optimal and Approximate Adaptive Stochastic Quantization,” with R. Basat, Y. Ben-Itzhak, and S. Vargaftik. In *Proceedings of the 38th Conference on Neural Information Processing Systems*, 2024.

“SkipPredict: When to Invest in Predictions for Scheduling,” with R. Shahout. In *Proceedings of the 38th Conference on Neural Information Processing Systems*, 2024.

“Beyond Throughput and Compression Ratios: Towards High End-to-end Utility of Gradient Compression,” with W. Han, S. Vargaftik, B. Karp, and R. Basat. In *Proceedings of the Twenty-Third ACM Workshop on Hot Topics in Networks*, pp. 186-194, 2024.

“Learning-Augmented Frequency Estimation in Sliding Windows,” with R. Shahout and I. Sabek. In *Proceedings of the Workshop on Intelligent Classification of High-Speed Network Traffic (at ICNP ’24)*.

“Learning-Based Heavy Hitters and Flow Frequency Estimation in Streams,” with R. Shahout. In *Proceedings of the 32nd IEEE International Conference on Network Protocols*, 2024.

“Tabula: Efficiently Computing Nonlinear Activation Functions for Secure Neural Network Inference,” with M. Lam, V. J. Reddi, G. Wei, and D. Brooks. In *Transactions on Machine Learning Research*, 2024.

“Accelerating Federated Learning with Quick Distributed Mean Estimation,” with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 3410-3442, 2024.

- “THC: Accelerating Distributed Deep Learning Using Tensor Homomorphic Compression,” with M. Li, R. Basat, S. Vargaftik, C. Lao, K. Xu, X. Tang, and, M. Yu. In *Proceedings of NSDI*, pp. 1191-1211, 2024.
- “Separating k -Player from t -Player One-Way Communication, with Applications to Data Streams,” with E. Du, D. Woodruff, and G. Yang. In *Theory of Computing*, vol. 19, Article 10, pp. 1-44, 2023.
- “Direct Telemetry Access,” with J. Langlet, R. Basat, G. Oliaro, M. Yu, and G. Antichi. In *Proceedings of ACM SIGCOMM*, pp. 832-849, 2023.
- “Can Learned Models Replace Hash Functions?” with I. Sabek, K. Vaidya, D. Horn, A. Kipf, and T. Kraska. *Proceedings of the VLDB Endowment*, vol. 16, no. 3, pp. 532-545, 2022.
- “Algorithmic Tools for Mining the Motif Structure of Networks,” with B. Matejek, T. Chen, and C. Tsourakakis. To appear in *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases*, 2022.
- “Edge-colored directed subgraph enumeration on the connectome,” with B. Matejek, D. Wei, T. Chen, C. Tsourakakis, and H. Pfister. *Scientific Reports*, 2022.
- “Algorithms with Predictions,” with S. Vassilvitskii. *Communications of the ACM*, Volume 65 Issue 7, pp. 33-35, 2022.
- “EDEN: Communication-Efficient and Robust Distributed Mean Estimation for Federated Learning,” with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In *Proceedings of the 39th International Conference on Machine Learning*, pp. 21984-22014, 2022.
- “SNARF: A Learning-Enhanced Range Filter,” with K. Vaidya, S. Chatterjee, E. Knorr, S. Idreos, and T. Kraska. In *Proceedings of the VLDB Endowment*, 15(8):1632-1664, 2022.
- “Proteus: A Self-Designing Range Filter,” with E. Knorr, B. Lemaire, A. Lim, S. Luo, H. Zhang, and S. Idreos. In *Proceedings of the SIGMOD International Conference on Management of Data*, pp. 1670-1684, 2022.
- “Designing Sensitive Viral Diagnostics with Machine Learning,” with H. Metsky, N. Welch, P. Pillai, N. Haradhvala, L. Rumker, S. Mantena, Y. Zhang, D. Yang, C. Ackerman, J. Weller, P. Blainey, C. Myhrvold, and P. Sabeti. *Nature Biotechnology*, 2022.
- “The Supermarket Model with Known and Predicted Service Times,” with M. Dell’Amico. *IEEE Transactions on Parallel and Distributed Systems*, 2022.
- “Uniform Bounds for Scheduling with Job Size Estimates,” with Z. Scully and I. Grosf. In *Proceedings of 13th Innovations in Theoretical Computer Science Conference*, 114:1-114:30, 2022.
- “DRIVE: One-bit Distributed Mean Estimation,” with S. Vargaftik, R. Basat, A. Portnoy, G. Mendelson, and Y. Ben-Itzhak. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, pp. 362-377, 2021.
- “Zero-CPU Collection with Direct Telemetry Access,” with J. Langlet, R. Basat, S.S. Ramanathan, G. Oliaro, M. Yu, and G. Antichi. In *Proceedings of the Twentieth ACM Workshop on Hot Topics in Networks*, pp. 108-115, 2021.

“Gradient Disaggregation: Breaking Privacy in Federated Learning by Reconstructing the User Participant Matrix,” with M. Lam, V. J. Reddi, D. Brooks, and G. Wei. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 5959-5968, 2021.

“Putting the ‘Learning’ into Learning-Augmented Algorithms for Frequency Estimation,” with E. Du and F. Wang. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 2860-2869, 2021.

“Queues with Small Advice.” In *Proceedings of the SIAM Conference on Applied and Computational Discrete Algorithms*, pp. 1-12, 2021.

“How to Send a Real Number Using a Single Bit (and Some Shared Randomness),” with R. Basat and S. Vargaftik. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming*, 25:1-25:20, 2021.

“SALSA: Self-Adjusting Lean Streaming Analytics,” with R. Basat, G. Einzinger, and S. Vargaftik. In *The 37th IEEE International Conference on Data Engineering*, pp. 864-875, 2021.

“Partitioned Learned Bloom Filters,” with K. Vaidya, E. Knorr, and T. Kraska. In *The Ninth International Conference on Learning Representations (ICLR)*, 2021.

“Improved Sublinear Time Algorithm for Longest Increasing Subsequence,” with S. Seddighin. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pp. 1934-1947, 2021.

“DISCOVERing the heavy hitters with disaggregated sketches,” with V. Bruschi, R. Basat, Z. Liu, G. Antichi, and G. Bianchi. In *Proceedings of the 16th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 536-537, 2020.

“Detecting Routing Loops in the Data Plane,” with J. Kucera, R. Basat, M. Kuka, G. Antichi, and M. Yu. In *Proceedings of the 16th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 466-473, 2020.

“Optimal Learning of Joint Alignments with a Faulty Oracle,” with K. Larsen and C. Tsourakakis. In *Proceedings of ISIT*, pp.2492-2497, 2020.

“PINT: Probabilistic In-band Network Telemetry,” with R. Basat, S. Ramanathan, Y. Li, G. Antichi, and M. Yu. In *Proceedings of ACM SIGCOMM*, pp. 662-680, 2020.

“Dynamic Algorithms for LIS and Distance to Monotonicity,” with S. Seddighin. In *Proceedings of the 52nd Annual Symposium on Theory of Computing*, pp. 671-684, 2020.

“Clustering with a Faulty Oracle,” with K. Larsen and C. Tsourakakis. In *Proceedings of the Web Conference*, pp. 2831-2834, 2020.

“Prophets, Secretaries, and Maximizing the Probability of Choosing the Best,” with H. Esfandiari, M. Hajiaghayi, and B. Lucier. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics*, pp. 3717-3727, 2020.

“Faster and More Accurate Measurement through Additive-Error Counters,” with R. Basat, G. Einzinger, and S. Vargaftik. In *Proceedings of IEEE INFOCOM*, pp. 1251-1260, 2020. o “Scheduling with Predictions and the Price of Misprediction.” In *Proceedings of 11th Innovations in Theoretical Computer Science Conference*, 14:1-14:18, 2020.

- “Equitability, Interval Estimation, and Statistical Power,” with Y. Reshef, D. Reshef, and P. Sabeti. *Statistical Science*, (35)2:202-217, 2020.
- “A Digital Fountain Retrospective,” with J. Byers and M. Luby. *ACM SIGCOMM Computer Communication Review*, 49(5), pp. 82-85, 2019.
- “Robust Set Reconciliation via Locality Sensitive Hashing,” with T. Morgan. In *Proceedings of the 38th Symposium on Principles of Database Systems*, pp. 164-181, 2019.
- “Online Pandora’s Boxes and Bandits,” with H. Esfandiari, M. Hajiaghayi, and B. Lucier. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, pp. 1885-1892, 2019.
- “Arithmetic Progression Hypergraphs: Examining the Second Moment Method.” In *Proceedings of ANALCO*, pp. 127-134, 2019.
- “A Bayesian Nonparametric View on Count-Min Sketch,” with R. Adams and D. Cai. In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pp. 8782-8791, 2018.
- “A Model for Learned Bloom Filters and Optimizing by Sandwiching.” In *Proceedings of the 31st Conference on Neural Information Processing Systems*, pp. 462-471, 2018.
- “Metric Sublinear Algorithms via Linear Sampling,” with H. Esfandiari. In *Proceedings of the 59th Annual IEEE Symposium on Foundations of Computer Science*, pp. 11-22, 2018.
- “EMOMA: Exact Match in One Memory Access,” with S. Pontarelli and P. Reviriego. *IEEE Transactions on Knowledge and Data Engineering*, 30(11):2120-2133, 2018.
- “Weightless: Lossy Weight Encoding for Deep Neural Network Compression,” with B. Reagen, U. Gupta, R. Adolf, A. Rush, G. Wei, and D. Brooks. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 4321-4330, 2018.
- “Load Thresholds for Cuckoo Hashing with Double Hashing,” with K. Panagiotou and S. Walzer. In *Proceedings of the 16th Scandinavian Symposium and Workshops on Algorithm Theory*, pp. 29:1-29:9, 2018.
- “Joint Alignment From Pairwise Differences with a Noisy Oracle,” with C. Tsourakakis. In *Proceedings of the 15th Workshop on Algorithms and Models for the Web Graph*, pp. 59-69, 2018. Journal version: *Internet Mathematics*, at <https://doi.org/10.24166/im.06.2019>, December 2020.
- “Reconciling Graphs and Sets of Sets,” with T. Morgan. In *Proceedings of the 37th Symposium on Principles of Database Systems*, pp. 33-47, 2018.
- “Simulated Annealing for JPEG Quantization,” with M. Hopkins and S. Wagner-Carena. In *Proceedings of the 2018 IEEE Data Compression Conference*, p. 414.
- “Adaptive Cuckoo Filters,” with S. Pontarelli and P. Reviriego. In *Proceedings of the ALENEX Conference*, pp. 36-47, 2018. Journal version: *ACM Journal of Experimental Algorithmics*, Article 1.1, March 2020.
- “An Empirical Comparison of the Maximal and Total Information Coefficients to Leading Measures of Dependence,” with D. Reshef, Y. Reshef, and P. Sabeti. *Annals of Applied Statistics*, 12(1):123-155, 2018.

- “Simple Multi-Party Set Reconciliation,” with R. Pagh. *Distributed Computing* 31(6):441-453, 2018.
- “Technical Perspective: Building a Better Hash Function.” *Communications of the ACM*, 60(7), p. 93, 2017.
- “Compresso: Efficient Compression of Segmentation Data For Connectomics,” with D. Haehn, F. Lekschas, B. Matejek, and H. Pfister. In *Proceedings of the 20th International Conference on Medical Image Computing and Computer Assisted Intervention*, pp. 781-788, 2017.
- “Auditable Data Structures,” with M. Goodrich, E. Kornaropoulos, and R. Tamassia. In *Proceedings of the European Symposium on Security and Privacy*, pp. 285-300, 2017.
- “2-3 Cuckoo Filters for Faster Triangle Listing and Set Intersection,” with D. Eppstein, M. Goodrich, and M. Torres. In *Proceedings of the 36th Symposium on Principles of Database Systems*, pp. 247-260, 2017.
- “Scalable motif-aware graph clustering,” with C. Tsourakakis and J. Pachocki. In *Proceedings of the 26th International Conference on the World Wide Web*, pp. 1451-1460, 2017.
- “Measuring Dependence Powerfully and Equitably,” with Y. Reshef, D. Reshef, H. Finucane, and P. Sabeti. *Journal of Machine Learning Research*, 17(212):1-63, 2016.
- “Quantized Random Projections and Non-Linear Estimation of Cosine Similarity,” with P. Li and M. Slawski. In *Advances In Neural Information Processing Systems 29*, pp.2748-2756, 2016.
- “Analyzing Distributed Join-Idle-Queue: A Fluid Limit Approach.” In *Proceedings of the 54th Annual Allerton Conference on Communication, Control, and Computing*, pp. 312-318, 2016.
- “Hardness of Peeling with Stashes,” with V. Nathan. *Information Processing Letters*, 116(11), pp. 682-688, 2016.
- “More Practical and Secure History-Independent Hash Tables,” with M. Goodrich, E. Kornaropoulos, and R. Tamassia. In *European Symposium on Research in Computer Security*, pp. 20-38, 2016.
- “Models and Algorithms for Graph Watermarking,” with D. Eppstein, M. Goodrich, J. Lam, N. Mamano, and M. Torres. In *Proceedings of the Information Security Conference*, pp. 283-301, 2016.
- “Better Bounds for Coalescing-Branching Random Walks on Graphs,” with R. Rajaraman and S. Roche. In *Proceedings of the 28th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 313-323, 2016. Journal version: *ACM Transactions on Parallel Computing*, 5(1), 2:1-2:23, 2018.
- “OMASS: One Memory Access Set Separation,” with S. Pontarelli and P. Reviriego. *IEEE Transactions on Knowledge and Data Engineering*, 28(7):1940-1943, 2016.
- “Voronoi Choice Games”, with M. Boppana, R. Hod, and T. Morgan. In *Proceedings of 43rd International Colloquium on Automata, Languages, and Programming*, 23:1-23:13, 2016.

- “Space Lower Bounds for Itemset Frequency Sketches,” with E. Liberty, J. Thaler, and J. Ullman. In *Proceedings of the 35th Symposium on Principles of Database Systems*, pp. 441-454, 2016.
- “Technical Perspective: Catching Lies (and Mistakes) in Offloaded Computation,” with J. Thaler. *Communications of the ACM*, 59(2), p. 102, 2016.
- “A New Approach to Analyzing Robin Hood Hashing.” In *Proceedings of ANALCO*, pp. 10-24, 2016.
- “More Analysis of Double Hashing for Balanced Allocations.” In *Proceedings of ANALCO*, pp. 1-9, 2016.
- “Theory Without Experiments: Have We Gone Too Far?” *Communications of the ACM*, 58(9), pp. 40-42, 2015.
- “Scaling Up Clustered Network Appliances with XBricks,” with D. Zhou, B. Fan, H. Lim, R. Wang, M. Kaminsky, D. Andersen, and A. Singh. In *Proceedings of ACM SIGCOMM*, 2015.
- “Scalable Large Near-Clique Detection in Large-Scale Networks via Sampling,” with J. Pachocki, R. Peng, C. Tsourakakis, and S. Chen Xu. In *Proceedings of the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2015.
- “Repeated Deletion Channels,” with B. Haeupler. In *Proceedings of the IEEE Information Theory Workshop*, pp. 152-156, 2014.
- “Multi-Party Set Reconciliation Using Characteristic Polynomials,” with A. Boral. In *Proceedings of the 52nd Annual Allerton Conference on Communication, Control, and Computing*, pp. 1182-1187, 2014.
- “Cleaning Up the Record on the Maximal Information Coefficient and Equitability,” with D. Reshef, Y. Reshef, and P. Sabeti. *Proceedings of the National Academy of Sciences* 111(33):E3362-3, 2014.
- “Cuckoo Filter: Practically Better than Bloom,” with B. Fan, D.G. Andersen, and M. Kaminsky. In *Proceedings of the 10th International Conference on Emerging Networks Experiments and Technologies (CoNEXT)*, pp. 75-88, 2014.
- “Coding for Random Projections,” with P. Li and A. Shrivastava. In *Proceedings of the 31st International Conference on Machine Learning (ICML)*, pp. 676–684, 2014.
- “Balanced Allocations and Double Hashing.” In *Proceedings of the 26th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 331-342, 2014.
- “Parallel Peeling Algorithms,” with J. Jiang and J. Thaler. In *Proceedings of the 26th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 319-330, 2014. Journal version: *ACM Transactions on Parallel Computing*, 3(1), 7, 2016.
- “Wear Minimization for Cuckoo Hashing: How Not to Throw a Lot of Eggs into One Basket”, with D. Eppstein, M. Goodrich, and P. Pszona. In *Proceedings of the 13th International Symposium on Experimental Algorithms*, pp. 162-173, 2014.

- “Efficient Estimation for High Similarities using Odd Sketches,” with R. Pagh and N. Pham. In *Proceedings of the 23rd International Conference on the World Wide Web*, pp. 109-118, 2014.
- “Improving the Performance of Invertible Bloom Lookup Tables,” with S. Pontarelli and P. Reviriego. *Information Processing Letters*, 114(4), pp. 185-191, 2014.
- “Special Section on the Forty-First Annual ACM Symposium on Theory of Computing (STOC 2009),” with N. Immerlica, J. Katz, R. Servedio, and C. Umans. *SIAM Journal on Computing*, 41(6):1591-1592, 2012.
- “The Daily Deals Marketplace: Empirical Observations and Managerial Implications,” with J. Byers and G. Zervas. *AC SIGecom Exchanges*, 11(2):29-31, 2012.
- “Peeling Arguments and Double Hashing,” with J. Thaler. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1118-1125, 2012.
- “The Complexity of Object Reconciliation, and Open Problems Related to Set Difference and Coding,” with G. Varghese. In *Proceedings of the 50th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1126-1132, 2012.
- “Cache-Oblivious Dictionaries and Multimaps with Negligible Failure Probability,” with M. Goodrich, D. Hirschberg, and J. Thaler. In *Proceedings of the First Mediterranean Conference on Algorithms*, pp. 203-218, 2012.
- “An Economic Analysis of User-Privacy Options in Ad-Supported Services,” with J. Feigenbaum and G. Zervas. In *Proceedings of the 8th Workshop on Internet and Network Economics*, pp. 30-43, 2012.
- “Verifiable Computation with Massively Parallel Interactive Proofs,” with H. Pfister, M. Roberts, and J. Thaler. In *Proceedings of the USENIX HotCloud Workshop, 2012*.
- “Biff (Bloom Filter) Codes: Fast Error Correction for Large Data Sets,” with G. Varghese. In *Proceedings of the International Symposium on Information Theory*, pp. 483-487, 2012.
- “Continuous Time Channels with Interference,” with I. Ivan, J. Thaler, and H. Yuen. In *Proceedings of the International Symposium on Information Theory*, pp. 860-864, 2012.
- “The Groupon Effect on Yelp Ratings: a Root Cause Analysis,” with J. Byers and G. Zervas. In *Proceedings of the ACM Conference on Electronic Commerce*, pp. 248-265, 2012.
- “Anonymous Card Shuffling and Its Applications to Parallel Mixnets,” with M. Goodrich. In *Proceedings of 39th International Colloquium on Automata, Languages, and Programming*, pp. 549-560, 2012.
- “Chernoff-Hoeffding Bounds for Markov Chains: Generalized and Simplified,” with K. Chung, H. Lam, and Z. Liu. In *Proceedings of the 29th International Symposium on Theoretical Aspects of Computer Science*, pp. 124-135, 2012.
- “Practical oblivious storage,” with M. Goodrich, O. Ohrimenko and R. Tamassia. In *Proceedings of the Second ACM Conference on Data and Application Security and Privacy*, pp. 13-24, 2012.

- “Hierarchical Heavy Hitters with the Space Saving Algorithm,” T. Steinke and J. Thaler. In *Proceedings of 2012 Meeting on Algorithm Engineering and Experiments*, pp. 160-174.
- “Daily Deals: Prediction, Social Diffusion, and Reputational Ramifications,” with J. Byers and G. Zervas. In *Proceedings of the 5th International Conference on Web Search and Web Data Mining*, pp. 543-552, 2012.
- “Practical Verified Computation with Streaming Interactive Proofs,” with G. Cormode and J. Thaler. In *Proceedings of Innovations in Theoretical Computer Science*, pp. 90-112, 2012.
- “Information Dissemination via Random Walks in d -Dimensional Space,” with H. Lam, Z. Liu, X. Sun, and Y. Wang. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1612-1622, 2012.
- “Privacy-Preserving Group Data Access via Stateless Oblivious RAM Simulation,” with M. Goodrich and R. Tamassia. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 157-167, 2012.
- “An Efficient Rigorous Approach for Identifying Statistically Significant Frequent Itemsets,” with A. Kirsch, A. Pietracaprina, G. Pucci, E. Upfal, and F. Vandin. *Journal of the ACM*, 59(3), 2012.
- “Detecting Novel Associations in Large Datasets,” with D. Reshef, Y. Reshef, H. Finucane, S. Grossman, G. McVean, P. Turnbaugh, E. Lander, and P. Sabeti. *Science*, Vol. 3004, No. 6602, pp. 1518-1524, December 2011.
- “External-Memory Multimaps,” with E. Angelino, M. Goodrich, and J. Thaler. In *Proceedings of the 22nd International Symposium on Algorithms and Computation*, pp. 384-394, 2011. Journal version: *Algorithmica*, 67(1), pp. 23-48, 2013.
- “Brief announcement: large-scale multimaps,” with M. Goodrich. In *Proceedings of the 23rd ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 259-260, 2011.
- “Oblivious RAM Simulation with Efficient Worst-Case Access Overhead,” with M. Goodrich, O. Ohrimenko, and R. Tamassia. In *Proceedings of 3rd ACM Cloud Computing Workshop*, pp. 95-100.
- “Invertible Bloom Lookup Tables,” with M. Goodrich. In *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*, pp. 792-799, 2011.
- “Cuckoo Hashing with Pages,” with M. Dietzfelbinger and M. Rink. In *Proceedings of the 19th Annual European Symposium on Algorithms*, pp. 615-627, 2011.
- “Network Coding Meets TCP: Theory and Implementation,” with J.K. Sundararajan, D. Shah, M. Medard, S. Jakubczak, and J. Barros. *Proceedings of the IEEE*, 99(3), pp. 490-510, 2011.
- “Privacy-Preserving Access of Outsourced Data via Oblivious RAM Simulation,” with M. Goodrich. In *Proceedings of 38th International Colloquium on Automata, Languages, and Programming*, pp. 576-587, 2011.

- “Graption: A Graph-based P2P Traffic Classification Framework for the Internet Backbone,” with M. Iliofotou, H. Kim, M. Faloutsos, P. Pappu, and G. Varghese. *Computer Networks*, 55(8), pp. 1909-1920, 2011.
- “On the Zero-Error Capacity Threshold for Deletion Channels,” with I. Kash, J. Thaler, J. Ullman. In *Proceedings of the Information Theory and Applications Workshop*, pp. 1-5, 2011.
- “Heapable Sequences and Subsequences,” with J. Byers, B. Heeringa, and G. Zervas. In *Proceedings of ANALCO*, pp. 33-44, 2011.
- “An Introduction to Human Guided Search.” *ACM Crossroads*, 17(2):34-35, 2010.
- “Popularity is Everything: A New Approach to Protecting Passwords from Statistical Guessing Attacks,” with S. Schechter and C. Herley. In *Proceedings of the 5th Usenix Workshop on Hot Topics in Security*, pp. 1-8, 2010.
- “Streaming Graph Computations with a Helpful Advisor,” with G. Cormode and J. Thaler. In *Proceedings of the European Symposium on Algorithms*, pp. 231-242, 2010. Journal version: *Algorithmica*, 65(2), pp. 409-442, 2013.
- “Tight Thresholds for Cuckoo Hashing via XORSAT,” with M. Dietzfelbinger, A. Goerdt, A. Montanari, R. Pagh, and M. Rink. In *Proceedings of 37th International Colloquium on Automata, Languages, and Programming*, pp. 213-225, 2010.
- “Local cluster aggregation models of explosive percolation,” with R.M. D’Souza. *Physical Review Letters*, 14(19), p. 104-107, 2010.
- “Tight Asymptotic Bounds for the Deletion Channel with Small Deletion Probabilities,” with A. Kalai and M. Sudan. In *Proceedings of the International Symposium on Information Theory*, pp. 997-1001, 2010.
- “An Improved Analysis of the Lossy Difference Aggregator,” with H. Finucane. *Computer Communication Review*, 40(2), pp. 4-11, 2010.
- “Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank,” with J. Byers and G. Zervas. In *Proceedings of the ACM Conference on Electronic Commerce*, pp. 1-12, 2010.
- “Carousel: Scalable Logging for Intrusion Prevention Systems”, with T. Lam and G. Varghese. In *Proceedings of NSDI 2010*, pp. 361-376.
- “AMS Without 4-Wise Independence on Product Domains,” with V. Braverman, K. Chung, Z. Liu, and R. Ostrovsky. In *Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science*, pp. 119-130, 2010.
- “Adaptive Weighing Designs for Keyword Value Computation,” with J. Byers and G. Zervas. In *Proceedings of the Third International Conference on Web Search and Web Data Mining* pp. 331-340, 2010.
- “Human-guided search,” with G. Klau, N. Lesh, and J. Marks. *Journal of Heuristics*, 16(3), pp. 289-310, 2010.

- “Real-Time Parallel Hashing on the GPU,” with D. Alcantara, A. Sharf, F. Abbasinejad, S. Sengupta, J. Owens, and N. Amenta. *ACM Transactions on Graphics*, 28(5), p. 154, 2009.
- “Exploiting Dynamicity in Graph-based Traffic Analysis: Techniques and Applications,” with M. Iliofotou and M. Faloutsos. In *Proceedings of ACM Co-NEXT*, pp. 241-252, 2009.
- “An Analysis of Random-Walk Cuckoo Hashing,” with A. Frieze and P. Mellsted. In *Proceedings of 2009 APPROX-RANDOM*, pp. 490-503, 2009. Journal version: *SIAM Journal on Computing*, 40(2), pp. 291-308, 2011.
- “Some Open Questions Related to Cuckoo Hashing.” In *Proceedings of the 17th Annual European Symposium on Algorithms*, pp. 1-10, 2009.
- “On Compressing Social Networks,” with F. Chierichetti, R. Kumar, S. Lattanzi, A. Panconesi, and P. Raghavan. In *Proceedings of the 15th ACM SIGKDD Int’l Conference on Knowledge Discovery and Data Mining*, pp. 219-228, 2009.
- “An Efficient Rigorous Approach for Identifying Statistically Significant Frequent Itemsets,” with A. Kirsch, A. Pietracaprina, G. Pucci, E. Upfal, and F. Vandin. In *Proceedings of the 28th Symposium on Principles of Database Systems*, pp. 117-126, 2009.
- “Using the Power of Two Choices to Improve Bloom Filters,” with Steven Lumetta. *Internet Mathematics*, vol. 4. no. 1, pp. 17-33, 2009.
- “Network coding meets TCP,” with J. Sundararajan, D. Shah, M. Medard, and J. Barros. In *Proceedings of IEEE INFOCOM*, pp. 280-288, 2009.
- “An Economically Principled Generative Model of AS Graph Connectivity,” with J. Corbo, S. Jain, and D. Parkes. In *Proceedings of IEEE INFOCOM (Mini-conference)*, pp. 2941-2945, 2009.
- “Designing floating codes for expected performance,” with H. Finucane and Z. Liu. In *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, 2008. Journal version (with F. Chierichetti, H. Finucane, and Z. Liu): *IEEE Transactions on Information Theory*, vol 56, Issue 3, pp. 968-978, 2010.
- “On the performance of multiple choice hash tables with moves on deletions and inserts,” with A. Kirsch. In *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, pp. 1284-1290, 2008.
- “More Robust Hashing: Cuckoo Hashing with a Stash,” with A. Kirsch and U. Wieder. In *Proceedings of the 16th Annual European Symposium on Algorithms*, pp. 611-622, 2008. Journal version: *SIAM Journal on Computing*, 39(4), pp. 1543-1561, 2009.
- “A survey of results for deletion channels and related synchronization channels.” In *Proceedings of the 2008 Scandinavian Workshop on Algorithm Theory*, pp. 1-3, 2008. Journal version: “New Results and Open Problem for Channels with Synchronization.” *Probability Surveys*, 2009.
- “The Power of One Move: Hashing Schemes for Hardware,” with A. Kirsch. In *Proceedings of IEEE INFOCOM*, pp. 106-110, 2008. Journal version: *IEEE/ACM Transactions on Networking*, 18(6):1752-1765, 2010.

- “Distributed Beamforming with Binary Signaling,” with M. Johnson and K. Ramchandran. In *Proceedings of the 2008 Int’l Symposium on Information Theory (ISIT)*, pp. 890-894, 2008.
- “The Hiring Problem and Lake Wobegon Strategies,” with A. Broder, A. Kirsch, R. Kumar, E. Upfal, and S. Vassilvitskii. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1184-1193, 2008. Journal version: *SIAM Journal on Computing*, 39(4), pp. 1233-1255, 2009.
- “Why Simple Hash Functions Work: Exploiting the Entropy in a Data Stream,” with S. Vadhan. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 746-755, 2008. Journal version: with K. Chung and S. Vadhan. *Theory of Computing*, 9, pp. 897-945-1255, 2013.
- “Trace reconstruction with constant deletion probability and related results,” with T. Holenstein, R. Panigrahy, and U. Wieder. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 389-398, 2008.
- “Capacity bounds for sticky channels.” *IEEE Transactions on Information Theory*, vol 54, Issue 1, pp. 72-77, 2008.
- “Network Monitoring Using Traffic Dispersion Graphs,” with M. Iliofotou, P. Pappu, G. Varghese, M. Faloutsos, and S. Singh. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, pp. 315-320, 2007.
- “HEXA: Compact Data Structures for Faster Packet Processing,” with S. Kumar, J. Turner, and P. Crowley. In *Proceedings of the IEEE International Conference on Network Protocols*, pp. 246-255, 2007.
- “Using a Queue to De-amortize Cuckoo Hashing in Hardware,” with A. Kirsch and M. Mitzenmacher. In *Proceedings of the 45th Annual Allerton Conference on Communication, Control, and Computing*, 2007.
- “Improved lower bounds for the capacity of i.i.d. deletion and duplication channels, with E. Drinea. *IEEE Transactions on Information Theory*, vol 53, Issue 8, pp. 2693-2714, 2007.
- “Wired Geometric Routing,” with J. Ledlie, P. Pietzuch, and M. Seltzer. In *Proceedings of the 6th International Workshop on Peer-to-Peer Systems (IPTPS)*, 2007.
- “Capacity Upper Bounds for Deletion Channels,” with S. Diggavi and H. Pfister. In *Proceedings of the 2007 Int’l Symposium on Information Theory (ISIT)*, pp. 1716-1720, 2007.
- “Codes for Deletion and Insertion Channels with Segmented Errors,” with Z. Liu. In *Proceedings of the 2007 Int’l Symposium on Information Theory (ISIT)*, pp. 846-850, 2007. Journal version: *IEEE Transactions on Information Theory*, vol 56, Issue 1, pp. 224-232, 2010.
- “Towards a theory of networked computation : Executive Summary,” with J. Feigenbaum. *ACM SIGACT News*, (4):22-26, 2006.
- “Bloom Filters via d-Left Hashing and Dynamic Bit Reassignment,” with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of the Allerton Conference*, pp. 877-883, 2006.

- “Beyond Bloom Filters: From Approximate Membership Checks to Approximate State Machines,” with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of ACM SIGCOMM*, pp. 315-326, 2006.
- “An Improved Construction for Counting Bloom Filters,” with F. Bonomi, R. Panigrahy, S. Singh, and G. Varghese. In *Proceedings of the European Symposium on Algorithms*, pp. 684-695, 2006.
- “Less Hashing, Same Performance: Building A Better Bloom Filter,” with A. Kirsch. In *Proceedings of the European Symposium on Algorithms*, pp. 456-467, 2006. Journal version: *Random Structures and Algorithms*, 33(2), pp. 187-218, 2008.
- “Stochastic Shortest Paths via Quasi-Convex Maximization,” with M. Brand, J. Kelner, and E. Nikolova. In *Proceedings of the European Symposium on Algorithms*, pp. 552-563, 2006.
- “Polynomial Time Low-Density Parity-Check Codes with Rates Very Close to the Capacity of the q -ary Random Deletion Channel for Large q .” *IEEE Transactions on Information Theory*, vol 52, Issue 12, pp. 5496-5501, 2006.
- “A Simple Lower Bound for the Capacity of the Deletion Channel,” with E. Drinea. *IEEE Transactions on Information Theory*, vol 52, Issue 10, pp. 4657-4660, 2006.
- “On Lower Bounds for the Capacity of Deletion Channels,” with E. Drinea. *IEEE Transactions on Information Theory*, vol 52, Issue 10, pp. 4648-4657, 2006.
- “On the Theory and Practice of Data Recovery with Multiple Versions.” In *Proceedings of the International Symposium on Information Theory*, pp. 982-986, 2006.
- “Fine-Grained Layered Multicast with STAIR,” with J. Byers, Gu-In Kwon, and M. Luby. *IEEE/ACM Transactions on Networking*, 14:1, pp. 81-93, 2006.
- “Distance-Sensitive Bloom Filters,” with A. Kirsch. In *Proceedings of the ALENEX Conference*, pp. 41-50, 2006.
- “BubbleSearch: A Simple Heuristic for Improving Priority-Based Greedy Algorithms,” with N. Lesh. *Information Processing Letters*, 97(4), pp. 161-169, 2006.
- “Network-Aware Overlays with Network Coordinates,” with P. Pietzuch, J. Ledlie, and M. Seltzer. In *Proceedings of the International Workshop on Dynamic Distributed Systems*, 2006.
- “Editorial: The Future of Power Law Research.” *Internet Mathematics*, vol. 2. no. 4, pp. 525-534, 2006.
- “X-Tolerant Test Response Compaction,” with S. Mitra, S. Lumetta, and N. Patil. *IEEE Design and Test of Computers*, vol. 22, no. 6, pp. 566-574, Nov/Dec 2005.
- “Privacy preserving keyword searches on remote encrypted data,” with Y.-C. Chang. In *Proceedings of the Third Intl. Conference on Applied Cryptography and Network Security*, pp. 442-455, 2005.
- “Simple Summaries for Multiple Choice Hashing,” with A. Kirsch. In *Proceedings of the 43rd Annual Allerton Conference on Communication, Control, and Computing*, 2005. Journal version: *IEEE/ACM Transactions on Networking*: 16(1), pp. 218-231, 2008.

- “A case study in large-scale interactive optimization,” with M. Chimani, N. Lesh, and C. Sidner. In *Proceedings of Artificial Intelligence and Applications 2005*, pp. 24-29.
- “New heuristic and interactive approaches to 2D rectangular strip packing,” with N. Lesh, A. McMahon, and J. Marks. *ACM Journal of Experimental Algorithms*, 10, 2005.
- “Multidimensional Balanced Allocations,” with A. Broder. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 195-196, 2005.
- “Digital Fountains: A Survey and Look Forward.” 2004 Information Theory Workshop.
- “Improved Lower Bounds on the Capacity of I.I.D. Deletion Channels,” with E. Drinea. In *Proceedings of the 42nd Annual Allerton Conference on Communication, Control, and Computing*, 2004.
- “Improved Lower Bounds on the Capacity of Deletion Channels,” with E. Drinea. In *Proceedings of the 2004 IEEE International Symposium on Information Theory*.
- “Geometric Generalizations of the Power of Two Choices,” with J. Byers and J. Conside. In *Proceedings of the 16th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 54-63, 2004.
- “Interactive Data Summarization: An Example Application,” with N. Lesh. In *Proceedings of the 2004 Workshop on Advanced Visual Interfaces*, pp. 183-187.
- “On Power Law Distributions for Monkeys Typing Randomly: The Case of Unequal Probabilities,” with B. Conrad. *IEEE Transactions on Information Theory*, 50:7, pp. 1403-1414, 2004.
- “Dynamic Models for File Sizes and Double Pareto Distributions.” *Internet Mathematics*, 1(3):305-334, 2003.
- “A Scaling Result for Explosive Processes,” with R. Oliveira and J. Spencer. *Electronic Journal of Combinatorics*, 11(1), R31, 2004.
- “An Exhaustive Approach to 2D Strip Packing,” with N. Lesh, A. McMahon, and J. Marks. *Information Processing Letters*, 90(1), pp. 7-14, 2004.
- “New Models and Methods for File Size Distributions,” with B. Tworetzky. In *Proceedings of the 41st Annual Allerton Conference on Communication, Control, and Computing*, pp. 603-612, 2003.
- “Capacity Approaching Signal Constellations for Channels with Memory,” with A. Kavčić, and X. Ma. *IEEE Transactions on Information Theory*, 49:7, pp. 1636-1652, 2003.
- “The Power Spectra of Good Codes for Partial Response Channels,” with A. Kavcic and X. Ma. In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 45.
- “Concatenated Codes for Deletion Channels,” with J. Chen, C. Ng, and N. Varnica. In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 218.
- “Verification Codes for Deletion Channels.” In *Proceedings of the 2003 IEEE International Symposium on Information Theory*, p. 217.

- “Estimating and Comparing Entropies Across Written Natural Languages Using PPM Compression,” with F. Behr, V. Fossum, and D. Xiao. In *Proceedings of the 2003 IEEE Data Compression Conference*, p. 416.
- “Simple Load Balancing for Distributed Hash Tables,” with J. Byers, and J. Considine. Appears in the 2nd International Workshop on Peer-to-Peer Systems (IPTPS), pp. 80-87, 2003.
- “A Complete and Effective Move Set for Simplified Protein Folding,” with N. Lesh, and S. Whitesides. In *Proceedings of the 7th Annual International Conference on Research in Computational Molecular Biology*, pp 188-195, 2003.
- “Network Applications of Bloom Filters: A Survey,” with A. Broder. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, pp. 636-646, 2002. Journal version: *Internet Mathematics*, 1(4):485-509, 2003.
- “Verification Codes,” with M. Luby. In *Proceedings of the 40th Annual Allerton Conference on Communication, Control, and Computing*, pp. 38-47, 2002. Journal version: *IEEE Transactions on Information Theory*: 51(1), pp. 120-127, 2005.
- “Balls and Bins with Memory,” with B. Prabhakar and D. Shah. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science*, pp. 799-808, 2002.
- “Informed Content Delivery Across Adaptive Overlay Networks,” with J. Byers, J. Considine, and S. Rost. In *Proceedings of ACM SIGCOMM '02*, available as *Computer Communications Review*, vol. 32:4, pp. 47-60, 2002. Journal version: *IEEE/ACM Transactions on Networking*, 12:5, pp. 767-780, 2004.
- “Optimal Plans for Aggregation,” with A. Broder. In *Proceedings of the 21st Annual ACM Symposium on Principles of Distributed Computing*, pp. 144-152, 2002.
- “Human-Guided Tabu Search,” with G. Klau, N. Lesh, and J. Marks. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pp. 41-47, 2002.
- “The HuGS Platform: A Toolkit for Interactive Optimization,” with G. Klau, N. Lesh, J. Marks, and G. Schafer. In *Proceedings of the Workshop on Advanced Visual Interfaces*, pp. 324-330, 2002.
- “Exact Sampling of TCP Window States,” with A. Goel. In *Proceedings of IEEE INFOCOM*, pp. 259-265, 2002.
- “Balls and Bins Models with Feedback,” with E. Drinea and A. Frieze. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 308-315, 2002.
- “A Brief History of Lognormal and Power Law Distributions.” In *Proceedings of the 39th Annual Allerton Conference on Communication, Control, and Computing*, pp. 182-191, 2001. Journal version: *Internet Mathematics*, 1(2):226-251, 2003.
- “Capacity Approaching Signal Constellations for Channels with Memory,” with A. Kavčić, X. Ma, and N. Varnica. In *Proceedings of the 39th Annual Allerton Conference on Communication, Control, and Computing*, pp. 311-320, 2001.
- “Compressed Bloom Filters.” In *Proceedings of the 20th Annual ACM Symposium on Principles of Distributed Computing*, pp. 144-150, 2001. Journal version: *IEEE/ACM Transactions on Networks*, 10:5, pp. 613-620, October 2002.

“Towards More Completed Models of TCP Throughput,” with R. Rajaraman. *Journal of Supercomputing*, 20(2), pp. 137-160, September 2001.

“Deriving Performance Bounds for ISI Channels Using Gallager Codes,” with A. Kavčić, B. Marcus, and B. Wilson. In *Proceedings of the 2001 International Symposium on Information Theory*, p. 345, 2001.

“Toward Compressing Web Graphs,” with M. Adler. In *Proceedings of the 2001 Data Compression Conference*, pp. 203-212, 2001.

“On the Hardness of Finding Multiple Preset Dictionaries.” In *Proceedings of the 2001 Data Compression Conference*, pp. 411-418, 2001. Journal version: *IEEE Transactions on Information Theory*, 50:7, pp. 1536-1539, 2004.

“Fine-Grained Layered Multicast,” with J. Byers and M. Luby. In *Proceedings of IEEE INFOCOM 2001*, pp. 1143-1151, 2001. Journal version: “Fine-Grained Layered Multicast with STAIR,” with J. Byers, Gu-In Kwon, and M. Luby. To appear in *IEEE/ACM Transactions on Networking*.

“Using Multiple Hash Functions to Improve IP Lookups,” with A. Broder. In *Proceedings of IEEE INFOCOM 2001*, pp. 1454-1463, 2001.

“Improved Results for Route Planning in Stochastic Transportation Networks,” with J. Boyan. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 895-902, 2001.

“Estimating Resemblance of MIDI Documents,” with S. Owen. In *Proceedings of the 3rd Workshop on Algorithm Engineering and Experiments*, available as *Lecture Notes in Computer Science 2153*, pp. 78-90, 2001.

“Linear Waste of Best Fit Bin Packing on Skewed Distributions,” with Claire Kenyon. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pp. 582-589, 2000. Journal version: invited to special issue of *Random Structures and Algorithms*, 20(3), pp. 441-464, 2002.

“FLID/DL: Congestion Control for Layered Multicast,” with J. Byers, M. Frumin, G. Horn, M. Luby, A. Roetter, and W. Shaver. In *Proceedings of the 2nd International Workshop on Networked Group Communication (NGC 2000)*, pp. 71-81, 2000. Journal version: with J. Byers, G. Horn, M. Luby, and W. Shaver. *IEEE Journal on Selected Areas in Communications*, 20:8, pp. 1558-1570, October 2002.

“On Near-Uniform URL Sampling,” with M. Henzinger, A. Heydon, and M. Najork. In *Proceedings of the Ninth International World Wide Web Conference*, pp. 295-308, 2000.

“An Extension of Path Coupling and Its Application to the Glauber Dynamics for Path Coupling,” with M. Dyer, L. Goldberg, C. Greenhill, and M. Jerrum. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 616-623, 2000. Journal version: *SIAM Journal on Computing*, 30(6), pp. 1962-1975, 2001.

“Improved Classification via Connectivity Information,” with A. Broder and R. Krauthgamer. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 576-585, 2000.

“The Asymptotics of Selecting the Shortest of Two, Improved,” with B. Vöcking. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, pp. 326-327, 1999. Extended version appears as a chapter in *Analytic Methods in Applied Probability: In Memory of Fridrikh Karpelevich*, American Mathematical Society, pp. 165-176, 2002.

“Completeness and Robustness Properties of Min-Wise Independent Permutations,” with A. Broder. In *Proceedings of Random '99*, available as *Lecture Notes in Computer Science 1671*, pp. 1-10, 1999. Journal version: *Random Structures and Algorithms*, 18(1), pp. 18-30, January 2001.

“Measuring Index Quality Using Random Walks on the Web,” with M. Henzinger, A. Heydon, and M. Najork. In *Proceedings of the Eighth International World Wide Web Conference*, pp. 213-225, Elsevier Science, 1999.

“Analysis of Timing Based Mutual Exclusion with Random Times,” with E. Gafni. In *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing*, pp. 13-21, 1999. Journal version: *SIAM Journal on Computing*, 31:3, pp. 816-837, 2001.

“Accessing Multiple Mirror Sites in Parallel: Using Tornado Codes to Speed Up Downloads,” with J. Byers and M. Luby. In *Proceedings of the 18th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '99)*, pp. 275-284, 1999.

“Unscrambling Address Lines (Abstract),” with A. Broder and L. Moll. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 870-871, 1999.

“A Derandomization Using Min-Wise Independent Permutations,” with A. Broder and M. Charikar. In *Proceedings of Random '98*, available as *Lecture Notes in Computer Science 1518*, pp. 15-24, October 1998. Journal version: *Journal of Discrete Algorithms*, 1:1, pp. 11-20, 2003.

“On Balls and Bins with Deletions,” with R. Cole, A. Frieze, B. M. Maggs, A. W. Richa, R. Sitaraman, and E. Upfal. In *Proceedings of Random '98*, available as *Lecture Notes in Computer Science 1518*, pp. 145-158, October 1998.

“Studying Balanced Allocations with Differential Equations.” *Combinatorics, Probability, and Computing*, vol. 8, pp. 473-482, 1999.

“Delayed Information and Action in On-Line Algorithms,” with S. Albers and M. Charikar. In *Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science (FOCS-98)*, pp. 71-80, November 1998. Journal version: *Information and Computation*, 170, pp. 135-152, 2001.

“A Digital Fountain Approach to the Reliable Distribution of Bulk Data,” with John Byers, Michael Luby, and Ashu Rege. In *Proceedings of ACM SIGCOMM '98*, available as *Computer Communications Review*, vol. 28:4, pp. 56-67, 1998. Journal version: *IEEE Journal on Selected Areas in Communication*, 20:8, pp. 1528-1540, October 2002.

“Analyses of Load Stealing Models Using Differential Equations.” In *Proceedings of the 10th Annual ACM Symposium on Parallel Algorithms and Architectures*, pp. 212-221, 1998. Journal version: *Theory of Computing Systems*, 34:1, pp. 77-98, January 2001.

“Load Adaptive Routing on Butterfly Networks,” with R. Cole, B. Maggs, F. Meyer auf der Heide, A. Richa, K. Schröder, R. Sitaraman, and B. Vöcking. In *Proceedings of the*

30th Annual ACM Symposium on Theory of Computing (STOC-98), pp. 378-388, ACM Press, May 23-26, 1998.

“Min-Wise Independent Permutations,” with Andrei Broder, Moses Charikar, and Alan Frieze. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pp. 327-336, ACM Press, May 23-26, 1998. Journal version: *Journal of Computer and System Sciences*, 60(3), pp. 630-659, 2000.

“Analysis of Low Density Codes and Improved Designs Using Irregular Graphs,” with Michael Luby, Amin Shokrollahi, and Dan Spielman. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC-98)*, pp. 249-258, ACM Press, May 23-26, 1998.

“Improved Low Density Parity Check Codes Using Irregular Graphs and Belief Propagation,” with Michael Luby, Amin Shokrollahi, and Dan Spielman. In *Proceedings of the 1998 International Symposium on Information Theory*. Extended version available as Compaq Systems Research Center Technical Note 1998-009. Journal version: “Improved Low Density Parity Check Codes Using Irregular Graphs,” *IEEE Transactions on Information Theory*, 47(2), pp. 585-598, 2001.

“Analysis of Random Processes via And-Or Tree Evaluation,” with Michael Luby and Amin Shokrollahi. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 364-373, 1998.

“Average Case Analyses of First Fit and Random Fit Bin Packing,” with Susanne Albers. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 290-299, 1998. Journal version: *Random Structures and Algorithms*, 16(3), pp. 240-259, 2000.

“How Useful is Old Information?” In *Proceedings of the 16th ACM Symposium on Principles of Distributed Computing*, pp. 83-91, 1997. Journal version: *IEEE Transactions on Parallel and Distributed Systems*, 11:1, pp. 6-20, January 2000.

“Revisiting the COUNTER Algorithms for List Update,” with Susanne Albers. *Inf. Proc. Letters*, 64, pp. 155-160, 1997.

“Practical Loss-Resilient Codes,” with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. In *Proceedings of the 29th ACM Symposium on the Theory of Computing*, pp. 150-159, 1997. Journal version: “Efficient Erasure Correcting Codes,” with Michael Luby, Amin Shokrollahi, and Dan Spielman. *IEEE Transactions on Information Theory*, 47(2), pp 569-584, 2001.

“On the Analysis of Randomized Load Balancing Schemes.” In *Proceedings of the 9th ACM Symposium on Parallel Algorithms and Architectures*, pp. 292-301, 1997. Journal version: special conference issue of *Theory of Computing Systems*, vol. 32, pp. 361-386, 1999.

“Load Balancing and Density Dependent Jump Markov Processes.” In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pp. 213-222, 1996. Journal version: “The Power of Two Choices in Randomized Load Balancing.” *IEEE Transactions on Parallel and Distributed Computing*, 12(10), pp. 1094-1104, 2001.

“Designing Stimulating Programming Assignments for an Algorithms Course: A Collection of Problems Based on Random Graphs.” In *SIGCSE Bulletin*, 28:3, pp. 29-36, September 1996.

“Constant Time per Edge is Optimal in Rooted Tree Networks.” In *Proceedings of the 8th ACM Symposium on Parallel Algorithms and Architectures*, pp. 162-169, 1996. Journal version: *Distributed Computing* vol. 10:4, 1997, pp. 189-197.

“Average Case Analyses of List Update Algorithms, with Applications to Data Compression,” with Susanne Albers. In *Proceedings of the 1996 International Colloquium on Automata, Languages, and Programming*, pp. 514-525. Journal version: *Algorithmica*, vol. 21:3, 1998, pp. 312-329.

“Pattern-based Compression of Text Images,” with Andrei Broder. In *1996 Data Compression Conference*, pp. 300-309.

“Parallel Randomized Load Balancing,” with Micah Adler, Soumen Chakrabarti, and Lars Rasmussen. In *Proceedings of the 27th ACM Symposium on the Theory of Computing*, pp. 238-247, 1995. Journal version: *Random Structures and Algorithms*, vol. 13:2, September 1998, pp. 159-188.

“Bounds on the Greedy Routing Algorithm in Array Networks.” In *Proceedings of the 6th ACM Symposium on Parallel Algorithms and Architectures*, pp. 346-353, 1994. Journal version: special conference issue of the *Journal of Computer and System Sciences* vol. 53:3, December 1996, pp. 317-327.

“Computational Complexity of Loss Networks,” with Graham Louth and Frank Kelly. *Theoretical Computer Science*, vol. 125, May 1994, pp. 45-59.

Other Publications

“Lyra: An Efficient and Expressive Subquadratic Architecture for Modeling Biological Sequences,” with Krithik Ramesh, Sameed Siddiqui, Albert Gu, and Pardis Sabeti. Available at <https://arxiv.org/abs/2503.16351>.

“Juniper: Reconstructing Transmission Events from Next-Generation Sequencing Data at Scale,” with Ivan Specht, Gage Moreno, Taylor Brock-Fisher, Lydia Kraslinkova, Brittany Petros, Jonathan Pekar, Mark Schifferli, Ben Fry, Catherine Brown, Lawrence Madoff, Megan Burns, Stephen Schaffner, Daniel Park, Bronwn MacInnis, Al Ozonoff, Patrick Varilly, and Pardis Sabeti. Available at <https://www.medrxiv.org/content/10.1101/2025.03.02.25323192v1>.

“Delphy: scalable, near-real-time Bayesian phylogenetics for outbreaks,” with Patrik Varilly, Mark Schifferli, Katherine Yang, Tim Burcham, Paul Cronan, Olivia Glennon, Olivia Jacks, Ellory Laning, Libby Marrs, Kyle Oba, Shannon Yeung, Edyth Parker, Ifeanyi Omah, Jonathan E. Pekar, Laura Luebbert, Kristian G. Andersen, Daniel J. Park, Stephen F. Schaffner, Bronwyn L. MacInnis, Christian Happi, Jacob E. Lemieux, Al Ozonoff, Michael D. Mitzenmacher, Ben Fry, and Pardis C. Sabeti. Available at <https://www.biorxiv.org/content/10.1101/2025.03.25.645253v1>.

“Use of popularity information to reduce risk posed by guessing attacks,” with Stuart Schechter and Cormac Herley. US Patent 08959644. Issued 2/17/2015.

“Method and apparatus providing network coding based flow control,” with Jay Kumar Sundararajan, Devavrat Shah, Muriel Medard, Joao Barros, and Szymon Jakubczak. US Patent 08526451. Issued 9/3/2013.

“Method and apparatus providing network coding based flow control,” with Jay Kumar Sundararajan, Devavrat Shah, Muriel Medard, Joao Barros, and Szymon Jakubczak. US Patent 08130776. Issued 3/6/2012.

“Method for finding optimal paths using a stochastic network model,” with Evdokia Nikolova and Matthew Brand. US Patent 07573866. Issued 8/11/2009.

“Ranking search engine results,” with Monika Henzinger. US Patent 07541388. Issued 11/11/2008.

“Distributed, compressed Bloom filter Web cache server.” US Patent 06920477. Issued 7/19/2005.

“Method for packing rectangular strips,” with Neal Lesh and Joe Marks. US Patent 06832129. Issued 12/14/2004.

“System and method for near-uniform sampling of web page addresses,” with Monika Henzinger, Allan Heydon, and Marc Najork. US Patent 06594694. Issued 7/15/2003.

“On demand encoding with a window,” with Michael Luby, Gavin Horn, Jeffrey Persch, John Byers, and Armin Haken. US Patent 06486803. Issued 11/26/2002.

“Generating high weight encoding symbols using a basis,” with Armin Haken, Michael Luby, Gavin Horn, Diane Hernek, and John Byers. US Patent 06411223. Issued 6/25/2002.

“Space-efficient multi-cycle barrel shifter circuit,” with Laurent Moll. US Patent 06314156. Issued 11/6/2001.

“Method for determining a random permutation of variables by applying a test function,” with Andrei Broder, Laurent Moll, and Mark Shand. US Patent 06292762. Issued 9/18/2001.

“Loss resilient code with double heavy tailed series of redundant layers,” with Michael Luby. US Patent 06195777. Issued 2/27/2001.

“Message encoding with irregular graphing,” with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06163870. Issued 12/19/2000.

“Compression of grey scale images of text,” with Andrei Broder. US Patent 06088309. Issued 7/11/2000.

“Irregularly graphed encoding technique,” with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06081909. Issued 6/27/2000.

“Loss resilient decoding technique,” with Michael Luby, Amin Shokrollahi, Dan Spielman, and Volker Stemann. US Patent 06073250. Issued 6/6/2000.

“Task processing optimization in a multiprocessor system,” with Andrei Broder. US Patent 05991808. Issued 11/23/1999.

“Compression protocol with multiple preset dictionaries,” with Andrei Broder and Jeff Mogul. US Patent 05953503. Issued 9/14/1999.

“When Simple Hash Functions Suffice,” with K. Chung and S. Vadhan. Appears as a chapter in *Beyond Worst Case Analysis*, edited by Tim Roughgarden, Cambridge University Press.

“Algorithms with Predictions,” with S. Vassilvitskii. Appears as a chapter in *Beyond Worst Case Analysis*, edited by Tim Roughgarden, Cambridge University Press.

“Some Practical Randomized Algorithms and Data Structures,” with E. Upfal. Appears as a chapter in *Computing Handbook, Third Edition: Computer Science and Software Engineering*, edited by T. Gonzalez, J. Diaz-Herrera, and A. Tucker. Chapman and Hall, 2014.

“Building an Efficient Hash Table on the GPU,” with D. Alcantara, V. Volkov, S. Sengupta, J. Owens, and N. Amenta. Appears as a chapter in *GPU Computing Gems: Jade Edition*, edited by Wen-mei W. Hwu. Morgan Kaufmann, 2011.

“Codes – Protecting Data Against Loss and Errors.” Appears as a chapter in *Algorithms Unplugged*, edited by B. Vöcking, M. Dietzfelbinger, R. Reischuk, and H. Vollmer. Springer, 2010.

“Hash-Based Techniques for High-Speed Packet Processing”, with A. Kirsch and G. Varghese. Appears as a chapter in *Algorithms for Next Generation Networks*, edited by G. Cormode and M. Thottan. Springer, 2010.

“Bloom Filters.” *Encyclopedia of Database Systems*, 2009.

“Unbiasing Random Bits.” *Dr. Dobbs Journal*, No. 335, pp. 101-104, April 2002.

“Good Hash Tables & Multiple Hash Functions.” *Dr. Dobbs Journal*, No. 336, May 2002.

“Challenging Students with Creative Assignments.” *SIGACT News*, 32(1):70-73, 2001.

“An Experimental Assignment on Random Processes.” *SIGACT News*, 32(1):74-78, 2001.

“The Power of Two Random Choices: A Survey of Techniques and Results,” with A. Richa and R. Sitaraman. Appears as a chapter in *Handbook of Randomized Computing: Volume 1*, edited by P. Pardalos, S. Rajasekaran, and J. Rolim, pp. 255-312.

“Tornado Codes for Forward Error Correction,” with Michael Luby, in *Digital Forefront*.

“Bounds and Improvements for BiBa Signature Schemes,” with A. Perrig. Harvard Computer Science Group Technical Report TR-02-02.

“A Note on Low Density Parity Check Codes for Erasures and Errors.” Digital SRC Technical Note 1998-017.

“Tight Thresholds for the Pure Literal Rule.” Digital SRC Technical Note 1997-011.

“The Power of Two Choices in Randomized Load Balancing.” Ph.D. thesis, U.C. Berkeley, August 1996.

“*Calculus: A Solution Manual*,” with others, Harcourt Brace Jovanovich, 1993.

“*Elliptic Curves in Computer Science: Primality Testing, Factoring, and Cryptography*,” senior honors thesis, Harvard University, 1991. Published as a Hoopes Prize winner.

Conference Presentations “Queueing, Predictions, and LLMs: Challenges and Open Problems.” Tutorial, SIGMETRICS. June 2025.

“SkipPredict: When to Invest in Predictions for Scheduling. Presented at the TTIC Workshop on Learning-Augmented Algorithms, August 2024.

“Algorithms with Predictions.” Invited talk presented at the INFORMS 2023 Annual Meeting, October 2023.

“Algorithms with Predictions.” Presented at the 2nd Vienna Workshop on Computational Optimization. March 2023.

“Algorithms with Predictions.” Presented at the Sublinear Algorithms Workshop, FODSI/MIT. August 2022.

“Algorithms with Predictions: State of the Area.” Presented at the Algorithms with Predictions Workshop at the 54th Annual ACM Symposium on Theory of Computing. June 2022. Also panel lead.

“Efficient Randomized Network Telemetry Algorithms.” Presented at the Data plane algorithms in programmable networks tutorial at ACM SIGMETRICS 2022.

“Scheduling with Predictions and the Price of Misprediction.” Presented at the 11th Annual Innovations in Theoretical Computer Science Conference. January 2020.

“Learning-Augmented Algorithms: How ML Can Lead to Provably Better Algorithms.” ALGO Keynote Talk, September 2019.

“Robust Set Reconciliation via Locality Sensitive Hashing.” Presented at the 38th Symposium on Principles of Database Systems, July 2019.

“Arithmetic Progression Hypergraphs: Examining the Second Moment Method.” Presented at the 2019 ANALCO Workshop, January 2019.

“A Model for Learned Bloom Filters and Optimizing by Sandwiching.” Presented at the NeurIPS 2018 Poster Session.

“Adaptive Cuckoo Filters.” Presented at the 2018 ALENEX Conference, January 2018.

“Analyzing Distributed Join-Idle-Queue: A Fluid Limit Approach.” Presented at the 54th Annual Allerton Conference on Communication, Control, and Computing, October 2016.

“Voronoi Choice Games.” Presented at the 2016 International Colloquium on Automata, Languages, and Programming, July 2016.

“A New Approach to Analyzing Robin Hood Hashing.” Presented at the 2016 ANALCO Workshop, January 2016.

“More Analysis of Double Hashing for Balanced Allocations.” Presented at the 2016 ANALCO Workshop, January 2016.

“Data exploration: equitability and the maximal information coefficient.” Presented at the Workshop for Data-driven Algorithmics, September 2015.

“Multi-Party Set Reconciliation Using Characteristic Polynomials.” Presented at the 2014 Allerton Conference, October 2014.

“Parallel Peeling Algorithms.” Presented at the 2014 Symposium on Parallelism in Algorithms and Architectures, June 2014.

“Balanced Allocations and Double Hashing.” Presented at the 2014 Symposium on Parallelism in Algorithms and Architectures, June 2014.

“Peeling Arguments and Double Hashing.” Presented at the 2012 Allerton Conference, October 2012.

“The Complexity of Object Reconciliation and Open Problems Related to Set Difference and Coding.” Presented at the 2012 Allerton conference, October 2012.

“The Ecology of Communication, Control, and Computing Research and Education (Panel).” Presented at the 2012 Allerton conference, October 2012.

“Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes.” Presented at the 2012 Workshop on Algorithms for Modern Massive Data Sets (MMDS), July 2012.

“Invertible Bloom Lookup Tables.” Presented at the 2011 Allerton conference, September 2011.

“Brief Announcement: Large-Scale Multimaps.” Presented at the 2011 Symposium on Parallel Algorithms and Architectures, June 2011.

“On the Zero-Error Capacity Threshold for Deletion Channels.” Presented at the 2011 ITA Workshop, February 2011.

“Some Recent Uses of Hashing for Network Applications.” Presented at the 2010 Workshop on Research and Analysis of Tail Phenomena, August 2010.

“Some Recent Uses of Hashing for Network Applications.” Presented at the 2010 Workshop on Random Graphs and Algorithms, June 2010.

“Tight Asymptotic Bounds for the Deletion Channel with Small Deletion Probabilities.” Presented at the 2010 International Symposium on Information Theory, June 2010.

“Some Open Questions on Cuckoo Hashing.” Presented at the 2009 European Symposium on Algorithms, August 2009.

“On the Performance of Multiple Choice Hash Tables with Moves on Deletions and Inserts.” Presented at the 2008 Allerton conference, September 2008.

“Designing Floating Codes for Expected Performance.” Presented at the 2008 Allerton conference, September 2008.

“A Survey of Results for Deletion Channels and Related Synchronization Channels.” Presented at the 11th Scandinavian Workshop on Algorithm Theory, July 2008.

“Why Simple Hash Functions Work: Exploiting the Entropy in a Data Stream.” Presented at the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2008.

“Using a Queue to De-amortize Cuckoo Hashing in Hardware.” Presented at the 2007 Allerton conference, September 2007.

“Capacity Upper Bounds for Deletion Channels.” Presented at the 2007 Int’l Symposium on Information Theory (ISIT), June 2007.

“Codes for Deletion and Insertion Channels with Segmented Errors.” Presented at the 2007 Int’l Symposium on Information Theory (ISIT), June 2007.

“Digital Fountains: Applications and Related Issues.” Presented at the INFORMS Annual Meeting, November 2006.

“New Directions for Power Law Research.” Presented at the Radcliffe Exploratory Seminar on Dynamic Networks, October 2006.

“Bloom Filters via d -left Hashing and Dynamic Bit Reassignment.” Presented at the 2006 Allerton conference, September 2006.

“Beyond Bloom Filters: From Approximate Membership Checks to Approximate State Machines.” Presented at the 2006 SIGCOMM conference, September 2006.

“An Improved Construction for Counting Bloom Filters.” Presented at the European Symposium on Algorithms, September 2006.

“Building a Better Bloom Filter.” Presented at the European Symposium on Algorithms, September 2006.

“On the Theory and Practice of Data Recovery with Multiple Versions.” Presented at the 2006 International Symposium on Information Theory, July 2006.

“New Results for an Old Data Structure: Bloom Filters and Beyond.” Presented at the 2006 Stochastic Networks conference, June 2006.

“New Results and Open Problems for Deletion Channels.” Presented at ANALCO 2006, January 2006 (plenary talk).

“Digital Fountains and their Application to Informed Content Delivery over Adaptive Overlay Networks.” Presented at 19th International Symposium on Distributed Computing, September 2005 (plenary talk).

“Simple Summaries for Multiple Choice Hashing.” Presented at the 2005 Allerton Conference.

“New Directions in Power Law Research.” MSRI Workshop on Models of Real-World Random Networks, April 2005.

“Multi-dimensional Balanced Allocations.” Presented at the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2005.

“Digital Fountains: A Survey and Look Forward.” Presented at the 2004 Information Theory Workshop. October 2004.

“Improved Lower Bounds on the Capacity of I.I.D. Deletion Channels.” Presented at the 2004 Allerton Conference.

“Geometric Generalizations of the Power of Two Choices.” Presented at the 16th ACM Symposium on Parallel Algorithms and Architectures, June 2004.

“New Models and Methods for File Size Distributions.” Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2003.

“Verification Codes for Deletion Channels.” Presented at the International Symposium on Information Theory, July 2003.

“New Exhaustive, Heuristic, and Interactive Algorithms for 2-Dimensional Bin-Packing.” Presented at the 2nd Annual Workshop on Random Graphs and Algorithms, June 2003.

“Network Applications of Bloom Filters: A Survey.” Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2002.

“Verification Codes.” Presented at the 40th Annual Allerton Conference on Communication, Control, and Computing, October 2002.

“New Directions in Balls and Bins Problems.” Invited talk for the 6th International Workshop on Randomization and Approximation Methods in Computer Science. September 2002.

“Optimal Plans for Aggregation.” Presented at the 21st Annual ACM Symposium on Principles of Distributed Computing, July 2002.

“How to Get Across Town by Bus.” Presented at the INFORMS Annual Meeting, November 2001.

“A Brief History of Lognormal and Power Law Distributions.” Presented at the 39th Annual Allerton Conference on Communication, Control, and Computing, October 2001.

“Compressed Bloom Filters.” Presented at the 20th Annual ACM Symposium on Principles of Distributed Computing, August 2001.

“How Useful Is Old Information?” Presented at the INFORMS Applied Probability Conference, July 2001.

“Using Multiple Hash Functions to Improve IP Lookups.” Presented at the INFOCOM, April 2001.

“On the Hardness of Finding Multiple Preset Dictionaries.” Presented at the 2001 Data Compression Conference, March 2001.

“Improved Results for Route Planning in Stochastic Transportation Networks.” Presented at the 12th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2001.

“Estimating Resemblance of MIDI Documents.” Presented at the 3rd Workshop on Algorithm Engineering and Experiments, January 2001.

“Linear Waste of Best Fit Bin Packing on Skewed Distributions.” Presented at the 41st Annual IEEE Symposium on Foundations of Computer Science, November 2000.

“Improved Classification via Connectivity Information.” Presented at the 11th Annual ACM Symposium on Discrete Algorithms, January 2000.

“The Asymptotics of Selecting the Shortest of Two, Improved.” Presented at the 37th Annual Allerton Conference on Communication, Control, and Computing, August 1999.

“Analysis of Timing Based Mutual Exclusion with Random Times.” Presented at the 18th Annual ACM Symposium on Principles of Distributed Computing, May 1999.

“Accessing Multiple Mirror Sites in Parallel: Using Tornado Codes to Speed Up Downloads.” Presented at INFOCOM '99, March 1999.

“Delayed Information and Action in On-Line Algorithms.” Presented at the 39th Annual IEEE Symposium on Foundations of Computer Science, November 1998.

“A Derandomization Using Min-Wise Independent Permutations.” Presented at Random ’98, October 1998.

“On Balls and Bins with Deletions.” Presented at Random ’98, October 1998.

“The Power of Two Choices and Other Examples of Using Differential Equations to Analyze Algorithms.” Invited presentation at the Warwick Randomized Algorithms and Stochastic Simulation Tutorial and Workshop. July 1988.

“Analysis of Low Density Codes and Improved Designs Using Irregular Graphs.” Presented at the 30th Annual ACM Symposium on Theory of Computing, May 1998.

“Average Case Analyses of First Fit and Random Fit Bin Packing.” Presented at the 9th Annual ACM Symposium on Discrete Algorithms, January 1998.

“How Useful Is Old Information?” Presented at the 16th Annual ACM Symposium on Principle of Distributed Computing, August 1997.

“On the Analysis of Randomized Load Balancing Schemes.” Presented at the 9th ACM Symposium on Parallel Algorithms and Architectures, Newport, Rhode Island, June 1997.

“Load Balancing and Density Dependent Jump Markov Processes.” Presented at the 37th Annual Symposium on Foundations of Computer Science, Burlington, Vermont, October 1996.

“Constant Time per Edge is Optimal in Rooted Tree Networks.” Presented at the 8th ACM Symposium on Parallel Algorithms and Architectures, Padua, Italy, June 1996.

“GRID: Pattern-based Compression of Text Images.” Presented at the 1996 Data Compression Conference, Snowbird, Utah, April 1996.

“Parallel Randomized Load Balancing.” Presented at the 27th ACM Symposium on the Theory of Computing, Las Vegas, Nevada, May 1995. Extended version presented at the Bay Area Theory Seminar, July 1995.

“Bounds on the Greedy Routing Algorithm on Array Networks.” Presented at the 6th ACM Symposium on Parallel Algorithms and Architectures, Cape May, New Jersey, June 1994.

**Other
Presentations**

“Algorithms with Predictions.” Georgia Tech. 2023.

“Modern Algorithms Lecture Series.” Lipari Summer School. 2022.

“Algorithms with Predictions.” Carnegie Mellon, UC Santa Barbara, Technical University of Munich. 2022.

“Locality Sensitive Hashing.” Broad Institute (Models, Inference, and Algorithms).

“Algorithms with Predictions.” Open Data Science Conference East 2020, Open Data Science Conference Europe 2020, Foundations of Data Science Institute Kickoff Workshop 2020.

“Learning-Augmented Algorithms: How ML Can Lead to Provably Better Algorithms.” Rice University. 2019.

“Metric Sublinear Algorithms via Linear Sampling.” Sublinear Algorithms Workshop (MIT), Baidu Research. 2018.

“Bloom Filters, Cuckoo Hashing, Cuckoo Filters, and Adaptive Cuckoo Filters.” Northeastern University, Univ. of Illinois at Urbana-Champaign, Georgetown University, University of California Irvine, Baidu Research, Georgia Tech. 2018.

“Hashing Lecture Series.” Swedish Summer School in Computer Science. 2016.

“More Analysis of Double Hashing for Balanced Allocations.” U.C. Irvine. 2016.

“Algorithm Engineering with Hashing,” U.C. Irvine, AMD. 2016.

“Bloom Filters and Such,” EADS Summer School on Hashing: Theory and Applications, July 2014.

“Cuckoo Hashing, Balanced Allocations, and Such,” EADS Summer School on Hashing: Theory and Applications, July 2014.

“Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes,” ICERM Workshop on Stochastic Graph Models, March 2014.

“Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes,” University of Copenhagen. 2013.

“Some Hash-Based Data Structures and Algorithms Everyone Should Know.” UC Davis, 2013.

“Peeling Algorithms,” Simons Institute. 2013.

“Verifying Computations in the Cloud (and Elsewhere),” TTI/Vanguard : Ginormous Systems. Microsoft Faculty Summit. 2013.

“The Groupon Effect on Yelp Ratings: a Root Cause Analysis,” Cornell, 2012.

“Peeling Arguments: Invertible Bloom Lookup Tables and Biff Codes,” Google. 2012.

“Detecting Novel Associations in Large Data Sets.” eHarmony (LA Machine Learning). 2012.

“Cuckoo Hashing: New Results and Open Questions.” University of Madison-Wisconsin. 2011.

“Practical Verified Computation with Streaming Interactive Proofs.” University of Madison-Wisconsin. 2011.

“The Power of Choice: Simple Hash-Based Data Structures Everyone Should Know.” Santa Fe Institute. 2011.

“Information Asymmetries in Pay-Per-Bid Auctions: How Swoopo Makes Bank.” Liverpool University, Microsoft UK. 2010.

“Some Open Questions on Cuckoo Hashing.” Stanford, Microsoft Research, Yale, Dartmouth, Princeton, Cambridge University, University College London . 2009-2010.

“An Efficient, Rigorous Approach for Identifying Statistically Significant Frequent Itemsets.” Google. 2009.

“Some Recent Results on Cuckoo Hashing.” AT&T. 2009.

“Some Results on Coding for Flash Memory.” University College London. 2009.

“Network Coding Meets TCP.” University College London. 2009.

“A Survey of Results for Deletion Channels and Related Synchronization Channels.” Cambridge, Georgia Tech, Microsoft Research, UCLA, 2009-2010.

“Bloom Filters, Related Data Structures, and their Applications.” Cambridge (Kuwait Lecture). 2009.

“A History of and New Directions for Power Law Research.” SUNY Buffalo. 2009.

“Why Simple Hash Functions Work.” Cisco. 2008.

“Cuckoo Hashing and CAMs.” Cisco, Google. 2008.

“New Results in Hash-Based Data Structures.” Cisco. 2007.

“The Hiring Problem and Lake Wobegon Strategies.” Microsoft, U.C. Berkeley, Princeton, Google, Yahoo. 2007-2008.

“A Brief History of Lognormal and Power Law Distributions.” Harvard Statistics. 2006.

“New Results and Open Problems for Deletion Channels and Related Synchronization Problems.” CMU, Brown, Harvard. 2006.

“Some New Results on an Old Data Structure, Bloom Filters .” Stanford, Cisco, Microsoft. 2005.

“Improved Lower Bounds for Deletion Channels.” MIT, EPFL (Lausanne). 2005.

“Building a Better Bloom Filter.” Google, Tufts, EPFL (Lausanne), MIT. 2004-05.

“Geometric Generalizations of the Power of Two Choices.” Columbia, Brown, MIT, Georgia Tech. 2004-05.

“New Lower Bounds on the Capacity of the Binary Deletion Channel.” MIT, 2004.

“Digital Fountains, and their Use for Informed Content Delivery across Overlay Networks.” Cornell, 2003.

“Verification Codes for Error and Deletion Channels.” U.C. Berkeley, 2003.

“New Exhaustive, Heuristic, and Interactive Algorithms for 2D Strip Packing.” U.C. Berkeley, 2003.

“Dynamic Models for File Sizes and Double Pareto Distributions.” IBM Research, Microsoft Silicon Valley, Stanford, Tufts, Boston University, Random Graphs and Algorithms Workshop. 2002.

“Verification Codes.” DIMACS workshop. 2001.

“Getting Across Town by Bus, Aggregation, and Reliability.” MIT (LIDS), Mitsubishi Electronic Research Labs. 2001.

“New Ideas for Layered Multicast Congestion Control.” International Computer Science Institute. 2001.

“Compressed Bloom Filters, and Towards Compressing Web Graphs.” U. of Washington, AT&T Labs, U. C. Berkeley, MIT, Brandeis, Carnegie Mellon University, Stanford, DIMACS workshop on compression, U.C. San Diego. 2001.

“Toward Compressing Web Graphs.” Compaq Systems Research Center, Altavista Corporation. 2000.

“Improved Results for Route Planning in Stochastic Transportation Networks.” Compaq Systems Research Center. 2000.

“Choosing the Shortest of Two, Improved, with Applications to IP Routing.” U. Washington, Princeton, Yale, U.C. Berkeley, Boston University, Stanford, M.I.T., Northeastern, Compaq Systems Research Center. 2000.

“Tornado Codes, with Applications to Reliable Multicast.” Harvard, CMU, Duke, AT&T Labs, Univ. of Illinois at Urbana-Champaign, University of Rhode Island, Boston University. 1998-2000.

“How Useful Is Old Information?” MIT, U.C. Berkeley, Stanford, Microsoft. 1997.

“Practical Loss-Resilient Codes.” U.C. Berkeley, Harvard. 1997.

“An Average Case Analysis for First-Fit Bin Packing.” Stanford, DIMACS workshop. 1997.

“The Power of Two Choices.” U.C. Berkeley, Brown, Harvard, Boston University, IBM Almaden, Digital Systems Research Center. 1997.

“The Power of Two Choices and Other Examples of Using Differential Equations to Analyze Algorithms.” U. Mass.