Differential Privacy Based Access Control

Nadia Metoui[1,2] and Michele Bezzi[3]

[1] Security & Trust Unit, FBK-Irst, Trento, Italy
[2] DISI, University of Trento, Italy
[3] SAP Product Security Research, Sophia-Antipolis, France

Abstract. The huge availability of data is giving organizations the opportunity to develop and consume new data-intensive applications (e.g., predictive analytics). However, data often contain personal and confidential information, and their usage and sharing come with security and legal risks; so there is the need of devising appropriate, task specific, data release mechanisms to find the balance between advantages of big data and the potential risks.
We propose a novel privacy-aware access control model, based on differential privacy. The model allows for data access at different privacy levels, generating an anonymized data set according to the *privacy clearance* of each request. The architecture also supports re-negotiation of the privacy level, in return of fulfilling a set of obligations. We also show, how the model can address the privacy and utility requirements, in an human-resource motivated use-case with a classification task. The model provides a flexible access control, improving data availability, while guaranteeing a certain level of privacy.

Keywords: privacy risk, differential privacy, privacy-aware access control

1 Introduction

Modern organizations have access to an increasing number of huge and diverse data sets (*big data*), and new data-usage or data-businesses are emerging, such as predictive analytics, benchmarking services, or data generation for testing. However, the usage and sharing of data come with security and legal risks, and it needs appropriate data release mechanisms. For example, part of data is often constituted by personal information, which are subject to strict regulations, and companies have to comply with them to use the data, to avoid major risks of fines or reputation. In short, organizations are becoming more concerned with setting up appropriate access control processes to better exploit data knowledge, balancing advantages of big data with the potential security and legal risks.

In this context, risk-based access control models [1,2] provide an important step towards a better management and exploitation of data. In fact, these models bring more flexibility, replacing (or integrating) pre-defined access control

policies, with access decisions based on the risk estimation of specific requests compared to a user/role dependent risk thresholds. However, most of existing risk-based access control models only support a binary access decision (i.e., the outcome is either fully allowed or denied access), which is not the most convenient approach to adopt, especially in a privacy context, where a third option could be allowing access to a partial and/or sanitized (anonymized) version of the data.

The anonymized data could still be useful *enough* for certain applications (or even required), and expose to lower privacy-risk. We recently proposed an extension to risk-aware access control systems, including data anonymization [3], based on k-anonymity. Although widely used in practice, k-anonymity framework (and the related family of syntactic privacy metrics [4]), is susceptible to various attacks (e.g., [5]), and, in the last 10 years, another formal approach has been proposed to provide strong privacy guarantee: differential privacy [6].

In this paper, we propose a privacy-aware access control model, which uses differential privacy to reduce the disclosure risk. The model, in case the access to raw data is not permitted, is able to provide a differential private data set, according to the *privacy clearance* of the user. This allows for a more flexible access, improving data availability, and at the same time, guaranteeing a formal level of privacy.

The main contributions of the paper are as follows:

- We propose a privacy-aware access control model that evaluates access and clearance decisions based on a privacy-preserving approach.
- We propose to use a differential private algorithm to enforce these decisions, respecting the adequate privacy level.
- We define an architecture for our access control system, which integrates a classical policy based access control, and also supports mechanisms for (temporally) increasing privacy clearance.
- We implement a proof-of-concept prototype and run preliminary experiments, to evaluate the utility of the data, using a simple classification task, and the performance of the system.

Structure of the Paper: In the next section, we provide a motivating use case for our work. In Section 3, we give a short overview on Differential Privacy. In Section 4, we introduce our privacy-aware access control model. Section 5 is dedicated to the description of the architecture of the access control framework. In Section 6, we describe the experimental evaluation and discuss the main results. In Section 7, we discuss the related work in terms of privacy- and risk-aware access control approaches, and, in Section 8, we conclude with some final remarks.

2 Use Case

Human Resource (HR) data are becoming increasingly important for the management of the company workforce. Whereas, traditionally, they were mostly

accessed in tabular form from the HR department and people managers, there is nowadays a large number of additional analytics and functionalities to improve HR key processes [7,8] (e.g., talent discovery, compensation process, trainings), and, correspondingly, there is an increased need for access to HR data, reports and analytics, involving multiple actors in the company. At the same time, HR data contain sensitive and personal information, which is subject to, often complex, data protection regulations, and data access should carefully manage.

For example, an HR manager can have a full view on the HR information for her/his department, but an aggregated view for the HR data from other departments. In some cases, for example employee survey results for collecting employee feedback, a certain level of anonymity is needed even for the data within the department. Legal framework, such as the European data protection regulations [9], can additionally impose geographical constraints on access and transfer of personal information.

HR data are also needed for the testing phase in the development of HR applications. In this case, the data should not contain personal information, but they should be realistic *enough* to allow for significant testing. So, an in-house developer may have access in a controlled environment to an anonymized version of the data. If the development task is outsourced to an external company, an even stronger anonymized is likely needed.

HR data (e.g., compensation and healthcare cost data) are also sometimes shared with external parties for benchmarking purpose (see e.g. Bureau of Labor Statistics (BLS) [10]). and, for that scope, they need an high level of privacy guarantees to be released.

The requirements of these illustrative examples can be summarized as in Table 1.

Table 1. Possible usage scenarios, comprising different devices and locations, and expected utility (i.e., type of reports needed) and security levels.

#	Role	Operation	Risk	Utility
1	HR manager	HR view (int.)	Low Risk	full access
2	HR manager	HR view (ext.)	Medium Risk	aggregated
3	HR developer	Testing data	Medium-High Risk	anonymized
4	HR Benchmarking	Benchmark	High Risk	anonymized

These scenarios show how a rather complex access control framework should be set up to address the privacy requirements. Currently, in most cases, these requirements are addressed with a mix of specific configurations of traditional access control systems (e.g, RBAC systems for the HR manager use-case), usage of specific anonymization tools (e.g., for releasing data for application testing or benchmarking services), and, often, relying on human-based processes. In the next sections, we will show how these scenarios can be realized in our framework.

3 Background

Differential privacy [6] is a privacy framework devised for providing a, formal, strong privacy guarantee. Whereas, traditionally, privacy preserving data publishing was based on syntactic privacy [4] mechanisms, where, for example, it is imposed as condition that a record being indistinguishable from k other records [11] (equivalence group), or the sensitive values to be *well* distributed within the equivalence groups [12,13], differential privacy takes another approach, requiring that the answer to any query being *probabilistically indistinguishable* if a particular record is present in the database or not. In other words, an adversary cannot learn (almost) anything about an individual record, since the output does not (almost) change, whether that specific record is present or absent in the data set. Following [14], we can define differential privacy, in the context of privacy-preserving data publishing, as:

Definition 1. *A randomized algorithm $\mathcal{K}$ satisfies ϵ-differential privacy if for all pairs of data sets D, D', differing for at most one record ($D \sim D'$), and for all possible anonymized data sets $\hat{D}$, we have that:*

$$Pr[\mathcal{K}(D) = \hat{D}] \leq e^{\epsilon} \times Pr[\mathcal{K}(D) = \hat{D}]$$

where the probability is computed over the randomness of $\mathcal{K}$, and the parameter $\epsilon > 0$ sets the bound of the privacy guarantee, with low values of ϵ providing stronger privacy.

The mechanism for providing differential privacy (called ϵ-differentially private sanitizer) is typically based on noise addition. There are two approaches: interactive (for privacy-preserving data mining) and non-interactive (for privacy preserving data publishing). Historically, differential privacy was devised for the interactive framework [6]: a user sends a set of queries to a data base, and the data base owner, to assure privacy, adds some random perturbation to the query answer (e.g., adding Laplace noise with variance related to ϵ parameter). Although the interactive framework is mostly used, it has some drawbacks [15], e.g., after a limited number of queries the noise level should be increased, highly impacting the utility.

In the non-interactive framework the database owner anonymizes the original raw data, and then releases the anonymized version, providing the user a greater flexibility for data analysis, and basically no limitation in terms of queries. In this paper, we use the non-interactive framework, since we are dealing with access control of tables, which is analogous to data publishing; it has been shown that differential privacy can be used for data publishing [16,15,14], although with some limitations (see [17]), and generic assumption on the utility needed (e.g., assuming that the analysis relies mostly on the counts of certain attributes).

In particular for deriving differential private data set for our evaluation (see Sect. 6), we follow the approach of [14]. The method considers the raw data, and it computes the contingency tables, counting the number of records sharing a a combination of attributes. Then, it probabilistically (using an exponential mechanism) generates a generalized contingency table (generalizing attribute values

in wider classes). Then, it applies Laplacian noise to the generalized contingency table. The generalization step allows to increase the counts for the cells, resulting in lowering the utility-impact of the noise addition. Synthetic data can be produced from the generalized and randomized contingency table. The resulting data set, generated by a ϵ-differential privacy mechanisms, can be safely used for any data analysis (we will test it on a classification task, as in [14]).

4 Model

In this section we provide a general description of our Differential Privacy-Based Access Control model for tabular data. The access control model we proposed can be considered as an extension of the traditional policy based model, such as XACML model [18], augmented with the possibility, once the access to the raw data is denied, to get access to an anonymized version of the data.

Generally speaking, the model proceeds as follows: whenever a user/role needs to access a data set, the access control model checks if the request can be fulfilled, comparing the user/role access rights with the access control policy of the data set. Differently from the classical policy based access control, the system, in addition of a allow or deny decision, can deny access to the data set in the raw version, but still provide the user with an anonymized version of the data.

More formally, the access evaluation can be represented by the function $Auth(obj, u)$ defined as follows. User u[1] is granted access to an object obj (say the HR data of a department) if the access control policy of the object P_{obj} includes user u (say the people manager of the department). The policy could also specify, that the access is only provided to an anonymized version of the data (say for the people managers of other departments), in this case, the system retrieves the *privacy clearance* value, T_ϵ, associated to the user/request, and it applies the differentially private sanitizer ($Sanitize(T_\epsilon)$) to the original data to attain a data set of differential privacy $\epsilon = T_\epsilon$[2]. Access is denied in the other cases (say for people outside the company), i.e.,

$$Auth(obj, u) = \begin{cases} \textbf{Allow} & \text{if } u \in P_{obj} \\ \textbf{Sanitize}(T_\epsilon) & \text{if } \{u, \epsilon\} \in P_{obj} \\ \textbf{Deny} & \text{otherwise} \end{cases} \quad (1)$$

Where $T_\epsilon(u, C)$ is the privacy clearance of the request, which depends on user u and context information C (e.g., within the corporate network users may have

[1] In most cases the dependency is mediated by roles and permissions. For the sake of simplicity, we do not consider roles, and focused only on read access, for an extension of this model including roles, we can follow the lines of access control risk models as described in [19,20].

[2] Note that the system may have already in the cache the anonymized data set, if it had received the same data request at the same privacy clearance. In this case, there is no need to re-anonymize the data, and it uses the already produced data set, improving performance and security.

a larger clearance).

Note that the privacy clearance parameter T_ϵ, here, plays a role similar to the privacy budget [21] typically used for differential privacy models. But, in our case, we only consider accessing disjoint sets of data, so each user/role can *spend* all his/her budget for a single request, and he/she has access to data at the same, or lower, level as the privacy clearance. This is similar to the security clearance parameter in multi-level security models.

Adding the option of providing anonymized data can increase the flexibility and, ultimately, the access to data. On the other hand, especially for small privacy clearance, T_ϵ, the high level of anonymization can severely impact the utility, making the data not usable. To this aim, we foresee mechanisms to (temporarily) increase the privacy clearance, for example asking the user to fulfill some obligations (as we proposed in [3]). The architecture described in Sect. 5 can support these Privacy Clearance Enhancement functionalities, but, we do not discuss them in details in this study, focusing on data sanitization.

5 Architecture

In this section we present an abstract architecture for our Privacy-Aware Access Control Framework. The architecture, depicted in Figure 1, is composed of three main modules, which are described in the remaining part of the section.

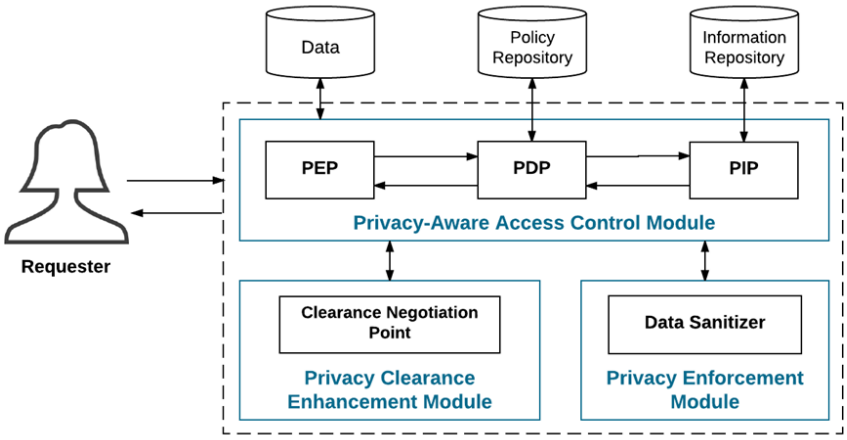

Fig. 1. Architecture of the Privacy-Aware Access Control framework

Privacy-Aware Access Control Module is the entry point of the system, through which users can submit requests to retrieve data from the underlying

database. This module evaluates the access request, and it grants access to (original or sanitized version of) the requested data or denies access.
For this scope, the Privacy-Aware Access Control Module assesses the data request against an access policy to determine whether the requester has the needed authorizations to access the resource (requested data-view) and, also, to evaluate the privacy clearance (as discussed in Sect. 4). Then, the decision is enforced by calling the Privacy Enforcement Module or renegotiate by calling the Privacy Clearance Enhancement Module.
The Privacy-Aware Access Control Module is based the on the XACML (eXtensible Access Control Markup Language) standard [18]. XACML is a declarative fine-grained, access control policy language. The standard also provides an access control architecture and a description of the access evaluation process (data-flows, access request, access decision etc.)
In this module Access Control is realized internally using a PEP-PDP [3] pair. A PIP (Policy Information Point) is used to provide additional information needed to evaluate the request and estimate its privacy clearance (e.g., in our use case if the requester is a manager, we would like to know her/his department in order to define her/his privacy clearance, if the requested data contains information about his department this clearance will be higher than about an other departments)

Privacy Enforcement Module. After evaluation of the access request, the Privacy Enforcement Module receives a data view (non-anonymized version) and a privacy clearance value. The role of this module is applying data sanitization algorithms, and generating an anonymized version of this data view, according to the privacy clearance.

Privacy Clearance Enhancement Module. The Privacy clearance defined by the Privacy-Aware Access Control Module can be re-negotiated to a higher level in some cases (e.g., for example if the utility of the anonymized data is not sufficient) to allow more flexibility. The user can ask (temporally) for higher clearance, in exchange, for example, of fulfilling some obligations to mitigate the additional risk. These operations are typically expressed as access and usage control obligations (see [3]), for example imposing deletion of a resource after that a retention period expires, or providing stronger authentication credentials.

6 Experimental Evaluation

In order to evaluate the practical feasibility of our approach, we developed a proof-of-concept implementation of the framework, to assess: *i)* the impact of

[3] In XACML the PDP is the point that evaluates an access request against an authorizations policy and issues an access decision and the PEP Policy Enforcement Point is the point that intercept user's request, it calls the PDP for an access decision then it enforces the decision by allowing or denying the access.

our privacy preserving access control on the data quality. To this aim, we defined a simple classification task, and test the performance using data sanitized at different privacy clearances. *ii)* to evaluate the impact of the enforcement of different privacy clearance levels (anonymization by applying differential privacy) on the performance of our access control system, in terms of response time.

To address these questions, we implemented a prototype of our Privacy Enforcement Module as described in Sect. 5. As data sanitizer we used *"DiffGen"* a Differentially-private anonymization algorithm based on Generalization, proposed and implemented by Mohammed et *al.* in [14].

DiffGen anonymizes the raw data by probabilistically generalizing the attributes. More in details, starting from the most general state (one-single group), a set of specializations are randomly selected, using an exponential mechanism with a predefined scoring function (e.g. a utility-based function assessing the information gain for each specialization). Then, the algorithm computes the contingency tables, counting the number of records sharing a combination of attributes, and, it applies Laplacian noise, with variance ϵ, to the generalized contingency table. Synthetic data can be then produced from the generalized and randomized contingency table (see in [14] for details). The resulting data set, generated by a ϵ-differential privacy mechanisms, can be safely used for any data analysis.

In our case, to represent the use case, described in Sect. 2, ϵ will be expressed on term of privacy clearance (T_ϵ) as in Table 2.

Table 2. Example of Privacy Clearance levels for use case in Sect. 2.

#	Role	Operation	Risk	Privacy Clearance
1	HR manager	HR view (int.)	Low Risk	$T_\epsilon > 1$
2	HR manager	HR view (ext.)	Medium Risk	$T_\epsilon \in]0.1, 1]$
3	HR developer	Testing data	Medium-High Risk	$T_\epsilon \in]0.05, 0.1]$
4	HR Benchmarking	Benchmark	High Risk	$T_\epsilon \leq 0.05$

For our test, we use the Adult Data Set [4] from the UCI Machine Learning Repository. This dataset contains $45K$ records from the US Census data set with 15 demographic and employment-related variables (6 numerical , 8 categorical, and 1 binary class column representing two income levels, $\leq 50K$ or $> 50K$). The Experiments were conducted on an Intel Core i5 2.6GHz PC with 8GB RAM.

Data Quality To evaluate the impact of anonymization on the Utility of data we propose to assess its impact on the accuracy of a simple classifier trained and tested using anonymized data at different clearance levels (shown in Table 2).

We use as (binary) class attribute the income level, $\leq 50K$ or $> 50K$, and as classifier the well-known C4.5 Algorithm [22]. Each anonymized data set is split

[4] Available at http://archive.ics.uci.edu/ml/datasets/Adult

in two. First part of the data (2/3) is used as training data to build a classifier, and the remaining data (1/3) is used as test data to measure the classification accuracy.

In Fig. 2, we report the accuracy of classifiers for different privacy clearances. We can observe that for small values of T_ϵ, the accuracy is highly impacted.

In fact with $T_\epsilon = 0.01$ the attributes are almost fully generalized, and the accuracy is close to the case where all the attributes (but the class attribute, of course) are removed. Still, the accuracy level of $\simeq 75\%$ could be enough for many benchmarking tasks.

The accuracy goes up, as expected, for higher privacy clearance values. Privacy clearance in the range of $[0.1, 0.5]$, still considered reasonably safe in practical cases, allows to produce data able to provide an accuracy close to 80%, which it could be sufficient as testing data for development, and to have a general view for a manager on other department analytics.

Privacy clearance > 1 (manager view on own team data, in our use case), gives levels of accuracy close to the raw data $\simeq 85\%$.

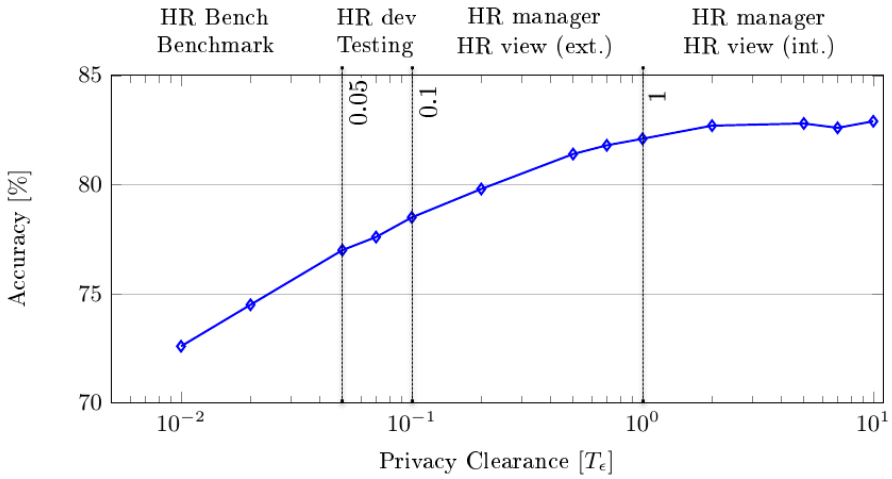

Fig. 2. Classifier Accuracy for different privacy clearance T_ϵ. Each data point represent the average over 100 runs (parameters of DiffGen: number of specialization of specialization $N_s = 10$, and scoring function $u = Max$).

Performance We estimate the computational overhead caused by the anonymization. From these experiments, we observe that the time for performing the anonymization can be easily of order of seconds, see Fig. 3. The effect of T_ϵ on the time is limited.

Despite being preliminary results, it is clear that for reaching real-time performance (as it is possible for k-anonymity algorithms, see [23]), it is needed

to include some optimization, for example in terms of caching or testing other algorithms for generating differential private data set.

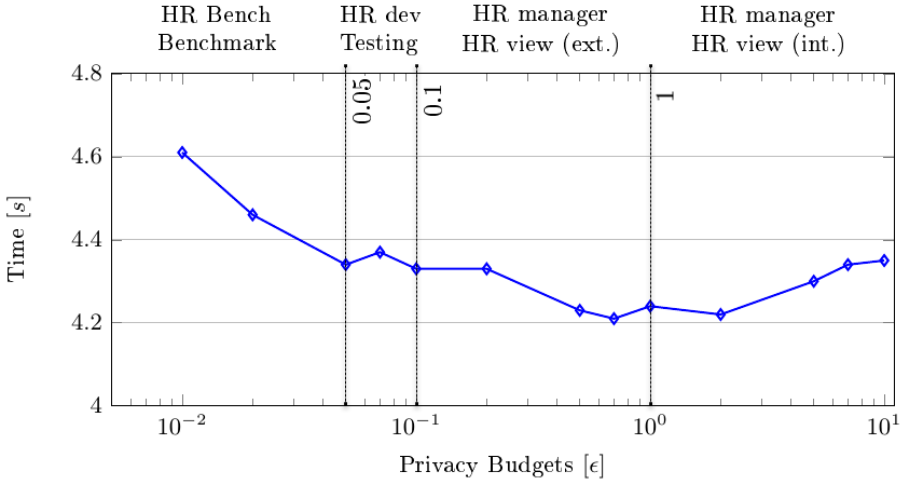

Fig. 3. Anonymization time for different privacy clearance T_ϵ. Each data point represent the average over 100 runs (parameters of DiffGen: number of specialization of specialization $N_s = 10$, and scoring function $u = Max$).

7 Related Work

Several privacy-aware access control approaches where proposed in the literature (see [24,25,26] for review), most of them rely on policies expressing privacy rules and preferences to be enforced during or after access evaluation. For instance, Purpose-based Access Control Systems [27] propose to evaluate access requests based on the purpose of the access, and they allow conditional or unconditional access only for specific proposes predefined in privacy policies. A related approach is based on the concept of *Sticky polices* [28], in this framework, privacy policies, expressing users preferences for data handling, are attached to the data, enabling to improve control over the usage of personal information and to define usage constraints and obligations as data travels across multiple parties (e.g., in the cloud). These policies based approach do not consider anonymization, nor other risk mitigation strategies.

The model proposed in this paper is inspired by the risk-aware access control models [1,29,2,30,31], in particular to the privacy risk aware access control model, which we introduced in [3]. Typically in these risk-aware models, for each access request or permission activation, the corresponding risk is estimated and, if the risk is lower than a threshold (e.g., related to trust) then the operation

is permitted, otherwise it is denied. Cheng et *al.* [2] estimate risk and trust thresholds from the sensitivity labels of the resource and clearance level of the users in a multi-level-security system. They also consider a trust enhancement mechanism (the authors call it risk mitigation strategy in their paper) that allows users to spend *tokens* to access resources with risk higher than their trust level. The details on how this mechanism can be applied in real cases are not provided. Compared to these studies, we introduce the privacy clearance concept, which plays a similar role of the clearance level in multi-level-security, but we focus on privacy, and we explicitly define a data sanitization mechanism.

Chen et *al.* [1] introduced an abstract model which allows role activation based on a risk evaluation compared to predefined risk thresholds. Trust values are considered, and they impact (decreasing) risk calculation. If risk is too high, the model includes mitigation strategies, indicated as system obligations. The paper does not specify how to compute the risk thresholds, trust, and the structure of obligations. In a derived model [29], mitigation strategies have been explicitly defined in terms of user obligations (actions that have to be fulfilled by the user). The model also introduces the concept of *diligence score*, which measures the diligence of the user to fulfill the obligations (as a behavioral trust model), and impacts the risk estimation.

Another extension has been proposed [32,23], focusing on re-identification risk and anonymization (based on k-anonymity) is used as mitigation strategy.

Although reminiscent to our approach, these models differ from our proposal since: 1) we use a formal guarantee for privacy (differential privacy) 2) we do not explicitly introduce the concept of risk, since it is very hard to estimate risk for a data set generated by a differentially private mechanisms [33]. In fact, unlike k-anonymity and related metrics, differential privacy is a property of the mechanism not of the data set and, consequently, the risk estimation cannot be immediately derived by the privacy parameter (ϵ) [34].

Our research is focused on the *usage* of differential privacy for access control, and as such it does not propose any novel algorithm for differential privacy. Indeed, our experiments are based on the DiffGen framework proposed in [14]. For the differential privacy literature, we refer the reader to [4,6], as well as the references introduced in Sec. 3.

8 Conclusions and future work

In this paper we proposed a novel privacy-aware access control model, based on differential privacy. The model allows for data access at different privacy levels, generating an anonymized data set according to the privacy clearance of the request. To evaluate our approach we developed a proof-of-concept prototype.

The first experimental assessment, considering a HR related use case, and a benchmarking data set, indicates that the model can address complex privacy and utility requirements. However, it also presents a number of open issues to be solved for a practical usage. For example, the performance of the current imple-

mentation, are not in real-time, therefore different algorithms and optimization strategies for the anonymization need to be investigated.

In addition, whereas in previous models we used the concept of privacy risk, which has a clear business interpretation, here we used the ϵ parameter of differential privacy. In future works, we would like to relate the two approaches, including explicitly privacy-risk assessment and adjustment mechanisms based on the concepts of differential identifiability [34] and interactive differential privacy [6].

We also presented an architecture, which supports mechanisms for increasing the privacy clearance of the user, we do not detail this part in the study, but we will investigate it in future works, in particular how it could be realized using obligations and how we can implement and enforce these obligations using the access control policies.

Acknowledgments The research leading to these results was supported by the EU-funded project TOREADOR (contract n. H2020-688797) and FP7 EU-funded project SECENTIS (FP7-PEOPLE-2012-ITN, grant no. 317387).

References

1. Chen, L., Crampton, J.: Risk-aware role-based access control. In Meadows, C., Fernandez-Gago, C., eds.: Security and Trust Management. Volume 7170 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2012) 140–156
2. Cheng, P.C., Rohatgi, P., Keser, C., Karger, P.A., Wagner, G.M., Reninger, A.S.: Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In: IEEE Symposium on Security and Privacy, IEEE Computer Society (2007) 222–230
3. Armando, A., Bezzi, M., Di Cerbo, F., Metoui, N.: Balancing trust and risk in access control. In: On the Move to Meaningful Internet Systems: OTM 2015 Conferences. Springer International Publishing (2015) 660–676
4. Clifton, C., Tassa, T.: On syntactic anonymity and differential privacy. Trans. Data Privacy **6**(2) (August 2013) 161–183
5. Ganta, S.R., Kasiviswanathan, S.P., Smith, A.: Composition attacks and auxiliary information in data privacy. In: Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '08, New York, NY, USA, ACM (2008) 265–273
6. Dwork, C.: Differential privacy. In Bugliesi, M., Preneel, B., Sassone, V., Wegener, I., eds.: Automata, Languages and Programming. Volume 4052 of Lecture Notes in Computer Science. Springer Berlin / Heidelberg (2006) 1–12 10.1007/11787006_1.
7. SERVICES, I.G.B.: Getting smart about your workforce: Why analytics matter. Technical report, IBM CANADA (March 2009)
8. Martin, L.: Getting smart about your workforce: Why analytics matter. Technical report, Oracle Corporation & CedarCrestone (September 2011)
9. of Europe, C.: Handbook on european data protection law. Technical report (2014)
10. Royster, S.: Working with big data. Technical report, U.S. Bureau of Labor Statistics (2013)

11. Samarati, P., Sweeney, L.: Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Technical report, Technical report, SRI International (1998)
12. Machanavajjhala, A., Kifer, D., Gehrke, J., Venkitasubramaniam, M.: l-diversity: Privacy beyond k-anonymity. ACM Transactions on Knowledge Discovery from Data (TKDD) **1**(1) (2007) 3
13. Li, N., Li, T., Venkatasubramanian, S.: t-closeness: Privacy beyond k-anonymity and l-diversity. In: 2007 IEEE 23rd International Conference on Data Engineering, IEEE (2007) 106–115
14. Mohammed, N., Chen, R., Fung, B.C., Yu, P.S.: Differentially private data release for data mining. In: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '11, New York, NY, USA, ACM (2011) 493–501
15. Soria-Comas, J., Domingo-Ferrer, J., Sánchez, D., Martínez, S.: Enhancing data utility in differential privacy via microaggregation-based $$k$$ k -anonymity. The VLDB Journal **23**(5) (2014) 771–794
16. Li, N., Qardaji, W., Su, D.: On sampling, anonymization, and differential privacy or, k-anonymization meets differential privacy. In: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security. ASIACCS '12, New York, NY, USA, ACM (2012) 32–33
17. Leoni, D.: Non-interactive differential privacy: a survey. In: Proceedings of the First International Workshop on Open Data, ACM (2012) 40–52
18. Moses, T., et al.: extensible access control markup language (xacml) version 2.0. Oasis Standard **200502** (2005)
19. Chen, L., Crampton, J.: Risk-aware role-based access control. In Meadows, C., Fernandez-Gago, C., eds.: Security and Trust Management. Volume 7170 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2012) 140–156
20. Baracaldo, N., Joshi, J.: An adaptive risk management and access control framework to mitigate insider threats. Computers and Security **39, Part B**(0) (2013) 237 – 254
21. Peng, S., Yang, Y., Zhang, Z., Winslett, M., Yu, Y.: Query optimization for differentially private data management systems. In: Data Engineering (ICDE), 2013 IEEE 29th International Conference on, IEEE (2013) 1093–1104
22. Salzberg, S.L.: C4. 5: Programs for machine learning by j. ross quinlan. morgan kaufmann publishers, inc., 1993. Machine Learning **16**(3) (1994) 235–240
23. Armando, A., Bezzi, M., Metoui, N., Sabetta, A.: Risk-based privacy-aware information disclosure. Int. J. Secur. Softw. Eng. **6**(2) (April 2015) 70–89
24. Ghani, N.A., Selamat, H., Sidek, Z.M.: Analysis of existing privacy-aware access control for e-commerce application. Global Journal of Computer Science and Technology **12**(4) (2012)
25. Ardagna, C., De Capitani di Vimercati, S., Paraboschi, S., Pedrini, E., Samarati, P., Verdicchio, M.: Expressive and deployable access control in open web service applications. IEEE Transactions on Service Computing (TSC) **4**(2) (April-June 2011) 96–109
26. Ardagna, C.A., Cremonini, M., De Capitani di Vimercati, S., Samarati, P.: A privacy-aware access control system. J. Comput. Secur. **16**(4) (December 2008) 369–397
27. Byun, J.W., Bertino, E., Li, N.: Purpose based access control of complex data for privacy protection. In: Proceedings of the Tenth ACM Symposium on Access Control Models and Technologies. SACMAT '05, New York, NY, USA, ACM (2005) 102–110

28. Pearson, S., Casassa-Mont, M.: Sticky policies: An approach for managing privacy across multiple parties. Computer **44**(9) (2011) 60–68
29. Chen, L., Crampton, J., Kollingbaum, M.J., Norman, T.J.: Obligations in risk-aware access control. In Cuppens-Boulahia, N., Fong, P., García-Alfaro, J., Marsh, S., Steghöfer, J., eds.: PST, IEEE (2012) 145–152
30. Dickens, L., Russo, A., Cheng, P.C., Lobo, J.: Towards learning risk estimation functions for access control. In: In Snowbird Learning Workshop. (2010)
31. Shaikh, R.A., Adi, K., Logrippo, L.: Dynamic risk-based decision methods for access control systems. Volume 31. (2012) 447–464
32. Armando, A., Bezzi, M., Metoui, N., Sabetta, A.: Risk-aware information disclosure. In: Data Privacy Management, Autonomous Spontaneous Security, and Security Assurance. Springer (2015) 266–276
33. McClure, D., Reiter, J.P.: Differential privacy and statistical disclosure risk measures: An investigation with binary synthetic data. Trans. Data Privacy **5**(3) (December 2012) 535–552
34. Lee, J., Clifton, C.: Differential identifiability. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '12, New York, NY, USA, ACM (2012) 1041–1049