A Privacy Enhancing Infrastructure for Context-Awareness

Patrik Osbakk
Computing Laboratory
University of Kent
Canterbury, Kent
CT2 7NF, UK
pjo2@kent.ac.uk

Nick Ryan
Computing Laboratory
University of Kent
Canterbury, Kent
CT2 7NF, UK
n.s.ryan@kent.ac.uk

ABSTRACT

Privacy protection remains a serious bar to the widespread deployment of context-aware and ubiquitous computing environments. Here, we outline an experimental privacy-enhancing infrastructure based on Role- Based Access Control and the use of P3P.

Keywords: Context, Privacy, RBAC, P3P

INTRODUCTION:

Rapid progress has been made in context-aware computing and projects including visitor guides [1][2], the active badge system and its successors [3], cooltown [4], and MusicFX [5] demonstrate that context-aware environments are set to become a reality. However, before they are ready to be introduced into our daily lives, a major challenge remains to be addressed, *privacy*. Surveys suggest that people do worry about security when sending personal data over the Internet [6] and fear of misuse affects their behaviour [7]. Context-aware environments are expected to make these issues more acute [8]. It is therefore essential that privacy protection mechanisms are in place from the start, so that a relationship of trust can be formed between technology users and providers. Here, we outline a privacy-enhancing infrastructure being developed as part of a project that aims to ease the development of privacy-friendly context-aware applications and to evaluate the practicality of privacy protection in ubiquitous environments.

CONTEXT

We intentionally employ a very broad definition: *context is information related to an entity, where the information may be an entity itself.* An entity can be anything from “people, places, and things” [4] to activities and concepts. The intention is to stress the existence of relationships between entities and between entities and data values. In this way, we allow application-specific definitions to coexist within the infrastructure. In the real world we will get an almost infinitely complex network defining the contextual relationships between entities.

PRIVACY DEFINED

We have asserted that privacy is desirable and needs to be protected, but what is privacy? In this paper we focus on the flow of information rather than physical privacy. The definition that has been adopted is “Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others” [9]. As such the ownership of information is with the subject, who alone should control the release of information. This leads to an important assumption. It is assumed that once information has been revealed no control over its use can exist. Therefore, a primary determinant for whether information is disclosed is the trust placed in the recipient.

REQUIREMENTS

Firstly, the infrastructure must implement a privacy protection mechanism. The level of protection deemed necessary should equate to what we enjoy if we are offline. Perfect privacy is, from a practical point of view, seen to be impossible to achieve. Using location as an example, we accept that anyone who wants to know our location could do so by utilising our environment, e.g. by using CCTV or asking a mutual friend. In most cases they would ask us, allowing an informed choice whether to disclose the location. As such we do not feel that our privacy is threatened. Additionally, a privacy protection mechanism needs to be able to support both known and previously unknown agents, where an agent can be either a context consumer (client) or a context producer (service).

Secondly, for the infrastructure to be useful it will need to provide support for common context-aware applications such as in-out boards, call forwarding, etc.

Thirdly, the infrastructure needs to run on heterogeneous limited devices with variable connectivity. This is essential to maximise the usefulness of the infrastructure and to support connected but offline groups.

INFRASTRUCTURE MODEL

A flexible decentralised infrastructure is under development [10]. It is distinct from, though closely related to, other ongoing work at Kent [11]. Each entity has at least one associated context manager (CM). Multiple, synchronised, CMs may be employed in a distributed environment. The CM is responsible for storing, processing, and protecting the entity’s context information. All context information owned by an entity flows through their CM. Requests from other agents are sent to the owner’s CM. The CM then fetches the context element from a service or storage. Hence the acquisition and use of context information is completely separated by the context manager, a requirement that others have found desirable [12].

To support such a distributed structure with unknown and variable connectivity, a component-based approach is taken. Plug-ins handle all communication with a CM, allowing any medium to be used. Similarly, cipher plug-ins will be used to secure communication. This will allow use of custom hardware accelerators as well as supporting differences in legislation and requirements. Together, these plug-ins will allow context to be communicated securely from services to the context manager and then to clients.

Given that all information flows through the CM, it can fully control what is disclosed to whom and hence protect the privacy of the subject. The base mechanism for authentication is a username-password combination. This ensures full compatibility with limited devices. An optional public key mechanism is currently under consideration. To minimise the administrative burden a Role Based Access Control (RBAC) [13] is used. RBAC is a well-tested and formal access control mechanism that uses the idea of roles to group permissions. The model used here is based on $RBAC_0$ [13], but with a important difference. Roles are automatically activated so that a user is given the best access possible in any session. This difference simplifies the model, making it easier to both manage and process on limited devices.

In our RBAC implementation the permissions, contained in the roles, are a list of access controls. An access control is always linked to one context element and grants access above default access denied level. Any combination of read, write, and history, including none of them, may be granted. A restriction can be placed upon how far back in time history access is granted. To allow further customisation each user can be assigned a personal permission that takes precedence over the access granted by any role. The use of negative permission and constraints to provide even greater flexibility is currently being investigated.

To allow access by unknown agents, as well as an addition to the RBAC for known agents, P3P policies may be used. A P3P policy is here seen to be a contract stating how the information, if any, might be used. A P3P policy accompanying a request is compared to the available rulesets specified by the entity. A ruleset defines what is required to proceed with an action and can be associated with any number of roles in the RBAC. Each ruleset matching a policy will temporarily, for one request only, add its associated roles to those already assigned to the requesting agent.

Access can thus be granted to both known and unknown agents. This access control mechanism forms the core of the CM and every request is filtered to remove unauthorised actions. After filtration the CM can evaluate the request and create an appropriate response.

FUTURE WORK

Future work includes investigating the possibility of integrating multiple infrastructures, use of pseudonyms, and anonymity. Being able to integrate the infrastructure with both a context-aware office and a smart home would enable utilisation of their existing sensor networks. It would also allow the privacy protection mechanisms to be applied to an unprotected infrastructure. Although information would still be free within such an environment any access from the outside can be controlled.

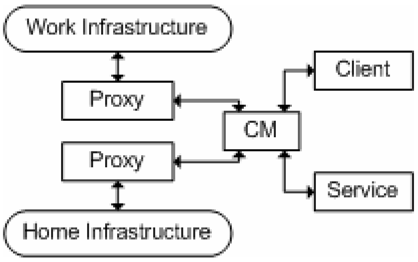

The use of pseudonyms and the possibility of being anonymous are classical privacy protection mechanisms that need to be investigated to see if they can improve privacy in ubiquitous environments. Is it useful to provide context information while being anonymous? What is being anonymous? Can a user known only by a pseudonym be trusted to use the information as agreed? Many issues need to be investigated before anonymity and pseudonymity can be introduced.

CONCLUSION

The privacy enhancing infrastructure described here provides an access control mechanism that enforces an entity's preferences. An early un-optimised implementation has shown these mechanisms to perform well even on limited devices. This progress shows that privacy-friendly context-aware applications are feasible and not an unrealistic desire.

REFERENCES

1. Abowd, G. D., C. G. Atkeson, et al. (1997). "Cyberguide: a mobile context-aware tour guide". *Wireless Networks* **3**(5): 421-433.

2. Cheverst, K., N. Davies, et al. (2000). "Developing a Context-aware Electronic Tourist Guide: Some Issues and Experiences". *Conference on Human Factors and Computing Systems*, The Hague, Netherlands.

3. Harter, A., A. Hopper, et al. (2002). "The Anatomy of Context-Aware Applications". *Wireless Networks* **8**(2-3): 187-197.

4. Kindberg, T., J. Barton, et al. (2002). "People, Places, Things: Web Presence for the Real World". *Mobile Networks and Applications* **7**(5): 365-376.

5. McCarthy, J. F. (1998). "MusicFX: An Arbiter of Group Preferences". *AAAI Spring Symposium on Intelligent Environments*, Palo Alto, USA.

6. ICM poll (2002). Big Brother part 1. *The Guardian*, 07/09/2002**:** 3.

7. European Commission (2002). "Questionnaire on the implementation of the Data Protection Directive (95/46/EC)". Your voice in Europe, *http://europa.eu.int/yourvoice/results/204/index_en.htm*.

8. Crowcroft, J. (2003). "Scalable Ubiquitous Computing Systems or just Ubiquitous Systems. A 15-year Grand Challenge for computer science". *http://www.nesc.ac.uk/esi/events/Grand_Challenges/proposals/US.pdf*

9. Westin, A. F. (1970). *Privacy and freedom*. London, Bodley Head.

10. Osbakk, P. and Ryan, N. (2003) A Privacy Enhacing Infrastructure, *http://www.cs.kent.ac.uk/projects/ubi/infra/privacy/*.

11. Ryan, N. and Osbakk, P. (2003) The MobiComp Infrastructure, *http://www.cs.kent.ac.uk/projects/ubi/infra/mobicomp/*.

12. Dey, A. K. and G. D. Abowd (2000). "The Context Toolkit: Aiding the development of Context-Aware Applications." *Workshop on Software Engineering for wearable and pervasive computing*.

13. Sandhu, R. S., E. J. Coyne, et al. (1996). "Role-Based Access Control Models." *IEEE Computer* **29**(2): 38-47.