Understanding the landscape of privacy technologies[1]

Claudia Diaz and Seda Gürses

1 Introduction

In the last decades, much effort has been devoted to research on privacy across different subfields in computer science (e.g., security engineering, data mining, HCI), resulting in a broad range of solutions for addressing the "privacy problem". Privacy is a multifaceted and complex concept that can be tackled from very different perspectives. Existing solutions rely on different definitions of privacy as well as on a variety of (often implicit) social and technical assumptions. As a result, it is hard for non-experts to understand the privacy research landscape and to put available technologies into context.

In this paper we provide an overview of the landscape of privacy technologies following the classification in three privacy research paradigms proposed by Gürses [11,14]. For each of these paradigms we describe their conception of privacy and the types of privacy threats that the technologies are meant to address, present representative examples of technologies that have been proposed to address these threats, and discuss their distinguishing characteristics.

2 Privacy as "control"

Electronic services often involve the collection and processing of personal information by the organizations providing the services. For example, hospitals keep health records of their patients, businesses keep purchase records of their customers, employers keep records on their employees, and tax authorities keep the tax records of the country's citizens. In this context, privacy violations are often associated with the disclosure of information to other parties or to the broader public. Typical examples include:

- User-generated content uploaded to online social networks that becomes available to people who were not supposed to have access to it. For example, a potential employer gets to see "drunk pictures".
- Data security breaches that lead to the disclosure of the personal records kept in an organization's database. Examples include hackers breaking into a system to steal credit card numbers, and health records becoming publicly accessible on the Internet.
- Personal records being shared with (or sold to) other organizations, e.g., for targeted advertising, without the consent of the individuals who provided the data.
- Personal information being abused by a malicious insider. For example, preventing a hospital employee from looking into the health records of a celebrity and then selling the information to gossip magazines.
- Privacy violations also arise when personal information is used for illegitimate purposes, such as unlawful discrimination based on racial or religious profiles.

Privacy technologies in the "control paradigm" aim at addressing these privacy concerns by: (1) providing individuals with means to control the disclosure of their information; and (2) providing organizations with means to define and enforce data security policies and prevent the abuse of personal information for unauthorized purposes. Representative examples of privacy research in this paradigm include:

- **Privacy settings**: the lack of usability and the complexity of configuring privacy settings have been identified as one of the main causes for unintended data disclosure in online social networks [15]. Various recent proposals aim at making privacy settings more usable by providing users with visual feedback on the visibility of their content [18], default "suites" of privacy settings which have been specified by friends or trusted experts [3], or privacy wizards that automate the configuration of the settings [13].
- **Purpose-based access control**: privacy violations often arise when data that has been collected for a stated purpose (e.g., billing) is used for a different purpose (e.g., profiling). Purpose-based access

[1] The content presented in this extended abstract reflects ongoing work in progress by Seda Gürses, Claudia Diaz and Nicola Zannone.

control (PBAC) [5] provides a mechanism for specifying which are the allowed uses of collected information, and for verifying that the purpose of a data access is compliant with the policy.

- **Auditing:** even if a privacy policy has been specified and it is enforced through access control mechanisms, additional measures need to be taken to verify that no abuse has taken place. In order to audit the use of data in a system, the system is required to log the data access and processing operations. Logging excessively however, imposes an overhead on the system and may even enable additional privacy violations (that involve exploiting the information in the logs). Some proposals to address this problem focus on generating auditing specifications that produce logs that are both minimal and sufficient to audit whether the policies have been respected [2].

The main objective of these technologies is to provide individuals with control and oversight over the collection, processing, and use of data. As such, the underlying conception of privacy relates to Westin's definition of privacy as "*the right of the individual to decide what information about himself should be communicated to others and under what circumstances*" [20].

Some of the main characteristics of the "privacy as control" approach are:

- Privacy is defined as the ability to determine acceptable data collection and usage, which is articulated through policies. Some of these policies are defined by users (e.g., through their privacy settings), while others are defined by the organization holding the data (e.g., policies that establish who can access which data and for which purposes). The latter are typically driven by compliance with existing regulations such as data protection law.
- The organization that holds the data is trusted to protect the privacy interests of individuals. This includes: respecting the preferences expressed by users in their privacy settings, defining appropriate data access and data processing policies, enforcing these policies through access control mechanisms, and auditing that no violations of the policies have occurred. These technologies do not offer protection guarantees towards organizations that want to violate user privacy by abusing the data that they hold.
- The focus of these technologies is on defining appropriate data usage rather than limiting data disclosure and collection. While limitations to data collection may be encoded in the policies, these technologies do not aim at technically preventing data collection.

3 Privacy as "confidentiality"

The solutions introduced in the previous section heavily rely on the assumption that organizations that collect and process user data are going to competently and honestly act in the best interests of their users. Furthermore, once the data is under the control of an organization it is very difficult for individuals to verify how their data is actually being used.

A different family of privacy technologies considers however that placing such high levels of trust in organizations should be avoided whenever possible, as they leave individuals vulnerable to incompetent or malicious organizations. The large number of reported privacy breaches (due to a lack of appropriate data security practices), and the incentives to amass and use personal data for financial gain (without regard for users' privacy concerns), support the rationale of this approach. The main privacy threats considered by these technologies are:

- The mass collection of user information in databases is considered in itself a privacy threat both at the individual as well as at the societal level. Furthermore, mass information collection, storage, and processing opens the door to a variety of privacy violations (e.g., surveillance, profiling, or manipulation). It is considered that the abuse of private information may be stealthy and never become evident to the individuals themselves.
- The linkability of user information across different contexts (e.g., shopping, location, health, browsing, and social network information) through unique identifiers aggravates the previous problem by enabling more powerful profiling capabilities and inferences on the interests, lifestyle, and behavior of the individuals to whom the information relates.

The privacy technologies in this paradigm thus aim to create an individual autonomous sphere free from intrusions from both an overbearing state and the pressure of social norms, and as such are inspired by the definition of privacy as "the right to be let alone" [19]. To achieve this, disclosure of information is by default prevented, or information is minimally disclosed in a way that cannot be linked back to the individual. Some of the privacy technologies in this paradigm include:

- **Anonymous authentication protocols**: authentication is often the first step of a transaction. In many cases however, all that is needed is to ensure that an entity has certain properties, and in these cases establishing the identity of the entity is unnecessary. A typical example is age verification trough an identity card when buying alcohol or cigarettes: the vendor may require the buyer to prove that she is of legal age to purchase the products, but no other personal information typically contained in an identity card (e.g., name, date and place of birth, address) is needed. In anonymous authentication protocols [4,6], the user first obtains a credential from an issuer (e.g., the government) certifying a set of attributes. Later, the user is able to selectively prove properties on these attributes to a verifying party (e.g., a vendor). The main property of these protocols is that a statement on the attributes can be proven without revealing any additional information besides the statement itself. The user remains anonymous even if the issuer and the verifier collude. Furthermore, it is not possible to link transactions as being performed by the same user.

- **Anonymous communication networks**: even if the content of communications is protected cryptographically, sensitive information may be leaked by traffic data, which includes the timing, order, frequency, and volume of communications, as well as the location and identities of the parties engaged in the communication. Anonymous communication systems [10] aim at protecting traffic data by concealing who talks to whom, and this is achieved by relying communications through a set of intermediate proxies. The earliest designs focused on implementing high-latency anonymous email services [7]. Tor [12] provides low-latency bi-directional anonymous channels that enable anonymous web browsing. With hundreds of millions of daily users Tor is currently the most widely used anonymous communication network.

- **Private Information Retrieval (PIR)**: in some cases users may want to access a database without revealing to the server hosting the database which records they retrieve. This can be achieved with PIR protocols [8]. Note that in this case, privacy is achieved not through anonymity but rather through concealing other details of the transaction. These techniques can be used for example to retrieve information on nearby points of interest from a location-based service without revealing location information to the server [17].

As we can see in the above examples, these technologies substantially reduce the disclosure of private information, thus diminishing the opportunities for data collection that can then be abused. The main characteristics of technologies in the "privacy as confidentiality" paradigm are:

- Privacy notions are predefined as properties that are hard-coded in the technology itself. Contrary to the previous case, in which the definition of what constitutes a privacy violation is dependent on a policy, the technologies in this paradigm typically provide a formal definition of the privacy property that they aim to achieve (e.g., in anonymous communication networks, finding who communicates with whom violates the privacy provided by the system). The goal of the technology is then to ensure that these formally defined privacy properties hold.

- Focus on preventing data disclosure. These technologies mitigate privacy risks by enabling services and functionalities while minimizing data disclosure beyond what seems intuitively possible. The rationale is that once data is disclosed it is hard to prevent privacy abuses, and that data cannot be abused for privacy invasive purposes if it has not been previously made available.

- Minimize the need to trust others with appropriately handling identifiable and linkable data. The goal of these technologies is to minimize privacy risks by reducing the need to trust other entities, while still guaranteeing other security properties such as service integrity (e.g., in anonymous authentication protocols users should not able to fake attributes). The proposal of these technologies is typically accompanied by a security analysis that considers strategic adversarial behavior and verifies that the system provides the claimed privacy and security properties.

4 Privacy as "practice"

The technologies in the two previous paradigms have a strong security focus. Their goals are to either allow individual users to prevent information disclosure, or organizations to enhance the security of the personal data that they hold and prevent its abuse for illegitimate purposes. Privacy is however not just an individual matter – it has an important social dimension, as users often make privacy decisions not in isolation but based on how their communities make privacy decisions. Technologies in the privacy as practice paradigm aim at making information flows more transparent through feedback and awareness thus enabling individuals as well as collectives to better understand how information is collected, aggregated, analyzed, and used for decision-

making. Making systems more transparent enhances their ability to question, intervene, and renegotiate information practices.

The main privacy concerns addressed by technologies in this paradigm include:

- It is hard for users to understand what information is available to others, how it is being used, and which inferences or decisions could be made based on it. This obscures the privacy risks that arise from their practices when using digital technologies.
- Statistical profiles can be constructed based on anonymized and aggregated data. These profiles can be in turn used to discriminate against certain population subgroups (without requiring individual identification).

In this paradigm, the focus is not only on concealing and controlling information but also on improving transparency and enabling identity construction. The underlying notion of privacy fits Agre's definition of the concept as "the freedom from unreasonable constraints on the construction of one's own identity" [1]. Two representative examples of technologies that fit in this paradigm are:

- **Platform for Privacy Preferences (P3P)**: P3P [9] is a protocol that provides information on the data collection and data use practices of web sites to users browsing those sites. The P3P policy provides a way for a Web site to encode its practices in a machine-readable XML format. This enables web users to understand what data will be collected by sites they visit, how that data will be used, and what data/uses they may "opt-out" of or "opt-in" to.
- **Privacy Mirrors**: Privacy Mirrors strive to give users a better understanding of the flow, state, and history of information throughout that system. With this better understanding the user can see if his personal comfort level for privacy fits within the current workings of the system Furthermore, users become aware of how they participate in the socio-technical system, how others participate with respect to them and their information, and how they can take action to change their own behavior as well as the socio-technical environment [16].

As illustrated by these example technologies, the distinguishing characteristics of technologies in this paradigm are:

- The technologies themselves do not aim at defining a particular notion of what is privacy-preserving or privacy-invasive. Rather, they provide transparency on actual practices so that users themselves can understand how information is being used and make their own mind as to whether they are comfortable with which data about them is being collected and how it is being used.
- While privacy as confidentiality technologies focus on (preventing) data disclosure and privacy as control technologies focus on (ensuring adequate) data usage, privacy as practice technologies focus on making transparent both data disclosure and data usage. The main objective is to make a range of mechanisms and actions visible in order to render them a resource for action.
- The technologies in this paradigm have the potential to uncover malicious behavior by organizations (or other individuals) that involves, for example, excessive collection of data or use of those data for unacceptable purposes. The technologies themselves however typically do not take into account possible adversarial behavior. For example, a P3P policy could be inaccurate in describing the practices of a web site, and the input to a privacy mirror could be corrupted to provide bad quality feedback.

5 Conclusions

Privacy is a broad and subjective concept that can be defined in many different ways. It is therefore not evident what sort of protection is offered by a specific technology. Some of the questions that need to be asked when considering a privacy technology include the following: What is the implicit or explicit definition of the "privacy problem" that is addressed by the technology? Who defines what privacy is (the technology designer, the organization's policy, or the users themselves)? What are the trust assumptions (both relating to the technical infrastructures available and to the behaviour of the entities involved in the system)? What are the adversary models that are considered (e.g., is the technology designed to protect against surveillance by the government, against profiling by an organization, or against disclosure to the broader public)? What is the level of assurance provided by the privacy technology – does it provide technical mechanisms that are verified through a security analysis? What is the scope of the solution and what are the orthogonal aspects that are left out of the scope?

Having clear answers to these questions is essential to ensure that the technologies that are selected and deployed to improve privacy protection in a given digital system actually address the privacy problems of interest.

Acknowledgments

This work was supported in part by the projects: GOA TENSE (GOA/11/007), IAP Programme P6/26 BCRYPT, EC ICT-2007-216676 ECRYPT II, IWT SBO SPION, FWO G.0360.11N, and FWO G.068611N. C. Diaz is funded by the Fund for Scientific Research in Flanders (FWO).

References

[1] P. Agre, M. Rotenberg: Technology and Privacy: The New Landscape. Cambridge, Massachusetts, MIT Press, 325 pages, 1998.

[2] D. Biswas, V. Niemi: Transforming Privacy Policies to Auditing Specifications. 13th *IEEE International High Assurance Systems Engineering Symposium*, IEEE CS, Boca Raton, FL, USA, pp. 368-375, 2011.

[3] J. Bonneau, J. Anderson, and L. Church: Privacy suites: shared privacy for social networks. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (SOUPS '09). ACM, New York, NY, USA, 2009.

[4] S. Brands: Rethinking Public Key Infrastructures and Digital Certificates. *MIT Press*, Cambridge, MA, USA, 2000.

[5] J. Byun, N. Li: Purpose based access control for privacy protection in relational database systems. The *VLDB Journal* 17(4), pp. 603-619, 2008.

[6] J. Camenisch, A. Lysyanskaya: An Efficient System for Non-transferable Anonymous Credentials with Optional Anonymity Revocation. In *Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques: Advances in Cryptology* (EUROCRYPT '01), Springer-Verlag, pp. 93-118, 2001.

[7] D. Chaum: Untraceable electronic mail, return addresses, and digital pseudonyms. In *Commun. ACM* 24(2) pp. 84-90, 1981.

[8] B. Chor, E. Kushilevitz, O. Goldreich, M. Sudan: Private information retrieval. J. ACM 45(6), pp. 965-981, 1998.

[9] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. Reagle: The Platform for Privacy Preferences 1.0 P3P 1.0 Specification. World Wide Web Consortium, Recommendation REC-P3P-20020416, 2002.

[10] G. Danezis, C. Diaz, P. Syverson: Systems for Anonymous Communication. In *CRC Handbook of Financial Cryptography and Security*, CRC Cryptography and Network Security Series, Chapman & Hall, pp. 341-390, 2010.

[11] G. Danezis, S. Gürses: A critical review of 10 years of Privacy Technology. In *Surveillance Cultures: A Global Surveillance Society?* 18 pages, 2010.

[12] R. Dingledine, N. Mathewson, P. Syverson: Tor: the second-generation onion router. In *Proceedings of the 13th conference on USENIX Security Symposium - Volume 13*, USENIX Association, Berkeley, CA, USA, pp. 303-320, 2004.

[13] L. Fang and K. LeFevre: Privacy wizards for social networking sites. In *Proceedings of the 19th international conference on World Wide Web (WWW '10)*. ACM, New York, NY, USA, pp. 351-360, 2010

[14] S Gürses: Multilateral Privacy Requirements Analysis in Online Social Networks. *PhD thesis*, *KU Leuven*, 312 pages, 2010.

[15] M. Madejski, M. Johnson, and S. Bellovin. The Failure of Online Social Network Privacy Settings. Tech Report CUCS-010-11, Columbia University, February 2011.

[16] D. Nguyen, E. Mynatt: Privacy Mirrors: Understanding and Shaping Socio-technical Ubiquitous Computing Systems. In Georgia Institute of Technology GVU Technical Report (GIT-GVU-02-16), 2002.

[17] F. Olumofin, P. Tysowski, I. Goldberg, U. Hengartner: Achieving efficient query privacy for location based services. In *Proceedings of the 10th international conference on Privacy enhancing technologies* (PETS'10), Springer-Verlag, pp. 93-110, 2010.

[18] T. Paul, D. Puscher, T. Strufe: Improving the Usability of Privacy Settings in Facebook. In CoRR abs/1109.6046, 2011.

[19] S. Warren, L. Brandeis: The right to privacy. *Harward Law Review* 4(5), pp. 193–220, 1890.

[20] A. Westin: Privacy and freedom. Atheneum, New York, 1970.