Privacy through Accountability: A Computer Science Perspective*

Anupam Datta

Computer Science Department
Electrical and Computer Engineering Department
CyLab
Carnegie Mellon University

Abstract. Privacy has become a significant concern in modern society as personal information about individuals is increasingly collected, used, and shared, often using digital technologies, by a wide range of organizations. To mitigate privacy concerns, organizations are required to respect privacy laws in regulated sectors (e.g., HIPAA in healthcare, GLBA in financial sector) and to adhere to self-declared privacy policies in self-regulated sectors (e.g., privacy policies of companies such as Google and Facebook in Web services). This article provides an overview of a body of work on formalizing and enforcing privacy policies. We formalize privacy policies that prescribe and proscribe *flows* of personal information as well as those that place restrictions on the *purposes* for which a governed entity may use personal information. Recognizing that traditional preventive access control and information flow control mechanisms are inadequate for enforcing such privacy policies, we develop principled *accountability* mechanisms that seek to encourage policy-compliant behavior by detecting policy violations, assigning blame, and punishing violators. We apply these techniques to several U.S. privacy laws and organizational privacy policies, in particular, producing the first complete logical specification and audit of all disclosure-related clauses of the HIPAA Privacy Rule.

1 Introduction

Privacy has become a significant concern in modern society as personal information about individuals is increasingly collected, used, and shared, often using digital technologies, by a wide range of organizations. Certain information handling practices of organizations that monitor individuals' activities on the Web, data aggregation companies that compile massive databases of personal information, cell phone companies that collect and use location data about individuals, online social networks and search engines—while enabling useful services—have aroused much indignation and protest in the name of privacy (see, for example, a series of articles in the Wall Street Journal [21]). Similarly, as healthcare organizations are embracing electronic health record systems and patient portals to enable patients, employees, and business affiliates more efficient access to personal health information, there is trepidation that the privacy of patients may not be adequately protected if information handling practices are not carefully designed and enforced [11, 14, 19]. To mitigate privacy concerns, organizations are required to respect privacy laws in regulated sectors (e.g., HIPAA in healthcare, GLBA in financial sector) and to adhere to self-declared privacy policies in self-regulated sectors (e.g., privacy policies of companies such as Google and Facebook in Web services).

This article provides an overview of a body of work on formalizing and enforcing practical privacy policies using computational techniques [1, 2, 4–9, 16–18] conducted jointly with my students, postdoctoral researchers, and colleagues at Carnegie Mellon, Stanford, and New York

* This work was partially supported by the NSF Science and Technology Center TRUST, the NSF Trustworthy Computing grant "Privacy Policy Specification and Enforcement: Information Use and Purpose," and HHS Grant no. HHS 90TR0003/01. 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.

University. We find that one significant difference from traditional security settings is that the enforcement mechanisms in privacy settings often have only *black-box access* to the programs and people who operate on personal information. For example, a class of privacy threats in hospitals arises from authorized insiders (e.g., doctors, nurses, administrative staff) who have a legitimate right to access personal information, but may abuse that right to inappropriately share and use that information; an enforcement mechanism employed by the hospital can observe the behavior of authorized insiders as recorded on audit logs, but does not have access to the programs (algorithms) running inside their minds. Similarly, a Web user or privacy advocacy group interested in checking if Google is using sensitive information, such as race, for advertising can interact with Google's program over the Web by supplying different kinds of information to it and observing the displayed ads, but will typically not have access to the code for Google's advertising program. Thus, my research program has focused on principled audit and accountability mechanisms for enforcing privacy properties by detecting policy violations, assigning blame and optimally managing risks stemming from privacy violations. These mechanisms operate with black-box models of the systems (programs and people) that operate over personal information.

The rest of the paper is organized as follows. Section 2 provides an overview of contextual integrity—a normative theory of privacy—and a logic of privacy that we developed informed by this theory. We used this logic to produce the first complete formalization of the HIPAA Privacy Rule and the Gramm-Leach-Bliley Act. Section 3 provides an overview of our algorithm for checking incomplete audit logs for compliance with policies expressed in the logic. This algorithm automatically checks some parts of privacy policies (e.g., pertaining to temporal conditions) and outputs other parts (e.g., pertaining to purposes and beliefs) in a residual policy that has to be checked by other means. Section 4 describes our work on formalizing and enforcing purpose restrictions in privacy policies. Finally, Section 5 describes our work on audit algorithms that prescribe effective resource allocation strategies for auditors interacting with byzantine and strategic adversaries.

2 Contextual Integrity and Logic of Privacy

The central thesis of contextual integrity is that *privacy is a right to appropriate flow of personal information* [13]. The building blocks of this theory are *social contexts* and *context-relative informational norms*. A context captures the idea that people act and transact in society not simply as individuals in an undifferentiated social world, but as individuals in certain capacities (roles) in distinctive social contexts, such as healthcare, education, friendship and employment. Norms prescribe the flow of personal information in a given context, e.g., in a healthcare context a norm might prescribe flow of personal health information from a patient to a doctor and proscribe flows from the doctor to other parties who are not involved in providing treatment. Norms are a function of the following parameters: the respective roles of the sender, the subject, and the recipient of the information, the type of information, and the principle under which the information is sent to the recipient. Examples of transmission principles include confidentiality (prohibiting agents receiving the information from sharing it with others), reciprocity (requiring bi-directional information flow, e.g., in a friendship context), consent (requiring permission from the information subject before transmission), and notice (informing the information subject that a transmission has occured). When norms are contravened, people experience a violation of privacy. This theory has been used to explain why a number of technology-based systems and practices threaten privacy by violating entrenched informational norms. In addition, it provides a prescriptive method for determining appropriate norms for a context (see [13]). This theory is now well known in the privacy community and has influenced privacy policy in the US (for example, 'respect for context' was included in the Consumer Privacy Bill of Rights released by the White House in 2012 [10]).

The idea that privacy expectations can be stated using context-relative informational norms is formalized in a *semantic model* and *logic of privacy* proposed with colleagues at Stanford and New York University [1] and developed further in follow-up work with my students and

postdoctoral researchers [8]. At a high-level, the model consists of a set of interacting agents in roles who perform actions involving personal information in a given context. For example, Alice (a patient) may send her personal health information to Bob (her doctor). Following the structure of context-relative informational norms, each transmission action is characterized by the roles of the sender, subject, receipient and the type of the information sent. Interactions among agents give rise to *traces* where each trace is an alternating sequence of states (capturing roles and knowledge of agents) and actions performed by agents that update state (e.g., an agent's knowledge may increase upon receiving a message or his role might change).

Transmission principles prescribe which traces respect privacy and which traces don't. While contextual integrity talks about transmission principles in the abstract, we require a precise logic for expressing them since our goal is to use information processing systems to check for violation of such principles. We were guided by two considerations in designing the logic: (a) *expressivity*—the logic should be able to represent practical privacy policies; and (b) *enforceability*—it should be possible to provide automated support for checking whether traces satisfy policies expressed in the logic.

A logic of privacy that meets these goals is presented in our recent work [9]. We arrive at this enforceable logic by restricting the syntax of the expressive first-order logic we used in our earlier work to develop the first complete formalization of two US privacy laws—the HIPAA Privacy Rule for healthcare organizations and the Gramm-Leach-Bliley Act for financial institutions [8][1]. These comprehensive case studies shed light on common concepts that arise in transmission principles in practice—data attributes, dynamic roles, notice and consent (formalized as temporal properties), purposes of uses and disclosures, and principals' beliefs—as well as how individual transmission principles are composed in privacy policies[2].

3 Policy Auditing over Incomplete Logs

We observe that traditional preventive access control and information flow control mechanisms are not sufficient for enforcing all privacy policies because at run-time there may not be sufficient information to decide whether certain policy concepts (e.g., future obligations, purposes of uses and disclosures, and principals' beliefs) are satisfied or not. We therefore take the position that audit mechanisms are essential for privacy policy enforcement. The importance of audits has been recognized in the computer security literature. For example, Lampson [12] takes the position that audit logs that record relevant evidence during system execution can be used to detect violations of policy, establish accountability and punish the violators. More recently, Weitzner et al. [22] also recognize the importance of audit and accountability, and the inadequacy of preventive access control mechanisms as the sole basis for privacy protection in today's open information environment. However, while the principles of access control and information flow control have been extensively studied, there is comparatively little work on the principles of audit. Our work is aimed at filling this gap.

Our first insight is that *incomplete audit logs* provide a suitable abstraction to model situations (commonly encountered in practice) in which the log does not contain sufficient information to determine whether a policy is satisfied or violated, e.g., because of the policy concepts alluded to earlier—future obligations, purposes of uses and disclosures, and principals' beliefs. We formalize incomplete logs as partial structures that map each atomic formula to true, false or unknown. We design an algorithm, which we name `reduce`, to operate iteratively over such incomplete logs that evolve over time. In each iteration, `reduce` provably checks as much of the policy as possible over the current log and outputs a residual policy that can only be checked when the log is extended with additional information. We implement `reduce` and use it to check simulated audit logs for compliance with the entire HIPAA Privacy Rule. Our experimental results

[1] This logic, in turn, generalizes the enforceable propositional temporal logic in [1].

[2] The model and logic supports information use actions in addition to transmission actions, so, strictly speaking, it can express policies that are more general than transmission principles.

demonstrate that the algorithm scales to realistic audit logs. This technical result is reported in a joint paper with my then postdoctoral researchers D. Garg and L. Jia [9].

4 Formalizing and Enforcing Purpose Restrictions

In recent work, we developed the first formal semantics for privacy policies that place restrictions on the purposes for which a governed entity may use personal information—an important and pervasive class of policies in practice (PhD thesis of M. C. Tschantz co-advised with J. M. Wing) [16, 18]. Purpose occupies a central place in numerous influential privacy guidelines and regulations, including OECDs Privacy Guidelines, the EU Privacy Directive, US privacy laws and organizational privacy policies in sectors as diverse as healthcare, finance, Web services, insurance, education, and government. For example, HIPAA requires that hospital employees use personal health information only for certain purposes (e.g. treatment). We argue that (a) an action is for a purpose if it is part of a plan for achieving that purpose and (b) a piece of information is used for a purpose if it affects the planning process. We model planning using (Partially Observable) Markov Decision Processes and design algorithms for auditing actions of agents by building on algorithms for plan recognition. The algorithms compare logged actions to a model of how an agent attempting to achieve the allowed purpose would plan to do so. If the logged actions differ from the model, the algorithm reports a potential violation.

5 Audit Games

Recognizing that audit mechanisms are constrained by available resources (it may not be possible to inspect every potential violation) and adversaries may adapt to beat them, I have also initiated a formal study of audit games (jointly with my students J. Blocki and A. Sinha, and colleagues N. Christin and A. Procaccia at Carnegie Mellon) that model the interaction between the auditor and auditees as a game. We have developed algorithms for computing optimal audit strategies that prescribe resource allocation for auditing Byzantine [5, 7] and rational auditees [4, 6]. The algorithms advance the state-of-the-art in online learning and algorithmic game theory to address these problems.

Let me highlight one result in this line of work. In [4], we model the audit process as a game between a defender (e.g, a hospital) and an adversary (e.g., an employee). The defender audits a given set of targets (e.g., health record accesses) and the adversary chooses a target to attack. The defender's action space in the audit game includes two components. First, the allocation of its inspection resources to targets; this component also exists in a standard model of physical security games [15]. Second, we introduce a continuous punishment rate parameter that the defender employs to deter the adversary from committing violations. However, punishments are not free and the defender incurs a cost for choosing a high punishment level. For instance, a negative work environment in a hospital with high fines for violations can lead to a loss of productivity (see [3] for a similar account of the cost of punishment). The adversary's utility includes the benefit from committing violations and the loss from being punished if caught by the defender. Our model is parametric in the utility functions. Thus, depending on the application, we can instantiate the model to either allocate resources for detecting violations or preventing them. This generality implies that our model can be used to study all the applications previously described in the security games literature [15]. To analyze the audit game, we use the Stackelberg equilibrium solution concept [20] in which the defender commits to a strategy, and the adversary plays an optimal response to that strategy. This concept captures situations in which the adversary learns the defender's audit strategy through surveillance or the defender publishes its audit algorithm. In addition to yielding a better payoff for the defender than any Nash equilibrium, the Stackelberg equilibrium makes the choice for the adversary simple, which leads to a more predictable outcome of the game. Furthermore, this equilibrium concept respects the computer security principle of avoiding "security through obscurity"— audit mechanisms like

cryptographic algorithms should provide security despite being publicly known. We view this work as a first step toward a computationally feasible model of audit games.

References

[1] Barth, A., Datta, A., Mitchell, J. C., and Nissenbaum, H. Privacy and contextual integrity: Framework and applications. In *Proceedings of the 27th IEEE Symposium on Security and Privacy (Oakland)* (2006), pp. 184–198.

[2] Barth, A., Datta, A., Mitchell, J. C., and Sundaram, S. Privacy and utility in business processes. In *Proceedings of the 20th IEEE Computer Security Foundations Symposium (CSF)* (2007), pp. 279–294.

[3] Becker, G. S. Crime and punishment: An economic approach. *Journal of Political Economy 76* (1968), 169.

[4] Blocki, J., Christin, N., Datta, A., Procaccia, A. D., and Sinha, A. Audit games. In *IJCAI* (2013).

[5] Blocki, J., Christin, N., Datta, A., and Sinha, A. Regret minimizing audits: A learning-theoretic basis for privacy protection. In *Proceedings of the 24th IEEE Computer Security Foundations Symposium (CSF)* (2011), pp. 312–327.

[6] Blocki, J., Christin, N., Datta, A., and Sinha, A. Audit mechanisms for provable risk management and accountable data governance. In *GameSec* (2012), pp. 38–59.

[7] Blocki, J., Christin, N., Datta, A., and Sinha, A. Adaptive regret minimization in bounded-memory games. In *GameSec* (2013). To appear.

[8] DeYoung, H., Garg, D., Jia, L., Kaynar, D., and Datta, A. Experiences in the logical specification of the HIPAA and GLBA privacy laws. In *Proceedings of the 9th Annual ACM Workshop on Privacy in the Electronic Society (WPES)* (2010). Full version: Carnegie Mellon University Technical Report CMU-CyLab-10-007.

[9] Garg, D., Jia, L., and Datta, A. Policy auditing over incomplete logs: Theory, implementation and applications. In *Proceedings of the 18th ACM Conference on Computer and Communications Security (CCS)* (2011).

[10] House, T. W. Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy, February 2012.

[11] Hulme, G. Steady Bleed: State of HealthCare Data Breaches. InformationWeek, September 2010. Available at http://www.informationweek.com/blog/healthcare/229200720.

[12] Lampson, B. W. Computer security in the real world. *IEEE Computer 37*, 6 (2004), 37–46.

[13] Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life.* Stanford University Press, 2010.

[14] Robertson, J. New data spill shows risk of online health records. Yahoo News, August 2011. Available at http://news.yahoo.com/data-spill-shows-risk-online-health-records-120743449.html.

[15] Tambe, M. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned.* Cambridge University Press, 2011.

[16] Tschantz, M. C., Datta, A., and Wing, J. M. Formalizing and enforcing purpose restrictions in privacy policies. In *IEEE Symposium on Security and Privacy* (2012), pp. 176–190.

[17] Tschantz, M. C., Datta, A., and Wing, J. M. Information flow investigations. Technical report cmu-cs-13-118, Carnegie Mellon University, 2013.

[18] Tschantz, M. C., Datta, A., and Wing, J. M. Purpose restrictions on information use. In *ESORICS* (2013), pp. 610–627.

[19] US Health and Human Services. HIPAA enforcement, Accessed September 24,2013. Available at http://www.hhs.gov/ocr/privacy/hipaa/enforcement/index.html.

[20] von Stackelberg, H. *Marktform und Gleichgewicht. - Wien & Berlin: Springer 1934. VI, 138 S. 8.* J. Springer, 1934.

[21] Wall Street Journal. What they know, Accessed on September 24, 2013. Available at http://online.wsj.com/public/page/what-they-know-digital-privacy.html.

[22] Weitzner, D. J., Abelson, H., Berners-Lee, T., Feigenbaum, J., Hendler, J. A., and Sussman, G. J. Information accountability. *Commun. ACM 51*, 6 (2008), 82–87.