CCC
Computing Community Consortium
Catalyst

Privacy by Design-Engineering Privacy

Workshop 3 Report

Executive Summary

Within the last few years, privacy experts in industry, government and academia have called for new thinking that enables engineering privacy into information technology. An engineering-based approach to privacy leads to improved measurement and prediction of the degree of privacy obtained by individuals through a system. This ability has traditionally been built upon strong fundamentals in expressive design languages, knowledge of environmental risks, and recurring patterns of design.

In order to identify a shared research vision to support these different facets of the practice of Privacy by Design, the Computing Community Consortium (CCC) is sponsoring a series of four workshops throughout 2015. The first workshop framed common conceptions of privacy, and the second workshop explored privacy and user experience design. The third workshop took place in late August in Pittsburgh to discuss the emerging challenges in engineering privacy. An interdisciplinary group of 65 participants had expertise spanning cryptography, software engineering, technology policy, and law. The workshop was structured in a single track with 14 sessions, most of which featured a panel of speakers followed by very active participant discussions. The session topics include "Requirements and Policy Languages," "Practical De-Identification," and "Design Patterns for Privacy," among others. Two sessions were dedicated to an interactive critique of three prototype privacy tools.

Key Insights

The discussions at the workshop underscored the point that privacy must be addressed at design time to enable the successful engineering of privacy-protective systems. Because "privacy can be broken or made in the details," having privacy designed from the very beginning is crucial; it further means that additions and changes to a system must continue to take privacy into account. A number of challenges and open questions, however, remain. Across sessions, workshop participants noted the difficulty of measuring different aspects of privacy. Many participants questioned whether there exist *any best practices* for precisely and systematically measuring privacy protection from a tool or system feature, gauging improvements in privacy as systems were modified, and quantifying privacy risk. The workshop highlighted the need for systematic and widely agreed-upon definitions and procedures for measuring these concepts.

Many discussions also centered on the great potential for even more comprehensive and widely agreed-upon design patterns for privacy, which could spur efforts towards effectively engineering privacy into widely deployed systems. Participants pointed out the potential benefits of making privacy practices, as well as the implications of information disclosure and big-data inferences, more transparent to users.

Below, we expand on the themes that cut across workshop sessions, most of which emerged organically in discussion among the participants. We conclude by identifying key open challenges to which solutions have the potential to push forward the nascent discipline of privacy engineering. The full agenda of the workshop is available at http://cra.org/ccc/events/pbd-engineering-privacy/#agenda

Formal specifications of systems must balance abstraction and realism, improve transparency and ensure humans are involved in privacy-critical decisions. Formalism was discussed as a means by which engineers can express information sharing scenarios and check these scenarios for privacy problems. Research on formal privacy models and languages must contend with *coherence*, which asks whether the formalism can reliably answer questions about privacy, and *correspondence*, which asks how well the formalism represents real-world scenarios, privacy problems and various context. Three contrasting examples from Barth et al., May et al., and Breaux et al., were presented to illustrate this continuum from strong coherence to strong correspondence. In addition, policy and law frequently prescribe non-computational decisions that at present can only be answered by human judgment. How to identify these situations and ensure that systems appropriately defer to human judgment when they arise (e.g., emergency situations where access rules may need to be violated, or other situations where what is reasonable depends upon knowledge of multiple contextual factors) is an important challenge. This is especially difficult, when there is no transparency into the programs that are performing privacy-sensitive operations on data. There is also concern about “drift” in runtime operations that can lead to deviations from well-formed, formalized policies. Finally, engineering also includes non-deterministic, interpretative and aesthetic activities that situate artifacts in the environment and questions were raised about whether formalism is compatible with those activities. For example, the “tags” that may be assigned to data with respect to a policy can be interpreted differently in different contexts, which may lead to non-determinism, particularly when engineers are focused on one use case at the expense of another.

Privacy-enhancing technologies (PETs) should clearly present their definition of privacy and how they support users and designers. Users and designers may not agree on what privacy means, therefore, it is important that designers understand how users define privacy to evaluate whether their tools improve privacy. One observation was that stakeholders are often willing to work on assumed definitions and engage in discussions where the definition of privacy-related terms may even be ambiguous. For example, there is no consensus on what is meant by ‘context’, yet everyone seems to agree that understanding context is fundamental to improving privacy. In addition, the target audience of PETs can vary: end-user tools help users manage their own privacy, whereas privacy engineering requires a new set of tools to help designers communicate among themselves how their systems achieve privacy by design.

Some argued for a unifying theory or conceptual framework that can unite users, designers, engineers, and other stakeholders that have an interest in promoting privacy preserving practices. Examples discussed at the workshop include the recent NIST Privacy Risk Management Framework and the OASIS Privacy By Design for Software Engineers effort. It is clear that a central lexicon and agreed-upon definitions of privacy would lead us closer to a privacy framework. Better aligning users’ goals with practice is equally crucial. One participant pointed out the disconnect between users’ privacy goals and privacy engineering, raising the question that “People might like a guarantee that their deep secrets won’t be shared - can we get this?”

Engineering privacy can devolve into improving security, yet important differences between privacy and security exist. Security and privacy overlap, particularly when speaking about confidentiality. It is true that some privacy problems arise from security attacks, in which designers may define privacy in terms of data breaches. However, it was noted at the workshop that measures that improve security can reduce privacy. For example, efforts to improve security can increase surveillance and data collection, which introduces potential privacy risks. Furthermore, composability of systems can introduce new security threats: once we put systems together, the resulting composition may introduce unexpected security vulnerabilities that did not exist in the separate systems prior to composition. Similar concerns can arise in privacy, for example, when systems are composed, thus allowing data sets to be combined in ways that reduce privacy by de-anonymizing or revealing private information.

When it comes to threat modeling, there is a distinction between security threat models and privacy threat models. Security analysts commonly think of vulnerabilities, adversaries, and threats surrounding confidentiality, integrity and availability. In privacy, however, it was highlighted that researchers deal with relentless attackers, especially in re-identification attacks, whereas industry must also contend with accidental "threats" due to designers who failed to perceive the privacy risk, or to programmers who made code-level errors. This underscores the emergent value of privacy threat modelling, such as the LINDDUN method. The LINDDUN method includes categories aimed at reducing privacy threats, such as ensuring unlinkability across datasets and unobservability.

Quantifying privacy and privacy risk, which remains elusive, can inform how to prioritize limited design resources. As one participant stated, "Nobody ever talks about how you identify risk in privacy." In security, risk has been defined as likelihood of a vulnerability being exploited, multiplied by the impact to the system of a successful exploit. Privacy risk, however, increases from routine and often authorized access, such as inferring a user's personality traits and behaviors from their shopping habits, or from the lack of transparency into decision making or lack of user control over one's personal information. In this context, adopting the security term "attack" to describe a threat to privacy introduces a malicious intent that does not exist. In response, an emerging definition of privacy risk cited by NIST that avoids this value-laden terminology is the *likelihood of a problematic data action* multiplied by *the impact of a problematic data action.* The NIST approach distinguishes between potential problems for users (e.g., stigmatization, loss of trust, etc.) and business impact factors (noncompliance cost, reputational cost, etc.) An observation from one participant was that some privacy is also a public good, such as benefits to democracy and autonomy, which are impacts to society and not exclusive to individuals or companies. Another participant noted that often no one is responsible in an organization for privacy concerns that cut across other organizations and supply chains, and that most organizations only focus on how their specific practices impact privacy.

One challenge in privacy risk analysis is that the individuals mostly affected are not traditionally providing input to the risk analysis. For example, a company can quantify and assess their own risk of data breach in a security risk assessment, but the same company often would not have access to data on the privacy impact to individuals, particularly when these individuals are not the company's customers. Ironically, to gain insight to individual privacy risk perceptions, those individuals may need to disclose more personal information.

Finally, the question of when to address risk in design was raised. At design time, the engineer introduces specific mitigations, possibly those described by privacy design patterns. An alternative is to allow the designer to introduce an unrestricted design, only to check for non-compliance with a privacy policy at runtime. The former approach may be more deterministic, depending on the mitigation, whereas the latter approach may be more flexible at the cost of allowing non-compliant events to occur. In many of these situations, a reliable map of where data is and how it moves through an organization is needed.

Privacy design patterns offer promise for sharing design knowledge and have emerged from both academia and industry. Privacy design patterns are defined as repeatable solutions to recurring problems within a given context. Participants noted that patterns are generally more useful in late-stage design. While the pattern offers a solution, such as ambient notice as opposed to an elaborate privacy policy, the designer must still select patterns and determine when to apply patterns. In addition, privacy may require multiple patterns to be composed together using a pattern language, which describes a network of related and interconnecting patterns. Anti-patterns describe what designers should avoid: for example, collecting personal information and re-purposing the data, or not allowing users to control access. During discussion, participants proposed several interesting patterns: private or unguessable URLs, privacy dashboards, and scoping disclosures per transaction versus per account. It was generally agreed that engineers could be trained to apply patterns, but that discovering new patterns requires significant experience across multiple contexts and systems.

Market incentives have made it difficult to achieve practical privacy standards. Participant experience with the IETF process is that such standards emerge based on the consensus of those who are present. For security standards, enforcement is driven by compliance and interoperability, whereas there are no commonly accepted compliance requirements for privacy across U.S. industries. If privacy is truly a public interest that affects individuals in different ways, then the incentives to adopt privacy standards may not yet exist. One participant noted that privacy standards committee meetings are often thwarted by simple matters, such as defining terminology, which are biased by how committee members define their business processes. However, the benefits of working with standards are clear: improved interoperability and predictability concerning data handling. The Flash language is an example where one company sought to push their own standard, which required developers to buy into the whole system as opposed to allowing design flexibility, and ultimately the system failed. Overall, standards making processes are believed to be more successful when they are motivated by legal requirements or regulation and when they cover interoperability design as opposed to user interface design.

De-identification techniques should be tailored to the privacy risk and legal context. The participants identified three basic categories of data release: direct access, dissemination-based access, and query-based access, that cover various forms of de-identification and differential privacy. Risk was defined as a function of multiple factors, including data volume, sensitivity, intended use, recipient and consent, among others. In light of linking datasets, it was proposed that there is no meaningful definition of personally identifiable information, since seemingly benign information can be combined with other data to identify a person. Sparsity in data sets makes it more difficult to anonymize, and there is an ontological challenge to determine when attributes can be inferred from data (e.g., whether cancer is prostate or ovarian, can leak gender).

Concerns were expressed that anonymization technology has not been sufficiently proven to reduce the privacy risk in light of the cost, which may explain the slow adoption. Furthermore, it’s argued that many disclosures are non-public, in which case the re-identification threat comes from the company’s deliberate practices or rogue employees. The challenge of assessing risk was distinguished in two ways: when disclosing data row-by-row, versus disclosing data through a combination of seemingly independent queries. The former can be assessed using statistical disclosure control, which provides various models under different assumptions for assessing risk. The latter was the motivation for differential privacy. Participants noted, however, that the epsilon value needed to increase or decrease privacy under differential privacy is not a substitute measure of privacy risk. Synthetic data, which generally represents only the statistical population of the original data, was viewed as unusable in certain fields, such as medicine. Moreover, increasing privacy by reducing quasi-identifiers has been shown to destroy the data mining utility of datasets (Brickell and Shmatikov, 2008). In similar and emerging fields, more work may be needed to improve the legal and ethical guidelines on appropriate uses of non-anonymized data.

Increased transparency, empowering users, and the need to recognize the liability of personal data. Participants bemoaned the lack of transparency in many privacy-impactful systems. While the idea of transparency, both informing users of what information is collected and for what purposes, as well as providing access to this data, has been a core tenet of privacy principles for decades, transparency is lacking in the current state of affairs. Key challenges remaining for increased transparency include how best to empower end users with this information, as well as how to support designers in crafting transparency mechanisms. “Data isn’t just the data, it’s lots of data and inference (emerging out of a sea of lots of data).”

A related tension is balancing the idea that, in one participant’s words, “data is an asset” with the idea that “data is a liability” inside a company or institution. In the former case, the application of machine learning to big data holds the promise of uncovering useful, otherwise hidden trends. Such an approach favors more data. In the latter case, however, the more data a company holds about its users, the greater increase in the risk that this data will be released in a breach or will reveal embarrassing latent patterns that the users did not wish to reveal. The actionable challenge for privacy engineers is to shift the conversation to focus on the liabilities of holding data without a concrete purpose, perhaps emphasizing the risks of doing so.

There is a need for consensus on how to evaluate privacy . One unique feature of the workshop was clinics of three privacy tools in which an author of each tool presented a demonstration and solicited feedback from participants. The clinics brought up a number of themes that reappeared in other workshop sessions throughout the two days. First, participants wondered how to judge a privacy tool’s success, or lack thereof. As one participant explained, “In applying the label ‘privacy technology’ or ‘privacy engineering,’ we ask, ‘What is it achieving?’” This point raised the difficulty of quantifying or systematically measuring the success of any privacy intervention.

A related theme was whether each tool was truly a “privacy tool,” and whether this distinction matters. For instance, tools that improve privacy by facilitating more secure data sharing or helping an organization define the requirements for its system precisely may lead to improved overall privacy, however, they do not explicitly address an individual’s perception of privacy. One participant noted that such distinctions might be meaningless, stating, “Often with privacy we can think of something as a privacy tool or privacy technology, but also realize that it needs something in the social world when you embed it in a situation to have it perform privacy.” Thus, to measure the success of tools, one may need to measure the effectiveness of the tool to achieve some larger goal, beyond mere usability.

Additionally, participants explored whether making systems more privacy-protective would, for realistic systems, impact utility to the degree that the systems were no longer practically useful. Other common discussion points centered on frameworks and approaches for improving the usability of privacy tools as part of the privacy-engineering design process, as well as exploring how to generalize tools beyond specific application domains to support privacy engineering as broadly as possible.

Summary of Key Future Research Topics:

1. What are the definitions of privacy, and how can we establish a unified lexicon of privacy-related terminology so that we can advance the state of the art?
2. How do we measure and quantify privacy?
    a. What are the dimensions of privacy risks? How do we establish a taxonomy of all the factors that contribute to harm when privacy is not upheld?
    b. What is privacy risk? How do we determine the likelihood and impact of privacy risks?
    c. How do we quantify the success or failure of privacy technologies, aside from the retroactive assay of the impact that privacy breaches have had?
3. How do we understand better what the relationship is between data in terms of its utility to business practices versus the liability it may carry when privacy breaches occur?
4. What is the difference, if any, between a privacy tool and a security tool?
    a. How do we measure how well a privacy tool is achieving its goals?
5. What is the extent of the relationship between privacy and security?
    a. How much does privacy and security intersect?
    b. Is there a shared lexicon of terms between the two domains?
    c. What can we learn from the state of the art in security, in order to further our understanding of its impact on privacy?

Workshop Participants

Annie	Antón	GA Tech
Eleanor	Birrell	Cornell University
Travis	Breaux	CMU
Koen	Buyens	Cigital
Bethan	Cantrell	Microsoft
Richard	Chow	Intel
Sandra	Corbett	CRA
Lorrie	Cranor	CMU
Anupam	Datta	CMU
Frank	Dawson	Nokia
Jose	del Alamo	Universidad Politecnica de Madrid
Damien	Desfontaines	Google
Nick	Doty	UC Berkeley
Ann	Drobnis	CCC
Khaled	El Emam	University of Ottawa
Robert	Ferguson	Automatic Labs
Matthew	Fredrikson	CMU
Gerald	Friedland	UC Berkeley
Simson	Garfinkel	NIST
Carmela	Gonzalez Troncoso	Gradiant
Nathan	Good	Good Research
Susan	Graham	Berkeley
Paul	Grassi	Connect.gov
Mohit	Gupta	Clever
Seda	Gurses	NYU
Greg	Hager	Johns Hopkins
Joseph	Hall	CDT
Peter	Harsha	CRA
Hanan	Hibshi	CMU
Jaap-Henk	Hoepman	Radboud University Nijmegen
Giles	Hogben	Google
Jason	Hong	CMU

Brian	Ince	DNI
Sabrina	Jacobs	CRA
Limin	Jia	CMU
Dawn	Jutla	St. Mary's University
Apu	Kapadia	Indiana
David	Kelts	MorphoTrust
Aleksandra	Korolova	USC in fall (had been at Google)
Susan	Landau	Worcester Polytechnic Institute
Naomi	Lefkovitz	NIST
Christopher	Lubinski	18f
Ashwin	Machanavajjhala	Duke
Keith	Marzullo	NITRD
Aaron	Massey	Georgie Tech
Ilya	Mironov	Google
Deirdre	Mulligan	Berkeley
Helen	Nissenbaum	NYU
Lake	Polan	PhD Student at UChicago
Sören	Preibusch	Google
Rebecca	Richards	NSA
Ira	Rubinstein	NYU
Norman	Sadeh	CMU
Tomas	Sander	HP
Stuart	Shapiro	MITRE
Katie	Shilton	UM College Park
Manya	Sleeper	CMU
Daniel	Smullen	CMU
Karen	Sollins	MIT/IETF
Michael	Tschantz	ICSI
Blase	Ur	CMU
Elizabeth	Van Couvering	Karlstad Unviersity

Richmond	Wong	Berkeley
Helen	Wright	CCC
Heng	Xu	NSF

This material is based upon work supported by the National Science Foundation under Grant No. (1136993).

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.