A Privacy Model to Support Minimal Disclosure in Virtual Organizations

Fabio Massacci
University of Trento
massacci@dit.unitn.it

John Mylopoulos
University of Trento
jm@dit.unitn.it

Nicola Zannone
University of Trento
zannone@dit.unitn.it

Abstract

The last years have seen an increasing attention on privacy-aware technologies and mechanisms for the negotiation of private information between customers and enterprises. Unfortunately, current proposals are still unsatisfactory since they do not cover the entire spectrum of privacy management. Moreover, they do not provide support for emerging business models such as the inter-organizational business process (also known as virtual organizations). In this paper we propose a privacy model complying with the minimal disclosure principle when a coalition of organizations integrate their efforts to provide services to customers.

1 Introduction

Data protection and privacy are becoming central parts of today's IT landscape in business contexts. This reflects the growing attention of customers to their personal data and the increasing number of laws, policies, and regulations that are intended to safeguard them. These define privacy as the right of data subjects to determine how their data can be used.

Many research efforts have proposed new privacy-aware technologies [1, 3, 6] that attempt to comply with the above definition. In particular, several languages, frameworks, and standards based on the notion of *purpose* (i.e., intended usage of information) have been proposed in the last years. They cope with different privacy aspects ranging from the specification of privacy policy and the enforcement of privacy promises and customer preferences across the entire organization to the negotiation of personal information. However, there are no proposals which deal with the entire spectrum of privacy management. Consequently, enterprises do not have technologies for fully managing and enforcing privacy.

Some proposals attempt to bridge the gap between the specification of privacy policies and their enforcement but they only address a small part of the problem [7]. Together with the notion of purpose, privacy legislations also define a number of privacy principles that organizations have to meet in order to guarantee customer privacy. At the basis of the exchange between enterprises and customers, there is the principle of *transparency*: enterprises should publish their privacy policy stating which data are collected and for what purpose. Another important principle is the notion of *minimal disclosure*: enterprises should maintain only information necessary to provide the service for which it was collected. Looking at the privacy policy, customers would be able to understand how their personal data will be used and, in case they agree, disclose them. Obviously, if a customer is not confident with the enterprise's privacy practices, they can get their services elsewhere.

In real business scenarios, enterprises offer customers with different alternatives for getting services. Customer than chooses one of them. However, this selection can be very difficult when an enterprise provides many alternatives. It gets worse when we do not have a single organization that provides the service, but rather a set of collaborating organizations participating in a single business process (*virtual organization* [4]).

Virtual organizations offer services that can be dynamically customized in at least two different ways. Different service components are dynamically chosen for fulfilling the same high level service, but possibly using different data. Moreover, different partners (subcontractors) are chosen for provide the same service, but possibly with different service level agreements. The selection of partners and the identification of a particular alternative have significant impact on the privacy of individual customers. Different customizations may require different data for which customers have its own privacy sensitivities. There might be also trusted and untrusted partners offering the same service. In this setting, customers may desire to get the service through the process that best protects their privacy based on their preferences.

Therefore, we are interested in solutions that support both customers and companies. On one side, virtual organizations should be able to publish comprehensive privacy policies spanning their partners. On the other side, the solution should allow customers to personalize services they want on the basis of their own preferences.

This position paper discusses a privacy model to automatically derive the "minimal" access control matrix governing inter-organizational business processes. In particular, our approach provides (1) a framework to model business processes spanning across multiple partners, (2) algorithms to select during the design phase the business plan and the business partners that fulfill the desired purpose while minimizing the private information requested from customers, (3) on-line algorithms that allow a customer to change its privacy preferences on-the-fly and to recompute the selection of the business plan and the business partners according the customer's own criteria for minimal disclosure.

2 Modeling Privacy Policies across Virtual Organizations

Since our reference business model is that of virtual organizations, we assume that there will often be more than one alternative to provide services to customers. In particular, enterprises might need to "decompose" a service into more specific ones since they are not completely able to entirely provide it by themselves, and so they may outsource the provisioning of sub-parts to third parties. These issues mainly affect the definition of the access control policy governing the business process since the privacy policy cannot be directly mapped to it without introducing unnecessary authorizations. Therefore, enterprises need a framework able to express complex business strategies and reason about them.

The common approach to overcome above issues is to decompose services into its subparts and consider them as "independent" services. Customers then can opt in or out the use of their data for each subservice. However, using such an approach we lose the logical relation between a service and its parts. In particular, this notation does not support the modeling of alternatives and, consequently, it does not allow reasoning about the effective delivery of services. For example, a customer might opt out a subservice that, however, is necessary for the delivery of the required service. Therefore, the enterprise may collect information that is altogether insufficient to provide the service.

We propose to model and analyze privacy policies using goal-oriented approaches. Accordingly, we analyze purposes and organize them into AND/OR tree hierarchies. In essence, AND-decomposition defines a plan for achieving a purpose, while OR-decomposition defines alternatives for achieving a purpose. Once we have built such a hierarchy, we need a data structure to represent it in order to design algorithms for reasoning about the fulfillment of root purposes. Our choice for this is hypergraphs [2], where AND- and OR-decompositions can be represented as hyperarcs.

Each alternative process is represented in our model as a decomposition path, so that different decomposition paths may exists to reach the high level service. Such paths may differ in an important aspect, notably they may require different private data items. Thus, depending on each customer's individual preferences, the same decomposition path might have a significantly different privacy "cost" for different customers. Accordingly, the model allows customers to express their preferences in the form of privacy penalties associated with each personal data item and each partner of the business process. The first represents the privacy penalty to disclose data items; the latter represents the privacy penalty of delegating information to sub-contractors. Both these assignments are given by customers based on their own preferences. A sketch of the procedure for representing the privacy policy governing a virtual organization through hypergraphs is given in Figure 1 (see [5] for more details).

1	For each supplier privacy policy table, purposes are analyzed through a goal refinement process.
2	Once a hypergraph for each supplier is defined, the hypergraph representing the privacy policies of the entire business process is built by merging them.
3	Each purpose is associated with the data items directly needed to achieve it.
4	A privacy penalty is associated with each decomposition arc.

Figure 1: Process for building privacy policy of virtual organization

Figure 2 presents the hypergraph representation of a simple privacy policy governing an e-Store and its sub-contractors. In such a figure broken lines partition the hypergraph into subgraphs, and each of them represents the privacy policy of single partner involved in the business process; arcs linking nodes across subgraphs represents delegation arcs; finally, the leafs represent data items.

3 Deriving an Access Control Policy Complying with Minimal Disclosure

Based on the purpose hierarchy derived through a goal refinement process, we propose algorithms for finding the process with the smallest privacy penalty [5]. Such a process can then be used to determine the minimum set of authorizations necessary to deliver a service according to customer preferences. In particular, the framework permits us to automatically

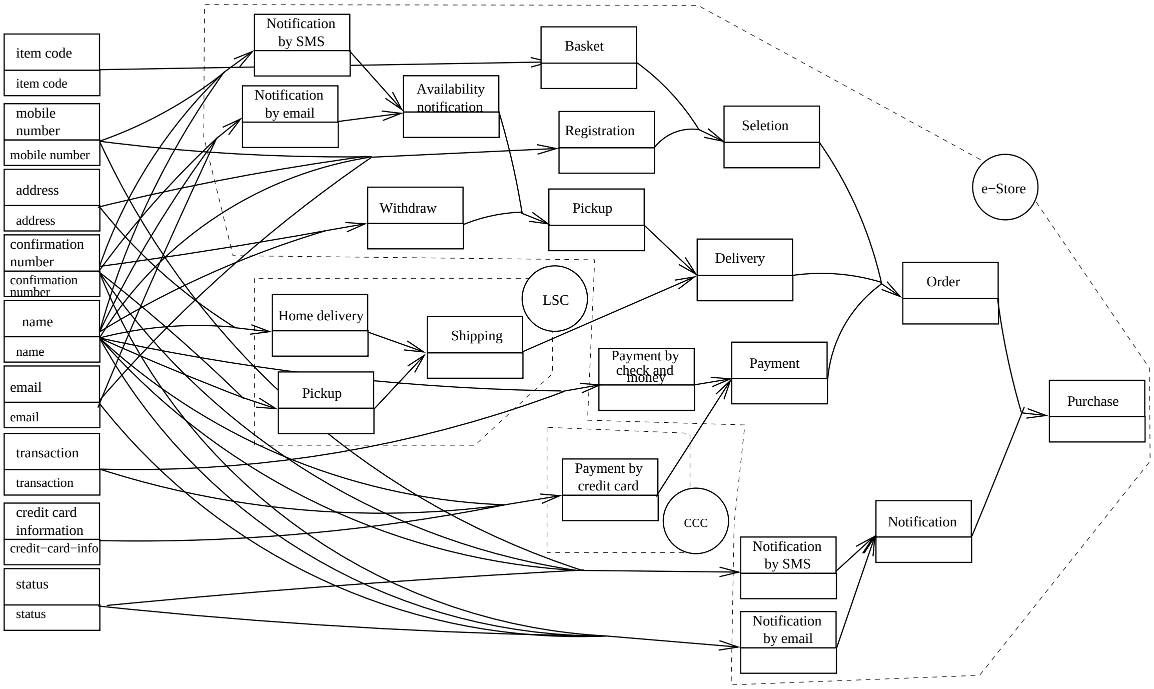

Figure 2: Purpose DAG

derive access control policies for an inter-organizational business process from the collection of privacy policies associated with different participating enterprises.

To support the dynamic customization of the process, we also give algorithms for efficiently updating it when customers change the privacy penalty of data items or choose among the alternatives that an enterprise offers for achieving a required service. By using effective on-line algorithms, the business process can be customized at run-time and guarantees maximal privacy protection because it was selected with criterion of the smallest privacy penalty.

4 Conclusion

We have proposed a privacy model for supporting the management of privacy-sensitive data within business processes provided by virtual organizations. Such a model does not pretends to substitute current technologies for privacy management but it aims at adding a piece to the privacy puzzle. In particular, our framework takes place between the specification and the enforcement of privacy promises.

There are clearly open issues. There is the need to extend current transaction models for virtual organizations in order to ensure that no information is disclosed until all partners have committed to the delivery of their part of the business plan. Our algorithms guarantee that the business process can be completed but a certificate may be desirable in some cases. The commitment should ensure customers that the virtual organization can actually provide the service once they have disclosed the requested personal information.

Acknowledgments

This work was partly supported by the projects RBNE0195K5 FIRB-ASTRO, 016004 IST-FP6-FET-IP-SENSORIA, 27587 IST-FP6-IP-SERENITY, 2003-S116-00018 PAT-MOSTRO, 1710SR-B/P PAT-STAMPS.

References

[1] P. Ashley, S. Hada, G. Karjoth, C. Powers, and M. Schunter. Enterprise Privacy Authorization Language (EPAL 1.2). W3C Recommendation, 2003. Available at http://www.w3.org/Submission/EPAL/.

[2] G. Ausiello, P. G. Franciosa, and D. Frigioni. Directed Hypergraphs: Problems, Algorithmic Results, and a Novel Decremental Approach. In *Proc. of ICTCS'01*, *LNCS 2202*, pages 312–327. Springer, 2001.

[3] L. Cranor, M. Langheinrich, M. Marchiori, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. W3C Recommendation, Apr. 2002.

[4] Y. Malhotra. *Knowledge Management and Virtual Organizations.* Idea Group, 2000.

[5] F. Massacci, J. Mylopoulos, and N. Zannone. Hierarchical Hippocratic Databases with Minimal Disclosure for Virtual Organizations. *VLDBJ*, 2006.

[6] OASIS. Privacy policy profile of XACML v2.0. OASIS Standard, 2005. Available at http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-privacy_profile-spec-os.pdf.

[7] C. S. Powers, P. Ashley, and M. Schunter. Privacy promises, access control, and privacy management. Enforcing privacy throughout an enterprise by extending access control. In *Proc. of ISEC'02*, pages 13–21. IEEE Press, 2002.