Dynamic Trust Models for Ubiquitous Computing Environments

Colin English, Paddy Nixon, Sotirios Terzis,
Andrew McGettrick and Helen Lowe.

Department of Computer and Information Sciences
University of Strathclyde, Glasgow, Scotland.
Colin.English@cis.strath.ac.uk

Abstract: A significant characteristic of ubiquitous computing is the need for interactions of highly mobile entities to be secure: secure both for the entity *and* the environment in which the entity operates. Moreover, ubiquitous computing is also characterised by partial views over the state of the global environment, implying that we cannot guarantee that an environment can always verify the properties of the mobile entity that it has just received. Secure in this context encompasses both the need for cryptographic security and the need for trust, on the part of both parties, that the interaction is functioning as expected. In this paper we make a broad assumption that trust and cryptographic security can be considered as orthogonal concerns (i.e. an entity might encrypt a deliberately incorrect answer to a legitimate request). We assume the existence of reliable encryption techniques and focus on the characteristics of a model that supports the management of the *trust relationships* between two entities during an interaction in a ubiquitous environment.

1 Introduction

Ubiquitous computing premises a massively networked world supporting a population of diverse but cooperating mobile entities where autonomous operation is necessary due to lack of central control. The composition and characteristics of this infrastructure will be both highly dynamic and unpredictable. Entities will have to deal with unforeseen circumstances ranging from unexpected interactions to disconnected operation with incomplete information about the environment.

The infrastructure that supports this ubiquitous computing system introduces new security challenges not addressed in existing security models; particularly in the domain of trust management. Humans use trust as a means to reason about and accept risk in situations of partial information and assign privileges accordingly. It is subjective and situation specific [2] in its nature as an individual's opinions are based on observations in a particular environment. Trust in one environment does not transfer to another environment and as a result, a notion of context is necessary [3]. This makes it very difficult to form a definition incorporating all views and types of trust identified by humans [4, 5].

The trend in trust management systems is to view trust implicitly through the delegation of privileges to trusted entities via the use of certificates, which can be chained to represent recommendations and the propagation of trust [6]. Trusted entities are decided by some central authority, be it the end user or system administrator. This coarse view of trust fails to capture the many intricacies of trust as intuitively viewed by humans.

1.1 Motivation

In decentralized ubiquitous systems, current trust models fail on a number of points. Firstly, only partial information may be available, as requests can come from unknown entities or environments may be unfamiliar or hostile. Secondly, mobile entities are likely to become disconnected from their home network and must be able to make fully autonomous security decisions without relying on a specific security infrastructure. Thus the use of certification authorities may not be possible. Third, the formation and evolution of trust, which are central to human intuition of the phenomenon, are neglected in current systems [7,8,9]. The only attempts at evolution are based around certificate revocation, which reduces options when a task must be carried out by the best of a bad bunch. Choosing between alternative collaborators is difficult in the case of implicit trust representation.

These issues must be resolved to be able to assign meaningful privileges and facilitate interaction in such a complex world and bring tremendous potential for new services. The aim of this work is to help create a user-intuitive Information Society where people have confidence in the systems they use everyday. Lack of trust in security mechanisms is evident in the reluctance to accept e-commerce, fuelled by a number of publicised attacks exposing weaknesses, which need addressed before users will adopt services provided by these systems.

Our position is that the ability to form and evolve explicit values for trust in other principles in an interaction allows autonomous computational entities to make better decisions in situations where only partial information is available.

2 Adopted Approach

2.1 Objectives

- Facilitate the ad-hoc interaction of unknown autonomous entities in situations of partial information by the definition of a trust model sufficiently detailed to allow entities to reason about and compare the trustworthiness of other entities for security related decisions.
- Capture the dynamic aspects of trust formation and trust evolution with fine granularity.
- The model must capture human intuitions about trust to ensure understanding by users, thus reducing security vulnerabilities in implementations.

2.2 Initial Ideas

Ad-hoc interaction between mutually unknown entities can take place only if there is an adequate level of trust between the parties. As mentioned above, the implicit, coarse and static view of trust in current systems fails to model the notion of trust, as human intuition understands it. A dynamic model of trust will provide the ability to operate and make decisions autonomously. While trust defies stringent definition, it is proposed that a model with explicit trust values can be realised in sufficient detail to be used either to augment other security mechanism or as a basis for unencrypted interactions. With a range of explicit values representing trust, a finer granularity of representation is achieved, providing entities with enhanced information on which to base decisions. Values may also be stored in memory, to represent historical information on the behavioural patterns of specific entities. It is also proposed that in situations where a task must be carried out by the 'best of a bad bunch', finer granularity of trust representation will facilitate comparisons between entities.

There are three main sources of trust information about another entity. Personal observations of the entity's behaviour are essential for the subjective evaluation of trustworthiness. Therefore the outcome of interactions is recorded and made available as evidence to all principals. Recommendations from trusted third parties provide the possibility for trust to be propagated between unknown entities in a similar manner to the deferment of trust as seen in current trust models, including the PGP 'web of trust' [10]. The reputation of an entity can be consulted in the absence of experience or recommendation, in effect, acting as an anonymous recommendation. Recommendations may take the form of signed credentials to be evaluated subjectively within a specific environment.

A downfall of most access control mechanisms on the Internet is the reliance on authenticated identity of the principal involved to provide access control. In the types of systems in the GCI vision, it may be impossible to establish the identity of unknown entities. Even when identity can be established, for example via intersecting certificate hierarchies in PKI [11], this conveys no *a priori* information about the likely behaviour of an entity. It is therefore proposed that all participants be assumed virtually anonymous, with consideration given to recognition of entities rather than identity. In this way, the necessity for prior configuration of collaborative entities is removed, allowing unforeseen circumstances to dealt with autonomously as they arise. Recognition through digital signatures is based on previous encounters whereas identity is established before interaction takes place. To allow this, auto-configuration measures must be in place for the formation of an initial level of trust when entities meet for the first time.

2.3 Dynamic Aspects of the Model

This paper proposes that the use of a range of explicit values for trust will provide finer granularity for the dynamic aspects of trust between agents in the systems described above. This will result in a more flexible model able to represent trust in a manner that captures human intuitions, such that positive outcomes of interactions will preserve or amplify trust, while trust erodes without periodic interactions or recommendations.

2.3.1 Trust Formation

The process of establishing the initial trustworthiness of each collaborator is referred to as trust formation. A summary of an entity's trustworthiness can be synthesized from the history of its past interactions to be used by other entities when allocating privileges with specific risks. Evidence relevant to the current context will carry the most weight. Initially new entities have no evidence of past behaviour to establish a base for interaction. To form an opinion of trustworthiness in this case requires the presence of some optimistic entities willing to take risks in unknown situations, allocating privileges judiciously until experience shows that it was unwise.

2.3.2 Trust Evolution

The evolution process can be regarded as iterating the process of trust formation as additional evidence becomes available. Accumulation of evidence with experience of new interactions must modify the level of trust to be placed in an entity, incrementing the summary information to maintain accuracy. The risk assessment for an entity performing an action in a particular context will change depending on how much is known about positively or negatively perceived actions in the past. A successful high-risk interaction results in greater increase of trust than a successful low risk interaction. Conversely, the lower the level of risk, the greater the penalty for a failed interaction.

This granularity of evolution is seen to be necessary when Byzantine behaviour is considered. The reason for a failure may be more important than the fact that the failure occurred. Most people would alter their level of trust in another more radically if a failure were intentional and malicious rather than accidental. Using historical information, patterns in previous behaviour may be analysed to help determine the reason behind failure. The only evidence of the outcome of interactions may be from dishonest sources, requiring measures to be in place to modify the reputation of certificate signatories and collaborators in cases of framing or collusion.

2.3.3 Trust Exploitation

The essential problem in exploitation is to determine behaviour on the basis of trust, which balances risk and benefit within the context appropriately. Security policy for access control is expressed in terms of trust and specifies the level of positive experiences required to allow access to a specific resource. Policy will determine whether an entity is optimistic or pessimistic about an interaction depending on the risks involved.

3 Status and Open Issues

As part of the SECURE project an initial formal trust model is being developed which addresses some of the issues that arise in using trust as part of a security mechanism, such as the representation of trust and of recorded evidence. The model will help determine exactly where the importance of context lies, what constitutes the context and how context-awareness can be achieved. Similarly, the model being developed will also lead to a better understanding of how risk can be estimated and the mechanisms to provide privileges based on trust.

The development of simulations to test aspects of the model is intended to address questions such as the auto-configuration mechanism and methods for entity recognition. The simulation work, currently based around an agent based file sharing facility, will also aid understanding of the temporal aspects of memory and, awareness and predictability of dishonest behaviour. It is hoped that privacy implications of displaying historical information will become clearer through these investigations.

It is envisaged that these insights will allow the development of a trust-risk based access control system for mobile entities and a supporting lifecycle management system for such interactions.

It is worth adding that the workshop presentation will outline the principles of the formal model and of the simulation system in their current state and will endeavour to present early results derived from them.

Acknowledgements

The work is this paper is supported by the EU project **SECURE:** Secure Environments for Collaboration among Ubiquitous Roaming Entities (IST-2001-32486) and the EU FET project **GLOSS:** Global Smart Spaces (IST-2000-26070).

References

[1] EU Future Emerging Technologies, Global Computing Initiative.

http://www.cordis.lu/ist/fetgc.htm

[2] D. McKnight and N. Chevany: “The Meanings of Trust”. Working paper, Carlson School of Management, University of Minnesota, 1996.

[3] S. Marsh: “Formalising Trust as a Computational Concept”. Ph.D. Thesis, University of Stirling, 1994.

[4] M. Deutsch: “Cooperation and Trust: Some Theoretical Notes”. In: M. Jones (ed), Nebraska Symposium on Motivation. Nebraska University Press, 1962.

[5] R. Golembiewski and M. McConkie: “The Centrality of Interpersonal Trust in Group Processes”. In: C. Cooper (ed), Theories of Group Processes, Wiley, 1975.

[6] A. Jøsang: “The right type of trust for distributed systems”. In: C. Meadows (ed), Proceedings of the 1996 New Security Paradigms Workshop. ACM, 1996.

[7] M. Blaze, J. Feigenbaum, and J. Lacy: “Decentralized trust management”. In Proceedings of the 1996 IEEE Symposium on Security and Privacy, pp.164-173, May 1996.

[8] M. Blaze, J. Feigenbaum, J. Ioannidis, and A. Keromytis: “The KeyNote Trust Management System - Version 2”. Internet Engineering Task Force, September 1999. RFC 2704.

[9] Y.-H. Chu, J. Feigenbaum, B. LaMacchia, P. Resnick, and M. Strauss: "REFEREE: Trust Management for Web Applications," World Wide Web Journal, 2 (1997), pp. 706--734.

[10] S. Garfinkel: “PGP: Pretty Good Privacy”, O'Reilly & Associates, Inc., 1995.

[11] U. Maurer: “Modelling a Public-Key Infrastructure”. In Proceedings of the 1996 European Symposium on Research in Computer Security, Lecture Notes in Computer Science, vol. 1146, pp. 325-350, 1996.