

Formal Methods for Privacy

Michael Carl Tschantz* **Jeannette M. Wing[†]**

September 2009
CMU-CS-09-154

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

This paper will appear in the Proceedings of Formal Methods 2009.

*mtschant@cs.cmu.edu

[†]wing@cs.cmu.edu

This research was sponsored in part by the US Army Research Office under contract no. DAAD19-02-1-0389 (“Perpetually Available and Secure Information Systems”) at Carnegie Mellon University’s CyLab. It was also partially supported by the National Science Foundation, while the second author was working at the Foundation. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government, or any other entity.

Keywords: privacy, formal methods

Abstract

Privacy means something different to everyone. Against a vast and rich canvas of diverse types of privacy rights and violations, we argue technology's dual role in privacy: new technologies raise new threats to privacy rights and new technologies can help preserve privacy. Formal methods, as just one class of technology, can be applied to privacy, but privacy raises new challenges, and thus new research opportunities, for the formal methods community.

1 Introduction

What is privacy? Today, the answer seems to be “It all depends on whom you ask.” There are philosophical, legal, societal, and technical notions of privacy. Cultures differ in their expectations regarding privacy. In some cultures, it is impolite to ask someone’s age or someone’s salary. Governments differ in their citizens’ rights to privacy; just witness the difference in privacy among the United States, the European Union, and China. What an adult thinks as private differs from what a teenager thinks, and vice versa [18].

New technologies give rise to new privacy concerns. Warren and Brandeis’s 1890 seminal paper, “The Right to Privacy,” was written after photographic and printing technologies made it easier to share and spread images and text in public [76]. Skipping ahead a century, with the explosion of the Internet, privacy is finally getting serious attention by the scientific community. More and more personal information about us is available online. It is by our choice that we give our credit card numbers to on-line retailers for the convenience of on-line shopping. Companies like Google, Yahoo, and Microsoft track our search queries to personalize the ads we see alongside the response to a query. With cloud computing, we further entrust in third parties the storage and management of private information in places unknown to us. We are making it easier for others to find out about our personal habits, tastes, and history. In some cases it is deliberate. The rise of social networks like Facebook, on-line community sites like Flickr, and communication tools like Twitter raises new questions about privacy, as people willingly give up some privacy to enhance social relationships or to share information easily with friends. At the same time, cyberattacks have increased in number and sophistication, making it more likely that unintentionally or not, personal information will fall into the wrong hands.

The National Academies study *Engaging Privacy and Information Technology in a Digital Age* [44] presents a compelling argument for the need for technology and policy experts to work together in addressing privacy, especially as new technology raises new privacy concerns. It is our responsibility as scientists and engineers to understand what can or cannot be done from a technical point of view on privacy: what is provably possible or impossible and what is practically possible or impossible. Otherwise, society may end up in a situation where privacy regulations put into place are technically infeasible to meet.

In this paper, we start in Section 2 by painting a broad picture of the diverse types of privacy. Against this canvas, we discuss the dual role of technology: how new technologies pose new threats to privacy (Section 3) and how technologies can help preserve privacy (Section 4). Finally, focusing on formal methods, as a specific class of technology, we identify some opportunities and challenges in using formal methods to protect privacy (Section 5).

2 Types of Privacy Rights and Violations

Philosophers justify the importance of privacy in different ways. Bloustein defends privacy as necessary for human dignity [7]. Others focus on privacy’s role in enabling intimate relations [24, 27, 28, 13] or interpersonal relations in general [54]. Gavison views privacy as a means of controlling access to the person [25].

Given the numerous philosophical justifications, legal scholars, starting with Prosser [53], have generally viewed privacy as a collection of related rights rather than a single concept. Solove in 2006 provided a taxonomy of possible privacy violations [59]. He collects these related violations into four groups: *invasions*, *information collection*, *information processing*, and *information dissemination*.

Invasions represent interference in what is traditionally considered the private sphere of life. Solove identifies two forms of invasions. The first involves physical *intrusions* either upon private property (such as trespassing in the home) or upon the body (such as blocking one's passage). The second is *decisional interference*, which is interfering with personal decisions. For example, the Supreme Court of the United States has used the right to privacy to justify limiting the government's ability to regulate contraceptives [61, 63], abortion [64], and sodomy [71] (cf. [67]). However, some view invasions as violations of other rights such as property and security rights in the case of intrusions [73], or the rights to autonomy and liberty in the case of decisional interference [51].

Solove's remaining three groupings of privacy rights are more difficult to reduce to other rights. They all involve a *data subject* about whom a *data holder* has information. The data holder may commit privacy violations in how he collects the information, how he processes it, or how he disseminates it to others.

Information collection includes making observations through *surveillance* and seeking information through *interrogation*. Information collection affects privacy by making people uneasy in how the collected information could be used. Thus, it is a violation of privacy even if the collected information is never used. Furthermore, interrogation can place people in the awkward position of having to refuse to answer questions. Even in the absence of these violations *per se*, information collection should be controlled to prevent other violations of privacy such as blackmail.

Even if information is collected in privacy-respecting ways, it can be processed in ways that violate privacy. Such information processing violations have the following forms. *Aggregation* is similar to surveillance in that it makes information available, but aggregation does so by combining diffuse pieces of information rather than collecting new information. Aggregation enables inferences that would be unavailable otherwise. *Identification*, linking information with a person by way of an identifier, also makes information more available and may alter how a person is treated. *Insecurity* makes information more available to those who should not be granted access such as identity thieves and can also lead to distortion of data if false data is entered. *Secondary uses* make information available for purposes for which it was not originally intended. *Exclusion* is the inability of a data subject to know what records are kept, to view them, to know how they are used, or to correct them. All these forms of information processing create uncertainty on the part of the data subject. Exclusion directly causes this uncertainty by keeping information about the information kept on the data subject secret. The other forms of information processing create this uncertainty by making information available in new, possibly unanticipated ways. Even in the absence of more material misuse of the information, such uncertainty can be a harm in of itself as it forces the data subject to live in fear of how his information may be used.

After information is processed, the data holder will typically disseminate it to others for use. Some forms of information dissemination can violate privacy by providing information to inappro-

prate entities. A breach of *confidentiality* occurs when a trusted data holder provides information about a data subject. An example would be a violation of patient-physician confidentiality. *Disclosure* involves not a violation of trust as with confidentiality, but rather the making of private information known outside the group of individuals who are expected to know it. *Exposure* occurs when embarrassing but trivial information is shared stripping the data subject of his dignity. *Distortion* is the presentation of false information about a person. Distortion harms not only the subject, whose reputation is damaged, but also third parties who are no longer able to accurately judge the subject's character. *Appropriation* is related to distortion. Appropriation associates a person with a cause or product that he did not agree to endorse. Appropriation adversely affects the ability of the person to present himself as he chooses. *Increased accessibility* occurs when a data holder makes previously available information more easily acquirable. It is a threat to privacy as it makes possible uses of the information that were previously too inefficient, and furthermore, potentially encourage unintended secondary uses. Rather than disseminating information, *black-mail* involves the threat of disseminating information unless some demand is met. It uses private information to create an inappropriate power relation with no social benefits.

These types of violations exist independent of technologies. However, technology plays a dual role in privacy. On the one hand, new technologies can create new ways of infringing upon privacy rights. On the other hand, new technologies can create new ways of preserving privacy.

3 Technology Raises New Privacy Concerns

Technological advances normally represent progress. The utility of these advances, however, must be balanced against any new privacy concerns they create. This tension forces society to examine how a new technology could affect privacy and how to mitigate any ill effects.

The courts often lead this examination. The first important U.S. law review article on privacy, Warren and Brandeis's "The Right to Privacy," was written in response to the ability of new cameras to take pictures quickly enough to capture images of unwilling subjects [76]. The advent of wire tapping technology led first to its acceptance [60] and then to its rejection [62] by the U.S. Supreme Court as its understanding of the technology, people's uses of phones, and government's obligations to privacy changed. Other new forms of surveillance including aerial observation [68, 69], tracking devices [65, 66], hidden video cameras [47], and thermal imaging [70] have all also been studied by courts in the U.S.

New technology has driven governments to create new regulations. The rise of large computer databases with new aggregation abilities led to the U.S. Federal Trade Commission's Fair Information Practice Principles requiring security and limiting secondary uses and exclusion [57]. In France, the public outcry over a proposal to create an aggregate government database, the System for Administrative Files Automation and the Registration of Individuals (SAFARI), forced the government to create the National Data Processing and Liberties Commission (CNIL), an independent regulatory agency. The rise of electronic commerce and the privacy concerns it created resulted in Canada's Personal Information Protection and Electronic Documents Act. Privacy concerns about electronic health records lead to the Privacy Rule under the Health Insurance Portability and Accountability Act (HIPPA) in the U.S. to mixed results [23]. Each of these regulations is designed

to allow new technologies to be used, but not in ways that could violate privacy.

Society is still forming its response to some new technologies. For example, data mining, one technique used for aggregation, has received a mixed reaction. In the U.S., the Total Information Awareness data mining program was largely shut down by Congress, only to be followed by the Analysis, Dissemination, Visualization, Insight and Semantic Enhancement (ADVISE) system, also shut down. However, rather than banning the practice, the Federal Agency Data Mining Reporting Act of 2007 requires agencies to report on their uses of data mining to Congress. Apparently, Congress has not come to a consensus on how to limit data mining and is still studying the concern on a case by case basis.

4 Technology Helps Preserve Privacy

Some of the new threats to privacy created by technology cannot efficiently or effectively be addressed by government action alone. Further technological advances can in some cases provide ways to mitigate these new threats.

In this section, we first give a quick tour through many different technical approaches used to complement or to reinforce non-technical approaches to preserving privacy (Section 4.1), and then focus in detail on two related classes of privacy violations, *disclosure and aggregation*, which have garnered the most attention recently from the computer science community (Section 4.2). We save till Section 5 our discussion of the role that formal methods, as a class of technology, can play in privacy.

4.1 A Diversity of Technical Approaches

While a government may legislate punishment for breaching the security of computer systems storing private records, such punishments can at best only dissuade criminals; they do not prevent privacy violations in any absolute sense. Cryptographic-based technologies with provably secure properties (e.g., one-time pads that guarantee perfect secrecy) or systems that have been formally verified with respect to a given security property (e.g., secure operating systems kernels [5, 58, 34]) can actually make some violations impossible. Likewise, identity theft laws might discourage the practice, but digital signatures can prevent appropriation [19, 55]. Even security technologies, such as intrusion detection systems and spam filters, which may not have provably secure properties, are indispensable in practice for mitigating attacks of intrusion.

In some cases, a data subject might not trust the government or third-party data holders to prevent a violation. For example, political bosses or coercive agents might attempt to learn for which candidate someone voted. In such cases, voting schemes that inherently prevent the disclosure of this information, even to election officials, would be more trustworthy; such schemes have been developed using cryptography (e.g., [8, 4]) or paper methods inspired by cryptography [56]. Political dissidents who wish to hide their online activities can use onion routing, based on repeated encryption, for anonymous Internet use [30]. Privacy preserving data mining (e.g., [75]) offers the government a way of finding suspicious activities without giving it access to private information [45, 6]. *Vanishing data* guarantees data subjects that their private data stored in the “cloud”

be permanently unreadable at a specific time; this recent work by Geambasu et al. [26] relies on public-key cryptography, Shamir’s secret sharing scheme, and the natural churn of distributed hash tables in the Internet.

Mathematical formulations of different notions of privacy are also useful for guiding the development of privacy preserving technologies and making it easier to identify privacy violations. Halpern and O’Neill formalize privacy relevant concepts such as secrecy and anonymity using logics of knowledge [31]. In response to Gavison’s desire for “protection from being brought to the attention of others” [25], Chawla et al. formalize a notion of an individual’s record being conspicuously different from the other records in a set [9]; they characterize this notion in terms of high-dimensional spaces over the reals.

4.2 A Heightened Focus on Disclosure and Aggregation

As Solove notes, aggregation can violate privacy [59]. The form of aggregation Solove describes is when the data holder combines data from multiple sources. Another form of aggregation occurs when the data holder publishes a seemingly harmless data set and an adversary combines this data set with others to find out information that the data holder did not intend to be learned. In this case, the adversary commits the violation of aggregation, but the data holder inadvertently commits the violation of disclosure. Thus, a responsible data holder must ensure that any data he releases cannot be aggregated by others to learn private information.

In the context of databases and anonymization, researchers have studied a special case of the above attack, called *linkage attacks*. In its simplest form, a collection of records, each about an individual, is anonymized by removing any explicit identifiers, such as names or IP addresses. After a data holder releases the anonymized database, an adversary compares it to another database that is not anonymized but holds information about some of the same people in the anonymized database. If one database holds a record r_1 and the second database holds a record r_2 such that r_1 and r_2 agree on values of attributes tracked by both databases, then the adversary can infer that the two records, r_1 and r_2 , refer to the same person with some probability. For example, suppose we know a person, Leslie, is in two databases: one lists him as the only person who has the zip code 15217 and who is male; the anonymized one contains only one person who has the zip code 15217 and is male, and furthermore this person has AIDS. We may conclude that Leslie has AIDS. This attack works despite the first database listing no private information (presuming that one’s zip code and gender are not private) and the second attempting to protect privacy by anonymization.

In light of the 2006 release of AOL search data, attempts to anonymize search query logs have shown they are prone to linkage and other attacks as well (e.g., see [32, 36]). In the same year Netflix released an anonymized database of rented movies for its Netflix Prize competition; Narayanan and Shmatikov showed how to use a linkage-based attack to identify subscriber records in the database, and thus discover people’s political preferences and other sensitive information [43].

A variety of attempts have been made to come up with anonymization approaches not subject to this weakness. One such approach, k -Anonymity, places additional syntactic requirements on the anonymized database [72]. However, for some databases, this approach failed to protect against slightly more complicated versions of the linkage attack. While further work has ruled out some of these attacks (e.g., [38, 37, 77]), no robust, compositional approach has been found.

A different approach comes from the statistics community. *Statistical disclosure limitation* attempts to preserve privacy despite releasing statistics. (For an overview see [22].) Two methods in this line of work are based on releasing tables of data, where entries in the table are either frequencies (counts), e.g., the number of respondents with the same combination of attributes, or magnitudes, the aggregate of individual counts. A third method uses microdata, a sanitization of individual responses. The public is most familiar with these statistical approaches since they are the basis for publishing census data, performing medical studies, and conducting consumer surveys. Surveyors collect information on a large number of individuals and only release aggregations of responses. These aggregations provide statistically significant results about the problem at hand (e.g., the efficacy of a new pharmaceutical) while not including information that an adversary may use to determine the responses of any of the individual respondents.

A more semantic approach originates with Dalenius. He proposed the requirement that an adversary with the aggregate information learns nothing about any of the data subjects that he could not have known without the aggregate information [17]. Unfortunately, Dwork proves that if a data holder provides the exact value of a “useful” aggregate (where “useful” is measured in terms of a utility function), it is impossible for Dalenius’s requirement to hold [20]. Fortunately, she with others showed that by adding noise to the value of the statistic, an adversary could be kept from learning much information about any one individual, leading to the formal definition of *differential privacy* [21]. This formal work on differential privacy inspired practical applications such as the Privacy Integrated Queries (PINQ) system, an API for querying SQL-like databases [42], and an algorithm for releasing query click graphs [35].

Differential privacy is theoretical work, complete with formal definitions, theorems explaining its power, and provable guarantees for systems developed to satisfy it [20]. While PINQ was developed with the specification of differential privacy in mind, the development exemplifies “formal methods light” with no attempt to verify formally that the resulting system satisfies the specification. This line of work on differential privacy could benefit from formal methods that enables such verification.

5 Opportunities and Challenges for Formal Methods

Formal methods can and should be applied to privacy; however, the nature of privacy offers new challenges, and thus new research opportunities, for the formal methods community.

We start in Section 5.1 with our traditional tools of the trade, and for each, hint at some new problems privacy raises. We then point out in Section 5.2 privacy-specific needs, exposing new territory for the formal methods community to explore.

5.1 Formal Methods Technology

All the machinery of the formal methods community can help us gain a more rigorous understanding of privacy rights, threats, and violations. We can use formal models, from state machines to process algebras to game theory, to model the behavior of the system and its threat environment.

We can use formal logics and formal languages to state different aspects of privacy, to state desired properties of these systems, to state privacy policies, to reason about when a model satisfies a property or policy, and to detect inconsistencies between different privacy policies. Automated analyses and tools enable us to scale the applicability of these foundational models and logics to realistic systems. Privacy does pose new challenges, requiring possibly new models, logics, languages, analyses, and tools.

Models

In formal methods, we traditionally model a system and its environment and the interactions between the two. Many methods may simply make assumptions about the environment in which the system operates, thus focusing primarily on modeling the system. To model failures, for example, due to natural disasters or unforeseen events, we usually can get away with abstracting from the different classes of failures and model a single failure action (that could occur at any state) or a single failure state.

Security already challenges this simplicity in modeling. We cannot make assumptions about an adversary the way we might about hardware failures or extreme events like hurricanes. On the other hand, it often suffices to include the adversary as part of the system's environment, and assume the worst case (e.g., treating an adversary's action as a Byzantine failure).

Privacy may require yet a new approach to or at least a new outlook on modeling. Privacy involves three entities: the data holder (system), an adversary (part of the environment), and the data subject. Consider this difference between security and privacy: In security, the entity in control of the system also has an inherent interest in its security. In privacy, the system is controlled by the data holder, but it is the data subject that benefits from privacy. Formal methods akin to proof-carrying code [46], which requires the data holder to provide an easy-to-check certificate to the data subject, might be one way to address this kind of difference.

Privacy requires modeling different relationships among the (minimally) three entities. Complications arise because relationships do not necessarily enjoy simple algebraic properties and because relationships change over time. For example if person X trusts Y and Y trusts Z that does not mean X trusts Z . X needs to trust that Y will not pass on any information about X to Z . Moreover, if X eventually breaks his trust relation with Y then X would like Y to forget all the information Y had about X . This problem is similar to revoking access rights in security except that instead of removing the right to access information (knowledge about X), it is the information itself that is removed.

Logics

The success of many formal methods rests on decades of work on defining and applying logics (e.g., temporal logics) for specifying and reasoning about system behavior. Properties of interest, which drive the underlying logics needed to express them, are often formulated as assertions over traces (e.g., sequences of states, sequences of state transitions, or sequences of alternating states and transitions).

McLean, however, shows that a class of information-flow properties cannot be expressed as trace properties [41]. In particular, *non-interference*, which characterizes when no information flows from a high-level (e.g., top secret) subject to a low-level (e.g., public) subject [29], cannot

be expressed as a property over a single trace. Non-interference formalizes the notion of keeping secure information secret from an adversary. Since secrecy is often a starting point for thinking about privacy, we will likely need new logics for specifying and reasoning about such non-trace properties and other privacy properties more generally.

Formal Policy Languages

The privacy right of exclusion requires that data subjects know how their information will be used. Thus, data holders must codify their practices into publicly available privacy policies. While most of these policies are written in natural language, some attempts have been made to express them in machine readable formats. For example, EPAL is a language for expressing policies with the intention of allowing automated enforcement [52]. Other policy languages such as P3P [15], which has a formal notation, inform website visitors of the site's privacy practices and enable automated methods for finding privacy-conscientious sites [16]. These languages, however, lack formal semantics.

Barth et al. do provide a formal language for specifying notions expressed in privacy policies such as HIPAA, the Children's Online Privacy Protection Act, and the Gramm-Leach-Bliley Act (about financial disclosures) [3]. Their language uses traditional linear temporal logic and its semantics is based on a formal model of *contextual integrity*, Nissenbaum's philosophical theory of information dissemination [50]. Much work remains in extending such formal languages to handle more forms of privacy.

Abstraction and Refinement

Formal methods have been particularly successful at reasoning above the level of code. That success, however, relies fundamentally on abstraction and/or refinement. Commuting diagrams allow us to abstract from the code and do formal reasoning at higher levels of description, but these diagrams rely on well-defined abstraction functions or refinement relations. Similarly, methods that successively refine a high-level specification to a lower-level one, until executable code is reached, rely on well-defined correctness-preserving transformations.

As discussed above, some privacy relevant properties, such as secrecy, are not trace properties. Furthermore, while a specification may satisfy a secrecy property, a refinement of the specification might not. Mantel [39], Jürjens [33], and Alur et al. [1] define specialized forms of refinement that preserve such secrecy properties. Similarly, Clarkson and Schneider [12] develop a theory of *hyperproperties* (sets of properties), which can express information-flow properties, and characterize a set of hyperproperties for which refinement is valid. These works just begin to address aspects of privacy; attention to other aspects may require new abstraction and/or refinement methods.

Policy Composition

Given that different components of a system might be governed by different policies or that one system might be governed by more than one policy, we must also provide methods of compositional reasoning: Given two components, A and B , and privacy policies, P_1 and P_2 , if A satisfies P_1 and B satisfies P_2 , what does that say about the composition of A and B with respect to P_1 , P_2 , and $P_1 \wedge P_2$? Privacy policies are likely in practice not to be compositional. For example, the National Science Foundation has a privacy policy that says reviewers of each grant proposal must remain anonymous to the grant proposers; the National Institutes of Health has a different review

policy where the names of the study (review) group members are known to the grant proposers. For NSF and NIH to have a joint program, therefore, some compromise between the policies needs to be made, while still preserving “to some degree” the spirit of both policies. This general challenge of composition already exists for other properties such as serializability in databases, feature interaction in telephone services, and noninterference in security. Privacy adds to this challenge.

Code-level Analysis

Formal methods, especially when combined with static analysis techniques, have been successful at finding correctness bugs (e.g., [2]) and security vulnerabilities (e.g., [14, 48]) at the code level. What kind of code-level reasoning could we do for privacy, either to prove that a privacy policy is preserved or to discover a privacy violation?

Automated Tools

One of the advantages of formal methods is that formal specifications are amenable to machine manipulation and machine analysis (e.g., finding bugs or proving properties). Automation not just helps us catch human errors, but also enables us to scale up pencil-and-paper techniques.

We need to explore the use of and extensions required for formal methods tools, such as theorem provers and models checkers, for verifying privacy policies or discovering privacy violations. While much foundational work in terms of models, logics, and languages remain, none will become of practical import unless our automated analysis tools scale to work for realistic systems.

5.2 Privacy-Specific Needs

Statistical/Quantitative Reasoning

The statistical nature of privacy raises a new challenge for formal methods. For example, aggregating the weights of a large number of individuals into the average weight is expected to make it difficult for an adversary to learn much about any one of the individuals. Thus, this form of aggregation can protect the private information (individual weights) while still providing a useful statistic (the average weight). In security, information flow is viewed as black and white: if a flow occurs from high to low, a violation has occurred. In privacy, a “small” amount of flow may be acceptable since we are unlikely to learn a lot about the weight of any one person from learning the average of many. While some work has been done on *quantitative* information flow (e.g., [11, 10, 40, 49]), even the tools developed from this work would consider the system as violating security (see [74] for why and an approach that does not), and thus would be inappropriate for a statistical notion of privacy.

More generally, formal methods may need to be extended to assure statistical guarantees rather than our traditional black-and-white correctness guarantees. A hybrid approach would be to combine traditional formal models with statistical models or formal methods with statistical methods.

Trustworthy Computing: Conflicting Requirements

While trade-offs are hardly new to computer science, privacy raises a new set of such trade-offs. Trustworthy computing requires balancing privacy with security, reliability, and usability. It would be good to have a formal understanding of the relationships among these properties. For example, we want auditability for security, to determine the source of a security breach. However,

auditability is at odds with anonymity, a desired aspect of privacy. Thus, to what degree can we provide auditability while providing some degree of anonymity? (This is not suggest that security and privacy are opposites: security is necessary for privacy.) To achieve reliability, especially availability, we often replicate data at different locations; replicas increase the likelihood that an attacker can access private data and make it harder for users to track and manage (e.g., delete) their data. Trade-offs between privacy and usability are similar to those between security and usability. We want to allow users to control how much of their information is released to others, but we want to make it easy for them to specify this control, and even more challenging, to understand the implications of what they specify and to be able to change the specifications over time.

6 Summary

Privacy touches the philosophy, legal, political, social science, and technical communities. Technical approaches to privacy must be part of the basis in creating privacy laws and in designing privacy regulations. Laws and policies need to be technically feasible to implement.

In this paper we focused on the dual role of technology in this vast privacy space: new technologies cause us to revisit old laws or create new ones; at the same time, advances in technology can help preserve privacy rights or mitigate consequences of privacy violations.

Formal methods is a technology that can help by providing everything from foundational formalizations of privacy to practical tools for checking for privacy violations. However, we have barely begun to use formal methods to study privacy in depth; we hope the community is ready to rise to the challenge.

Acknowledgments

Our understanding of privacy has benefited from conversations with Anupam Datta, Dilsun Kaynar, and Jennifer Tam.

References

- [1] Rajeev Alur, Pavol Černý, and Steve Zdancewic. Preserving secrecy under refinement. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming*, 2006.
- [2] Tom Ball and Sriram Rajamani. Automatically validating temporal safety properties of interfaces. In *Proceedings of the SPIN 2001 Workshop on Model Checking of Software*, volume 2057 of *Lecture Notes in Computer Science*, May 2001.
- [3] Adam Barth, Anupam Datta, John C. Mitchell, and Helen Nissenbaum. Privacy and contextual integrity: Framework and applications. In *SP '06: Proceedings of the 2006 IEEE Symposium on Security and Privacy*, pages 184–198, Washington, DC, USA, 2006. IEEE Computer Society.

- [4] Josh Benaloh and Dwight Tuinstra. Receipt-free secret ballot elections. In *Proceedings of the 26th Annual ACM symposium on Theory of Computing*, Montreal, Canada, 1994.
- [5] Terry V. Benzel. Analysis of a kernel verification. In *Proceedings of the IEEE Symposium on Security and Privacy*, 1984.
- [6] Brian Bergstein. Research explores data mining, privacy. *USA Today*, June 18 2008.
- [7] E. Bloustein. Privacy as an aspect of human dignity: An answer to dean prosser. *New York University Law Review*, 39:962, 1964.
- [8] D. Chaum. Secret ballot receipts: True voter-verifiable elections. *IEEE J. Security and Privacy*, pages 38–47, 2004.
- [9] Shuchi Chawla, Cynthia Dwork, Frank McSherry, Adam Smith, and Hoeteck Wee. Toward privacy in public databases. In *2nd Theory of Cryptography Conference (TCC 2005)*, pages 363–385, 2005.
- [10] David Clark, Sebastian Hunt, and Pasquale Malacaria. A static analysis for quantifying information flow in a simple imperative language. *Journal of Computer Security*, 15:321–371, 2007.
- [11] Michael R. Clarkson, Andrew C. Myers, and Fred B. Schneider. Belief in information flow. In *CSFW '05: Proceedings of the 18th IEEE workshop on Computer Security Foundations*, pages 31–45, Washington, DC, USA, 2005. IEEE Computer Society.
- [12] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. In *Proceedings of IEEE Computer Security Foundations Symposium*, June 2008.
- [13] J. Cohen. *Regulating Intimacy: A New Legal Paradigm*. Princeton University Press, Princeton, 2002.
- [14] Crispin Cowan, Perry Wagle, Calton Pu, Steve Beattie, and Jonathan Walpole. Buffer overflows: Attacks and defenses for the vulnerability of the decade. In *SANS 2000*, 1999.
- [15] Lorrie Faith Cranor. *Web Privacy with P3P*. O'Reilly, September 2002.
- [16] Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. User interfaces for privacy agents. *ACM Trans. Comput.-Hum. Interact.*, 13(2):135–178, 2006.
- [17] T. Dalenius. Towards a methodology for statistical disclosure control. *Statistik Tidskrift*, 15:429–444, 1977.
- [18] danah boyd. Why youth (heart) social network sites: The role of networked publics in teenage social life. In David Buckingham, editor, *MacArthur Foundation Series on Digital Learning—Youth, Identity, Digital Media Volume*, Cambridge, MA, 2007. MIT Press.

- [19] W. Diffie and M. Hellman. New directions in cryptography. *Information Theory, IEEE Transactions on*, 22(6):644–654, Nov 1976.
- [20] Cynthia Dwork. Differential privacy. In *33rd International Colloquium on Automata, Languages and Programming (ICALP 2006)*, volume 2, pages 1–12, 2006.
- [21] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Proceedings of the 3rd Theory of Cryptography Conference*, pages 265–284. Springer, 2006.
- [22] Federal Committee on Statistical Methodology. Statistical disclosure limitation methodology. Statistical Policy Working Paper 22, 2005.
- [23] Theo Francis. Spread of records stirs fears of privacy erosion. *The Wall Street Journal*, December 28 2006.
- [24] C. Fried. *An Anatomy of Values*. Harvard University Press, Cambridge, Mass., 1970.
- [25] Ruth Gavison. Privacy and the limits of law. *Yale Law Journal*, 89(3):421–471, January 1980.
- [26] Roxana Geambasu, Tadayoshi Kohno, Amit Levy, and Henry M. Levy. Vanish: Increasing data privacy with self-destructing data. In *Proceedings of the USENIX Security Symposium*, Montreal, Canada, August 2009.
- [27] T. Gerety. Redefining privacy. *Harvard Civil Rights-Civil Liberties Law Review*, 12:233–296, 1977.
- [28] R. Gerstein. Intimacy and privacy. *Ethics*, 89:76–81, 1978.
- [29] Joseph A. Goguen and Jose Meseguer. Security policies and security models. In *Proceedings of the IEEE Symposium on Security and Privacy*, 1982.
- [30] David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Onion routing. *Commun. ACM*, 42(2):39–41, 1999.
- [31] Joseph Halpern and Kevin O’Neill. Secrecy in multiagent systems. In *CSFW ’02: Proceedings of the 15th IEEE workshop on Computer Security Foundations*, pages 32–46, Washington, DC, USA, 2002. IEEE Computer Society. A longer version available at <http://www.kevinoneill.org/papers/secrecy.pdf>.
- [32] Rosie Jones, Ravi Kumar, Bo Pang, and Andrew Tomkins. ”I Know What You Did Last Summer”: Query Logs and User Privacy. In *Proceedings of the Sixteenth ACM Conference on Information and Knowledge Management*, Lisbon, Portugal, 2007.
- [33] Jan Jürjens. Secrecy-preserving refinement. In *FME ’01: Proceedings of the International Symposium of Formal Methods Europe on Formal Methods for Increasing Software Productivity*, pages 135–152, London, UK, 2001. Springer-Verlag.

- [34] Gerwin Klein, Kevin Elphinstone, Gernot Heiser, June Andronick, David Cock, Philip Derin, Dhammika Elkaduwe, Kai Engelhardt, Rafal Kolanski, Michael Norrish, Thomas Sewell, Harvey Tuch, and Simon Winwood. sel4: Formal verification of an os kernel. In *Proceedings of the 22nd ACM Symposium on Operating Systems Principles*, Big Sky, Montana, October 2009.
- [35] Aleksandra Korolova, Krishnaram Kenthapadi, Nina Mishra, and Alexandros Ntoulas. Releasing search queries and clicks privately. In *Proceedings of the 2009 International World Wide Web Conference*, Madrid, Spain, 2009.
- [36] Ravi Kumar, Jasmine Novak, Bo Pang, and Andrew Tomkins. On anonymizing query logs via token-based hashing. In *Proceedings of the 16th International Conference on World Wide Web*, Banff, Alberta, Canada, 2007.
- [37] Ninghui Li, Tiancheng Li, and S. Venkatasubramanian. t -closeness: Privacy beyond k -anonymity and l -diversity. *Data Engineering, 2007. ICDE 2007. IEEE 23rd International Conference on*, pages 106–115, 15-20 April 2007.
- [38] Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke, and Muthuramakrishnan Venkatasubramanian. ℓ -Diversity: Privacy beyond k -anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1):3, 2007.
- [39] Heiko Mantel. Preserving information flow properties under refinement. In *SP '01: Proceedings of the 2001 IEEE Symposium on Security and Privacy*, page 78, Washington, DC, USA, 2001. IEEE Computer Society.
- [40] Stephen McCamant and Michael D. Ernst. A simulation-based proof technique for dynamic information flow. In *PLAS '07: Proceedings of the 2007 workshop on Programming languages and analysis for security*, pages 41–46, New York, NY, USA, 2007. ACM.
- [41] John McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *SP '94: Proceedings of the 1994 IEEE Symposium on Security and Privacy*, page 79, Washington, DC, USA, 1994. IEEE Computer Society.
- [42] Frank McSherry. Privacy integrated queries: An extensible platform for privacy-preserving data analysis. In *SIGMOD '09: Proceedings of the 2009 ACM SIGMOD international conference on Management of data*, New York, NY, USA, 2009. ACM. To appear. Available at <http://research.microsoft.com/apps/pubs/?id=80218>.
- [43] Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *SP '08: Proceedings of the 2008 IEEE Symposium on Security and Privacy*, pages 111–125, Washington, DC, USA, 2008. IEEE Computer Society.
- [44] National Research Council. Engaging privacy and information technology in a digital age. In James Waldo, Herbert S. Lin, and Lynette I. Millett, editors, *National Research Council of the National Academies*, Washington, D.C., 2007. The National Academies Press.

- [45] National Research Council. *Protecting Individual Privacy in the Struggle Against Terrorists*. The National Academies Press, Washington, D.C., 2008.
- [46] George C. Necula and Peter Lee. Safe kernel extensions without run-time checking. *SIGOPS Oper. Syst. Rev.*, 30(SI):229–243, 1996.
- [47] New Hampshire Supreme Court. Hamberger v. Eastman. *Atlantic Reporter*, 206:239, 1964.
- [48] James Newsome and Dawn Song. Dynamic taint analysis: Automatic detection, analysis, and signature generation of exploit attacks on commodity software. In *Network and Distributed Systems Security Symposium*, February 2005.
- [49] James Newsome and Dawn Song. Influence: A quantitative approach for data integrity. Technical Report CMU-CyLab-08-005, CyLab, Carnegie Mellon University, February 2008.
- [50] H. Nissenbaum. Privacy as contextual integrity. *Washington Law Review*, 79(1), 2004.
- [51] W. Parent. Privacy, morality and the law. *Philosophy and Public Affairs*, 12:269–288, 1983.
- [52] Calvin Powers and Matthias Schunter. Enterprise privacy authorization language (EPAL 1.2). W3C Member Submission, November 2003.
- [53] William L. Prosser. Privacy. *California Law Review*, 48:383, 1960.
- [54] J. Rachels. Why privacy is important. *Philosophy and Public Affairs*, 4:323–333, 1975.
- [55] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM*, 21(2):120–126, 1978.
- [56] Ronald L. Rivest and Warren D. Smith. Three voting protocols: Threeballot, vav, and twin. In *EVT’07: Proceedings of the USENIX Workshop on Accurate Electronic Voting Technology*, pages 16–16, Berkeley, CA, USA, 2007. USENIX Association.
- [57] Secretary’s Advisory Committee on Automated Personal Data Systems. Records, computers, and the rights of citizens. Technical report, U.S. Department of Health, Education, and Welfare, July 1973.
- [58] Jonathan Silverman. Reflections on the verification of the security of an operating system kernel. In *Proceedings of the Ninth ACM Symposium on Operating Systems Principles*, Bretton Woods, New Hampshire, 1983.
- [59] Daniel J. Solove. A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3):477–560, January 2006.
- [60] Supreme Court of the United States. Olmstead v. United States. *United States Reports*, 277:438, 1928.

- [61] Supreme Court of the United States. *Griswold v. Connecticut*. *United States Reports*, 381:479, 1965.
- [62] Supreme Court of the United States. *Katz v. United States*. *United States Reports*, 389:347, 1967.
- [63] Supreme Court of the United States. *Eisenstadt v. Baird*. *United States Reports*, 405:438, 1972.
- [64] Supreme Court of the United States. *Roe v. Wade*. *United States Reports*, 410:113, 1973.
- [65] Supreme Court of the United States. *United States v. Knotts*. *United States Reports*, 460:276, 1983.
- [66] Supreme Court of the United States. *United States v. Karo*. *United States Reports*, 468:705, 1984.
- [67] Supreme Court of the United States. *Bowers v. Hardwick*. *United States Reports*, 478:186, 1986.
- [68] Supreme Court of the United States. *Dow Chemical Co. v. United States*. *United States Reports*, 476:227, 1986.
- [69] Supreme Court of the United States. *Florida v. Riley*. *United States Reports*, 488:455, 1989.
- [70] Supreme Court of the United States. *Kyllo v. United States*. *United States Reports*, 533:27, 2001.
- [71] Supreme Court of the United States. *Lawrence v. Texas*. *United States Reports*, 538:918, 2003.
- [72] Latanya Sweeney. *k*-Anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, 2002.
- [73] J. Thomson. The right to privacy. *Philosophy and Public Affairs*, 4:295–314, 1975.
- [74] Michael Carl Tschantz and Aditya V. Nori. Measuring the loss of privacy from statistics. In Sumit Gulwani and Sanjit A. Seshia, editors, *Proceedings of the 1st Workshop on Quantitative Analysis of Software (QA'09)*, Technical Report UCB/EECS-2009-93, Electrical Engineering and Computer Sciences, University of California at Berkeley, pages 27–36, June 2009.
- [75] V. S. Verykios, E. Bertino, I. N. Fovino, L. P. Provenza, Y. Saygin, and Y. Theodoridis. State-of-the-art in privacy preserving data mining. *ACM SIGMOD Record*, 3(1):50–57, March 2004.
- [76] Warren and Brandeis. The right to privacy. *Harvard Law Review*, IV(5), December 1890.

- [77] Xiaokui Xiao and Yufei Tao. *m*-Invariance: Towards privacy preserving re-publication of dynamic datasets. In *SIGMOD '07: Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, pages 689–700, New York, NY, USA, 2007. ACM Press.