UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

AMERICAN CIVIL LIBERTIES UNION; AMERICAN CIVIL LIBERTIES UNION FOUNDATION; NEW YORK CIVIL LIBERTIES UNION; and NEW YORK CIVIL LIBERTIES UNION FOUNDATION, Plaintiffs, v. JAMES R. CLAPPER, in his official capacity as Director of National Intelligence; KEITH B. ALEXANDER, in his official capacity as Director of the National Security Agency and Chief of the Central Security Service; CHARLES T. HAGEL, in his official capacity as Secretary of Defense; ERIC H. HOLDER, in his official capacity as Attorney General of the United States; and ROBERT S. MUELLER III, in his official capacity as Director of the Federal Bureau of Investigation, Defendants.	**DECLARATION OF PROFESSOR EDWARD W. FELTEN** Case No. 13-cv-03994 (WHP) **ECF CASE**

DECLARATION OF PROFESSOR EDWARD W. FELTEN

I, Edward W. Felten, declare under penalty of perjury, pursuant to 28 U.S.C. § 1746, that the following is true and correct:

1. The plaintiffs in this lawsuit have challenged what they term the "mass call-tracking" program of the National Security Agency, and they have asked me to explain the sensitive nature of metadata, particularly when obtained in the aggregate. Below, I discuss how advances in technology and the proliferation of metadata-producing devices, such as phones, have produced rich metadata trails. Many details of our lives can be gleaned by examining those trails, which often yield information more easily than do the actual content of our communications.

Superimposing our metadata trails onto the trails of everyone within our social group and those of everyone within our contacts' social groups, paints a picture that can be startlingly detailed.

2. I emphasize that I do not in this declaration pass judgment on the use of metadata analysis in the abstract. It can be an extraordinarily valuable tool. But because it can also be an unexpectedly revealing one—especially when turned to the communications of virtually everyone in the country—I write in the hope that courts will appreciate its power and control its use appropriately.

Biography

3. My name is Edward W. Felten. I am Professor of Computer Science and Public Affairs, as well as Director of the Center for Information Technology Policy, at Princeton University.

4. I received a Bachelor of Science degree in Physics from the California Institute of Technology in 1985, a Master's degree in Computer Science and Engineering from the University of Washington in 1991, and a Ph.D. in the same field from the University of Washington in 1993. I was appointed as an Assistant Professor of Computer Science at Princeton University in 1993, and was promoted to Associate Professor in 1999 and to full Professor in 2003. In 2006, I received an additional faculty appointment to Princeton's Woodrow Wilson School of Public and International Affairs.

5. I have served as a consultant or technology advisor in the field of computer science for numerous companies, including Bell Communications Research, International Creative Technologies, Finjan Software, Sun Microsystems, FullComm and Cigital. I have authored numerous books, book chapters, journal articles, symposium articles, and other publications relating to computer science. Among my peer-reviewed publications are papers on the inference

of personal behavior from large data sets[1] and everyday objects,[2] as well as work on the extraction of supposedly protected information from personal devices.[3]

6. I have testified several times before the United States Congress on computer technology issues.

7. In 2011 and 2012, I served as the first Chief Technologist at the U.S. Federal Trade Commission ("FTC"). In that capacity, I served as a senior policy advisor to the FTC Chairman, participated in numerous civil law enforcement investigations, many of which involved privacy issues, and acted as a liaison to the technology community and industry. My privacy-related work at the FTC included participating in the creation of the FTC's major privacy report issued in March 2012,[4] as well as advising agency leadership and staff on rulemaking, law enforcement, negotiation of consent orders, and preparation of testimony.

8. Among my professional honors are memberships in the National Academy of Engineering and the American Academy of Arts and Sciences. I am also a Fellow of the Association of Computing Machinery. A copy of my curriculum vitae is attached as Exhibit 1 to this declaration.

---

[1] Joseph A. Calandrino, Ann Kilzer, Arvind Narayanan, Edward W. Felten & Vitaly Shmatikov, *"You Might Also Like:" Privacy Risks of Collaborative Filtering*, Proceedings of IEEE Symposium on Security and Privacy (May 2011), http://bit.ly/kUNh4c.

[2] William Clarkson, Tim Weyrich, Adam Finkelstein, Nadia Heninger, J. Alex Halderman & Edward W. Felten, *Fingerprinting Blank Paper Using Commodity Scanners*, Proceedings of IEEE Symposium on Security and Privacy (May 2009), http://bit.ly/19AoMej.

[3] J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum & Edward W. Felten, *Lest We Remember: Cold Boot Attacks on Encryption Keys*, Proceedings of USENIX Security Symposium (August 2008), http://bit.ly/13Ux38w.

[4] Federal Trade Commission, *Protecting Consumer Privacy in an Era of Rapid Change: Recommendations for Businesses and Policymakers* (March 2012), http://1.usa.gov/HbhCzA.

The Mass Call Tracking Program

9. On June 5, 2013, *The Guardian* disclosed an order issued by the Foreign Intelligence Surveillance Court ("FISC") pursuant to Section 215 of the Patriot Act (the "Verizon Order").[5] This order compelled a Verizon subsidiary, Verizon Business Network Services ("Verizon"), to produce to the National Security Agency ("NSA") on "an ongoing daily basis . . . all *call detail records* or 'telephony metadata' created by Verizon for communications (i) between the United States and abroad; or (ii) wholly within the United States, including local telephone calls."[6] The Director of National Intelligence subsequently acknowledged the authenticity of the Verizon Order.[7]

10. Following the disclosure of the Verizon Order, government officials indicated that the NSA's acquisition of call detail records is not limited to customers or subscribers of Verizon. In particular, the NSA's collection of this data encompasses telephone calls carried by the country's three largest phone companies: Verizon, AT&T, and Sprint.[8] Because these companies provide at least one end of the vast majority of telecommunications connectivity in the country, these

[5] Secondary Order, *In re Application of the FBI for an Order Requiring the Production of Tangible Things from Verizon Bus. Network Servs., Inc. on Behalf of MCI Commc'n Servs., Inc. d/b/a Verizon Bus. Servs.*, No. BR 13-80 at 2 (FISA Ct. Apr. 25, 2013), *available at* http://bit.ly/11FY393.

[6] *Id.* at 2 (emphasis added).

[7] James R. Clapper, *DNI Statement on Recent Unauthorized Disclosures of Classified Information*, Office of the Director of National Intelligence (June 6, 2013), http://1.usa.gov/13jwuFc.

[8] *See* Siobhan Gorman et al., *U.S. Collects Vast Data Trove*, Wall St. J., June 7, 2013, http://on.wsj.com/11uD0ue ("The arrangement with Verizon, AT&T and Sprint, the country's three largest phone companies means, that every time the majority of Americans makes a call, NSA gets a record of the location, the number called, the time of the call and the length of the conversation, according to people familiar with the matter. . . . AT&T has 107.3 million wireless customers and 31.2 million landline customers. Verizon has 98.9 million wireless customers and 22.2 million landline customers while Sprint has 55 million customers in total.").

statements suggest that the NSA is maintaining a record of the metadata associated with nearly every telephone call originating or terminating in the United States.

11. Assuming that there are approximately 3 billion calls made every day in the United States, and also assuming conservatively that each call record takes approximately 50 bytes to store, the mass call tracking program generates approximately 140 gigabytes of data every day, or about 50 terabytes of data each year.

12. Assuming (again conservatively) that a page of text takes 2 kilobytes of storage, the program generates the equivalent of about 70 million pages of information every day, and about 25 billion pages of information every year.

13. Members of Congress have disclosed that this mass call tracking program has been in place for at least seven years, since 2006.[9]

14. On July 19, 2013, the day that the Verizon Order was set to expire, the Director of National Intelligence disclosed that the FISC had renewed the NSA's authority to collect telephony metadata in bulk.[10]

15. As noted above, the Verizon Order requires the production of "call detail records" or "telephony metadata." According to the order itself, that term encompasses, among other things, the originating and terminating telephone number and the time and duration of any call. Call detail records also typically include information about the location of the parties to the call. *See* 47 C.F.R. § 64.2003 (2012) (defining "call detail information" as "[a]ny information that

[9] *See* Dan Roberts & Spencer Ackerman, *Senator Feinstein: NSA Phone Call Data Collection in Place 'Since 2006,'* Guardian, June 6, 2013, http://bit.ly/13rfxdu; *id.* (Senator Saxby Chambliss: "This has been going on for seven years."); *see also* ST-09-0002 Working Draft – Office of the Inspector General, National Security Agency & Central Security Service (Mar. 24, 2009), http://bit.ly/14HDGuL.

[10] Press Release, Foreign Intelligence Surveillance Court Renews Authority to Collect Telephony Metadata, Office of the Director of National Intelligence (July 19, 2013), http://1.usa.gov/12ThYlT.

pertains to the transmission of specific telephone calls, including, for outbound calls, the number called, and the time, location, or duration of any call and, for inbound calls, the number from which the call was placed and the time, location, or duration of any call").

16. Although this latter definition of "call detail information" includes data identifying the location where calls are made or received, I will not address mobile phone location information in this declaration. While senior intelligence officials have insisted that they have the legal authority under Section 215 to collect mobile phone location information, they have stated that the NSA is not collecting phone location information "under this program."[11]

17. The information sought from Verizon also includes "session identifying information"—*e.g.*, originating and terminating telephone number, International Mobile Subscriber Identity (IMSI) number, International Mobile station Equipment Identity (IMEI) number, etc. These are unique numbers that identify the user or device that is making or receiving a call. Although users who want to evade surveillance can make it difficult to connect these numbers to their individual identities, for the vast majority of ordinary users these numbers can be connected to the specific identity of the user and/or device.

18. The information sought from Verizon also includes the "trunk identifier" of telephone calls. This provides information about how a call was routed through the phone network, which naturally reveals information about the location of the parties. For example, even if the government never obtains cell site location information about a call,[12] trunk identifier

[11] *See* Siobhan Gorman & Julian E. Barnes, *Officials: NSA Doesn't Collect Cellphone-Location Records*, Wall St. J., June 16, 2013, http://on.wsj.com/13MnSsp; Pema Levy, *NSA FISA Metadata Surveillance: Is The Government Using Cell Phones To Gather Location Data?*, Int'l Bus. Times, Aug. 2, 2013, http://bit.ly/18WKXOV.

[12] Cell site location information ("CSLI") reflects the cell tower and antenna sector a phone is connected to when communicating with a wireless carrier's network. Most carriers log and retain CSLI for the start and end of each call made or received by a phone, and some carriers log CSLI

information revealing that a domestic call was carried by a cable from Hawaii to the mainland United States will reveal that the caller was in the state of Hawaii at the time the call was placed.

19. In the present case, government officials have stated that the NSA retains telephony metadata gathered under the Verizon Order, and others similar to it, for five years.[13] Although officials have insisted that the orders issued under the telephony metadata program do not compel the production of customers' names, it would be trivial for the government to correlate many telephone numbers with subscriber names using publicly available sources. The government also has available to it a number of legal tools to compel service providers to produce their customer's information, including their names.[14]

Metadata Is Easy to Analyze

20. Telephony metadata is easy to aggregate and analyze. Telephony metadata is, by its nature, *structured data*. Telephone numbers are standardized, and are expressed in a predictable format: In the United States, a three digit area code, followed by a three digit central office exchange code, and then a four digit subscriber number. Likewise, the time and date information

---

for text messages and data connections as well. Wireless carriers can also obtain CSLI by "pinging" a phone whenever it is turned on, even if it is not engaged in an active call. The precision of CSLI varies according to several factors, and "[f]or a typical user, over time, some of that data will inevitably reveal locational precision approaching that of GPS." *The Electronic Communications Privacy Act (ECPA), Part 2: Geolocation Privacy and Surveillance: Hearing Before the Subcomm. on Crime, Terrorism, Homeland Sec. & Investigations of the H. Comm. On the Judiciary*, 113th Cong. (2013) (statement of Matt Blaze, Associate Professor, University of Pennsylvania), http://1.usa.gov/1awvgOa.

[13] *See* Letter from Ronald Weich, Assistant Attorney General, to Hon. Dianne Feinstein & Hon. Saxby Chambliss, Feb. 2, 2011, http://1.usa.gov/1cdFJ1G (enclosing *Report on the National Security Agency's Bulk Collection Programs for USA PATRIOT Act Reauthorization*); Siobhan Gorman & Julian E. Barnes, *Officials: NSA Doesn't Collect Cellphone-Location Records*, Wall St. J., June 16, 2013, http://on.wsj.com/13MnSsp.

[14] *See* 18 U.S.C. § 2709 (national security letter); 18 U.S.C. § 2703(c), (d) (court order for records concerning electronic communication service).

associated with the beginning and end of each call will be stored in a predictable, standardized format.

21. By contrast, the contents of telephone calls are not structured. Some people speak English, others Spanish, French, Mandarin, or Arabic. Some people speak using street slang or in a pidgin dialect, which can be difficult for others to understand. Conversations also lack a common structure: Some people get straight to the point, others engage in lengthy small talk. Speakers have different accents, exhibit verbal stutters and disfluencies. Although automated transcription of speech has advanced, it is still a difficult and error-prone process.

22. In contrast, the structured nature of metadata makes it very easy to analyze massive datasets using sophisticated data-mining and link-analysis programs. That analysis is greatly facilitated by technological advances over the past 35 years in computing, electronic data storage, and digital data mining. Those advances have radically increased our ability to collect, store, and analyze personal communications, including metadata.

23. Innovations in electronic storage today permit us to maintain, cheaply and efficiently, vast amounts of data. The ability to preserve data on this scale is, by itself, an unprecedented development—making possible the maintenance of a digital history that was not previously within the easy reach of any individual, corporation, or government.

24. This newfound data storage capacity has led to new ways of exploiting the digital record. Sophisticated computing tools permit the analysis of large datasets to identify embedded patterns and relationships, including personal details, habits, and behaviors. As a result, individual pieces of data that previously carried less potential to expose private information may now, in the aggregate, reveal sensitive details about our everyday lives—details that we had no intent or expectation of sharing.

25. IBM's Analyst's Notebook and Pen-Link are two such computing tools. Both are widely used by law enforcement and intelligence agencies for this purpose.[15]

26. IBM's Analyst Notebook product is a multi-purpose intelligence analysis tool that includes specific telephony metadata analysis features, which are "routinely" used to analyze large amounts of telephony metadata.[16] IBM even offers training courses entirely focused on using Analyst's Notebook to analyze telephone call records.[17]

27. Pen-Link is a tool that is purpose-built for processing and analyzing surveillance data. It is capable of importing subscriber Call Detail Record ("CDR") data from the proprietary formats

[15] *Public Safety & Law Enforcement Operations*, International Business Machines (last visited Aug. 22, 2013), http://ibm.co/1avGItq ("IBM® i2® solutions help law enforcers to turn huge volumes of crime data into actionable insights by delivering tools for tactical lead generation, intelligence analysis, crime analysis and predictive analysis."); *see also Defense and National Security Operations*, International Business Machines (last visited Aug. 22, 2013), http://ibm.co/18nateN ("IBM i2 solutions for military and national security organizations have been used across the world to process and analyze the vast quantities of information that they collect, to generate actionable intelligence and to share insights that help identify, predict and prevent hostile threats."); *see also* Pen-Link, *Unique Features of Pen-Link v8* at 16 (April 17, 2008), http://bit.ly/153ee9g ("Many U.S. Federal Law Enforcement and Intelligence agencies have acquired agency-wide site license contracts for the use of Pen-Link in their operations throughout the United States…Pen-Link systems are also becoming more frequently used by U.S. intelligence efforts operating in several other countries.").

[16] *Case Studies: Edith Cowan University, IBM i2 Solutions Help University Researchers Catch a Group of Would-Be Hackers*, International Business Machines (Mar. 27, 2013), http://ibm.co/13J2o36 ("Analyzing this volume of data is nothing new to many law enforcement users who routinely analyze tens of thousands of telephone records using IBM® i2® Analyst's Notebook®.").

[17] *Course Description: Telephone Analysis Using i2 Analyst's Notebook*, International Business Machines (last visited Aug. 22, 2013), http://ibm.co/1d5QlB8 ("This intermediate hands-on 3-day workshop focuses on the techniques of utilizing i2 Analyst's Notebook to conduct telephone toll analysis…Learn to import volumes of call detail records from various phone carriers, analyze those records and identify clusters and patterns in the data. Using both association and temporal charts, discover how to use different layouts and more advanced tools to analyze telephonic data quickly and effectively.").

used by the major telephone companies,[18] it can import and export call data to several federal surveillance databases,[19] as well as interact with commercial providers of public records databases such as ChoicePoint and LexisNexis. Pen-Link can perform automated "call pattern analysis," which "automatically identifies instances where particular sequences of calls occur, when they occur, how often they occur, and between which numbers and names."[20] As the company notes in its own marketing materials, this feature "would help the analyst determine how many times Joe paged Steve, then Steve called Barbara, then Steve called Joe back."[21]

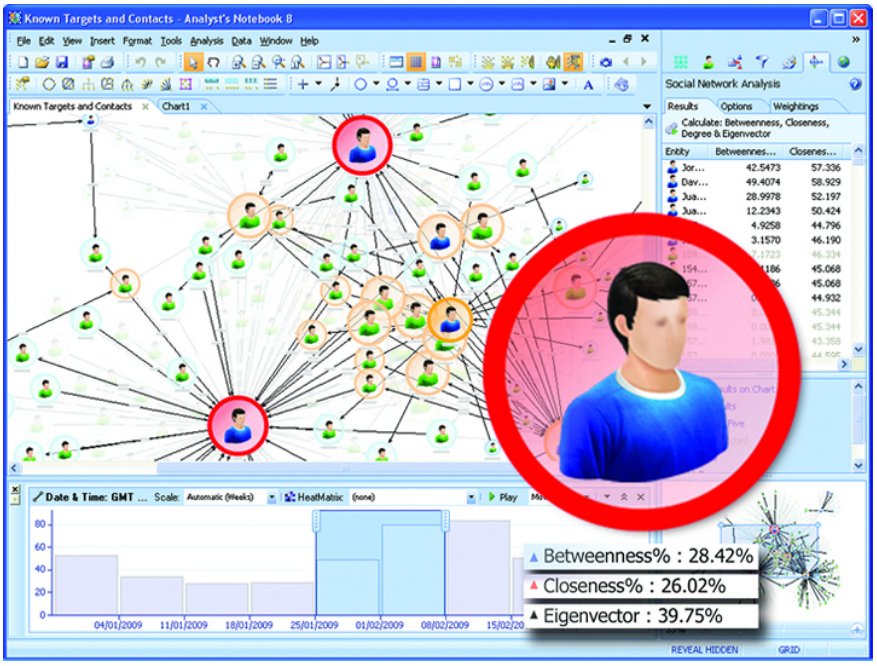

Figure 1: Screenshot of IBM's Analyst Notebook.[22]

[18] *See* Pen-Link, *Unique Features of Pen-Link v8* at 4 (Apr. 17, 2008), http://bit.ly/153ee9g (describing the capability to import 170 different data formats, used by phone companies to provide call detail records).

[19] *Id.* at 4.

[20] *Id.* at 7.

[21] *Id.*

[22] Image taken from *Data Analysis and Visualization for Effective Intelligence Analysis*, International Business Machines (last visited Aug. 22, 2013), http://ibm.co/16qT3hw.

28. The contents of calls are far more difficult to analyze in an automated fashion due to their unstructured nature. The government would first have to transcribe the calls and then determine which parts of the conversation are interesting and relevant. Assuming that a call is transcribed correctly, the government must still try to determine the meaning of the conversation: When a surveillance target is recorded saying "the package will be delivered next week," are they talking about an order they placed from an online retailer, a shipment of drugs being sent through the mail, or a terrorist attack? Parsing and interpreting such information, even when performed manually, is exceptionally difficult. To do so in an automated way, transcribing and data-mining the contents of hundreds of millions of telephone calls per day is an even more difficult task.

29. It is not surprising, then, that intelligence and law enforcement agencies often turn first to metadata. Examining metadata is generally more cost-effective than analyzing content. Of course, the government will likely still have analysts listen to every call made by the highest-value surveillance targets, but the resources available to the government do not permit it to do this for all of the calls of 300 million Americans.

The Creation of Metadata Is Unavoidable

30. As a general matter, it is practically impossible for individuals to avoid leaving a metadata trail when engaging in real-time communications, such as telephone calls or Internet voice chats.

31. After decades of research (much of it supported by the U.S. government), there now exist many tools that individuals and organizations can use to protect the confidentiality of their communications content. Smartphone applications are available that let individuals make encrypted telephone calls and send secure text messages.[23] Freely available software can be used

[23] Somini Sengupta, *Digital Tools to Curb Snooping,* N.Y. Times, July 17, 2013, http://nyti.ms/12JKz1s (describing RedPhone and Silent Circle).

to encrypt email messages and instant messages sent between computers, which can frustrate government surveillance efforts traditionally performed by intercepting communications as they are transmitted over the Internet.

32. However, these secure communication technologies protect only the content of the conversation and do not protect the metadata. Government agents that intercept an encrypted email may not know what was said, but they will be able to learn the email address that sent the message and the address that received it as well as the size of the message and when it was sent. Likewise, Internet metadata can reveal the parties making an encrypted audio call and the time and duration of the call, even if the voice contents of the call are beyond the reach of a wiretap.

33. There also exist security technologies specifically designed to hide metadata trails, but those technologies do not work quickly enough to allow real-time communication. The general technique for hiding the origin and destination information for an internet communication involves sending data through a series of intermediaries before it reaches the destination, thus making it more difficult for an entity such as a government agency to learn both the source and destination of the communication. (Such data is conventionally encrypted so that the intermediaries cannot capture it; and a series of intermediaries is used so that no one intermediary knows the identities of both endpoints.)

34. The most popular and well-studied of these metadata hiding systems is The Tor Project, which was originally created by the U.S. Naval Research Lab, and has since received significant funding from the State Department. One significant and widely acknowledged limitation of Tor is the noticeable delay introduced by using the tool. Web browsing conducted through Tor is much slower than through a direct connection to the Internet, as all data must be sent through a series of Tor relays, located in different parts of the world. These volunteer-run relays are

oversubscribed—that is, the demands on the few relays from hundreds of thousands of Tor users are greater than the relays can supply, leading to slowdowns due to "traffic jams" at the relay.

35. Browsing the web using Tor can be painfully slow, in some cases requiring several seconds or longer to load a page. Real-time audio and video communications require a connection with minimal delay, which Tor cannot deliver. Internet telephony and video conferencing services are simply unusable over metadata-protecting systems like Tor.

36. As a result, although individuals can use security technologies to protect the contents of their communications, there exist significant technical barriers that make it difficult, if not impossible, to hide communications metadata, particularly for real-time communications services like Internet telephony and video conferencing.

37. Over the last three decades, and especially with the widespread adoption of mobile phones in the past decade, our reliance on telecommunications has significantly increased. Mobile phones are today ubiquitous, and their use necessarily requires reliance on a service provider to transmit telephone calls, text messages, and other data to and fro. These communications inevitably produce telephony metadata, which is created whenever a person places a call. There is no practical way to prevent the creation of telephony metadata, or to erase it after the fact. The only reliable way to avoid creating such metadata is to avoid telephonic communication altogether.

Telephony Metadata Reveals Content

38. Telephony metadata can be extremely revealing, both at the level of individual calls and, especially, in the aggregate.

39. Although this metadata might, on first impression, seem to be little more than "information concerning the numbers dialed,"[24] analysis of telephony metadata often reveals information that could traditionally only be obtained by examining the contents of communications. That is, metadata is often a proxy for content.

40. In the simplest example, certain telephone numbers are used for a single purpose, such that any contact reveals basic and often sensitive information about the caller. Examples include support hotlines for victims of domestic violence[25] and rape,[26] including a specific hotline for rape victims in the armed services.[27] Similarly, numerous hotlines exist for people considering suicide,[28] including specific services for first responders,[29] veterans,[30] and gay and lesbian teenagers.[31] Hotlines exist for suffers of various forms of addiction, such as alcohol,[32] drugs, and gambling.[33]

---

[24] Administration White Paper, *Bulk Collection of Telephony Metadata Under Section 215 of the USA Patriot Act* 15 (Aug. 9, 2013), http://huff.to/1ey9ua5.

[25] *National Domestic Violence Hotline*, The Hotline (last visited Aug. 22, 2013), http://www.thehotline.org.

[26] *National Sexual Assault Hotline*, RAINN: Rape, Abuse & Incest National Network (last visited Aug. 22, 2013), http://www.rainn.org/get-help/national-sexual-assault-hotline.

[27] *About the Telephone Helpline*, DOD Safe Helpline (last visited Aug. 22, 2013), https://www.safehelpline.org/about-safe-helpline.

[28] *District of Columbia/Washington D.C. Suicide & Crisis Hotlines*, National Suicide Hotlines (last visited Aug. 22, 2013), http://www.suicidehotlines.com/distcolum.html.

[29] *Get Help Now! Contact us to Get Confidential Help via Phone or Email*, Safe Call Now (last visited Aug. 22, 2013), http://safecallnow.org.

[30] *About the Veterans Crisis Line*, Veterans Crisis Line (last visited Aug. 22, 2013), http://www.veteranscrisisline.net/About/AboutVeteransCrisisLine.aspx.

[31] *We Provide Crisis Intervention and Suicide Prevention for LGBTQ Youth*, The Trevor Project (last visited Aug. 22, 2013), thttp://www.thetrevorproject.org.

[32] *Alcohol Addiction Helpline*, Alcohol Hotline (last visited Aug. 22, 2013), http://www.alcoholhotline.com.

[33] *What is Problem Gambling*?, National Council on Problem Gambling (last visited Aug. 22, 2013), http://bit.ly/cyosu.

41. Similarly, inspectors general at practically every federal agency—including the NSA[34]—have hotlines through which misconduct, waste, and fraud can be reported, while numerous state tax agencies have dedicated hotlines for reporting tax fraud.[35] Hotlines have also been established to report hate crimes,[36] arson,[37] illegal firearms[38] and child abuse.[39] In all these cases, the metadata alone conveys a great deal about the content of the call, even without any further information.

42. The phone records indicating that someone called a sexual assault hotline or a tax fraud reporting hotline will of course not reveal the exact words that were spoken during those calls, but phone records indicating a 30-minute call to one of these numbers will still reveal information that virtually everyone would consider extremely private.

43. In some cases, telephony metadata can reveal information that is even more sensitive than the contents of the communication. In recent years, wireless telephone carriers have partnered with non-profit organizations in order to permit wireless subscribers to donate to charities by sending a text message from their telephones. These systems require the subscriber to send a specific text message to a special number, which will then cause the wireless carrier to add that

---

[34] Barton Gellman, *NSA Statements to the Post*, Wash. Post, Aug. 15, 2013, http://wapo.st/15LliAB.

[35] *Report Tax Fraud – Tax Fraud Hotline*, North Carolina Department of Revenue (last visited Aug. 22, 2013), http://www.dor.state.nc.us/taxes/reportfraud.html.

[36] *Report Hate Crimes*, LAMBDA GLBT Community Services (last visited Aug. 22, 2013), http://www.lambda.org/hatecr2.htm.

[37] *ATF Hotlines – Arson Hotline*, Bureau of Alcohol, Tobacco, Firearms and Explosives (last visited Aug. 22, 2013), http://www.atf.gov/contact/hotlines/index.html.

[38] *ATF Hotlines – Report Illegal Firearms Activity*, Bureau of Alcohol, Tobacco, Firearms and Explosives (last visited Aug. 22, 2013), http://www.atf.gov/contact/hotlines/index.html.

[39] *Childhelp National Child Abuse Hotline*, Childhelp (last visited Aug. 22, 2013), http://www.childhelp.org/pages/hotline-home.

donation to the subscriber's monthly telephone bill. For example, by sending the word HAITI to 90999, a wireless subscriber can donate $10 to the American Red Cross.

44. Such text message donation services have proven to be extremely popular. Today, wireless subscribers can use text messages to donate to churches,[40] to support breast cancer research,[41] and to support reproductive services organizations like Planned Parenthood.[42] Similarly, after a policy change in 2012 by the Federal Election Commission, political candidates like Barack Obama and Mitt Romney were able to raise money directly via text message.[43]

45. In all these cases, the most significant information—the recipient of the donation—is captured in the metadata, while the content of the message itself is less important. The metadata alone reveals the fact that the sender was donating money to their church, to Planned Parenthood, or to a particular political campaign.

46. Although it is difficult to summarize the sensitive information that telephony metadata about a single person can reveal, suffice it to say that it can expose an extraordinary amount about our habits and our associations. Calling patterns can reveal when we are awake and asleep; our religion, if a person regularly makes no calls on the Sabbath, or makes a large number of calls on Christmas Day; our work habits and our social aptitude; the number of friends we have; and even our civil and political affiliations.

---

[40] *Several Ways to Give*, The Simple Church (2013), http://bit.ly/1508Mgw; *Other Ways to Give*, North Point Church (last visited Aug. 22, 2013), http://bit.ly/16S3IkO.

[41] *Donate by Text*, Susan G. Komen for the Cure (last visited Aug. 22, 2013), http://sgk.mn/19AjGP7.

[42] *Help Support a New Future for Illinois Women and Families*, Planned Parenthood of Illinois (last visited Aug. 22, 2013), http://bit.ly/1bXI2TX.

[43] Dan Eggen, *Text to 'GIVE' to Obama: President's Campaign Launches Cellphone Donation Drive*, Wash. Post, Aug. 23, 2012, http://bit.ly/16ibjCZ.

Aggregated Telephony Metadata Is Even More Revealing

47. When call metadata is aggregated and mined for information across time, it can be an even richer repository of personal and associational details.

48. Analysis of metadata on this scale can reveal the network of individuals with whom we communicate—commonly called a *social graph*. By building a social graph that maps all of an organization's telephone calls over time, one could obtain a set of contacts that includes a substantial portion of the group's membership, donors, political supporters, confidential sources, and so on. Analysis of the metadata belonging to these individual callers, by moving one "hop" further out, could help to classify each one, eventually yielding a detailed breakdown of the organization's associational relationships.

49. For instance, metadata can help identify our closest relationships. Two people in an intimate relationship may regularly call each other, often late in the evening. If those calls become less frequent or end altogether, metadata will tell us that the relationship has likely ended as well—and it will tell us when a new relationship gets underway. More generally, someone you speak to once a year is less likely to be a close friend than someone you talk to once a week.

50. Even our relative power and social status can be determined by calling patterns. As *The Economist* observed in 2010, "People at the top of the office or social pecking order often receive quick callbacks, do not worry about calling other people late at night and tend to get more calls at times when social events are most often organized (sic), such as Friday afternoons."[44]

[44] *Mining Social Networks: Untangling the Social Web*, Economist, Sep. 2, 2010, http://econ.st/9iH1P7.

51. At times, by placing multiple calls in context, metadata analysis can even reveal patterns and sensitive information that would not be discoverable by intercepting the content of an individual communication.

52. Consider the following hypothetical example: A young woman calls her gynecologist; then immediately calls her mother; then a man who, during the past few months, she had repeatedly spoken to on the telephone after 11pm; followed by a call to a family planning center that also offers abortions. A likely storyline emerges that would not be as evident by examining the record of a single telephone call.

53. Likewise, although metadata revealing a single telephone call to a bookie may suggest that a surveillance target is placing a bet, analysis of metadata *over time* could reveal that the target has a gambling problem, particularly if the call records also reveal a number of calls made to payday loan services.

54. With a database of telephony metadata reaching back five years, many of these kinds of patterns will emerge once the collected phone records are subjected to even the most basic analytic techniques.

55. With an organization such as the ACLU, aggregated metadata can reveal sensitive information about the internal workings of the organization and about its external associations and affiliations. The ACLU's metadata trail reflects its relationships with its clients, its legislative contacts, its members, and the prospective whistleblowers who call the organization. Second-order analysis of the telephony metadata of the ACLU's contacts would then reveal even greater details about each of those contacts. For example, if a government employee suddenly begins contacting phone numbers associated with a number of news organizations and then the ACLU and then, perhaps, a criminal defense lawyer, that person's identity as a prospective

whistleblower could be surmised. Or, if the government studied the calling habits of the ACLU's members, it could assemble a detailed profile of the sorts of individuals who support the ACLU's mission.

56. I understand from the plaintiffs that they sometimes represent individuals in so-called "John Doe" lawsuits, where the individuals filing suit request anonymity—and are granted it by the courts—because they are juveniles or because they wish to conceal sensitive medical or psychiatric conditions. In such cases, analysis of aggregated metadata might reveal the anonymous litigant. If, for example, the lawyers in the case have only a handful of contacts in common other than mutual co-workers, and one or more of the lawyers generally call the same one of those common contacts shortly before or after hearings or deadlines in the lawsuit, this would imply the identity of the anonymous litigant. If the attorneys' calling patterns suggest more than one possible identity for the "John Doe," metadata analysis of the candidate individuals could verify the identity of the "John Doe," by correlating facts about the individuals with facts detailed in the lawsuit—for example, that he lives in a particular area (based on the area code of his phone or those of the majority of his contacts), that he has a particular job (based on calls made during work hours), that he has a particular medical condition (based on calls to medical clinics or specialists), or that he holds particular religious or political views (based on telephone donations, calls to political campaigns, or contact with religious organizations).

57. Metadata analysis could even expose litigation strategies of the plaintiffs. Review of the ACLU's telephony metadata might reveal, for example, that lawyers of the organization contacted, for example, an unusually high number of individuals registered as sex offenders in a particular state; or a seemingly random sample of parents of students of color in a racially

segregated school district; or individuals associated with a protest movement in a particular city or region.

58. In short, aggregated telephony metadata allows the government to construct social graphs and to study their evolution and communications patterns over days, weeks, months, or even years. Metadata analysis can reveal the rise and fall of intimate relationships, the diagnosis of a life-threatening disease, the telltale signs of a corporate merger or acquisition, the identity of a prospective government whistleblower, the social dynamics of a group of associates, or even the name of an anonymous litigant.

Mass Collection of Metadata and Data-Mining Across Many Individuals

59. Advances in the area of "Big Data" over the past few decades have enabled researchers to observe even deeper patterns by mining large pools of metadata that span many telephone subscribers.

60. Researchers have studied databases of call records to analyze the communications reciprocity in relationships,[45] the differences in calling patterns between mobile and landline subscribers,[46] and the social affinity and social groups of callers.[47]

61. Researchers have discovered that individuals have unique calling patterns, regardless of which telephone they are using,[48] they have figured out how to predict the kind of device that is

[45] Lauri Kovanen, Jari Saramaki & Kimmo Kaski, *Reciprocity of Mobile Phone Calls*, Dynamics of Socio-Economic Systems (Feb. 3, 2010), http://arxiv.org/pdf/1002.0763.pdf.

[46] Heath Hohwald, Enrique Frias-Martinez & Nuria Oliver, *User Modeling for Telecommunication Applications: Experiences and Practical Implications* 8, (Data Mining and User Modeling Group, Telefonica Research, 2013), http://bit.ly/1d7WkUU ("Interestingly, Monday is the day with most calls for landline users, while Friday is the day with most calls for mobile users. . . Mobile users spend less time on the phone than landline users.").

[47] Sara Motahari, Ole J. Mengshoel, Phyllis Reuther, Sandeep Appala, Luca Zoia & Jay Shah, *The Impact of Social Affinity on Phone Calling Patterns: Categorizing Social Ties from Call Data Records*, The 6th SNA-KDD Workshop (Aug. 12, 2012), http://b.gatech.edu/1d6i4RY.

making the calls (a telephone or a fax machine),[49] developed algorithms capable of predicting whether the phone line is used by a business or for personal use,[50] identified callers by social group (workers, commuters, and students) based on their calling patterns,[51] and even estimated the personality traits of individual subscribers.[52]

62. The work of these researchers suggests that the power of metadata analysis and its potential impact upon the privacy of individuals increases with the scale of the data collected and analyzed. It is only through access to massive datasets that researchers have been able to identify or infer new and previously private facts about the individuals whose calling records make up the telephone databases. Just as multiple calls by the same person reveal more than a single call, so too does a database containing calling data about millions of people reveal more information about the individuals contained within it than a database with calling data about just one person. As such, a universal database containing records about all Americans' communications will reveal vastly more information, including new observable facts not currently known to the

---

[48] Corrina Cortes, Daryl Pregibon & Chris Volinsky, *Communities of Interest*, AT&T Shannon Research Labs, http://www.research.att.com/~volinsky/papers/portugal.ps.

[49] Haim Kaplan, Maria Strauss & Mario Szegedy, *Just the Fax – Differentiating Voice and Fax Phone Lines Using Call Billing Data*, AT&T Labs, http://bit.ly/19Aa8Ua.

[50] Corinna Cortes & Daryl Pregibon, *Giga-Mining*, AT&T Labs-Research, http://bit.ly/153pMcI.

[51] Richard A. Becker, Ramon Caceres, Karrie Hanson, Ji Meng Loh, Simon Urbanek, Alexander Varshavsky & Chris Volinsky, *Clustering Anonymized Mobile Call Detail Records to Find Usage Groups*, AT&T Labs-Research, http://soc.att.com/16jmKdz.

[52] Rodrigo de Oliveira, Alexandros Karatzoglou, Pedro Concejero, Ana Armenta & Nuria Oliver, *Towards a Psychographic User Model from Mobile Phone Usage*, CHI 2011 Work-in-Progress (May 7–12, 2011), http://bit.ly/1f51mOy; *see also* Yves-Alexandre de Montjoye, Jordi Quoidbach, Florent Robic & Alex (Sandy) Pentland, *Predicting People Personality Using Novel Mobile Phone-Based Metrics*. Social Computing, Behavioral-Cultural Modeling and Prediction (2013), http://bit.ly/1867vWU.

research community, because no researcher has access to the kind of dataset that the government is presumed to have.

63. A common theme is seen in many of these examples of "big data" analysis of metadata. The analyst uses metadata about many individuals to discover patterns of behavior that are indicative of some attribute of an individual. The analyst can then apply these patterns to the metadata of an individual user, to infer the likely attributes of that user. In this way, the effect of collecting metadata about one individual is magnified when information is collected across the whole population.

64. The privacy impact of collecting all communications metadata about a single person for long periods of time is qualitatively different than doing so over a period of days. Similarly, the privacy impact of assembling the call records of every American is vastly greater than the impact of collecting data about a single person or even groups of people. Mass collection not only allows the government to learn information about more people, but it also enables the government to learn new, previously private facts that it could not have learned simply by collecting the information about a few, specific individuals.

Edward W. Felten

Dated: August 23, 2013

EXHIBIT 1

Edward W. Felten

Professor of Computer Science and Public Affairs
Director, Center for Information Technology Policy
Princeton University
Sherrerd Hall, Room 302
Princeton NJ 08544
(609) 258-5906
(609) 964-1855 fax
felten@cs.princeton.edu

Education

Ph.D. in Computer Science and Engineering, University of Washington, 1993.
Dissertation title: "Protocol Compilation: High-Performance Communication for Parallel Programs." Advisors: Edward D. Lazowska and John Zahorjan.
M.S. in Computer Science and Engineering, University of Washington, 1991.
B.S. in Physics, with Honors, California Institute of Technology, 1985.

Employment

Professor of Computer Science and Public Affairs, Princeton University, 2006-present.

Chief Technologist, U.S. Federal Trade Commission, 2011-2012.

Professor of Computer Science, Princeton University, 2003-2006.
Associate Professor of Computer Science, Princeton University, 1999-2003.
Assistant Professor of Computer Science, Princeton University, 1993-99.
Senior Computing Analyst, Caltech Concurrent Computing Project, California Institute of Technology, 1986-1989.

Director, Center for Information Technology Policy, Princeton University, 2005-present.

Elysium Digital LLC and various law firms. Consulting and expert testimony in technology litigation, 1998-present
U.S. Federal Trade Commission: consulting regarding spam policy and investigation, 2004, 2006.
U.S. Dept. of Justice, Antitrust Division: consulting and testimony in Microsoft antitrust case, 1998-2002..
Electronic Frontier Foundation. Consulting in intellectual property / free speech lawsuits, 2001-2010.
Certus Ltd.: consultant in product design and analysis, 2000-2002.
Cigital Inc.: Technical Advisory Board member, 2000-2007.

Cloakware Ltd.: Technical Advisory Board member, 2000-2003.
Propel.com: Technical Advisory Board member, 2000-2002.
NetCertainty.com: Technical Advisory Board member, 1999-2002.
FullComm LLC: Scientific Advisory Board member, 1999-2001.
Sun Microsystems: Java Security Advisory Board member, 1997-2001.
Finjan Software: Technical Advisory Board member, 1997-2002.
International Creative Technologies: consultant in product design and analysis, 1997-98.
Bell Communications Research: consultant in computer security research, 1996-97.

Honors and Awards

National Academy of Engineering, 2013.
American Academy of Arts and Sciences, 2011
ACM Fellow, 2007.
EFF Pioneer Award, 2005.
Scientific American Fifty Award, 2003.
Alfred P. Sloan Fellowship, 1997.
Emerson Electric, E. Lawrence Keyes Faculty Advancement Award, Princeton University School of Engineering, 1996.
NSF National Young Investigator award, 1994.
Outstanding Paper award, 1997 Symposium on Operating Systems Principles.
Best Paper award, 1995 ACM SIGMETRICS Conference.
AT&T Ph.D. Fellowship, 1991-93.
Mercury Seven Foundation Fellowship, 1991-93.

Research Interests

Information security. Privacy. Technology law and policy. Internet software. Intellectual property policy. Using technology to improve government. Operating systems. Interaction of security with programming languages and operating systems. Distributed computing. Parallel computing architecture and software.

Professional Service

Professional Societies and Advisory Groups

ACM U.S. Public Policy Committee, Vice Chair, 2008-2010, 2012-present.
DARPA Privacy Panel, 2010-2012.
Transportation Security Administration, Secure Flight Privacy Working Group, 2005.
National Academies study committee on Air Force Information Science and Technology Research, 2004-present.
Electronic Frontier Foundation, Advisory Board, 2004-2007.
ACM U.S. Public Policy Committee, 2004-present (Executive Committee, 2005-present)

ACM Advisory Committee on Security and Privacy, 2002-2003.
DARPA Information Science and Technology (ISAT) study group, 2002-2004.
Co-chair, ISAT study committee on "Reconciling Security with Privacy," 2001-2002.
National Academy study committee on Foundations of Computer Science, 2001-2004.

Program Committees

World Wide Web Conference, 2006.
USENIX General Conference, 2004.
Workshop on Foundations of Computer Security, 2003.
ACM Workshop on Digital Rights Management, 2001.
ACM Conference on Computer and Communications Security, 2001.
ACM Conference on Electronic Commerce, 2001.
Workshop on Security and Privacy in Digital Rights Management, 2001.
Internet Society Symposium on Network and Distributed System Security, 2001.
IEEE Symposium on Security and Privacy, 2000.
USENIX Technical Conference, 2000.
USENIX Windows Systems Conference, 2000.
Internet Society Symposium on Network and Distributed System Security, 2000.
IEEE Symposium on Security and Privacy, 1998.
ACM Conference on Computer and Communications Security, 1998.
USENIX Security Symposium, 1998.
USENIX Technical Conference, 1998.
Symposium on Operating Systems Design and Implementation, 1996.

Boards

Electronic Frontier Foundation, Board of Directors, 2007-2010.
DARPA Information Science and Technology study board, 2001-2003.
Cigital Inc.: Technical Advisory Board.
Sun Microsystems, Java Security Advisory Council.
Cloakware Ltd.: Technical Advisory Board.
Propel.com: Technical Advisory Board.
Finjan Software: Technical Advisory Board.
Netcertainty: Technical Advisory Board.
FullComm LLC: Scientific Advisory Board.

University and Departmental Service

Committee on Online Courses, 2012-present
Director, Center for Information Technology Policy, 2005-present.
Committee on the Course of Study, 2009-present.
SEAS Strategic Planning, 2004.
- Member, Executive Committee
- Co-Chair, Interactions with Industry area.
- Co-Chair, Engineering, Policy, and Society area.

Faculty Advisory Committee on Policy, 2002-present.
Council of the Princeton University Community, 2002-present (Executive Committee)
Faculty Advisory Committee on Athletics, 1998-2000.

Computer Science Academic Advisor, B.S.E. program, class of 1998 (approx. 25 students)
Faculty-Student Committee on Discipline, 1996-98.
Faculty-Student Committee on Discipline, Subcommittee on Sexual Assault and Harrassment, 1996-98.

Students Advised

Ph.D. Advisees:

Harlan Yu (Ph.D. 2012). Dissertation: Designing Software to Shape Open Government Policy.
Ariel J. Feldman (Ph.D. 2012). Dissertation: Privacy and Integrity in the Untrusted Cloud.
Joseph A. Calandrino (Ph.D. 2012). Dissertation: Control of Sensitive Data in Systems with Novel Functionality.
William B. Clarkson (Ph.D. 2012). Dissertation: Breaking Assumptions: Distinguishing Between Seemingly Identical Items Using Cheap Sensors. Technical staff member at Google.
Matthias Jacob (Ph.D. 2009). Technical staff member at Nokia.
J. Alex Halderman (Ph.D. 2009). Dissertation: Security Failures in Non-traditional Computing Environments. Assistant Professor of Computer Science, University of Michigan.
Shirley Gaw (Ph.D. 2009). Dissertation: Ideals and Reality: Adopting Secure Technologies and Developing Secure Habits to Prevent Message Disclosure. Technical staff member at Google.
Brent Waters (Ph.D. 2004). Dissertation: Security in a World of Ubiquitous Recording Devices. Assistant Professor of Computer Science, University of Texas.
Robert A. Shillingsburg (Ph.D. 2004). Dissertation: Improving Distributed File Systems using a Shared Logical Disk. Retired; previously a technical staff member at Google.
Michael Schneider (Ph.D. 2004). Dissertation: Network Defenses against Denial of Service Attacks. Researcher, Supercomputing Research Center, Institute for Defense Analyses.
Minwen Ji (Ph.D. 2001). Dissertation: Data Distribution for Dynamic Web Content. Researcher, HP Labs.
Dirk Balfanz (Ph.D. 2000). Dissertation: Access Control for Ad Hoc Collaboration. Technical staff member at Google.
Dan S. Wallach (Ph.D. 1998). Dissertation: A New Approach to Mobile Code Security. Associate Professor of Computer Science, Rice University.

Significant Advisory Role:

Drew Dean (Ph.D. 1998). Advisor: Andrew Appel. Program Manager at DARPA.
Stefanos Damianakis (Ph.D. 1998). Advisor: Kai Li. President and CEO, Netrics, Inc.

Pei Cao (Ph.D. 1996). Advisor: Kai Li. Staff technologist at Facedbook.
Lujo Bauer (Ph.D. 2003). Advisor: Andrew Appel. Research Scientist, School of Computer Science, Carnegie Mellon University.

Publications

Books and Book Chapters

[1] Enabling Innovation for Civic Engagement. David G. Robinson, Harlan Yu, and Edward W. Felten. In Open Government, Daniel Lathrop and Laurel Ruma, eds., O'Reilly, 2010.

[2] Securing Java: Getting Down to Business with Mobile Code. Gary McGraw and Edward W. Felten. John Wiley and Sons, New York 1999.

[3] Java Security: Web Browsers and Beyond. Drew Dean, Edward W. Felten, Dan S. Wallach, and Dirk Balfanz. In "Internet Besieged: Countering Cyberspace Scofflaws," Dorothy E. Denning and Peter J. Denning, eds. ACM Press, New York, 1997.

[4] Java Security: Hostile Applets, Holes and Antidotes. Gary McGraw and Edward Felten. John Wiley and Sons, New York, 1996

[5] Dynamic Tree Searching. Steve W. Otto and Edward W. Felten. In "High Performance Computing", Gary W. Sabot, ed., Addison Wesley, 1995.

Journal Articles

[6] Government Data and the Invisible Hand. David Robinson, Harlan Yu, William Zeller, and Edward W. Felten. Yale Journal of Law and Technology, vol. 11, 2009.

[7] Mechanisms for Secure Modular Programming in Java. Lujo Bauer, Andrew W. Appel, and Edward W. Felten. Software – Practice and Experience, 33:461-480, 2003.

[8] The Digital Millennium Copyright Act and its Legacy: A View from the Trenches. Illinois Journal of Law, Technology and Policy, Fall 2002.

[9] The Security Architecture Formerly Known as Stack Inspection: A Security Mechanism for Language-based Systems. Dan S. Wallach, Edward W. Felten, and Andrew W. Appel. ACM Transactions on Software Engineering and Methodology, 9:4, October 2000.

[10] Statically Scanning Java Code: Finding Security Vulnerabilities. John Viega, Tom Mutdosch, Gary McGraw, and Edward W. Felten. IEEE Software, 17(5), Sept./Oct. 2000.

[11] Client-Server Computing on the SHRIMP Multicomputer. Stefanos N. Damianakis, Angelos Bilas, Cezary Dubnicki, and Edward W. Felten. IEEE Micro 17(1):8-18, February 1997.

[12] Fast RPC on the SHRIMP Virtual Memory Mapped Network Interface. Angelos Bilas and Edward W. Felten. IEEE Transactions on Parallel and Distributed Computing, February 1997.

[13] Implementation and Performance of Integrated Application-Controlled File Caching, Prefetching and Disk Scheduling. Pei Cao, Edward W. Felten, Anna R. Karlin, and Kai Li. ACM Transactions on Computer Systems, Nov 1996.

[14] Virtual Memory Mapped Network Interface Designs. Matthias A. Blumrich, Cezary Dubnicki, Edward W. Felten, Kai Li, and Malena Mesarina. IEEE Micro, 15(1):21-28, February 1995.

Selected Symposium Articles

[15] Social Networking with Frientegrity: Privacy and Integrity with an Untrusted Provider. Ariel J. Feldman, Aaron Blankstein, Michael J. Freedman, and Edward W. Felten. Proc. USENIX Security Symposium, Aug. 2012.

[16] Bubble Trouble: Off-Line De-Anonymization of Bubble Forms. Joseph A. Calandrino, William Clarkson, and Edward W. Felten. Proc. USENIX Security Symposium, Aug. 2011

[17] You Might Also Like: Privacy Risks of Collaborative Filtering. Joseph A. Calandrino, Ann Kilzer, Arvind Narayanan, Edward W. Felten, and Vitaly Shmatikov. Proc. IEEE Symposium on Security and Privacy, May 2011.

[18] SPORC: Group Collaboration Using Untrusted Cloud Resources. Ariel J. Feldman, William P. Zeller, Michael J. Freedman, and Edward W. Felten. Proc. Symposium on Operating Systems Design and Implementation, 2010.

[19] SVC: Selector-Based View Composition for Web Frameworks. William Zeller and Edward W. Felten. Proc. USENIX Conference on Web Application Development, 2010.

[20] Defeating Vanish with Low-Cost Sybil Attacks Against Large DHTs. Scott Wolchok, Owen S. Hofmann, Nadia Heninger, Edward W. Felten, J. Alex Halderman, Christopher J. Rossbach, Brent Waters, and Emmet Witchel. Proc. 17th Network and Distributed System Security Symposium, 2010.

[21] Can DREs Provide Long-Lasting Security? The Case of Return-Oriented Programming and the AVC Advantage. Stephen Checkoway, Ariel J. Feldman, Brian Kantor, J. Alex Halderman, Edward W. Felten, and Hovav Shacham, Proc. Electronic Voting Technology Workshop, 2009.

[22] Some Consequences of Paper Fingerprinting for Elections. Joseph A. Calandrino, William Clarkson, and Edward W. Felten. Proc. Electronic Voting Technology Workshop, 2009.

[23] Software Support for Software-Independent Auditing. Gabrielle A. Gianelli, Jennifer D. King, Edward W. Felten, and William P. Zeller. Proc. Electronic Voting Technology Workshop, 2009.

[24] Fingerprinting Blank Paper Using Commodity Scanners. William Clarkson, Tim Weyrich, Adam Finkelstein, Nadia Heninger, J. Alex Halderman, and Edward W. Felten. Proc. ACM Symposium on Security and Privacy, May 2009.

[25] Lest We Remember: Cold Boot Attacks on Encryption Keys. J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. Proc. Usenix Security Symposium, 2008.

[26] In Defense of Pseudorandom Sample Selection. Joseph A. Calandrino, J. Alex Halderman, and Edward W. Felten. Proc. Electronic Voting Technology Workshop, 2008.

[27] Security Analysis of the Diebold AccuVote-TS Voting Machine. Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten. Proc. Electronic Voting Technology Workshop, 2007.

[28] Machine-Assisted Election Auditing. Joseph A. Calandrino, J. Alex Halderman, and Edward W. Felten. Proc. Electronic Voting Technology Workshop, 2007.

[29] Lessons from the Sony CD DRM Episode. J. Alex Halderman and Edward W. Felten. Proc. Usenix Security Symposium, 2006.

[30] A Convenient Method for Securely Managing Passwords. J. Alex Halderman, Brent R. Waters, and Edward W. Felten. Proc. 14th World Wide Web Conference, 2005.

[31] New Client Puzzle Outsourcing Techniques for DoS Resistance. Brent R. Waters, Ari Juels, J. Alex Halderman, and Edward W. Felten. ACM Conference on Computer and Communications Security. November 2004.

[32] Privacy Management for Portable Recording Devices. J. Alex Halderman, Brent R. Waters, and Edward W. Felten. 3rd Workshop on Privacy in Electronic Society. November 2004.

[33] Receiver Anonymity via Incomparable Public Keys. Brent R. Waters, Edward W. Felten, and Amit Sahai. ACM Conference on Computer and Communications Security. November 2003.

[34] Attacking an Obfuscated Cipher by Injecting Faults. Matthias Jacob, Dan Boneh, and Edward W. Felten. ACM Workshop on Digital Rights Management, November 2002.

[35] A General and Flexible Access-Control System for the Web. Lujo Bauer, Michael A. Schneider, and Edward W. Felten. 11th USENIX Security Symposium, August 2002.

[36] Informed Consent in the Mozilla Browser: Implementing Value-Sensitive Design. Batya Friedman, Daniel C. Howe, and Edward W. Felten. Hawaii International Conference on System Sciences, January 2002. (Best Paper award, organizational systems track.)

[37] Reading Between the Lines: Lessons from the SDMI Challenge. Scott A. Craver, John P. McGregor, Min Wu, Bede Liu, Adam Stubblefield, Ben Swartzlander, Dan S. Wallach, Drew Dean, and Edward W. Felten. USENIX Security Symposium, August 2001.

[38] Cookies and Web Browser Design: Toward Realizing Informed Consent Online. Lynette I. Millett, Batya Friedman, and Edward W. Felten. Proc. of CHI 2001 Conference on Human Factors in Computing Systems, April 2001.

[39] Timing Attacks on Web Privacy. Edward W. Felten and Michael A. Schneider. Proc. of 7th ACM Conference on Computer and Communications Security, Nov. 2000.

[40] Archipelago: An Island-Based File System for Highly Available and Scalable Internet Services. USENIX Windows Systems Symposium, August 2000.

[41] Proof-Carrying Authentication. Andrew W. Appel and Edward W. Felten. Proc. of 6th ACM Conference on Computer and Communications Security, Nov. 1999.

[42] An Empirical Study of the SHRIMP System. Matthias A. Blumrich, Richard D. Alpert, Yuqun Chen, Douglas W. Clark, Stefanos, N. Damianakis, Cezary Dubnicki, Edward W. Felten, Liviu Iftode, Margaret Martonosi, Robert A. Shillner, and Kai Li. Proc. of 25th International Symposium on Computer Architecture, June 1998.

[43] Performance Measurements for Multithreaded Programs. Minwen Ji, Edward W. Felten, and Kai Li. Proc. of 1998 SIGMETRICS Conference, June 1998.

[44] Understanding Java Stack Inspection. Dan S. Wallach and Edward W. Felten. Proc. of 1998 IEEE Symposium on Security and Privacy, May 1998.

[45] Extensible Security Architectures for Java. Dan S. Wallach, Dirk Balfanz, Drew Dean, and Edward W. Felten. Proc. of 16th ACM Symposium on Operating Systems Principles, Oct. 1997. Outstanding Paper Award.

[46] Web Spoofing: An Internet Con Game. Edward W. Felten, Dirk Balfanz, Drew Dean, and Dan S. Wallach. Proc. of 20th National Information Systems Security Conference, Oct. 1997.

[47] Reducing Waiting Costs in User-Level Communication. Stefanos N. Damianakis, Yuqun Chen, and Edward W. Felten. Proc. of 11th Intl. Parallel Processing Symposium, April 1997.

[48] Stream Sockets on SHRIMP. Stefanos N. Damianakis, Cezary Dubnicki, and Edward W. Felten. Proc. of 1st Intl. Workshop on Communication and Architectural Support for Network-Based Parallel Computing, February 1997. (Proceedings available as Lecture Notes in Computer Science #1199.)

[49] Early Experience with Message-Passing on the SHRIMP Multicomputer. Richard D. Alpert, Angelos Bilas, Matthias A. Blumrich, Douglas W. Clark, Stefanos Damianakis, Cezary Dubnicki, Edward W. Felten, Liviu Iftode, and Kai Li. Proc. of 23rd Intl. Symposium on Computer Architecture, 1996.

[50] A Trace-Driven Comparison of Algorithms for Parallel Prefetching and Caching. Tracy Kimbrel, Andrew Tomkins, R. Hugo Patterson, Brian N. Bershad, Pei Cao, Edward W. Felten, Garth A. Gibson, Anna R. Karlin, and Kai Li. Proc. of 1996 Symposium on Operating Systems Design and Implementation.

[51] Java Security: From HotJava to Netscape and Beyond. Drew Dean, Edward W. Felten, and Dan S. Wallach. Proc. of 1996 IEEE Symposium on Security and Privacy.

[52] Integrated Parallel Prefetching and Caching. Tracy Kimbrel, Pei Cao, Edward W. Felten, Anna R. Karlin, and Kai Li. Proc. of 1996 SIGMETRICS Conference.

[53] Software Support for Virtual Memory-Mapped Communication. Cezary Dubnicki, Liviu Iftode, Edward W. Felten, and Kai Li. Proc. of Intl. Parallel Processing Symposium, April 1996.

[54] Protected, User-Level DMA for the SHRIMP Network Interface. Matthias A. Blumrich, Cezary Dubnicki, Edward W. Felten, and Kai Li. Proc. of 2nd Intl. Symposium on High-Performance Computer Architecture, Feb. 1996

[55] Improving Release-Consistent Shared Virtual Memory using Automatic Update . Liviu Iftode, Cezary Dubnicki, Edward W. Felten, and Kai Li. Proc. of 2nd Intl. Symposium on High-Performance Computer Architecture, Feb. 1996

[56] Synchronization for a Multi-Port Frame Buffer on a Mesh-Connected Multicomputer. Bin Wei, Gordon Stoll, Douglas W. Clark, Edward W. Felten, and Kai Li. Parallel Rendering Symposium, Oct. 1995.

[57] A Study of Integrated Prefetching and Caching Strategies. Pei Cao, Edward W. Felten, Anna R. Karlin, and Kai Li. Proc. of 1995 ACM SIGMETRICS Conference. Best Paper award.

[58] Evaluating Multi-Port Frame Buffer Designs for a Mesh-Connected Multicomputer. Gordon Stoll, Bin Wei, Douglas W. Clark, Edward W. Felten, Kai Li, and Patrick Hanrahan. Proc. of 22nd Intl. Symposium on Computer Architecture.

[59] Implementation and Performance of Application-Controlled File Caching. Pei Cao, Edward W. Felten, and Kai Li. Proc. of 1st Symposium on Operating Systems Design and Implementation, pages 165-178, November 1994.

[60] Application-Controlled File Caching Policies. Pei Cao, Edward W. Felten, and Kai Li. Proc. of USENIX Summer 1994 Technical Conference, pages 171-182, 1994.

[61] Virtual Memory Mapped Network Interface for the SHRIMP Multicomputer. Matthias A. Blumrich, Kai Li, Richard D. Alpert, Cezary Dubnicki, Edward W. Felten, and Jonathan S. Sandberg. Proc. of Intl. Symposium on Computer Architecture, 1994.

[62] Performance Issues in Non-Blocking Synchronization on Shared-Memory Multiprocessors. Juan Alemany and Edward W. Felten. Proceedings of Symposium on Principles of Distributed Computing, 1992.

[63] Improving the Performance of Message-Passing Applications by Multithreading. Edward W. Felten and Dylan McNamee. Proceedings of Scalable High-Performance Computing Conference (SHPCC), 1992.

[64] A Highly Parallel Chess Program. Edward W. Felten and Steve W. Otto. 1988 Conference on Fifth Generation Computer Systems.

Selected Other Publications

[65] Strangers in a Strange Land. Review of *Blown to Bits: Your Life, Liberty, and Happiness after the Digital Explosion*, by Abelson, Ledeen, and Lewis. American Scientist, 97:4. July/August 2009.

[66] Lest We Remember: Cold-Boot Attacks on Encryption Keys. J. Alex Halderman, Seth D. Schoen, Nadia Heninger, William Clarkson, William Paul, Joseph A. Calandrino, Ariel J. Feldman, Jacob Appelbaum, and Edward W. Felten. *Communications of the ACM*, 52(5):91-98. May 2009.

[67] Security Analysis of the Diebold AccuVote-TS Voting Machine. Ariel J. Feldman, J. Alex Halderman, and Edward W. Felten. Sept. 2006.

[68] Digital Rights Management, Spyware, and Security. Edward W. Felten and J. Alex Halderman, *IEEE Security and Privacy,* Jan./Feb. 2006.

[69] Inside RISKS: DRM and Public Policy. Edward W. Felten. Communications of the ACM, 48:7, July 2005.

[70] Understanding Trusted Computing: Will its Benefits Outweigh its Drawbacks? Edward W. Felten. IEEE Security and Privacy, May 2003.

[71] A Skeptical View of DRM and Fair Use. Edward W. Felten. Communications of the ACM 46(4):56-61, April 2003.

[72] Consumer Privacy and Government Technology Mandates in the Digital Media Marketplace. Testimony before U.S. Senate Commerce Committee. September 2003.

[73] Secure, Private Proofs of Location. Brent R. Waters and Edward W. Felten. Submitted for publication, 2003.

[74] An Efficient Heuristic for Defense Against Distributed Denial of Service Attacks using Route-Based Distributed Packet Filtering. Michael A. Schneider and Edward W. Felten. Submitted for publication, 2003.

[75] Written testimony to House Commerce Committee, Subcommittee on Courts, the Internet, and Intellectual Property, oversight hearing on "Piracy of Intellectual Property on Peer to Peer Networks." September 2002.

[76] Written testimony to Senate Judiciary Committee hearings on "Competition, Innovation, and Public Policy in the Digital Age: Is the Marketplace Working to Protect Digital Creativity?" March 2002.

[77] Informed Consent Online: A Conceptual Model and Design Principles. Batya Friedman, Edward W. Felten, and Lynette I. Millett. Technical Report 2000-12-2, Dept. of Computer Science and Engineering, University of Washington, Dec. 2000.

[78] Mechanisms for Secure Modular Programming in Java. Lujo Bauer, Andrew W. Appel, and Edward W. Felten. Technical Report CS-TR-603-99, Department of Computer Science, Princeton University, July 1999.

[79] A Java Filter. Dirk Balfanz and Edward W. Felten. Technical Report 567-97, Dept. of Computer Science, Princeton University, October 1997.

[80] Inside RISKS: Webware Security. Edward W. Felten. Communications of the ACM, 40(4):130, 1997.

[81] Simplifying Distributed File Systems Using a Shared Logical Disk.Robert A. Shillner and Edward W. Felten. Princeton University technical report TR-524-96.

[82] Contention and Queueing in an Experimental Multicomputer: Analytical and Simulation-based Results. Wenjia Fang, Edward W. Felten, and Margaret Martonosi. Princeton University technical report TR-508-96.

[83] Design and Implementation of NX Message Passing Using SHRIMP Virtual Memory Mapped Communication. Richard D. Alpert, Cezary Dubnicki, Edward W. Felten, and Kai Li. Princeton University technical report TR-507-96.

[84] Protocol Compilation: High-Performance Communication for Parallel Programs. Edward W. Felten. Ph.D. dissertation, Dept. of Computer Science and Engineering, University of Washington, August 1993.

[85] Building Counting Networks from Larger Balancers. Edward W. Felten, Anthony LaMarca, and Richard Ladner. Univ. of Washington technical report UW-CSE-93-04-09.

[86] The Case for Application-Specific Communication Protocols. Edward W. Felten. Univ. of Washington technical report TR-92-03-11.

[87] A Centralized Token-Based Algorithm for Distributed Mutual Exclusion. Edward W. Felten and Michael Rabinovich. Univ. of Washington technical report TR-92-02-02.

[88] Issues in the Implementation of a Remote Memory Paging System. Edward W. Felten and John Zahorjan. Univ. of Washington technical report TR-91-03-09.