Can you engineer privacy? On the potentials and challenges of applying privacy research in engineering practice

Seda Gürses
New York University
firstname@nyu.edu

August 14, 2014

Industrial-size data spills, leaks about large-scale secret surveillance programs, and personal tragedies due to inappropriate flows of information are guaranteed to have at least one consequence: engineers will be increasingly expected to integrate privacy solutions into the systems they are building or maintaining (see, e.g., [1]. Yet, the task of engineering systems to address privacy concerns can be complex.

The seeming unwieldiness of the engineering task becomes evident in the concept of privacy itself and how this concept is negotiated. As a legal concept, privacy is defined rather vaguely. That vagueness, some argue, is part of its protective function. The open ended definition allows people to invoke privacy as a category to protect their personal lives and autonomy from intrusions by others – including the state that endows them with citizenship rights and runs surveillance programs. European Data Protection Directive (DPD) or Fair Information Practice Principles (FIPPs) on the other hand are procedural measures, e.g., like notice and choice, data retention limitation, subject access rights. These principles are seen to be instrumental to making the collection and processing activities of organizations transparent. Although less ambiguous, data protection principles still need to be translated into technical requirements and are vulnerable to narrow interpretations. Moreover, FIPPs fall short of mitigating all the privacy concerns of the users' towards a given organization. They also do not address privacy concerns users may have with respect to other users, with people in their social environments and towards a greater public.

Scholars from various fields have stepped up to the challenge of clearing the murky waters of privacy. Legal scholars and philosophers have proposed taxonomies of privacy violations [2] and a holistic framework for evaluating appropriate flows of information based on contextual social norms [3]. Social scientists and ethnographers have studied groups of people, online and offline, to develop better informed understandings of users' needs. But, how are engineers supposed to integrate and translate these frameworks into existing engineering

practice? It is in answering this question, privacy research conducted within computer science is valuable.

Over the years, privacy research in computer science has led to a whole palette of privacy solutions. The solutions originate from diverse sub-fields of computer science, e.g., security engineering, software engineering, HCI, AI. From a birds eye view, all of these researchers are studying privacy problems and solutions, and yet a closer look reveals that they also have their differences. In the following, I introduce a small taxonomy of prominent approaches to privacy within computer science. The categories of this taxonomy are not comprehensive or definitive, but they provide a way of distinguishing the guiding principles of the different approaches that privacy researchers are taking. In practice, engineers may mix and match these principles; pulling them apart allows us to think about the different ways in which we can engineer systems with privacy in mind.

1 Privacy as Confidentiality

The most prominent conception of privacy relies on the binary that exposure of information leads to a loss of privacy, while guaranteeing the confidentiality of information is a way to preserve or enhance privacy. This binary can be linked to Warren and Brandeis' "right to be let alone" [4], which was a response to the novel ways in which innovation in technology makes it increasingly possible to collect information about matters that would have previously been regarded as private.

Many privacy researchers work on privacy solutions that rely on this binary understanding of unwanted disclosures as privacy violations. These researchers rely on three important principles: data minimization, avoidance of a single point of failure and openness to scrutiny. The first principle, data minimization is about designing systems (and computational mechanisms) to only collect private information that is absolutely necessary for a given functionality. In its bare bones, this means that, by default, the users should be able to use the system anonymously. If the users have to be identified, the different interactions of the user throughout time should remain unlinkable. For example, a user of a service may utilize zero knowledge proofs to prove that she is above 18 without revealing her birthdate or any further private information, thus guaranteeing that her transactions are unlinkable. Communications as well as traffic data must also be kept confidential from unauthorized parties, sometimes including the service provider itself. The capabilities of these unauthorized parties, also called adversaries, is an important driver of threat models, and hence the design of privacy solutions in this line of research.

The second principle is to avoid designing architectures for information collection and processing with a "single point of failure". In other words, introducing a distributed trust model so that users do not need to rely on a single entity to protect their privacy. The third principle requires that the protocols, code and processes of development that underlie the privacy tools are open to

public scrutiny in order to increase trust in the privacy solution itself. Tor here is a popular example of privacy as confidentiality solution that is built using all three principles.

2 Privacy as Control

Another approach to privacy starts from the assumption that information will have to be disclosed in an increasingly networked world. Hence, Westin writes, privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others [5]. Westins work, which he further butressed with dozens of large-scale surveys, is fundamental to most legal and organizational measures introduced to protect personal data across the globe.

Based on this conception of information privacy, DPD and FIPPs list procedural mechanisms through which organizations can make their personal data collection and processing practices transparent to their data subjects, regulators and the general public. If these mechanisms are in place, ideally users can make informed decisions about and have greater control over the collection and flows of their personal information. Further, abuses can be detected or mitigated.

A good portion of privacy research focuses on developing methods and mechanisms for data protection compliance. These protection mechanisms are expected to compliment organizational measures, like privacy training for employees working with personal data or procedures for database breach notifications. Another type of notification is the privacy policy, through which organizations can inform data subjects about the purpose for which they are collecting personal data, which data, and for how long. Privacy researchers have studied ways in which to ease the burden of reading privacy policies often weighted with legal jargon. Proposals have included usable representations of the content through labeling mechanisms comparable to those in the food industry, as well as the design of machine readable privacy policies that can be matched to user preferences.

Privacy policy languages, as the latter are often called, and policy enforcement mechanisms like purpose based access control models can be applied in combination to provide organizations with mechanisms to ensure that the internal use of personal data adheres to the collection purpose. Data minimization also pops up here, but rather than aspiring to achieve anonymity or unlinkability through computational methods, the principle is about limiting the collection of personal data for the given purpose only. These researchers often assume that service providers are the main party trusted with the protection of personal data. This means that distributed trust models are rare in these proposals. Subject access rights, the ability of the users to access, rectify or delete their data collected by a service provider is another challenging requirement which has attracted research and has found popular implementations like Google dashboard.

3 Privacy as Practice

In this third approach, privacy is seen as the negotiation of social boundaries through a set of actions that users collectively or individually take with respect to disclosure, identity and temporality in environments that are mediated by technology. Hence, privacy is not seen as something that users can delegate to the machine. Rather, engineers explore how design mechanisms and principles mediate users' privacy practices [6]. These researchers dispense with the binary understanding of privacy as exposure and privacy as concealment, since interactions inform negotiation of privacy in unexpected ways, e.g., a user may signal to her peers through chat that she wants to be left alone, a disclosure that allows her to negotiate her privacy in a public space. As evident in this example, the distinction between online and offline privacy is also undone in the pursuit of understanding privacy practices.

Privacy researchers here emphasize that privacy is negotiated through collective dynamics, e.g., if unlucky, the same user might have very disrespectful peers that do not respect her request for privacy. Transparency and feedback mechanisms for raising awareness of the socio-technical systems workings are often proposed as central to establishing privacy practices. This understanding of transparency is greater than providing information about what data is collected by an organization as proposed in privacy as control. Instead, the objective is to make information systems, their affects, as well as the responses from (non-)user communities part of what needs to be made transparent.

For example, if possible consequences of a user's actions can be made transparent to her, she may be able to make better decisions about her interactions. The user may learn from her past interactions, so feedback on past practices may be used to inform future ones. For instance, information about how many friends have visited a user's profile may inform how much she wants to post on her profile in the future. Similarly, users may learn from their (mediated) social surroundings: information about how other friends manage their privacy settings may provide guidance to the user. In some cases, based on studies about good privacy norms, users can be nudged to develop better privacy practices [7], e.g., users may be opportunistically encouraged to review their privacy settings. Similarly, if a user is provided with feedback on the algorithms underlying a recommender system, she may better assess whether and how she wants to participate in such a system.

4 Future prospects

There are many more proposals for addressing privacy in systems than listed above. Some proposals fall in-between the three categories. For example, database anonymization and differential privacy both propose elaborate computational methods for data minimization comparable to solutions in privacy as confidentiality. However, the mechanisms are not intended to minimize data collection but to anonymize or obfuscate later disclosure. Further, database

anonymization is a way to exit the legal compliance regime, making it difficult to identify it as an organizational transparency mechanism typical of privacy as control. Furthermore, there are a number of proposals for addressing discrimination and fairness issues in the context of data mining like discrimination aware data mining as well as fairness in classification. However, it is open to discussion whether discrimination and fairness are privacy issues. Generally, many concerns we discuss under privacy may in fact be related to greater issues of social justice that require more elaborate rethinking of our societies as well as technological futures.

The taxonomy above shows that engineering decisions, be it when architecting infrastructures, designing organizational systems or crafting particular applications, co-determine the way in which people may negotiate their privacy. Yet, challenges are abound. Can we integrate these three approaches given their fundamental differences? For example, while privacy as confidentiality assumes a world in which trust in organizations that process private data should be minimized, privacy as control assumes trust in those organizations can be established through transparency. In contrast, most privacy as practice proposals assume that the service provider is honest and has a genuine interest in accommodating users' privacy practices above its own organizational and market interests. Hence, a skeptic could ask, given the slipperiness of the concept, the political and market contestations of privacy, as well as the differences between the solution sets, can we even speak of a privacy engineering project?

Engineering privacy may in fact only be another ideal like engineering security, dependability, or usability, and one that misleadingly suggests that we can engineer social and legal concepts. We cannot engineer society, but neither are our societies independent of the systems we engineer. Hence, as practitioners and researchers we have the responsibility to engineer systems that address privacy concerns. The above taxonomy attempts to provide an overview of existing approaches to privacy in computer science research. The robustness of these approaches will only grow through further engagement in all of them when we engineer systems.

References

[1] NIST, Privacy Engineering Workshop, April, 2014, `http://www.nist.gov/itl/csd/privacy-engineering-workshop.cfm` URL last visited: May 5, 2014.

[2] Daniel Solove, *A Taxonomy of Privacy*, University of Pennsylvania Law Review, 154 (3), 2006.

[3] Helen Nissenbaum, *Privacy as Contextual Integrity*, Washington Law Review, 79 (1), 2004

[4] Samuel Warren and Louis Brandeis, *The Right to Privacy*, Harvard Law Review, 4 (5), 1890.

[5] Alan Westin, *Privacy and Freedom*, Altheneum, New York, NY, 1967.

[6] Leysia Palen and Paul Dourish, *Unpacking "privacy" for a networked world*, pp. 129 – 136, CHI, 2003.

[7] Yang Wang, Pedro Giovanni Leon, Alessandro Acquisti, Lorrie Faith Cranor, Alain Forget, Norman M. Sadeh: *A field trial of privacy nudges for facebook*, CHI (2014), 2367-2376.