Using P3P in a web services-based context-aware application platform[1]

Martijn Zuidweg, José Gonçalves Pereira Filho[2], Marten van Sinderen
Centre for Telematics and Information Technology, University of Twente
PO Box 217, 7500 AE Enschede, The Netherlands
{zuidweg, filho, sinderen}@cs.utwente.nl

Abstract

This paper describes a proposal for a privacy control architecture to be applied in the WASP project. The WASP project aims to develop a context-aware service platform on top of 3G networks, using web services technology. The proposed privacy control architecture is based on the P3P privacy policy description standard defined by W3C. The paper identifies extensions to P3P and its associated preference expression language APPEL that are needed to operate in a context-aware environment.

1 Introduction

Context-aware computing is an emerging computing paradigm that tries to exploit information about the context of its users to provide new or improved services.

Dey and Abowd [4] have defined context as "*any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves*". This definition is widely used in literature today.

Schilit [12] identified four classes of services that can be provided using contextual information:

- **Proximate selection**: selection of service providers based on their location. An example would be an application that can provide a user with a list of museums within walking distance.
- **Automatic contextual reconfiguration**: application behaviour changes automatically depending on the context. For example, a mobile phone would know that a user is currently in a meeting and thus not ring when it is called.
- **Contextual information and commands**: the same user request returns different results when executed in different contexts. For example, a print request would always, automatically, be routed to the nearest printer in the building.
- **Context-triggered actions**: commands are automatically executed when certain contextual conditions are met. For example, a user could set a trigger to receive a reminder to buy bread when he passes by a bakery.

Context-aware computing environments may use information provided by many sensors to acquire knowledge about the context. These sensors can be invisible to users. It is obvious that these sensors, gathering information about people without being noticed, can be a threat to privacy. If the risks of privacy violation when using a context-aware application cannot be estimated, users may be unwilling to use such a system. *This is why privacy control is essential in the design of a context-aware computing platform.*

[1] The work described in this paper has been sponsored by Freeband Knowledge Impulse, a joint initiative of Dutch Government, knowledge institutes and industry.

[2] On leave from Departamento de Informática, Universidade Federal do Espírito Santo, Brasil. E-mail: zegonc@inf.ufes.br

This paper aims at providing a privacy control architecture for the context-aware application platform developed in the WASP project (see section 2).

The rest of this paper is structured as follows. Section 2 presents a quick overview of the WASP platform. In section 3, general privacy issues are explained and the P3P standard is introduced. This standard, originally developed for web sites, will be used to provide privacy control in WASP. The applicability of P3P in a context-aware platform such as WASP is explained in section 4. Section 5 identifies the extensions needed to P3P to be used for context-aware web services. Section 6 depicts the privacy control architecture of WASP, using P3P. Section 7 discusses related work. Finally, section 8 contains conclusions and presents future work to be done.

2 WASP

In the WASP (Web Architectures for Services Platforms) project [15], the University of Twente, Ericsson and the Telematica Instituut cooperate in developing a platform to support context-aware applications based on web services. The WASP platform operates on top of 3G networks, using Parlay X [11] as a web services interface to 3G network functions.

Typically, users interact with service providers offering context-aware applications through their mobile device. A context-aware application uses context information available from the WASP platform to provide its services (Figure 1).

Initially, the WASP project will focus on tourist applications using location-based services. Service providers, called Points of Interest (POI) in WASP, like museums and restaurants, provide a web services interface. The descriptions of these services are stored in a registry, the POI registry. Users can look up services of interest through this registry.

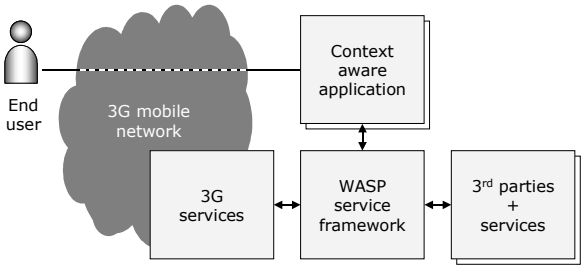

Figure 1: Interaction between a user, a service provider (context-aware application) and the WASP platform.

WASP users will mainly access WASP applications using mobile devices, such as PDAs and smart phones. WASP will provide tight integration with the underlying 3G technology by offering seamless invocation of network services, such as directly making a voice call to the POI from a WASP application.

3 Privacy and P3P

Whereas several architectures supporting context-aware services have been proposed and implemented, none of these architectures has fully integrated privacy control. Nonetheless, research on privacy has been recognized as an important topic, and results have been published on both privacy control in context-aware systems (see section 7) and privacy in general. This research has led to a number of widely adopted principles, both by researchers [1, 7] and legislation (US Privacy Act of 1974 and the EU Directive 95/46/EC):

- **Notice**: People should be informed when data is collected about them.
- **Choice and consent**: People must explicitly agree to data collection, and be able to opt out of using a service if they do not agree with its practises. Several privacy practise options should be available if the service can also be (partially) provided with less intrusion on user privacy, to avoid that people with stronger privacy concerns cannot use a particular service at all.
- **Access**: An individual should have access to any data that is gathered about him.

- **Anonymity and pseudonymity**: If not necessary, no personal information should be stored. If services can be used anonymously or using a pseudonym, this possibility should be provided. Legislation lays no restriction on the collection of unidentifiable data.

To provide privacy control in WASP, we want to develop and apply an adapted version of P3P. P3P (Platform for Privacy Preferences Project, [14]) is a protocol for web sites to inform web users of their data-collection practices. It was developed by the World Wide Web Consortium (W3C). P3P enables web sites to express their privacy practises in an XML document.

Figure 2 shows a part of such a P3P policy. This statement explains that some contact information about the user will be collected to be able to contact him for marketing purposes. As indicated by the <RECIPIENT> element, the information will used by the owner of the web site and, if the user explicitly chooses for it, spread to comparable companies.

```
<STATEMENT>
  <CONSEQUENCE>
     At your request, we will send you
     carefully selected marketing
     solicitations that we think you will be
     interested in.
  </CONSEQUENCE>
  <PURPOSE>
    <contact required="opt-in"/>
    <individual-decision required="opt-in"/>
    <tailoring required="opt-in"/>
  </PURPOSE>
  <RECIPIENT>
    <ours/>
    <same required="opt-in"/>
  </RECIPIENT>
  <RETENTION><stated-purpose/></RETENTION>
  <DATA-GROUP>
    <DATA ref="#user.name" optional="yes"/>
    <DATA ref="#user.home-info.postal"
          optional="yes"/>
    <DATA
       ref="#user.home-info.telecom.telephone"
       optional="yes"/>
    <DATA ref="#user.business-info.postal"
          optional="yes"/>
    <DATA  ref="#user.business-
                      info.telecom.telephone"
           optional="yes"/>
    <DATA ref="#user.home-info.online.email"
          optional="yes"/>
  </DATA-GROUP>
</STATEMENT>
```

Figure 2: Example P3P statement. *Extracted from [14].*

P3P provides a standardized way of associating these policies with web sites or parts of web sites, and includes a mechanism for transporting the policies over HTTP.

The intention of P3P is to automatically negotiate about a web site's privacy practice, addressing the principles of *notice* and *choice and (automatic) consent*. To support this, a user must have previously defined his privacy preferences and stored these in a machine-readable format. W3C has developed the privacy preference description language APPEL [14] for this purpose. When a user wants to access a web site, a user agent (embedded in the user's browser) first retrieves the web site's P3P policy. It then compares the policy with the user preferences. If the policy complies with the user's preferences, the web site is retrieved. If not, the user may be prompted for further evaluation, or the request may be cancelled. This behaviour is depicted in Figure 3.

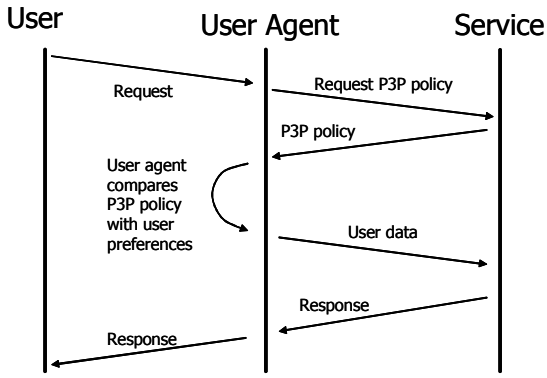

Figure 3: P3P automatic negotiation

The scope of P3P is limited to the concepts of notice, and choice and consent. Users are informed of a web site's privacy policy, and based on this policy they can choose to use the service or decline it. P3P does not intend to enforce privacy by technical means. It is up to the user to make a decision whether or not to trust the service provider. No guarantee is given that the service provider will actually conform to its privacy policy. This should be regulated by law.

4 Applicability of P3P in WASP

P3P was developed as a standard for web sites. Since the main purpose of P3P is to simply describe services, its applicability is actually much wider. Web services, used as the enabling technology in the WASP project, follow a client-server paradigm that is comparable to interaction on the World Wide Web. Furthermore, P3P is based on XML, as are web services. Integration of P3P into the domain of web services will be quite straightforward. The requirements for this integration are identified in section 5.

In a context-aware platform, contextual information is generally acquired through various sensors. Some of these sensors may reside on the user's device, such as a GPS receiver integrated in a mobile phone, while other sensors may be installed in walls or ceilings of the building close to or surrounding the user. It may be simple for users to control the release of contextual information that is stored on their own device, but it may be hard to control information gathered at different places in the platform. As P3P does not provide mechanisms for controlling data release to interested parties, but simply describes service behaviour, it is irrelevant to P3P where the data is coming from. So, P3P is suitable for an environment where the data is distributed.

Research has shown that the "inquirer", i.e., the service provider using contextual information, is an important determinant for people's privacy preferences [8]. This means that users usually have the same preferences for the same data collector, no matter where they are or what they are doing. So, privacy preferences are strongly influenced by a description of the data collecting service, which is exactly what P3P provides.

Other researchers have also proposed to use (adapted versions of) P3P in context-aware systems. The works of Langheinrich [6], Myles et al. [9] and Nilsson et al. [10] all propose privacy control mechanisms based on P3P for context-aware or location-based environments. These systems are discussed in more detail in section 7.

5 Extensions to P3P

Whereas the suitability for P3P in a context-aware web services-based environment such as WASP has been argued in the previous section, some extensions have to be implemented before we can actually use P3P in such an environment. This section discusses the extensions needed to P3P itself as well as to the privacy preference expression language APPEL.

5.1 P3P in web services-based context-aware environments

Two extensions to P3P are identified to make it suitable for a web services-based context-aware environment. First, P3P will need to be adapted so that it works with web services instead of just web sites. Next, P3P will need some mechanism to reason about contextual information. Finally, some adaptation is necessary to integrate P3P with the context-aware environment. This will involve several issues:

Policy discovery. Web sites can reference their policies in multiple ways. Either they can publish it at a well-know location (/w3c/p3p.xml relative to the site's address), they can link it from the (X)HTML source file, or they can reference it in HTTP headers.

For web services, we propose to publish the policy in either the WSDL file describing the service, or in a registry such as UDDI, or the WASP POI registry.

It is also possible to use a well-know function call (analogous to the well-known location for web sites), but this means that a web service has to be invoked before the policy is retrieved. This problem is analogous to the problem of linking to P3P files from an HTML document, or sending a reference to the P3P file in HTTP headers. In these cases, the web site has to be accessed before its policy can be evaluated.

Contextual information. The P3P specification describes many predefined data types. These data types are categorised, and have associated descriptions of their meaning. This allows also semantic agreement on the data collected, so

that there is no ambiguity between a user and a service provider about what is exactly collected. However, as P3P was developed for web sites, there are no data types addressing the kind of data that would be gathered in context-aware environments.

P3P provides extensibility in its specification, i.e., a service provider can mention other types of data (for example user location) than the ones specified by P3P in a policy. There are two problems with using such an extension. First, there is no semantic agreement on the actual meaning of this data, so it may mean one thing for one service provider and a completely different thing for another service provider. Second, if all contextual information has to be redefined in every P3P policy, policies may become quite large. This is especially problematic in environments like the WASP platform, where users typically use small devices with limited bandwidth connections. So, contextual information should be added to the list of predefined data types.

Trust. Privacy control using P3P relies on trust. P3P does not provide any technical means to enforce privacy. This is left to law and other regulations. To increase this trust, service providers could be screened by the context-aware platform provider, and receive a certificate [16].

5.2 User privacy preferences

P3P is intended to provide privacy control for simple request-reply web interaction. In this case, one request leads to one reply. So, a user can evaluate the site's privacy policy before each request, and decide whether or not to use the service. After the invocation of the service, no further interactions take place. A new interaction can only start on the initiative of the user.

In a context-aware system, it is common for a user to register with a service once, after which the service may contact the user many times. For example, a user may register to a location tracking service, which will provide the user with information every time a certain condition is met.

A problem that arises in this situation is that user preferences may also be constrained by context. For example, a user may allow his employer (in fact: a service acting on the employer's behalf) to track his location, but *only during work hours*. While a condition like this may still be checked by the employer before requesting a user's location (since both the user and the employer's service can have the same notion of time), it is not difficult to think of more complicated situations. Suppose the user will only allow his employer to track his location *while he is in the office.* This condition can no longer be checked by the employer before actually retrieving the user's location. In other words, a service cannot publish this kind of behaviour in its P3P policy and take care of complying with the policy itself. So, user preferences containing constraints on the contextual data that is being collected must be checked somewhere else, as they cannot be checked by the service provider.

There is a good argument not to include these context-dependent conditions in a privacy policy of a service at all. Context-dependent conditions are not part of the actual service behaviour. They describe user preferences, not service characteristics.

To support the inclusion of this kind of conditions in user preferences, we are currently developing an extension to APPEL. This extension will support the expression of constraints on the values of the contextual data specified by our extension to P3P. For example, it enables a user to express that a particular service provider may only retrieve data between nine and five, or only retrieve his location while the user is in his office building.

6 WASP privacy architecture

The proposed privacy architecture for the WASP platform is depicted in Figure 4. A user interacts with context-aware services through a user agent, which can automatically retrieve P3P policies for the requested services and compare them to the user's preferences. This user agent can reside on the user's device, or, in case of small user devices with limited band-

width and processing power, somewhere in the wired network.

Contextual information is shielded from service providers by the privacy control layer. This layer is responsible for checking context-dependent privacy preferences, as explained below.

Contextual information is available from the contextual information interpreter. This component aggregates all data from sensors and other context providers (such as for example a user's calendar).

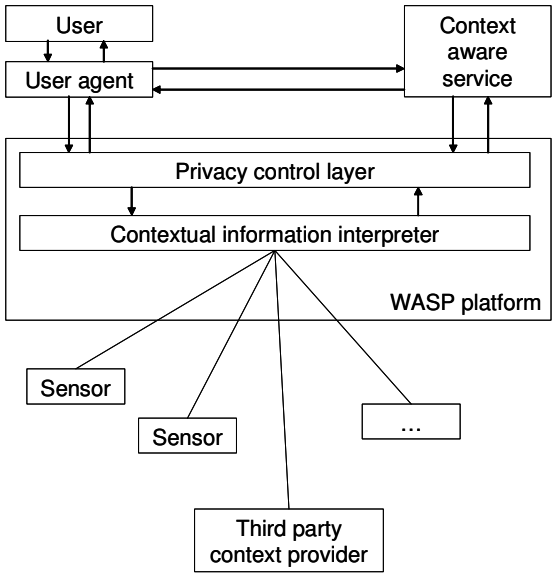

Figure 4: WASP privacy architecture overview.

In WASP, we separate the evaluation of the privacy policy of a service and the evaluation of a user's context-dependent preferences. Before a user invokes a service, his user agent will retrieve the P3P policy of the service. This will be compared against the user's preferences, expressed in the extended version of APPEL. If the service's policy is acceptable (evaluated automatically from the APPEL preferences, or after interaction with the user if no suitable rule is available for automatic evaluation), the user agent will store the association between the user, the service and the agreed privacy policy in the platform's privacy control layer.

Together with this association, the user's context-dependent part of the preferences will be stored. Every time a service needs contextual information, it will contact the WASP platform. The privacy control layer will then check for an association record between the service, its policy and the user it wants information about. If this association exists, the privacy control layer will evaluate the context-dependent preferences expressed by the user. If the context-dependent constraints are satisfied, the privacy control layer will release the requested contextual information to the service.

The separation between the checking of the privacy policy against the user's preferences and the checking of the request against the user's context-dependent preferences is made for two reasons. First, it expresses the conceptual difference between a privacy policy, describing a service's permanent characteristics, and context-dependent preferences, which are of a frequently changing nature. Second, it provides a performance improvement. P3P policies, which can be quite complex, need to be compared to the user's preferences only once. For further requests for context by a context-aware service, only the context-dependent constraints need to be evaluated.

In practise, this separation enables to establish the willingness of a user to use a certain service in principle, even though there are some conditions under which the user may temporarily wish to stop being tracked by the service.

7 Related work

The use of P3P in context-aware systems has been proposed by several researchers.

Marc Langheinrich has developed *pawS* [6], a privacy-awareness system using P3P. In this system, data collecting services, which may also include sensors such as a camera system, announce their presence to a user and provide the user with a P3P policy describing their behaviour. The user may then choose whether or not to use the service. Langheinrich expresses the need to extend P3P with the capability to describe contextual information. The system focuses on tracking services that are on by default. The main difference to the WASP approach is that in the latter, users take the

initiative to request (web) services from service providers.

Myles, Friday and Davies have also built a privacy control system based on P3P, called *LocServ* [9]. It addresses requests initiated completely by a service provider. Users of the context-aware system express their preferences in so-called validators. These validators contain APPEL-like rules. Whenever a service provider wants to collect information about a user, the context-aware platform will look for a validator containing a rule that allows data collection by this particular service provider. Only if a validator is found that accepts the request from the service provider, context information is released. The validators support context-dependent constraints.

Other approaches for controlling privacy in context-aware and ubiquitous computing environments include anonymous and pseudonymous communication [2, 3] and privacy control of documents by attaching metadata describing permissions on the documents [5].

8 Conclusions

In this paper, we identified the requirements for applying a P3P based privacy control mechanism in our context-aware WASP platform. A privacy architecture has been designed based on an extended version of P3P and its associated preference expression language APPEL.

We are currently implementing the proposed extensions to P3P and APPEL. The next step in our research will be the actual design and implementation of a prototype of the proposed privacy architecture. This prototype will be integrated in the WASP project.

References

[1] Ackerman, A. et al., Privacy in context, *Human-Computer Interaction*, 16, 2001, pp. 167-176.

[2] Al-Muhtadi, J. et al., Routing through the Mist: Privacy Preserving Communication in Ubiquitous Computing Environments, in *Intl. Conf. of Distributed Computing Systems* (ICDCS 2002) ,Vienna, July 2002, pp. 65-74.

[3] Beresford, A.R. and F. Stajano, Location Privacy in Pervasive Computing, *IEEE Pervasive Computing*, Jan.-Mar. 2003, pp. 46-55.

[4] Dey, A. K. and G. D. Abowd, *Towards a Better Understanding of Context and Context-Awareness*, Technical Report 99-22, Georgia Institute of Technology, 1999.

[5] Jiang, X. and J.A. Landay, Modeling Privacy Control in Context-Aware Systems, in *IEEE Pervasive Computing*, 2002.

[6] Langheinrich, M., A Privacy Awareness System for Ubiquitous Computing Environments, in *4th Intl. Conf. on Ubiquitous Computing* (UbiComp 2002), Springer-Verlag LNCS 2498, September 2002, pp. 237-245.

[7] Langheinrich, M., Privacy by Design - Principles of Privacy-Aware Ubiquitous Systems, in *3th Intl. Conf. on Ubiquitous Computing* (UbiComp 2001), Springer-Verlag LNCS 2201, 2001, pp. 273-291.

[8] Lederer, S. et al., Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing, to appear in Extended Abstracts of the *2003 Conf. on Human Factors in Computing Systems* (CHI 2003), April 2003.

[9] Myles, G. et al., Preserving Privacy in Environments with Location-Based Applications, *IEEE Pervasive Computing*, Jan.-Mar. 2003, pp. 56-64.

[10] Nilsson, M. et al., Privacy Enhancements in the Mobile Internet, in *IFIP WG 9.6/11.7 Working Conf. on Security and Control of IT in Society*, Bratislava, June 2001.

[11] Parlay Group, *Parlay X Web Services White Paper*, http://www.parlay.org/specs/library/ParlayX-WhitePaper-1.0.pdf, 2002.

[12] Schilit, B. et al., Context-Aware Computing Applications, in *1st Intl. Workshop on Mobile Computing Systems and Applications*, 1994, pp.85-90.

[13] W3C, *A P3P Preference Exchange Language 1.0 (APPEL1.0)*, available on http://www.w3.org/TR/P3P-preferences/, April 2002.

[14] W3C, *The Platform for Privacy Preferences 1.0 (P3P1.0) Specification*, available on http://www.w3.org/TR/P3P/, April 2002.
[15] WASP Project, http://www.freeband.nl/projecten/wasp/ENindex.html, 2002.
[16] Wu, M. and A. Friday, Integrating Privacy Enhancing Services in Ubiquitous Computing Environments*, in Workshop on Security in Ubiquitous Computing, 4th Intl. UbiComp Conf.*, 2002.