

POLICY REFINEMENT CHECKING (EXTENDED ABSTRACT)

NIKOLAOS PAPANIKOLAOU, SADIE CREESE, MICHAEL GOLDSMITH

ABSTRACT. We introduce refinement checking for privacy policies expressed in P3P and XACML. Our method involves a translation of privacy policies to a set of process specifications in CSP, which describe how the privacy policy is enforced. The technique is described through an example involving medical data collected by a biobank.

1. INTRODUCTION

We are interested in the use of privacy policies as a means to protecting digital assets. It can be necessary to compare two related policies, particularly when these policies are expressed using different languages. However, even when two policies are expressed in the same language, it is possible that they will differ syntactically but supposedly express the same intention. If mistakes are made in the comparison of privacy policies this could result in the wrong policy being implemented and exposure to risk.

The Privacy Preferences Platform Project (P3P) at the W3C is focussed on developing machine readable XML for expressing websites' privacy policies and users' privacy preferences; this is intended to enable the use of privacy-aware browsers and to allow websites to collect and process information they may require in a fashion that respects user privacy. The policy languages developed within P3P [1] so far are lacking a formal semantics and hence are prone to inconsistency and ambiguity [3]. XACML is a language for expressing role-based access control, and has been extended with a profile for expressing privacy policies, but also lacks a formal semantics. The lack of a widely accepted semantics for privacy policies is the main source of difficulty in policy comparison.

2. OUR APPROACH

Policy refinement [5] is a term used to refer to the process of synthesising lower level policies from policies expressing higher level concerns. In the study of policy refinement one sees policies which achieve the same overall goal, but show different levels of abstraction in how the goal is achieved. The checking of refinements for concurrent processes, however, has been studied extensively, and there is a clear opportunity to unify the two notions of refinement, since policies are readily expressible as sets of processes that implement policy rules.

Of interest is how any two related policies can be compared. Such comparisons become necessary in practical applications, especially when the policies in question are expressed using different languages. Even when two policies are expressed in the same language, it is possible that they will differ syntactically but supposedly express the same intention. The lack of a formal and hence unambiguous, and commonly accepted semantics for privacy policies is the main source of difficulty in policy comparison. What we are trying to address is not whether two given policies have one and the same meaning; we wish to determine whether one policy *refines* another, and we use the well-established technique of process refinement, commonly used for model checking CSP [6, 2]. Our focus on refinement as opposed to equality of policies is justified by the need to check that one policy is a valid implementation of another, as might arise in the setting of a supply chain, for example.

The key step is to model the intention of a privacy policy using a set of interconnected CSP processes which can be generated automatically from a P3P or XACML policy. We are developing a tool to perform this translation. A top-level CSP process expresses a policy as a whole, and links together a sequence of smaller processes, each of which checks the conditions contained in policy rules. We are working on the development of a tool that converts P3P policies to CSP models and interacts with the FDR model checker to enable refinement checking for these. Extending this to accept OASIS XACML is a direction for future work. We believe this is a fruitful avenue of investigation with important practical applications.

Related work includes the definition of ‘policy relations’ by May et al. [4], who proposed a high-level semantic framework for comparing policy outcomes. Their approach is mathematically elegant as it avoids comparing specific actions permitted by a policy, and focuses only on the outcomes of applying it. Yu et al. [7] have proposed a relational semantics for P3P, in an effort to give unambiguous meaning to syntactically different expressions of a single policy. Their ideas are complementary to the approach we propose here.

3. CASE STUDY: VERIFYING PRIVACY POLICIES FOR A BIOBANK

A biobank is a database in which is stored a vast amount of medical information (and even genetic material) for a large number of individuals; such a database is intended for research purposes, namely for the prevention of disease in future generations¹. Participation in a biobank is entirely voluntary, and individuals are expected to disclose personally identifiable information. We assume that individuals disclose such information online.

Suppose that a fictitious biobank requires the following data from each participant: name, date of birth, address, ethnic origin, history of family diseases, and height. These data are subject to the following privacy policy:

¹See for example the UK Biobank: <http://www.ukbiobank.ac.uk>

- the name, date of birth, address, and history of family diseases will be used by a research group working from within the biobank for internal research;
- the name, ethnic origin, and height will be disclosed by the biobank to a third party, an independent agency working on a sociological experiment; and
- the name, date of birth and height will be retained by the biobank for a maximum period of 30 days.

How can the biobank compare this policy (which is readily expressible in P3P or XACML) with a different policy, recommended by their legal team, which requires that all data be outsourced to the same third party, given that they are able to carry out all the research needed by the biobank? Through policy refinement checking we hope to be able to show formally the relationship between the levels of privacy provided to individuals by these two policies.

4. FUTURE WORK AND CONCLUSION

We are working on the development of a tool that converts P3P policies to CSP models and interacts with the FDR model checker to enable refinement checking for these. Extending this to accept OASIS XACML is a direction for future work. We believe this is a fruitful avenue of investigation with important practical applications.

REFERENCES

- [1] L. Cranor, B. Dobbs, S. Egelman, G. Hogben, J. Humphrey, M. Langheinrich, M. Marchiori, M. Presler-Marshall, J. M. Reagle, M. Schunter, D. A. Stampley, and R. Welling. The platform for privacy preferences 1.1 (p3p1.1) specification. World Wide Web Consortium, Note NOTE-P3P11-20061113, November 2006.
- [2] M. Goldsmith. *FDR2 User's Manual version 2.82*, June 2005.
- [3] G. Hogben. A technical analysis of problems with p3p v1.0 and possible solutions. In *Proceedings of W3C Workshop on the Future of P3P*, November 2002. Available at <http://www.w3.org/2002/p3p-ws/pp/jrc.html>.
- [4] M. J. May, C. A. Gunter, and I. Lee. Strong and weak policy relations. In *IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY'09)*, London, UK, July 2009.
- [5] J. D. Moffett and M. S. Sloman. Policy hierarchies for distributed systems management. *IEEE Journal on Selected Areas in Communications*, 11:1404–1414, 1993.
- [6] A. W. Roscoe, C. A. R. Hoare, and R. Bird. *The Theory and Practice of Concurrency*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1997.
- [7] T. Yu, N. Li, and A. I. Antón. A formal semantics for P3P. In *Proceedings of ACM Workshop on Secure Web Services*, Fairfax VA, USA, October 29 2004.