APPROVAL SHEET

Title of Thesis: Enhancing Web Privacy with Policy Language and Trust

Name of Candidate: Pranam Kolari Master of Science, 2004

Thesis and Abstract Approved: ________________________

Dr. Anupam Joshi
Associate Professor
Department of Computer Science and
Electrical Engineering

Date Approved: ______________________________

Curriculum Vitae

Name: Pranam Kolari.

Permanent Address: 5001I Westland Blvd., Baltimore, MD 21227.

Degree and date to be conferred: Master of Science, May 2004.

Date of Birth: September 28, 1979.

Place of Birth: Mangalore, India.

Secondary Education: M.E.S College, Bangalore, 1997.

Collegiate institutions attended:
University of Maryland, Baltimore County, M.S. Computer Science, 2004.
University Visvesvaraiah College of Engg., Bangalore, India,
B.E. Computer Science, 2001.

Major: Computer Science.

Minor(s):

Professional publications:

Pranam Kolari, Anupam Joshi
Web Mining - Research and Practice
to appear, IEEE Computing in Science and Engineering, Web Engineering Issue - 2004
Li Ding, Pranam Kolari, Shashidhara Ganjugunte, Tim Finin, Anupam Joshi
On Modeling and Evaluating Trust Network Inference
to be published, Seventh International Workshop on Trust in Agent Societies, AAMAS 2004
Pranam Kolari, Li Ding, Lalana Kagal, Shashidhara Ganjugunte, Anupam Joshi, Tim Finin
Enhancing P3P with Policy Language and Trust
submitted to the International Semantic Web Conference, 2004

Professional positions held:

Research Assistant, CSEE Department, UMBC. (Jan. 04 - Jun. 04)

Teaching Assistant, CSEE Department, UMBC. (Aug. 03 - Dec. 03)

Teaching Assistant, CSEE Department, UMBC. (Aug. ’02 - Jun. ’03).

Software Engineer, IBM Global Services, IBM India. (Oct. ’01 - Jul. ’02).

ABSTRACT

Title of Thesis:

Enhancing Web Privacy with Policy Language and Trust

Author: Pranam Kolari, Master of Science, 2004

Thesis directed by: Dr. Anupam Joshi, Associate Professor
Department of Computer Science and
Electrical Engineering

The Platform for Privacy Preferences (P3P) is a W3C standard that web sites can use to describe their privacy practices. The presence of P3P policies enables users to configure web browsers to constrain what they can and cannot do when visiting sites. It's a good idea that unfortunately is rarely used. We identify two reasons: (i) the languages available to define a user's privacy preferences are not expressive enough and (ii) most web sites do not have published P3P policies. We present enhancements to P3P that use semantic web languages and models of trust to help solve both of these problems.

We propose the use of the RDF-based Rei policy language to specify user privacy preferences through an ontological representation of user requirements. We also introduce a new trust model to capture trust between users and websites, as it relates to privacy practices. This model incorporates attributes of a website, which we term as web evaluation statements as they provide a metric for quantifying the trust with the website. This trust can also be used in making privacy decisions.

We show how our proposed architecture is effective even in the absence of published P3P policies. Finally, we present use cases to demonstrate the relevance of our work to the current web privacy landscape and offer it as a powerful enhancement that can promote P3P's adoption and use.

Enhancing Web Privacy with Policy Language and Trust

by

Pranam Kolari

Thesis submitted to the Faculty of the Graduate School
of the University of Maryland in partial fulfillment
of the requirements for the degree of
Master of Science
2004

Dedicated to my parents

ACKNOWLEDGMENTS

I would like to express my sincere gratitude to my graduate advisor Dr. Anupam Joshi for his constant support and guidance. I am also grateful to Dr. Tim Finin for his pointers and guidance and my committee members Dr. Vipul Kashyap and Dr. Tim Oates.

I would also like to thank Li Ding and Lalana Kagal. Li Ding was very enthusiastic about my work and has been a constant help in working on the aspects of Web of Trust. I would also like to thank my labmates and friends at UMBC, who have either directly or indirectly been helpful in my work.

TABLE OF CONTENTS

. . . . **i**

. . . . **ii**

ACKNOWLEDGMENTS . . . . **iii**

LIST OF FIGURES . . . . **vii**

LIST OF TABLES . . . . **viii**

Chapter 1 INTRODUCTION **1**

Chapter 2 THE CURRENT WEB PRIVACY LANDSCAPE . . . **3**

2.1 Web Privacy Protection 3

2.2 Limitations with User Privacy Preference Languages 4

2.3 Limitations with the P3P Trust Model 9

2.4 Related Work 10

Chapter 3 ENHANCEMENTS TO THE P3P FRAMEWORK . . **12**

3.1 Intelligent Privacy Proxy 12

3.2 Website Recommender Network 13

3.3 Privacy Expert 13

3.4 Rei Policy Engine 14

Chapter 4 USER PRIVACY PREFERENCE LANGUAGE 15

4.1 Rei Policy Language 15

4.1.1 Ontologies for policy modeling 16

4.1.2 Scope for future extensions 16

4.1.3 Rule engine 16

4.2 Subset of Rei for Web privacy 17

4.3 Comparing APPEL with Rei and XPref 19

4.3.1 Rule priority. 20

4.3.2 Constraint. 21

4.3.3 Precondition and Postcondition 22

4.3.4 Logical connectives. 22

Chapter 5 ENHANCEMENTS TO THE P3P TRUST MODEL . 23

5.1 Website Evaluation Ontology 23

5.2 Incorporating our trust model into Privacy decisions 27

5.3 Trust Based Recommender System 29

5.3.1 Web information sources 30

5.3.2 Personal social agent (PSA) 31

5.3.3 An example walkthrough 31

5.4 Enhancing the Recommender Network 32

Chapter 6 USE CASES FOR THE ENHANCED P3P FRAMEWORK 33

6.1 A comprehensive example 33

6.2 Examples of policy preferences 37

Chapter 7 FUTURE WORK **39**

Chapter 8 CONCLUSION **40**

REFERENCES **41**

LIST OF FIGURES

2.1 P3P-APPEL matching 6

2.2 Problem with P3P-APPEL matching 7

2.3 XPref Solution to APPEL 8

2.4 Internet explorer Privacy Tab 10

3.1 The enhanced P3P Web privacy framework 13

4.1 Rei concepts used in Web Privacy 17

4.2 Depiction of Rule Priority 20

5.1 Instance of Website Evaluation Ontology 28

5.2 The trust based recommender system 29

6.1 A Rei rule using trust and context 35

LIST OF TABLES

4.1 A language feature comparison of APPEL, Rei and XPref 19

Chapter 1

INTRODUCTION

The issue of “web privacy” is increasingly important for users. While accessing an online vendor or even just browsing a website, users’ private information is often collected explicitly or implicitly for their tracking or targeting. An example of implicit information is the click-stream, which collects the sequence of pages visited by a user. Many sites, as a requirement of use, also explicitly ask users to register and provide personal data. Moreover, distributed data mining techniques [8] can track a user across websites, user sessions, physical locations etc. Hence it is very important for users to be aware of the potential privacy hazards and to have better control over disclosing private information in their online activities.

The first step towards protecting web privacy is for a user to read the privacy policies published by websites and then decide how to interact with these sites. This is often the only step available today. Manual perusal is a time consuming task; therefore, it is not practical for a normal user who might visit tens of sites each day. Motivated by this manual limitation, W3C proposed the P3P system[1] to automate the privacy policy verification process. The P3P framework (i) requires that websites publish XML based privacy policies using the P3P vocabulary, (ii) lets users spec-

[1] http://www.w3.org/P3P/

ify their privacy preference profile using a recommended language (e.g. APPEL[2]), and (iii) lets P3P user agents (such as a web browser) manage users' privacy preference by automatically verifying if a website's P3P policy conforms to user preferences.

The P3P framework is useful; however, it has not been widely adopted. Cranor et al.[2] reported that only 538 of the top 5856 websites were P3P enabled (published valid P3P policies) till May 2003. A report from Ernst and Young [5] shows that P3P adoption in the top 500 sites increased from 16% (August 2002) to 23% (January 2004). Moreover, P3P user agents are not popular among users. Users are discouraged by the limited expressiveness of APPEL, and the limited website adoption of P3P to specify privacy policies.

We believe the key to making machine interpretable privacy policies more widely used is by improving the user side privacy decision; therefore, we propose a two-step enhancement to P3P using the Semantic Web technologies and models of trust. (i) Using an RDF based policy language, Rei [7], for more effective modeling of user privacy preference. Rei uses an ontology based approach for policy preference modeling. (ii) Better accommodating trust in the web privacy domain. The current P3P trust model builds users' trust in websites based only on the existence of P3P policies and privacy certifiers (e.g. TRUSTe.com). We argue that this trust model is insufficient and propose a new model of trust. In this model, trust is derived not only from the conformance of the site's stated policies to user preferences, but also from the existence of website evaluation statements, that can be obtained by consulting trusted recommenders. In both these approaches we show the utility of the ontology based approach.

[2] http://www.w3.org/TR/P3P-preferences/

Chapter 2

THE CURRENT WEB PRIVACY LANDSCAPE

2.1 Web Privacy Protection

The issue of web privacy protection presents many challenges. Browsing the Web for information as well as online shopping has become part of a daily routine for many users. During these browsing activities personal information of users is shared with websites either through explicit data collection or through implicit mechanisms.

In the explicit data collection process, a user is presented a form to be filled out as a prerequisite for the service offered by a website. The collected data might be shared with third parties who could use it in ways not agreeable to the user. On the other hand in the implicit data collection process, which is a more common practice, a user is tracked through stored cookies on client machines. A cookie is a small piece of software (file) on the client machine which allows statefulness in web browsing activities. Such implicit tracking through cookies lets advertisers like www.doubleclick.com track users across websites and search engines like www.google.com store a history of all past searches of a particular user. For example the cookie stored by www.google.com can keep state information of a particular user for a period of 35 years. Hence users should have some control over such data collection and usage. Web privacy protection aims to solve this problem.

Web privacy protection is however different from non-Web like scenarios in which user information is not shared often and users can afford to read and understand privacy policies prior to sharing their private information. To address the concern of automatic privacy protection, many tools and techniques currently exist which provide some level of semi-automated privacy protection to users. Two popular mechanisms are cookie cutters and anonymizing proxies. Cookie cutters are software programs that manage cookies on a host. Anonymizing proxies on the other hand provide anonymous browsing capabilities by shielding users from directly accessing a website. These techniques concentrate on privacy protection as it relates to implicit data collection. They are also semi-automatic in that the user explicitly employs anonymizing proxies when browsing certain websites or specifies cookies of websites that are to be blocked.

Given the need for a comprehensive and automated privacy protection mechanism, the W3C initiated the P3P project. Websites publish machine readable P3P policies which are matched with user privacy preferences to provide automatic user privacy protection. Further P3P policies cover both implicit and explicit data collection practices by websites. We introduce P3P, APPEL and their related limitations in the next section.

2.2 Limitations with User Privacy Preference Languages

The P3P policy is used by websites to publish information about their data collection, usage, retention and distribution policies using XML. Such a policy, if published, can specify data practices for every *http* request by a web browser. The P3P 1.0 Specification defines an XML-based vocabulary for websites to publish such machine interpretable privacy policies. The important entities of this vocabulary are:

1. <DATA> - The collected data like name, click-stream etc.
2. <PURPOSE> - The purpose of data collection like administration, tailoring etc.
3. <RECIPIENT> - The recipients of collected data like ours, third-party etc.
4. <RETENTION> - The period for which this data is retained like no-retention, indefinitely etc.

These entities are bounded by a <STATEMENT> element which groups together data elements with their usage specifications. More information about these elements and their allowed values are available in the P3P specification[11].

APPEL is A Privacy Preference Exchange Language. APPEL-based policy is the P3P counterpart on the client side, used by user agents (P3P user agents) to automatically make privacy decisions for users. The APPEL policy consists of multiple RULE elements which specify user requirements for matching with P3P policies. Figure 2.1 shows a simple example of matching a P3P policy and an APPEL user policy. For clarity the namespaces of these XML elements are not shown in the example. A rule in the APPEL policy matches with a STATEMENT in the website published P3P policy. The match is depicted by arrows connecting corresponding elements. The *behavior* attribute specifies the action to be taken on a rule match, which in this case is *request* i.e. allow access to a website.

Privacy Bird by AT&T [1] and P3P Proxy by JRC[2] are P3P user agents based on the above mentioned approach. However both these implementations are rarely used. Though one reason is the low adoption rate of P3P policy by websites, the other more direct issue is the inadequacy of APPEL, which is used by these user agents for user preference specification.

[1] http://privacybird.com

[2] http://p3p.jrc.it

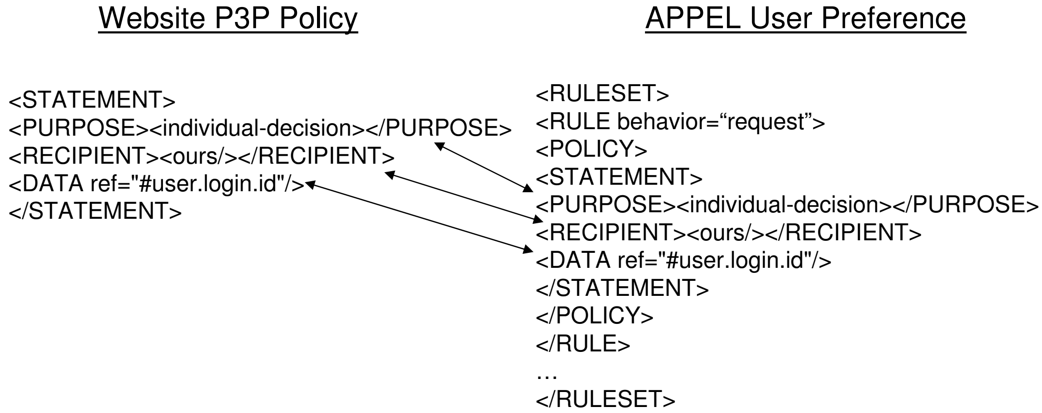

```
<STATEMENT>
<PURPOSE><individual-decision></PURPOSE>
<RECIPIENT><ours/></RECIPIENT>
<DATA ref="#user.login.id"/>
</STATEMENT>
```

```
<RULESET>
<RULE behavior="request">
<POLICY>
<STATEMENT>
<PURPOSE><individual-decision></PURPOSE>
<RECIPIENT><ours/></RECIPIENT>
<DATA ref="#user.login.id"/>
</STATEMENT>
</POLICY>
</RULE>
...
</RULESET>
```

Fig. 2.1. P3P-APPEL matching

Agrawal et al. [1] have detailed problems with APPEL, namely its notion of logical connectives to group elements, rule ordering and matching criteria. An example of such a preference match is depicted in Figure 2.2. In this example, a user wishes to express that she is willing to give away all of her data to the websites she visits for purpose of delivery of purchased products. A simple rule in APPEL is shown in Figure 2.2, using the "or-exact" connective, which means that the user agrees to only these two usages. But, APPEL's rule matching semantics will accept the P3P policy depicted in the same figure 2.2. Here "unrelated" means third-party and the user has implicitly agreed to such a usage of her data. Hence this rule has to be expressed in other ways in APPEL which complicates user rule specification.

The above mentioned issues have been put forward and a solution through the use of Xpref is suggested by [1]. XPref is a subset of XPATH 1.0 and additionally

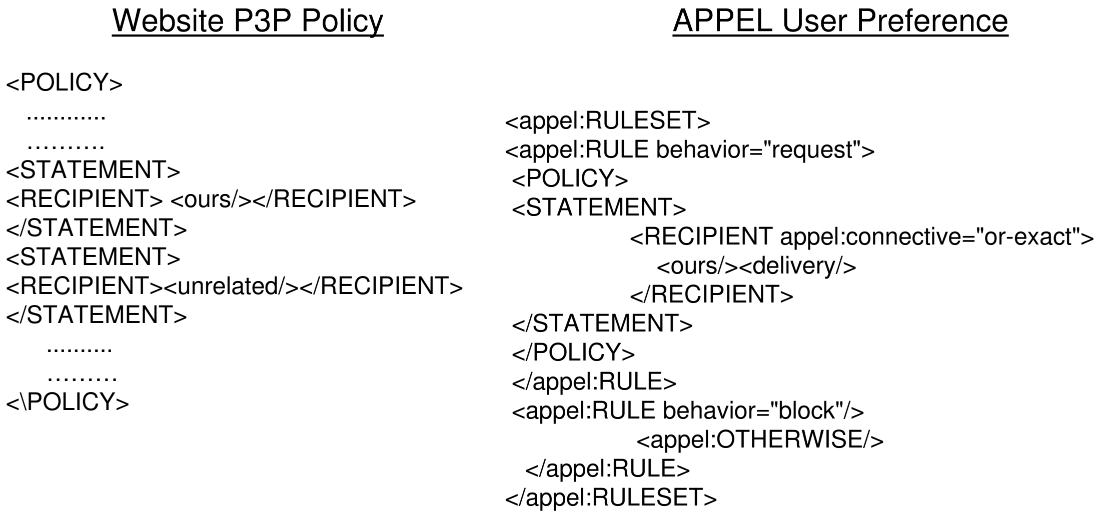

Website P3P Policy

```
<POLICY>
   ............
   ..........
<STATEMENT>
<RECIPIENT> <ours/></RECIPIENT>
</STATEMENT>
<STATEMENT>
<RECIPIENT><unrelated/></RECIPIENT>
</STATEMENT>
      ..........
      .........
<\POLICY>
```

APPEL User Preference

```
<appel:RULESET>
<appel:RULE behavior="request">
 <POLICY>
 <STATEMENT>
                <RECIPIENT appel:connective="or-exact">
                   <ours/><delivery/>
                </RECIPIENT>
 </STATEMENT>
 </POLICY>
 </appel:RULE>
 <appel:RULE behavior="block"/>
                 <appel:OTHERWISE/>
   </appel:RULE>
</appel:RULESET>
```

Fig. 2.2. Problem with P3P-APPEL matching

makes use of the expression “every” from XPATH 2.0, an XML query language, which uses path notations to query for content. XML documents can be represented as a tree and XPATH allows navigation and querying of this document. The user requirement can now be represented in Xpref as shown in figure 2.3. Note that the APPEL rule body has been replaced by an equivalent statement in Xpref, which is the value for the attribute *condition* of the RULE. This forces all the STATEMENT elements to be checked for the value of RECIPIENT. The first rule will fire only if all of the RECIPIENT values satisfy one of “ours” or “delivery”. A comprehensive example showing the advantage of Xpref over APPEL is described in [1] .

Though XPref solves the problems with rule matching in APPEL, it does not solve the other problem inherent in APPEL, that of privacy policy expression. This is because it is still based on APPEL which is in XML. XML does not provide enough

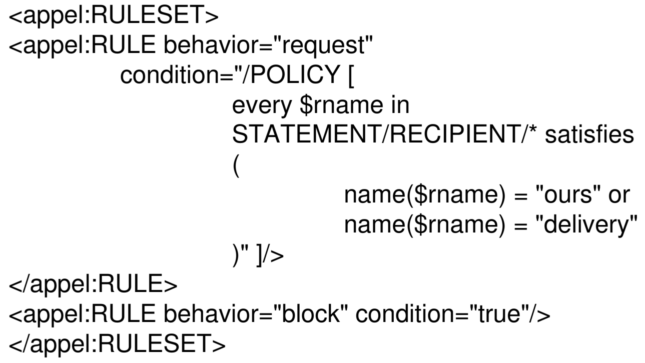

```
<appel:RULESET>
<appel:RULE behavior="request"
        condition="/POLICY [
                every $rname in
                STATEMENT/RECIPIENT/* satisfies
                (
                        name($rname) = "ours" or
                        name($rname) = "delivery"
                )" ]/>
</appel:RULE>
<appel:RULE behavior="block" condition="true"/>
</appel:RULESET>
```

Fig. 2.3. XPref Solution to APPEL

machine-understandable semantics. On the other hand semantic web languages like RDFS[3] and OWL[4] provide schemas and ontologies that help the programmatic understanding of the data. This limitation of XML also leads to another problem. Adding other domain specific information or creating conditions in APPEL or XPref is not simple. It would be difficult to extend APPEL to reason over other aspects of a website and not only its P3P policy.

We propose the use of a new user preference (exchange) language based on the Rei policy language, which solves the above mentioned problems with user privacy preference languages. Through Rei we also propose a language capable of matching of P3P policies specified in RDFS[12].

[3] http://www.w3.org/TR/rdf-schema

[4] http://www.w3.org/TR/owl-features

2.3 Limitations with the P3P Trust Model

Maintaining and building customer trust is a very important criteria for the growth of online business. A recent survey [3] by Ernst and Young suggests that 56% of online consumers believe that websites do not adhere to their privacy policies. Further 90% of these consumers believe that a single most important way of increasing their trust with a website is through independent verification of privacy policies and its subsequent enforcement. Websites have resorted to different mechanisms to build and maintain this trust, like customer service, better handling of user data, text privacy policy, certification, etc.

The P3P framework also attempts to build and maintain trust of consumers in websites. This is through publishing of P3P based policies and a legal entity that is accountable for a specified P3P policy, which is the counterpart of certifiers (e.g. Trust-E) used in human readable privacy policies. We argue that this model does not sufficiently incorporate trust. First, it is highly tied to the presence of a privacy certifier. If the adoption of P3P by websites is low, the certification of their privacy practices is even lower. Second, in the absence of a privacy certifier the model makes implicit assumptions that the presence of P3P policies is sufficient for building trust. Other factors like website popularity which also lead to trusting a website are not sufficiently incorporated. These factors are to some extent equivalent to privacy certifiers and can act as alternative trust certifiers. To gather this information, our approach uses a social recommender network for privacy related knowledge sharing, drawing from the popularity of similar systems like Epinions[5] and Bizrate[6]. These websites let users recommend websites and products to other users through a knowl-

[5] http://www.epinions.com

[6] http://www.bizrate.com

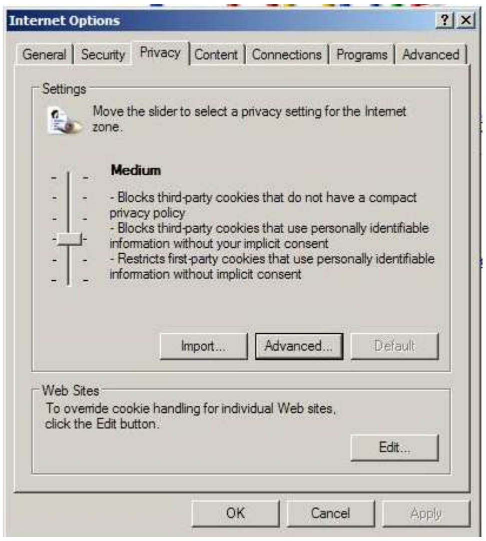

Move the slider to select a privacy setting for the Internet zone.

Medium

- Blocks third-party cookies that do not have a compact privacy policy
- Blocks third-party cookies that use personally identifiable information without your implicit consent
- Restricts first-party cookies that use personally identifiable information without implicit consent

To override cookie handling for individual Web sites, click the Edit button.

Fig. 2.4. Internet explorer Privacy Tab

edge sharing network, which generates the overall rating.Our system makes such a recommendation system machine understandable.

2.4 Related Work

Popular web browsers like Mozilla and Internet Explorer have inbuilt user agents which implement a simplification of P3P. Figure 2.4 shows the preference specification provided by Internet Explorer, a set of different levels of privacy protection based on cookies. Users select a particular level of protection based on their privacy preferences.

Current web browsers concentrate only on cookie handling heuristics. Though Mozilla does not formalize internal representation of user privacy requirements, Internet Explorer provides its own language for user privacy specification. We attribute the simplified approach of these browsers to the issues with APPEL and the low P3P adoption rate.

In the rest of the chapters we detail our enhancements to the P3P architecture, through our work in building a powerful user agent.

Chapter 3

ENHANCEMENTS TO THE P3P FRAMEWORK

Figure 3.1 provides a broad overview of our enhancements. Only key components of the system are depicted. We have prototyped these enhancements by building an intelligent privacy proxy, as an extension to the P3P JRC proxy.

3.1 Intelligent Privacy Proxy

The JRC Proxy was one of the first implementations of a P3P user agent. Users register with the proxy by publishing their user preference in APPEL. This is followed by all http requests (every html page can have many http requests) of the user being managed by the proxy. The proxy fetches the P3P XML policy of a website and matches it with user specified APPEL policy. The entity in the proxy which does the matching is known as the APPEL evaluator. Our prototype replaces this evaluator with an enhanced privacy evaluator, which we term the Privacy Expert. All user preference matching is now handled by the Privacy Expert. Registered users are also required to publish their user preferences in Rei.

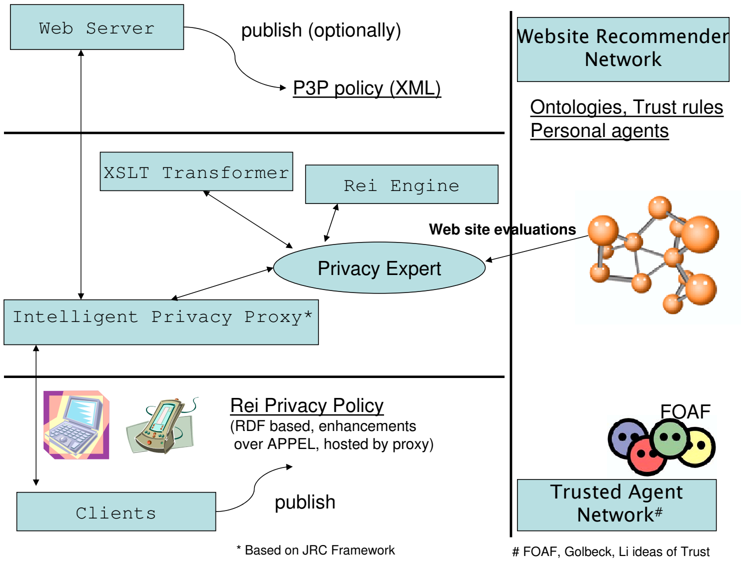

FIG. 3.1. The enhanced P3P Web privacy framework

3.2 Website Recommender Network

We incorporate a new trust model by using privacy recommendations/assertions provided by the Website Recommender Network . This recommender network is a social network of trusted agents which uses ontologies and rules of trust for knowledge sharing. Every user registered with the proxy has an associated Personal Social Agent(PSA) in such a network, which is responsible for gathering web evaluation statements from other PSA's as well as reputation servers. Trust between agents on this network can use FOAF(Friend of a Friend)[1] or other approaches[4, 6, 9].

3.3 Privacy Expert

The privacy expert is implemented as a web service, so that not only the modified JRC proxy, but other privacy agents can also query it for privacy decisions. The main

[1] www.foaf-project.org/

task of the privacy expert is to make privacy decisions for the user - given the P3P policy in XML, the Rei user preference, the user for whom a decision has to be taken and the website to be accessed. The privacy expert uses other services to complete its task. It converts P3P in XML, which is currently used by websites to P3P in RDF using the XSLT Transformer service[14]. It also queries the users Personal Social Agent (PSA) in a recommender network to obtain website evaluations.

3.4 Rei Policy Engine

The Rei policy engine is the final decision-making module of our framework and is also implemented as a web service. It is queried by the Privacy Expert, which in turn passes its response to the intelligent proxy to either block or allow access to a website. A notable point in the entire architecture is that no changes are required from web servers, making our scheme backward compatible, as it were.

In the following chapters, we detail individual components of our framework – the user preference language and our trust model. We recognize that coming up with privacy policies expressed in Rei/RDFS is not something that an average user will do. However, there are ongoing efforts by researchers to either learn user preferences from observing their behavior (web mining on the client side), or at least provide graphical interfaces and templates for policy specification. We also note that such a problem is inherent in APPEL as well.

Chapter 4

USER PRIVACY PREFERENCE LANGUAGE

4.1 Rei Policy Language

Rei is a declarative policy language, represented in RDFS [1], which includes notions of logic like variables for describing different kinds of conditions. It is modeled on deontic concepts of rights, prohibitions, obligations and dispensations [7]. Rei is based on the fact that most policies can be expressed as what an entity can/cannot (right/prohibit) do and what it should/should (obligation/dispensation) not do in terms of actions, services, conversations etc. Hence Rei is capable of describing a large variety of policies ranging from security policies to conversation and privacy policies.

In this chapter we show how Rei solves the issues related to user privacy preference languages through its ontological modeling approach. In the rest of this section we briefly describe some of the important features of Rei which make it a powerful language in the Web privacy domain. For a more detailed explanation readers are referred to [7].

[1] More recent versions support OWL

4.1.1 Ontologies for policy modeling

Rei policies model user privacy policy preferences using ontologies. This ontology-based approach provides rich semantics for specification of highly expressive policies. The policy language has domain independent ontologies but will also require specific domain dependent ontologies. The former includes concepts for permissions, obligations, actions, speech acts, operators, rule priorities etc. which are used to represent policies. The latter is a set of ontologies, which defines domain classes (trust, website, user context etc.) and properties associated with the classes (reputation, operating-system-used, time-of-day, etc.). Constraints and actions are instances of domain specific ontologies. Policies govern actions using a set of constraints.

4.1.2 Scope for future extensions

Rei also provides scope for rule specification involving obligation and delegation management. Obligations are future promises made by a website (e.g. e-mail notification on privacy policy updates) and delegations (e.g. “we share information with our trusted partners who do not have an independent right to further share this data”) are policies regarding distribution of data. Since websites cannot publish such information using P3P, we do not detail them here.

4.1.3 Rule engine

Associated with the policy language is a policy engine that interprets and reasons over the policies, related speech acts and domain information expressed in RDF-S to make decisions about applicable rights, prohibitions, obligations and dispensations.

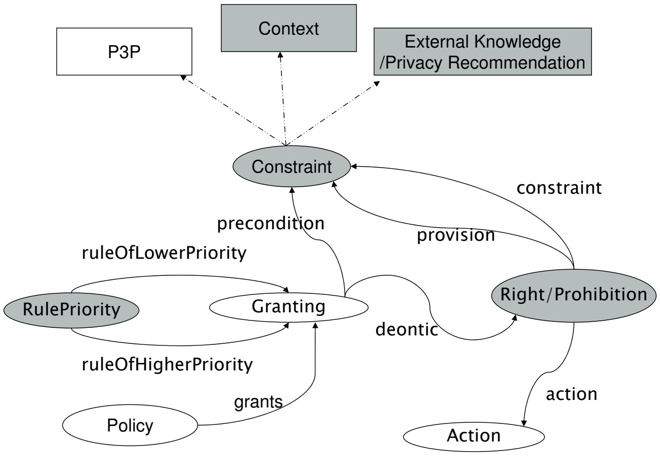

Fig. 4.1. Rei concepts used in Web Privacy

4.2 Subset of Rei for Web privacy

Rei has specific domain independent ontologies which mandate the policy syntax. This includes RDF-S ontologies for specifying entities, objects, policies, credentials, and other Rei specific ontologies[2].

The subset of the Rei ontologies used in the Web Privacy domain is depicted in figure 4.1. Oval nodes represent classes, while properties are defined by directed edges from the domain to range of the property. Dashed edges associate the class *Constraint* to different possible types of constraints. All entities with a white background have a counterpart in APPEL. Shaded entities are enhancements provided by Rei. The concepts used in Figure 4.1 are listed below:

- Policy

[2] http://www.cs.umbc.edu/ lkagal1/rei/ontologies

The root element of the ontology is *Policy*, the equivalent of APPEL *RULELIST* element. The user privacy preference for a particular user is an instance of this root element. This concept has a property *grants* which links a *Policy* instance to one or more rules which is of type *Granting*.

- Granting

  Instances of *Granting* are equivalent to individual rules specified in APPEL using the *RULE* element. *Granting* has an associated property *precondition* which has to be true for the constraints of the rule to be checked. *Granting* also has a property *deontic* which specifies the particular *deontic* object for which this rule applies.

- Right/Prohibition

  Right and *Prohibition* are deontic objects which encapsulate an *Action*, associated constraints and a postcondition. Properties *postcondition* and *constraint* have the range *Constraint*.

- Action

  Instances of *Action* are the domain specific actions which are constrained by the the policy. *Actions* as they relate to Web Privacy can be categorized into five types - request (allow), block, limited, request-prompt, limited-prompt. This provides a direct mapping with actions in APPEL. *Actions* are associated with constraints and postconditions, which decide if an *Action* can be performed or not.

- Constraint

  Constraints are the range of the *constraint* property. A constraint is of the form “x is a type of website” or more generally “x has an attribute y with value z” which in encapsulated by the *Constraint* object. *Constraint* is a triple with *subject*, *object* and *predicate* properties and their associated values. Such constraints are specified over domain specific ontologies like user context, browser capabilities, website[3] etc. *Constraint* can also be a *BooleanConstraint* like *AndConstraint, OrConstraint, NotConstraint* which can be used to group constraints together.

- RulePriority

  RulePriority has properties whose values decide the priority of Granting instances(rules). The two properties *ruleOfHigherPriority* and *RuleOfLowerPriority* decide the ordering between rules.

4.3 Comparing APPEL with Rei and XPref

Though XPref handles the issues involved with rule priority and logical connectives, the privacy rules it can express is still limited as it is still based on APPEL. We compare Rei, APPEL, and XPref in table 4.1. Based on this comparison, we detail specific advantages of Rei.

Table 4.1. A language feature comparison of APPEL, Rei and XPref

	APPEL	Rei	Xpref
RootElement	RULELIST	Policy	RULELIST
RuleElement	RULE	Granting	RULE
Actions	request,block, limited	Any domain action	request,block, limited
Ontologies	Not supported	RDFS	Not supported
Constraints	P3P Specific	Domain ontology	P3P Specific
Rule Priority	Not supported	specified by RulePriority	handled using “XPATH”

[3] http://www.cs.umbc.edu/ kolari1/ontologies/

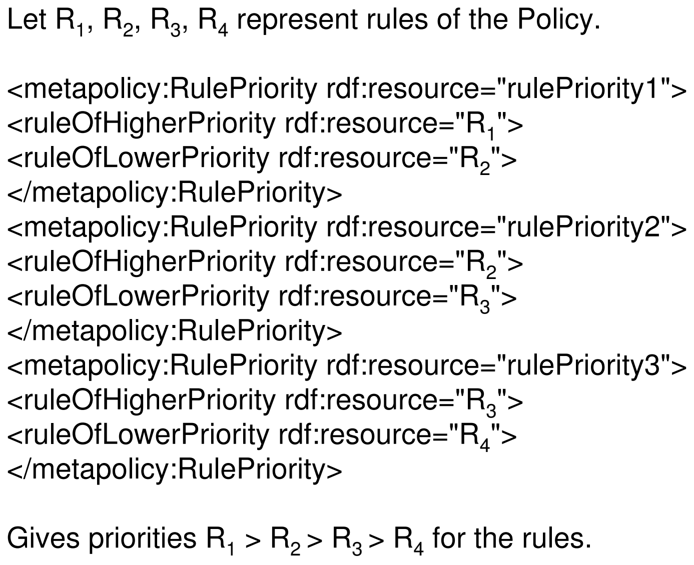

Let R_1, R_2, R_3, R_4 represent rules of the Policy.

```
<metapolicy:RulePriority rdf:resource="rulePriority1">
<ruleOfHigherPriority rdf:resource="R1">
<ruleOfLowerPriority rdf:resource="R2">
</metapolicy:RulePriority>
<metapolicy:RulePriority rdf:resource="rulePriority2">
<ruleOfHigherPriority rdf:resource="R2">
<ruleOfLowerPriority rdf:resource="R3">
</metapolicy:RulePriority>
<metapolicy:RulePriority rdf:resource="rulePriority3">
<ruleOfHigherPriority rdf:resource="R3">
<ruleOfLowerPriority rdf:resource="R4">
</metapolicy:RulePriority>
```

Gives priorities $R_1 > R_2 > R_3 > R_4$ for the rules.

FIG. 4.2. Depiction of Rule Priority

4.3.1 Rule priority.

This feature solves the problem of rule priority in APPEL. The properties of *RulePriority* class, namely *ruleOfHigherPriority* and *ruleOfLowerPriority* can be used to specify order in rules. The range of these properties are instances of *Granting*, which are the actual rules specified by the Policy. When there are more than two policy rules, multiple instances of *RulePriority* specify order and can be cascaded to obtain rule priorities governing all the rules. Figure 4.2 depicts such a usage. Rei also provides modality preference, for example a preference for Prohibition over a Right, which can also be used to specify that *block* action has higher preference than a *request* action for a website.

4.3.2 Constraint.

Rei can specify a wide range of constraints through the easy inclusion of domain specific ontologies. As shown in fig. 4.1 APPEL/XPref lets users specify constraints only on P3P policies published by websites. Rei allows constraints on other information, like context information (e.g. IP address of the client) and other privacy related assertions (e.g. popularity of the site). This capability is important, as it lets each user specify constraints based on personal preferences. It also allows the incorporation of our trust model which will be detailed in later sections.
The possible types of constraints are as follows:

- P3P constraint

  Constraints can be on P3P policies specified in RDF format. For example in a constraint as (statement s_1, recipient, Recipient-Ours), the subject s_1 has a property *recipient* with value “Recipient-Ours”.

- Context constraint

  By the inclusion of a domain ontology to model user context, constraints can now be specified on the context of a user. A simple ontology to model user context has been specified for this purpose. For example (user x, browsingFrom, home) specifies the constraint that the user is browsing from his home.

- Website evaluation statement constraint

  Website evaluation statements include additional knowledge about a website that can be useful for making privacy decisions. This lets the inclusion of our trust model, which will be detailed in the next chapter. For example (website x, isBasedOutOf, India) specifies the constraint that a website should be hosted in India.

4.3.3 Precondition and Postcondition

Rei policies can also be constrained by preconditions and postconditions. Preconditions for a rule (e.g. web resource being accessed is a website or a webservice) allows specifying order in the constraints themselves. Rei also provides the ability to specify postconditions which can be used to specify constraints for a particular *Action*, after its execution. One such example is the obligation on the part of the environment which can be queried by an external entity. For example a browser can be set an obligation that cookies from a website have to be deleted after a particular action.

4.3.4 Logical connectives.

Since Rei allows the use of operators *and, not* and *or* through *BooleanConstraint*, it can represent all kinds of logical connectives and unambiguously reason over them. Rei provides constraints of the types - *AndConstraint*, *OrConstraint* and *NotConstraint* which are all subclasses of *BooleanConstraint*

Representation of rules specified by APPEL is trivial using the Rei preference language[4]. Rei allows the specification of complex user preferences through the features discussed above. We present these use cases upon detailing our trust model in the next chapter.

[4] http://www.cs.umbc.edu/ kolari1/ontologies/wwwpolicy.rdf

Chapter 5

ENHANCEMENTS TO THE P3P TRUST MODEL

In order to solve the limitation of the current P3P trust model, we consider factors that lead a user to trust a website and formalize them through the Web Evaluation Ontology. The ability of our user preference language to make use of domain specific information for privacy decisions, over and above the published P3P policies, also allows easy integration of this trust model into the P3P framework.

In the rest of this chapter, we elaborate our trust model, detailing the web evaluation ontology and its usefulness for privacy/trust related decisions with a website. We also introduce a mechanism for instantiation of this ontology using a trusted recommender system for website evaluations. Such a network also allows verification of instances through collaborative consent.

5.1 Website Evaluation Ontology

Trusting a website's privacy practice includes two aspects: trusting if the website's privacy policy satisfies user privacy preference, and trusting if the website adheres to its privacy policy. A trust judgment requires statements characterizing

websites: either the statements in the website's P3P policy, or *website evaluation statements* gathered from external sources. A website evaluation statement provides evaluation of a website's privacy policies and practices and its reputation, and is independent of the P3P policy. Therefore, a user can make trust judgments even in absence of P3P policy, for e.g., "most of the .edu websites usually have an acceptable privacy policy". These trust judgments can in turn be used to make privacy decisions. We have developed an ontology [1] that models two categories of website evaluation statements:
(i) meta privacy evaluation statements and
(ii) implicit privacy evaluation statements.

The main focus of the meta privacy evaluation statements is the building of trust in a website's privacy practices itself. We identify the following properties as being in this category.

- hasPrivacyPolicy

  represents the URI where the website's human readable privacy policy is located. The presence of a text privacy policy has to be incorporated into automatic decision making and is currently not part of P3P. This information can be obtained by a service which crawls a website's homepage and checks for a link to the privacy policy.

- hasP3PPolicy

  represents the URI of the website's P3P policy. This information can be used to calculate the overall reputation of a website, as it shows commitment to automatic privacy protection of users.

[1] http://www.cs.umbc.edu/˜kolari1/ontologies/Website.rdfs

- privacyPolicyCertifiedBy

  links a website to another website/organization that certifies if the linking website adheres to its privacy policy. If a website publishes P3P, this information can be obtained directly, otherwise scraping of the text policy page will provide the required information.

- privacyPolicyEnforcedBy

  is the URI for the internal system that enforces a website's privacy policy. Though websites will be unwilling to give out such information, the existence of it can help in improving the trust of users.

In the above mentioned properties, P3P has support for hasP3PPolicy (the existence of P3P itself) and privacyPolicyCertifiedBy properties.

The main focus of implicit privacy evaluation statements is building of trust in a website itself; which to some extent implies trust in the website's privacy practices. We identify the following properties in this category.

- domainSuffix

  is the suffix of a website's domain name, such as ".com", ".gov", and ".edu". For example an educational website with domainSuffix .edu, would rarely set cookies, and generally not use it in ways that might breach user privacy.

- owner

  could be a person, organization etc. A website owned by a highly reputed company (such as a well-known bank) can be trusted to have good privacy protection mechanism.

- reputation

  represents the overall rating of a website offered by online reputation services, such as Google PageRank, Bizrate.com rating, and Epinion.com rating. These ratings are computed in different ways by agents on the recommender network. For example Google might consider this the page rank, Epinions might consider it the aggregation of user rating and a personal social agent (PSA) as the rating given by a particular user.

- popularity

  refers to the number of users who review the website with rating. Intuitively, it shows the confidence about the reputation of a website, either positive or negative. Link popularity can be obtained by web scraping sites like http://www.linkpopularity.com, http://www.popdex.com etc. which give the number of inlinks to a particular page. On the same lines the bizrate website rating is associated with the number of reviewers rating that site.

- subDomainOf

  provides information about the parent website to which this site is associated with. e.g. images.yahoo.com is associated with www.yahoo.com, a highly trusted site. Hence it can be trusted implicitly.

- isBasedOutOf

  links a website to a URI of a country in which the host machine of the website is located. This location information can be obtained from online services that maps IP address to spatial location. Different countries have different social value systems; hence they view user privacy protection in different ways.

- lawAccountability

  lawAccountability gives information about the privacy laws which are in effect for a particular website. This could be because of a particular website being hosted at a particular country or state. A user might be confident of privacy laws in U.S and not very certain about those in Asia.

- policySimilarTo

  specifies the similarity between two sites as it relates to their privacy practices. For example www.slashdot.org specifies that its privacy policies are based on those of its parent company, namely the OSDN network. A user trusting OSDN's privacy policy will in turn trust that of slashdot's.

- domainOfService

  represents website classification based on the type of service provided like informational site, advertising site, search site etc. This information can be obtained from online directory servers to make privacy decisions. It provides higher granularity over domainSuffix introduced earlier. The kind of service a website offers can give valuable information about data that has to be protected and released. For example, mailing information and clickstream information can be provided to a bookselling site (if the user is a book enthusiast), whereas for an advertising site (e.g. doubleclick) the release of clickstream information should be controlled as this might lead to user tracking across multiple websites.

Figure 5.1 gives an example of an instance of the web evaluation ontology, for the website http://www.slashdot.org. Blank nodes are those for which the value of a property is not known or irrelevant.

5.2 Incorporating our trust model into Privacy decisions

Instances of the web evaluation ontology enhance the privacy decision model by improving the trust model. They can be incorporated as constraints into the user privacy preferences, and can be used to make privacy decisions even in the absence of P3P, resulting in better user protection. For example, a rule could state that if

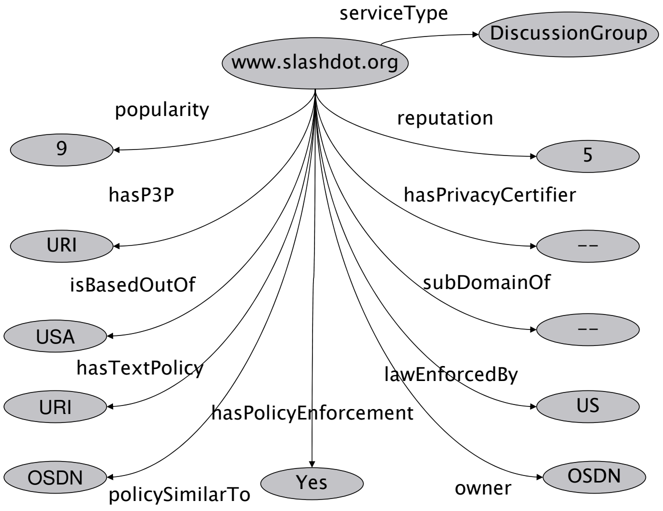

Fig. 5.1. Instance of Website Evaluation Ontology

the reputation of the site is 5 (where reputation ranges between 1 and 5), privacy practices of the website can be ignored by the user agent, since a user might be willing to give out private information in return to the services offered by a highly rated site.

Though instances of the web evaluation ontology are extremely useful for making privacy decision, obtaining such knowledge is not straight-forward. We identify two issues which are central to such knowledge acquisition:

- The process of data acquisition

  Assuming that every user has enough information about all websites is not feasible. Hence we propose the use of a Recommender system in which users share information about websites and also obtain knowledge from reputation servers like www.bizrate.com, www.stumbleupon.com, Google etc.

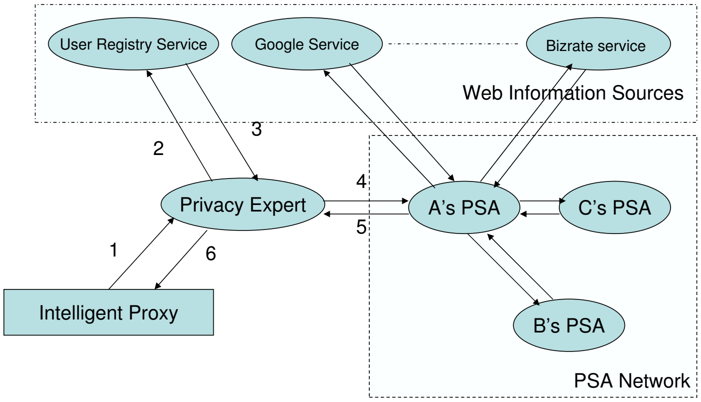

FIG. 5.2. The trust based recommender system

- The process of data aggregation

  The recommender system allows PSA's to query for information from multiple sources. However the question of weighing information provided by them based on their trust values and their merging using a computation model have to be solved. To address this issue we suggest a trust based recommender system, which incorporates trust between users for knowledge aggregation.

5.3 Trust Based Recommender System

Our trust based recommender system[2] is based on online recommendation systems. Currently, online review/rating systems are geared towards customer satisfaction in the online shopping domain, but in the future, we believe that privacy will also be an important issue discussed in online communities. We build such a website rating network. An important feature of our recommender system is that all agents are implemented as web services, which lets entities outside our system easily query

[2] Joint work with Li Ding - dingli1@cs.umbc.edu

services for information. As shown in figure 5.2, we are currently implementing a trust based recommender system that allows users to exchange their privacy evaluation statements based on their own experiences as well as obtaining recommendations from online reputation services. Our system consists of two categories of agents: *web information sources* and *personal social agents*.

5.3.1 Web information sources

Web information sources collect information from the Web, and encode this information using the Semantic Web technologies, which includes the web evaluation ontology. The lack of widespread adoption of the Semantic Web necessitates that these information sources be implemented by website scarping. Currently, we envision three types of web information sources and have implemented a subset of them.
(i) *Reputation systems* (such as Epinions.com, Bizrate.com, and eBay.com) collect customer satisfaction reports and generate overall rating about commercial websites[3].
(ii) *Trusted third parties* provide "objective" evaluation about websites, which is trusted by most online users. For example, Google reports statistics about the inlinks to a web page based on a significant amount of observations, and TRUSTe.com provides privacy certification based on its privacy verification mechanism.
(iii) *User registry service* can automatically discover users' personal web services by analyzing their online FOAF profiles, and users' relation through the "knows" relation. This gives the user's PSA a set of other PSA's and reputation servers that can be queried for information.

[3] We have implemented information extraction agents for Bizrate.com

5.3.2 Personal social agent (PSA)

A PSA interfaces a user with the online community. It provides a personal web service to the online community, so that the user's knowledge can be published online and shared under the user control. However the most important function of the PSA is that is allows a user to consult knowledge from the other PSAs and web information sources. By maintaining and evolving trust relations between agents, the PSAs can form a peer to peer (P2P) social network in a distributed environment. We assume a PSA can discover the other users' PSAs' address by using user registry service. We further assume trust can be derived between any pair of PSAs through trust network inference [4, 6, 9, 13].

5.3.3 An example walkthrough

The message flow in figure 5.2 depicts how the trust based recommender system helps the intelligent proxy to make privacy decision. When the intelligent proxy requires external privacy assertions (i.e., website evaluation statements) to satisfy user A's policy constraints for allowing access to website X, it sends message 1 to the privacy expert. The privacy expert then queries "user registry service" for user A's PSA address with A's name in message 2. When the privacy expert obtains the address successfully, it forwards the query to A's PSA in message 4. Since A's PSA does not have enough information to populate the website evaluation ontology, it consults two trusted peers B and C and two trusted web information sources for website rating and inlink statistics. A's PSA combines the decisions from B, C and web information sources and passes the result back to the privacy expert which passes it on to the intelligent proxy in messages 7 and 8. The returned result has values for all the properties of the website evaluation ontology, for which information could be obtained using the recommender network.

5.4 Enhancing the Recommender Network

As of now, our implementation assumes that trust between a two users and trust between a user and a web information source is available through complex trust network inference. Experimentation with these inference mechanisms that yields the best possible knowledge sharing network is not part of current work. We also assume that there exists no conflict between information provided by different trusted sources. These sources provide knowledge which populates different properties of the web evaluation ontology. Experimentation with the actual computation model to handle inconsistent and redundant knowledge is hence not part of current work. Further since Rei does not allow policies with probabilistic constraints, we do not consider such constraints and their propagation in the trust network. The use of probabilistic web evaluation statements, which is the case for most statements in a social knowledge sharing network would be the possible next step. The overall trust based recommender system presents many challenges for future research.

With the use of our proposed website recommender, privacy (P3P) user agents will be capable of making decisions even in the absence of P3P policies published by websites. We believe that this will increase the adoption of privacy user agents, primarily if they are incorporated into web browsers. The increase in user agent adoption will in turn force websites to publish policies in P3P, and in many cases to certify their privacy practices if they are to continue to build and maintain trust with users.

Chapter 6

USE CASES FOR THE ENHANCED P3P FRAMEWORK

6.1 A comprehensive example

Cookies set on client machines are very useful for online vendors, from being enablers of online shopping to tracking of user browsing behavior. Though cookies are browser dependent, they are independent of internet service provider and client location. As an implication of this, a notebook user's activities can be tracked across locations based on her IP address and past browsing preferences. Commonly accessed websites like e-mail and advertising can reconstruct the entire travel history of a particular user. Further an e-mail service like GMail[1] can associate this information with the e-mails being read by the user to identify the purpose of the visit.

A option in the above mentioned scenario would be to use cookie cutters to shield the user from such implicit data collection. However this requires the user to manually deny cookies to certain websites. To automate control so that a user can decide when cookies can be released to websites, we list the criteria (constraints) of a typical user as follows. Similar constraints can be specified for other implicit and

[1] www.gmail.com

explicit data collected.

(i) Trustable P3P constraint - The website publishes a P3P policy and specifies that data collected is user IP and click-stream, its usage is website tailoring, and retention period of this data is no-retention(deleted after current session) and the P3P policy is certified by Trust-E. If it is not certified, the overall rating from the web evaluation ontology is 5 (assuming that rating has integral value and is between 0 and 5). This allows trusting a website through either a certifier or a web evaluation statement.

(ii) Trusted website constraint - The website is one of a set of websites which is highy trusted by the user, be it her bank, the company she works for or her home page. For a bank or the employer, cookies might be an important way of accountability should there be a conflict in transactions.

(iii) Trusted domain constraint - The website is either of type(domainSuffix) "edu" or" gov" and is based out of USA. The user is willing to share information as she is confident of privacy practices of such websites.

(iv) DomainOfService Negative constraint - The website is neither of type (domainOfService) portlet(e.g. Yahoo, AOL) nor of type advertising. Such websites if allowed direct access might create and store a history of users travel habits.

(v) DomainOfService constraint - The domainOfService of the website is "travel". A travel related site could give her hints about her travel plans and save her both time and money. When the user is in a new city, allowing cookies (user preference) to a site providing local information or any other travel related site might be useful for the user, which in this case might present a user a free ticket to a concert when in Baltimore. Further, information that the website caters to a particular location for example, the state Maryland in which Baltimore is located can also be used if available, so that cookies are allowed to only such websites. This will allow specification of more specific and stricter constraints. All of the above mentioned constraints

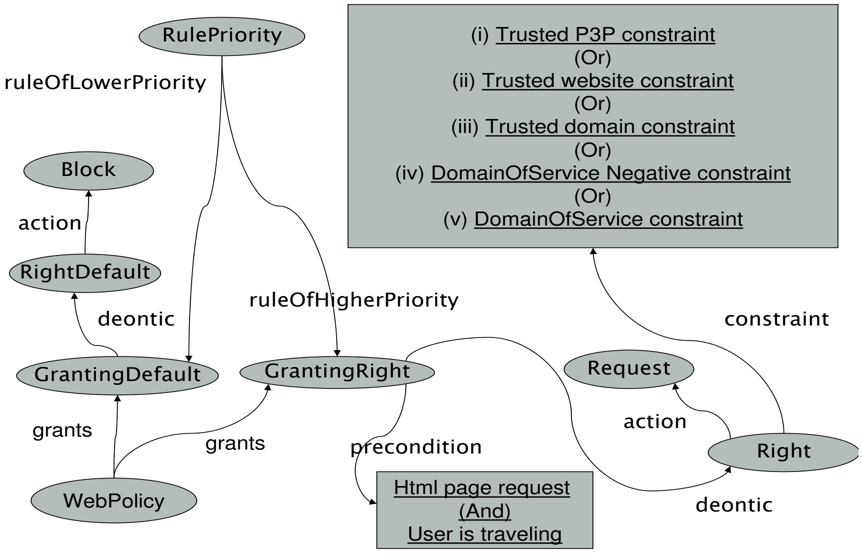

Fig. 6.1. A Rei rule using trust and context

are based on the assumptions that - user privacy is being protected when the user is travelling i.e context information browsingFrom has value “away”, cookies are not exchanged with a site during the process of querying for P3P policy itself and the website is not one where a user is required to login to access provided services(e.g. Gmail).

The ability of our user preference language to incorporate external knowledge, which in the above scenario is context knowledge and trust evaluation statements lets a user specify preferences of the above mentioned complexity with ease. Figure 6.1 depicts the actual policy specification of the required rules using Rei. One of the rules incorporates all the constraints of a user and another rule is the default. In the case of typical user preferences many such rules are specified with rule priorites deciding the order of rule matching.

WebPolicy is an instance of *Policy* and specifies two rules for actions related to

Web Privacy, namely *Block* and *Request*. *GrantingDefault* specifies the default rule, and *GrantingRight* specifies the rule for the constraints governing *Request* action. *RulePriority* sets higher priority to the *GrantingRight* rule. The default rule has a lower priority and blocks access to a website. Note that in this example, there are preconditions that the resource being accessed is a html page and the user's context is "away". This shows how we can use different policy rules when accessing web services [2], images, active-X controls or just when browsing from home.

The action associated with the deontic concept *Right* is *Request* i.e. allow access to a website. The *Right* has constraints. *Constraints* are underlined, to point to the fact that they are represented by English-like statements for clarity. In reality they are asserted as triples into our knowledge base, based on the domain ontology. The specified constraints are the one's listed above as a typical user preference in such a scenario. All constraints are specified in rectangular blocks to clearly delineate them from other parts of the user preference. Constraints are grouped together using *BooleanConstraint* provided by Rei. The complete example is available at http://www.cs.umbc.edu/~kolari1/ontologies/examples/comprehensive.rdf. Other examples which depict user preference examples for low, high and medium level privacy protection along with their descriptions are available at http://www.cs.umbc.edu/~kolari1/ontologies/examples/.

Given such a Rei policy, the Rei engine makes decisions based on satisfying constraints. Since *GrantingRight* has higher priority, constraints specific to it are first checked. If these constraints fail, the action *Block* specified by *GrantingDefault* is fired, which specifies the action "block". The action "block" could be interpreted in

[2] P3P for web services is a currently evolving specification – http://www.w3c.org/

different ways by the client software. It could mean denying access to site, blocking cookies, or requiring obligations from browsers or websites that cookies and information that they provide are deleted after every session. Enforcing a particular action depends on the capability of the browser and websites. The JRC proxy denies access to a website.

6.2 Examples of policy preferences

The following examples are given in English-like statements for clear understanding. The RDF equivalents are available online. We provide three example policies which can be represented.

- We categorize the policy rules in this example as “low” as it is low in terms of user privacy protection. The user is willing to give out more information in return to the services offered by websites.

  The set of rules for this example are:

  R$_1$ - Allow access to a website if its reputation is 5

  R$_2$ - Allow access to all sites whose domain suffix is not “.com”

  R$_3$ - Allow access to all sites who publish P3P and usage of collected data is not individual-analysis

  R$_4$ - Prompt for user action

  Priority set as R$_2$ > R$_1$ > R$_3$ > R$_4$

  Available at http://www.cs.umbc.edu/ kolari1/ontologies/examples/low.rdf

- This policy example provides medium privacy protection.

  The set of rules for this example are:

  R$_1$ - Allow access to a website if its reputation is 5 and has a privacy certifier

  R$_2$ - Allow access to all sites which collects clickstream data for pseudo-analysis

R_3 - Block access to sites which collect health related information, store indefinitely and use of individual analysis

R_4 - Allow access to all sites when browsing from home and websites donot have domainSuffix of “.com”

R_5 - Prompt for user action

Priority set as $R_3 > R_1 > R_2 > R_4 > R_5$

Available at http://www.cs.umbc.edu/ kolari1/ontologies/examples/medium.rdf

- This policy example provides high privacy protection.

  The set of rules for this example are:

  R_1 - Allow access to websites W, X, Y, Z

  R_2 - Allow access if the collected data is not used for individual analysis and the website is not a advertiser and is privacy certified and basedOutOf Europe

  R_3 - Block access if Website is an advertiser, identifiable data is collected and context is work

  R_4 - Allow access to all websites when browsing websites having domainSuffix of “.edu”

  R_5 - Allow access to websites with reputation of 5 and basedOutOf U.S and providing service “searchengine”

  R_6 - Prompt for user action Priority set as $R_1 > R_4 > R_3 > R_2 > R_5 > R_6$

  Available at http://www.cs.umbc.edu/ kolari1/ontologies/examples/high.rdf

Chapter 7

FUTURE WORK

One of the main issues in a recommender network is for users to share their knowledge. We are in the process of collecting such information by providing web forms to registered users. Since the data collection process is time consuming, our experiments do not take them into account. Currently real data exists only as part of the web scraping services. The knowledge base can also be augmented by the availability of additional entities/services which can collect data. For example http://www.stumbleupon.com gives user recommendation about websites. IP mapping and other services which all aid in populating the web evaluation ontology can be built.

Trust network inference is currently a widely researched area on the Semantic Web. To showcase our approach, we limit our system to simplistic trust inferences. Experimentation with various inference approaches to test their effectiveness in a real world application like ours is the next step in integrating a complex inference system to our trust recommender system.

Chapter 8

CONCLUSION

The ultimate goal of Web privacy protection is to safeguard private information automatically in online activities. The key to this goal is to have intelligent user agents automate privacy decision. Towards this end, our enhancements to P3P make the following contributions. In comparison with APPEL, we show that the RDF based policy language Rei is more expressive and suitable in the Web privacy domain. Through Rei, we contribute a user privacy preference language in RDFS for P3P. To make the P3P trust model effective in absence of P3P policy, we introduce a *website evaluation ontology* and a trust based recommender system for users to share privacy related information.

The overall system also shows the effectiveness of the Semantic Web in a new domain. The main motivation of the W3C P3P project is to let machines automatically make privacy decisions for users. For such a system to work, machine readable information from multiple sources have to be aggregated and used. The widespread deployment of the Semantic Web would have made such privacy decisions easier, without the need to implement web scraping services on a per-site basis and resulted in turn resulted in the widespread adoption of P3P.

REFERENCES

[1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. An xpath-based preference language for p3p. In *Proceedings of the twelfth international conference on World Wide Web*, pages 629–639. ACM Press, 2003.

[2] L. F. Cranor, S. Byers, and D. Kormann. An analysis of p3p deployment on commercial, government, and children's web sites as of may 2003. Technical report, AT &T Labs-Research, 2003. prepared for the 14 May 2003 Federal Trade Commission Workshop on Technologies for Protecting Personal Information.

[3] J. R. DeVault, D. Roque, jay Rosenblum, and K. Valente. Privacy promises are not enough. Technical report, Ernst&Young, 2001.

[4] L. Ding, L. Zhou, and T. Finin. Trust based knowledge outsourcing for semantic web agents. In *Proceedings of IEEE/WIC International Conference on Web Intelligence*, 2003.

[5] Ernst&Young. P3p dashboard report: Top 500 p3p dashboard, 2004.

[6] J. Golbeck, B. Parsia, and J. Hendler. Trust networks on the semantic web. In *Proceedings of Cooperative Intelligent Agents*, 2003.

[7] L. Kagal, T. Finin, and A. Joshi. A policy language for a pervasive computing environment. In *Proceedings of IEEE 4th International Workshop on Policies for Distributed Systems and Networks*, June 2003.

[8] H. Kargupta and P. Chan, editors. *Advances in Distributed and Parallel Knowledge Discovery*. MIT/AAAI Press, 2000.

[9] M. Richardson, R. Agrawal, and P. Domingos. Trust management for the semantic web. In *Proceedings of the Second International Semantic Web Conference*, 2003.

[10] B. Clark, and S. DeRose. XML Path Language, W3C Recommendation 16 November 1999. *http://www.w3.org/TR/xpath*

[11] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation 16 April 2002. *http://www.w3.org/TR/P3P/*

[12] B. McBride, R. Wenning, and L. Cranor. An RDF Schema for P3P, W3C Note 25 January 2002. *http://www.w3.org/TR/p3p-rdfschema/*

[13] B. Yu, M. Venkatraman, and M. Singh. An Adaptive Social Network for Information Access: Theoretical and Experimental Results *Journal of the Applied Artificial Intelligence*, 2003, 17,1.

[14] R. Wenning, W3C Privacy Activity Lead, Personal Communication *XSLT for P3P RDF, March 10, 2004*