

Security Models for Trusting Network Appliances

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 pervasive computing is the need for secure interactions between highly mobile entities and the services in their environment. Moreover, these decentralised systems are also characterised by partial views over the state of the global environment, implying that we cannot guarantee verification of the properties of the mobile entity entering an unfamiliar domain. 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. cryptographic measures do not ensure transmission of correct information). 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 devices during ad-hoc interactions.

1 Introduction

Ubiquitous and pervasive computing premises a massively networked world supporting a population of diverse but cooperating mobile entities [1], ranging from mobile computational agents to mobile devices such as handheld PDAs and mobile phones. Many of these devices stand to benefit from the ability to interact and cooperate with other entities and services, whether static or mobile, to allow successful execution of allocated tasks even in unfamiliar surroundings. The capability of PDAs to form an ad-hoc connection to a network printer on an unfamiliar LAN is an example of this. Within such an infrastructure with highly dynamic and unpredictable characteristics and composition, autonomous operation is

necessary due to lack of central control. Entities will have to deal with unforeseen circumstances ranging from unexpected interactions to disconnected operation with incomplete information about the environment. Security plays an important role in this infrastructure, as the risks inherent in interacting with services and other mobile entities are many and varied.

The infrastructure that supports this pervasive computing system introduces new security challenges not addressed in existing security models, including 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 therefore reasonable to consider trust as a mechanism to facilitate interaction between mobile devices and the facilities within the environment. Trust is subjective [2], being a personal opinion based primarily on first hand observations or carefully considered advice from others if available, allowing decisions to be made with only partial knowledge. Trust is also 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 necessarily transfer to another environment and as a result, a notion of context is necessary [3]. It is also a highly dynamic phenomenon, which evolves dependent on new evidence as it becomes available. These factors, while providing great flexibility, make it very difficult to form a definition incorporating all views and types of trust identified by humans [4, 5].

2 Current Trust Management

Matt Blaze et al. [6] define trust management as “a unified approach to specifying and interpreting security policies, credentials and relationships that allow direct authorisation of security-critical actions” In such trust

management systems, trust is viewed implicitly through the delegation of privileges to trusted entities via the use of credentials or certificates, which can be chained to represent recommendations and the propagation of trust between entities [7]. This implicit coarse view of trust fails to capture the many intricacies of trust as intuitively viewed by humans. Many of these aspects are essential for a trust model that must operate without central control, to allow security decisions to be made by autonomous entities or devices in situations where no specific security infrastructure can be relied upon.

Although current trust management systems as defined above provide many useful and valid insights, the lack of explicit trust evaluation precludes many of the aspects deemed necessary for autonomous entities to reason about trust for flexible security paradigms. In the systems considered here, these approaches fail on a number of other important, general points. Firstly, many rely on complete information, where 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 security decisions without relying on a specific security infrastructure or certification authority. The user or some central authority such as the system administrator currently often decides which entities are trustworthy and as a result, entities cannot dynamically reconfigure themselves to cope with unforeseen circumstances or requests for service from unknown devices entering their administrative domain. Thirdly, the dynamic aspects of trust formation, evolution and exploitation, which are central to human intuition of the phenomenon, are largely neglected in current systems [6], [8], [9]. Formation of implicit trust relationships generally requires some form of prior configuration, which may be impossible in situations where the device is disconnected from its home network. Most attempts at evolution are based around certificate revocation, which is a very negative and coarse view, making the choosing between alternative collaborators difficult via implicit trust representation.

These issues must be resolved to be able to assign meaningful privileges and facilitate interaction between devices 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, by removing the need for the user to consider or even understand the security implications of actions they take or have taken on their behalf. The lack of trust in current security mechanisms is evident in the reluctance to accept e-commerce, fuelled by a number of publicised attacks

exposing weaknesses that need addressed before users will adopt services provided by these systems.

The view taken here is that the ability to form and evolve explicit values for trust in other principles in an interaction allows autonomous computational entities within devices to make better decisions on the user's behalf in situations where only partial information is available.

3 Adopted Approach

3.1 Objectives

- Facilitate the ad-hoc interaction of unknown autonomous devices 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, thus reducing security vulnerabilities in implementations.

3.2 Characteristics of the Trust Model

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 devices with the ability to operate and make security related 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 mechanisms 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. The reputation of an entity can be consulted in the absence of experience or recommendation, in effect, acting as an anonymous recommendation. Further information on which to base trusting decisions can be extracted from the environment or domain in which the entity is operating and thus a notion of context is necessary to incorporate the situational nature of trust. A strong basis for trust is established through an entity's subjective observations and the collection of such evidence, although exactly what properties of the interaction should be recorded and how must be established. A means of recognition should be included, together with trust values, entity state before and after the encounter, desired state after the encounter and some notion of context. Recommendations may take the form of signed credentials or evidence to be evaluated subjectively within an environment similar to the context of the recommendation.

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 be dealt with as they arise. Auto-configuration measures must therefore be in place to remove the reliance on centralised certification authorities and allow the formation of an initial level of trust when entities meet for the first time, even when devices roaming between administrative domains may be disconnected from their home network.

3.3 Dynamic Aspects of the Model

The dynamic aspects of how trust is formed, how trust evolves over time due to available information and how trust can be exploited are collectively referred to as the trust lifecycle. A model of trust incorporating the lifecycle will provide an entity with the ability to reason about and make security related decisions autonomously. 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. The temporal aspect of memory must be addressed if trust is to be modelled realistically with a sense of history. Trust information or values may be stored in memory, to represent historical information on the behavioural patterns of specific entities. Before any new interaction an entity will choose what fraction of its past to reveal, affecting the awareness and predictability of dishonest behaviour, based on patterns in the available evidence. This dynamic view of trust 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.

3.3.1 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, in particular subjective observations made by the entity itself about previous interactions. 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.

Recommendations may be used to establish collaboration between entities that have never met, but who trust a common third party. Recommendation chains can be used as a recursive version of this principle. Recommendation and reputation concepts are similar to the concept of the web of trust in the literature, except that all aspects of trust are dealt with in this way. Reputation can be consulted in the absence of experience or recommendation.

3.3.2 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 is 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.

3.3.3 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. Trustworthiness is interpreted through historical information before deciding to interact with another entity, with evidence relevant to the current context carrying the most weight. The risk assessment for an entity performing an action in a particular context will also change depending on how much is known about positively or negatively perceived actions in the past.

Security policy for granting requests is expressed in terms of trust and specifies the level of positive experiences required to allow access to a specific resource or service. If we consider trust as a mechanism for expressing the amount of risk an entity will accept in a particular context, the entity must evaluate the level and type of risk associated with the context. Policy will determine whether an entity is optimistic or pessimistic about an interaction depending on the scale of the adverse consequences associated with the risks. An optimistic approach is appropriate when the risks are commensurate with the possible benefits, while a pessimistic one is likely to be adopted whenever the potential risk is high. Optimistic behaviour allows new entities to take the first steps towards establishing their trustworthiness. Pessimistic behaviour is essential when a great deal is at risk; we must be sure of past good behaviour in similar contexts.

4 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 to confer privileges based on trust. The formal trust model is being developed based on a lattice of trust values, using the trust policies of an entity

as functions for least fixed-point calculations. The model currently uses trust delegation rather than recommendations, and it is in the early stages of development.

A risk model is also being developed, considering the interaction between trust and risk, the properties of which are uncertain at present. It is unclear whether the fact that a particular entity is trusted affects the perception of risk or affects the willingness to accept risk in an interaction with that entity. Currently the assumption is the former perceptive notion, where willingness to accept risk is captured by some utility function, although this decision may be reconsidered as the model is developed further. Whether trust values can be incorporated into the risk model in both ways remains to be seen, although it seems that this is likely to introduce issues of double counting of trust values, resulting in unrealistic trusting decisions being made. Other current assumptions are that for a specific action, all outcomes are known and the costs or benefits associated with each one can be calculated.

Entity recognition is being addressed within SECURE through the use of a pluggable recognition module. The mechanism being developed should have two important properties. Firstly, a change of identity will be possible, but is discouraged by the penalty of loss of past history and privileges and secondly, spoofing of identities should not be possible.

Our priority is currently the development of a model of the dynamic aspects of the trust lifecycle, to determine the feasibility of such an approach and examine the issues of evidence collection and use in the formation, evolution and exploitation of trust values. The concepts of recommendation and reputation will be studied to determine how these can be best represented. Studying the dynamic aspects introduces the concept of time and memory to the trust model; therefore these aspects must be examined in relation to the evidence stored on trust from previous interactions, enhancing the awareness and predictability of dishonest behaviour. Issues of the exact scope of trust within the model will have to be addressed. It is envisaged that these insights will allow the development of a trust based security model for mobile entities, where risk is a central component, and a supporting lifecycle management system for such interactions.

The investigation of how to model and combine the three forms of trust evidence are important as the scalability of the system is affected by the decision on how to represent the information gained from an interaction. Evidence based on personal observations clearly has more value than that from other sources, but

evidence available from other entities is also important and can be passed as a recommendation, for subjective evaluation when, for example, they have vastly more experience than you. Reputation is less important in that it is less reliable than the other forms of evidence. It may be represented in various ways, being weaker by conveyance of less information or information source anonymity. This could mean that only the trust value is passed rather than the evidence, or that the information is an amalgamation of recommendations received and passed on by an entity with no experience of its own. Our work in this area contemplates a similar approach to the work of Catholijn Jonker and Jan Treur [12], using sequences of positive or negative past experiences of various degrees to evolve trust from all previous experience, or update an existing value. In the initial stages, we assume the absence of Byzantine behaviour and non-cooperative scenarios and leave these for longer-term future work.

We are developing a simulation framework, where entities are represented by agents, for the investigation of trust lifecycle issues. The model will be tested using simulations rather than implementation scenarios, as this allows control over independent variables and a range of complex behaviours to be studied. We are unlikely to be able to run “real-life” experiments of more than a few cases even if these were desirable in the first instance. In real life we cannot control independent variables so failure (of our model to live up to expectations) would tell us little. Also, we could only test very benign scenarios where no one was really going to get hurt as a result of their behaviour. Simulation is, therefore, an important weapon in our armoury. The simulations will test the applicability and scalability of all aspects of the model and address issues such as the auto-configuration mechanism and methods for entity recognition. It may be possible to examine the correctness of assumptions such as agents always behave in a rational manner and examine the effects on the system when such assumptions are removed. Work started with an implementation of two specific scenarios, an agent-based file sharing facility and trust based dynamic routing in ad-hoc networks, which we are now generalizing to produce the simulation framework. It is also hoped that privacy implications of displaying historical information will become clearer through these investigations.

5 Conclusion

Traditional hard coded security models lack the flexibility to be of use in pervasive systems consisting of ad-hoc interactions between autonomous mobile devices, where only incomplete information is available on which to base security decisions. A weaker, but more flexible

model is the application of trust based security models to cope with the risk inherent in interactions in this environment. Current trust management solutions fail to capture the notion of trust and its relation to risk in a manner suited to these systems with no form of central control. The requirement for pre-configured trust information highlights the lack of flexibility. This paper discusses the characteristics necessary to provide a basis for reasoning about trust in security related decisions for these systems. Although it is clear there are many open issues, these will only be fully determined by the continuing work, which is expected to address these problems.

Acknowledgements

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

[12] C. M. Jonker, J Treur: "Formal Analysis of Models for the Dynamics of Trust Based on Experiences" *Modelling Autonomous Agents in a Multi-Agent World 1999 European Workshop on Multi-Agent Systems*. pp. 221-231, 1999.