

Policy Refinement Checking

Dr Nick Papanikolaou

Research Fellow, e-Security Group
International Digital Laboratory
University of Warwick
<http://go.warwick.ac.uk/nikos>

Joint Work with **Sadie Creese** and **Michael Goldsmith**

Ninth International Workshop on
Automated Verification of Critical Systems
Gregynog, Wales
24 September 2009

Policy Refinement
Checking

N. Papanikolaou

Introduction

Platform for
Privacy
Preferences

Policy Refinement

Policy Refinement
Checking

EnCoRe and Case
Studies

Conclusions and
Future Work

Introduction: The Ubiquity of Policies

Policies arise in diverse computing applications:

- ▶ security policies esp. access control policies
- ▶ digital rights policies
- ▶ privacy policies
- ▶ legal policies
- ▶ policies for business practice
- ▶ ...

A policy is generally understood as a set of required behaviours (of an individual, of an enterprise, ...) under specified conditions.

A **privacy policy** specifies how a business collects and processes personal data.

Policy comparison

The **central question** of interest to us is how policies can be compared.

There is existing work on **equality** or compatibility of policies [Backes, Karjoth, Bagga & Schunter 04; Lin, Rao, Bertino & Lobo 07]

The Platform for Privacy Preferences - P3P
[<http://www.w3.org/P3P>]

- ▶ this has a **policy description language** (the "P3P" XML variant)
- ▶ and a **preferences description language** ("APPEL")
- ▶ the syntax of APPEL is designed to match elements of P3P policies for direct comparison
- ▶ Example - AT&T Privacy Bird browser plugin

P3P: Problems and Alternatives

The P3P specification (actually W3C Recommendation) defines a specific vocabulary for **purposes, data types, retention, recipients** of data that is collected online.

Problems:

- ▶ No formal semantics!
- ▶ Ambiguity - two policies may mean the same thing but be written different ways
- ▶ Semantic problems and ambiguities studied by [Hogben 02]

Interesting alternative approaches:

- ▶ OASIS XACML [Moses 05] + Privacy Profile
- ▶ Privacy RBAC [Ni, Bertino, Lobo, & Calo 09]
- ▶ Policy relations/bisimulation [May, Gunter, & Lee 09]

P3P Example: A Policy Rule (Statement)

```
<data-group base="">
<data ref="#user.name.given" />
</data-group>
<purpose>
<contact /> <tailoring /> <pseudo-analysis />
</purpose>
<recipient><ours /></recipient>
<retention>
<business-practices />
</retention>
```

- ▶ Consistency issues:
 - ▶ What if >1 rules are given for the same data item?
 - ▶ Conflicting rules could emerge!

Policy Hierarchies and Refinement

Of broader interest is the case where you have a **policy hierarchy**.

Policy refinement [Moffett & Sloman 93] is concerned with deriving lower-level policies from higher-level ones.

An enterprise has to fulfil:

- ▶ Legal and regulatory requirements for handling personal data (cf. Sarbanes-Oxley Act, Data Protection Act, Gramm-Leach-Bliley Act, ...)
- ▶ Good practice guidance (cf. Information Privacy Commissioner)
- ▶ Corporate policies specific to the enterprise
- ▶ Low-level IT policies about who can do what on which machine

Our Approach: Checking for **process refinement** between policies

Model checking the process algebra **CSP** using the FDR checker¹ involves checking for refinements between processes.

Example (traces refinement):

$A(x) = \text{collect}.x \rightarrow \text{process}_U.x \rightarrow \text{delete}.x \rightarrow A$

$B(x) = \text{collect}.x \rightarrow \text{store}.x \rightarrow \text{process}_U.x \rightarrow \text{delete}.x \rightarrow B$

$C(x) = \text{collect}.x \rightarrow \text{store}.x \rightarrow \text{delete}.x \rightarrow C$

We can see that $\forall x, B(x) \sqsubseteq_T A(x)$ and $B(x) \sqsubseteq_T C(x)$.

In this context **traces refinement** is intuitively inappropriate. Failures-divergences refinement allows us to have finer control.

¹FDR is © of Formal Systems Europe [www.fse1.com].

Modelling a P3P policy in CSP

The main thing about policies is that they prescribe behaviours that cannot happen, or rather conditional behaviours.

A privacy policy specifies:

- ▶ **conditions:** what? to whom?
- ▶ **obligations:** how long to retain the policy? obligation to delete the data
- ▶ **actions:** (only implicitly) purpose for data collection

To compare two policies, we need to check that they prevent the same set of actions from occurring.

Thus they should have the same set of **failures**.

We are still experimenting with different ways of expressing policies for FDR. (This is still **work in progress**.)

Case studies - The EnCoRe Project

EnCoRe [www.encore-project.info] is a multidisciplinary UK-wide research project ("Ensuring Consent and Revocation").

Aims:

- ▶ To enable business to adopt scalable, cost effective and robust consent and revocation methods for controlling the **use, storing, locating and sharing of personal data**.
- ▶ To benefit individuals by providing **meaningful, intuitive mechanisms which will allow them to control the use** of their personal information held by others.
- ▶ To help restore individual confidence in participating in the digital economy and so, in turn, benefit the wider society.

Policy Refinement
Checking

N. Papanikolaou

Introduction

Platform for
Privacy
Preferences

Policy Refinement

Policy Refinement
Checking

EnCoRe and Case
Studies

Conclusions and
Future Work

Case study

EnCoRe case studies:

- ▶ Employee data within an organisation
- ▶ Consent-based marketing
- ▶ Social Networking
- ▶ Government Agency
- ▶ Biobank
- ▶ Cloud computing

Policies arise in each of these contexts and there is interest in proving that one policy implements/refines another.

Also policy refinement checking will help identifying **policy classes** - general policy templates with features common to each case study

Case study: Biobank example

A user consents to sharing personal data (esp. health data) with a **biobank** for research purposes esp. for disease prevention in society in future. e.g. in UK:

[www.ukbiobank.ac.uk]

- ▶ Policy A: some data processed internally by biobank
- ▶ Policy B: some data outsourced to third party

End result of processing is the same, but policy A provides more privacy

- ▶ One can use **traces refinement** to show that end result is the same, of course.


Interesting question is how to cleverly encode policies so that subtle differentiations can also be made.

- ▶ Showing that one policy refines another is a simple goal in principle, but we should be able to use **counterexamples** cleverly to indicate *subtleties* and **implicit assumptions** in policies.

Future work

- ▶ autogeneration of CSP code for P3P and XACML policies
- ▶ experimentation with several example privacy policies - real(istic?) website privacy policies
- ▶ lattice-based models of C&R
- ▶ policies for revoking data
- ▶ conformance of lower-level policies to common law and regulation ie. **compliance automation**

More details in selected papers

 N. Papanikolaou, S. Creese, and M. Goldsmith.
Policy Refinement Checking (Extended Abstract)
Proceedings of AVoCS 09, Gregynog, University of
Swansea, 23-25 September 2009.

 I. Agrafiotis, S. Creese, M. Goldsmith and
N. Papanikolaou.
Reaching for Informed Revocation: Shutting Off the
Tap on Personal Data
Proceedings of Fifth International Summer School on
Privacy and Identity Management for Life, Nice,
France, 7th - 11th September 2009.

Policy Refinement
Checking

N. Papanikolaou

Introduction

Platform for
Privacy
Preferences

Policy Refinement

Policy Refinement
Checking

EnCoRe and Case
Studies

Conclusions and
Future Work




Review and Conclusions

- ▶ Introduced **refinement checking** for privacy policies and policies in general
- ▶ presented the issues specific to privacy policies, esp. with regards to the W3C P3P platform
- ▶ presented a formal verification approach using CSP/FDR
- ▶ discussed range of applications
- ▶ ...esp. in the context of EnCoRe project.

For more information visit

<http://go.warwick.ac.uk/nikos>.

For Further Reading

-  M. Backes, G. Karjoth, W. Bagga and M. Schunter.
Efficient comparison of enterprise privacy policies.
In *SAC '04*, pp. 375–382, 2004.
-  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.
-  M. J. May, C. A. Gunter, I. Lee.
Strong and Weak Policy Relations.
In *IEEE International Symposium on Policies for Distributed Systems and Networks (POLICY'09)*.
London, UK, July 2009.

Policy Refinement
Checking

N. Papanikolaou

Introduction

Platform for
Privacy
Preferences




Policy Refinement

Policy Refinement
Checking

EnCoRe and Case
Studies

Conclusions and
Future Work

For Further Reading (2)

-  J. D. Moffett, M. S. Sloman.
Policy Hierarchies for Distributed Systems Management.
IEEE Journal on Selected Areas in Communications
11:1404–1414, 1993.
-  T. Moses.
eXtensible Access Control Markup Language (XACML).
Standard Version 2.0, OASIS, February 2005.
-  Q. Ni, E. Bertino, J. Lobo, S.B. Calo.
Privacy-Aware Role-Based Access Control.
In IEEE Security and Privacy 7 (4), July/August 2009.

Policy Refinement
Checking

N. Papanikolaou

Introduction

Platform for
Privacy
Preferences

Policy Refinement

Policy Refinement
Checking

EnCoRe and Case
Studies

Conclusions and
Future Work