Quantitative Approaches to the Protection of Private Information: State of the Art and Some Open Challenges

Catuscia Palamidessi

INRIA Saclay and LIX, École Polytechnique

Privacy is a broad concept affecting a variety of modern-life activities. As a consequence, during the last decade there has been a vast amount of research on techniques to protect privacy, such as communication anonymizers [9], electronic voting systems [8], Radio-Frequency Identification (RFID) protocols [13] and private information retrieval schemes [7], to name a few.

Until some years ago, the prevailing technology for privacy protection was k-anonymity [17,16]. Similarly to other techniques like ℓ-diversity, k-anonymity is based on the principle of modifying opportunely the so-called quasi-identifier attributes so that for every combination of quasi-identifier values in the data set, there are at least k individuals with these values. The idea is that in this way, each individual would be concealed in a group of at least k individuals with the same characteristics. The problem is that this technique requires the set of the quasi-identifier to be static, fixed in advance, and to be the same for all the individuals. However, as the amount of publicly available information about individuals grows, the distinction between quasi-identifier and non-quasi-identifier attributes fades away: Any information that distinguishes one person from another can be used to re-identify the person. For instance, any behavioral or transactional profile like movie viewing histories, consumption preferences, shopping habits, browsing patterns, etc. Furthermore, while many attributes may not be uniquely identifying on their own, in combination with others any attribute can be identifying. Due to these shortcomings, anonymity techniques, and more in general, the privacy protection paradigm based on de-identifying the data, have proved mostly ineffective against the emergence of powerful re-identification algorithms based on background knowledge, cross-correlation between databases, and analysis of the network structure. For instance, Narayanan and Shmatikov [15] conducted research on two large social networks, Flickr and Twitter, and demonstrated that, by using their algorithm, one third of the users who were members of both networks could be recognized in the completely anonymous Twitter graph with only 12% error rate!

In recent years, a new framework for privacy, called *differential privacy* (DP) has become increasingly popular in the area of statistical databases [10,12,11]. The idea is that, first, the access to the data should be allowed only through a query-based interface. Second, it should not be possible for the adversary to *distinguish*, from the answer to the query, whether a *certain individual is present or not* in the database. Formally, the *likelihood* of obtaining a certain answer should not change too much (i.e., more than a factor e^{ϵ}, where ϵ is a parameter)

R. Focardi and A. Myers (Eds.): POST 2015, LNCS 9036, pp. 3–7, 2015.
DOI: 10.1007/978-3-662-46666-7_1

when the individual joins (or leaves) the database. This is achieved by adding *random noise* to the answer, resulting in a trade-off between the privacy of the mechanism and the utility of the answer: the stronger privacy we wish to achieve, the more the answer needs to be perturbed, thus the less useful it is. One of the important features of DP is that it does not depend on the side information available to the adversary. Related to this, another important advantage is that DP is robust with respect to composition attacks: by combining the results of several queries, the level of privacy of every mechanism necessarily decreases, but with DP it declines in a controlled way. This is a feature that can only be achieved with randomized mechanisms: With deterministic methods, such as *k-anonymity*, composition attacks may be catastrophic.

DP has proved to be a solid foundation for privacy in statistical databases. Various people have also tried to extend it to other domains. However, there are some inherent limitations that make it inadequate in several practical cases. First, DP assumes that the disclosed information is produced by aggregating the data of multiple individuals. However, many privacy applications involve only a single individual, making differential privacy inapplicable. Second, even when multiple individuals are involved, DP assumes that full range of possible values of an individual needs to be *completely protected*. In applications where perturbations in an individual's value lead to a non-negligible change in the result, this requirement is impractical since the noise that we need to add is so big that the result becomes useless. In such cases, we wish to adapt our privacy definition, to only partially protect the user's data (which is often sufficient), while lowering the noise to obtain an acceptable level of utility. Third, DP focuses on the worst-case, since it requires the likelihood property to be satisfied for every possible database and every possible result. There are situations, however, where an average notion (weighed with the probabilities) would be more suitable for measuring the risk. For example, an insurance company protecting credit cards will be interested in knowing the probability that a card is compromised (and the corresponding expected loss) in order to decide what fee to apply. And an individual user may want to know the probability of a privacy breach in order to decide whether it is worth employing some costly counter-measures.

Finally, differential privacy needs some care when handling correlated data. In such situation, in fact, the adversary can filter out some of the noise by statistical reasoning. The best solution offered so far to this problem is to assign a *privacy budget*, and subtract from this budget a certain amount at every release of information.

In our team, we have addressed some of these issues by defining an extended DP framework in which the indistinguishability requirement is based on an arbitrary notion of distance (d_x-privacy, [5]). In this way we can naturally express (protection against) privacy threats that cannot be represented with the standard notion, leading to new applications of the differential privacy framework. In particular, we have explored applications in geolocation [3,4] and smart metering [5]. In the context of geolocation, the problem of the correlated data becomes particularly relevant when we consider traces, which usually are composed of a

large amount of highly related points. We addressed this issue using *prediction functions* [6], obtaining encouraging results.

A different approach for measuring privacy is to employ techniques from *quantitative information flow* (QIF). Such techniques aim at quantifying the *leakage of secret information* through the observation of some public event, and have been successfully applied in several application contexts, such as programs, anonymity protocols, side channels attacks, etc. A common approach in this area is to use notions from information theory to measure the correlation between secret and observable information, the most prominent examples being Shannon and min-entropy. In contrast to differential privacy, these approaches typically provide average measures.

An important limitation of these entropy-based measures is that they treat secrets as atomic data, ignoring the structure and the relationship between the secrets. This structure is crucial for privacy applications: in geo-location systems, for instance, the *distance* between secrets (which are locations) plays a crucial role in defining the concept of privacy within a certain area.

Another limitation, common to both QIF and DP measures, is that they ignore several parameters that should play a fundamental role in an effective and realistic analysis of the privacy risk. First of all they ignore that the inference of the confidential information may have a cost for the adversary, which may considerably reduce the risk of an attack. Such cost can be, for example, in terms of computational resources or of some kind of deterrent. Furthermore, they ignore the gain that the adversary may obtain by acquiring the confidential data, and which is not necessarily the same for all the data (the credit card of Bill Gates is probably more worth than that of an average person...). Dually, they ignore the amount of damage that the user may suffer from the privacy breach, and that can be different depending on the data, or on the user.

The recently developed g-leakage framework [2,14,1] proposes a unified solution to the above issues by introducing the notion of *gain function*, which allows to express the gain of the adversary when guessing a secret. This richer definition of leakage has opened the way to new research directions.

Another shortcoming of the current approaches to privacy is that they are only applicable when the public information is well delimited and acquired in finite in time. Unfortunately, in most situation the source of public information is not necessarily bound, and some additional information can always be revealed in the future. At present, there are no techniques to verify privacy guarantees in situations in which the revelation of public information is not bound in time. This is a serious limitation, especially given that most of the systems which we use nowadays have an interactive nature, and usually are not under the control of the user.

In our team, we have started exploring a possible approach to this problem by defining a generalized version of the bisimulation distance based on the Kantorovich metric. In contrast to the standard bisimulation distance, which is additive and therefore not suitable to capture properties such as differential privacy, our framework considers the Kantorovich lifting on arbitrary metrics.

We have applied this framework to the particular case of the d_x-privacy, and provided an efficient method to compute it based on a dual form of the Kantorovich lifting. However, for other notions of leakage the quest for an efficient implementation remains open, as well as that of a generalized dual form.

References

1. Alvim, M.S., Chatzikokolakis, K., McIver, A., Morgan, C., Palamidessi, C., Smith, G.: Additive and multiplicative notions of leakage, and their capacities. In: IEEE 27th Computer Security Foundations Symposium, CSF 2014, Vienna, Austria, July 19-22, pp. 308–322. IEEE (2014)
2. Alvim, M.S., Chatzikokolakis, K., Palamidessi, C., Smith, G.: Measuring information leakage using generalized gain functions. In: Proceedings of the 25th IEEE Computer Security Foundations Symposium (CSF), pp. 265–279 (2012)
3. Andrés, M.E., Bordenabe, N.E., Chatzikokolakis, K., Palamidessi, C.: Geo-indistinguishability: differential privacy for location-based systems. In: Proceedings of the 20th ACM Conference on Computer and Communications Security (CCS 2013), pp. 901–914. ACM (2013)
4. Bordenabe, N.E., Chatzikokolakis, K., Palamidessi, C.: Optimal geo-indistinguishable mechanisms for location privacy. In: Proceedings of the 21th ACM Conference on Computer and Communications Security, CCS 2014 (2014)
5. Chatzikokolakis, K., Andrés, M.E., Bordenabe, N.E., Palamidessi, C.: Broadening the Scope of Differential Privacy Using Metrics. In: De Cristofaro, E., Wright, M. (eds.) PETS 2013. LNCS, vol. 7981, pp. 82–102. Springer, Heidelberg (2013)
6. Chatzikokolakis, K., Palamidessi, C., Stronati, M.: A Predictive Differentially-Private Mechanism for Mobility Traces. In: De Cristofaro, E., Murdoch, S.J. (eds.) PETS 2014. LNCS, vol. 8555, pp. 21–41. Springer, Heidelberg (2014)
7. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. In: Proceedings of the 36th Annual Symposium on Foundations of Computer Science, pp. 41–50. IEEE (1995)
8. Delaune, S., Kremer, S., Ryan, M.: Verifying privacy-type properties of electronic voting protocols. Journal of Computer Security 17(4), 435–487 (2009)
9. Dingledine, R., Mathewson, N., Syverson, P.F.: Tor: The second-generation onion router. In: Proceedings of the 13th USENIX Security Symposium, pp. 303–320. USENIX (2004)
10. Dwork, C.: Differential Privacy. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) ICALP 2006. LNCS, vol. 4052, pp. 1–12. Springer, Heidelberg (2006)
11. Dwork, C.: A firm foundation for private data analysis. Communications of the ACM 54(1), 86–96 (2011)
12. Dwork, C., Lei, J.: Differential privacy and robust statistics. In: Mitzenmacher, M. (ed.) Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), Bethesda, MD, USA, May 31-June 2, pp. 371–380. ACM (2009)
13. Juels, A.: Rfid security and privacy: A research survey. IEEE Journal on Selected Areas in Communications 24(2), 381–394 (2006)
14. McIver, A., Morgan, C., Smith, G., Espinoza, B., Meinicke, L.: Abstract Channels and Their Robust Information-Leakage Ordering. In: Abadi, M., Kremer, S. (eds.) POST 2014 (ETAPS 2014). LNCS, vol. 8414, pp. 83–102. Springer, Heidelberg (2014)

15. Narayanan, A., Shmatikov, V.: De-anonymizing social networks. In: Proceedings of the 30th IEEE Symposium on Security and Privacy, pp. 173–187. IEEE Computer Society (2009)
16. Samarati, P.: Protecting respondents' identities in microdata release. IEEE Trans. Knowl. Data. Eng. 13(6), 1010–1027 (2001)
17. Samarati, P., Sweeney, L.: Generalizing data to provide anonymity when disclosing information (abstract). In: ACM (ed.) PODS 1998. Proceedings of the ACM SIGACT–SIGMOD–SIGART Symposium on Principles of Database Systems, Seattle, Washington, June 1-3, pp. 188–188. ACM Press (1998)