

Trust for Ubiquitous, Transparent Collaboration*

Brian Shand, Nathan Dimmock, Jean Bacon

University of Cambridge Computer Laboratory

JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

{Firstname.Lastname}@cl.cam.ac.uk

1 Introduction


Ubiquitous computing needs mutual trust between participants in order to support collaborative tasks, such as arranging meetings, while protecting sensitive information used in the collaboration. At the same time, security measures must be proportional to the risk involved to allow the interaction between devices to be as automated as possible.

We believe that many security systems fail because of their high administrative overhead — passwords on post-it notes attached to monitors for example, because proper user-account administration is considered to be too much work — and so we aim to create a security mechanism suitable for use in the pervasive computing environment where human intervention is a valuable resource, yet due to the nature of the data involved, security remains important.

The full version of this paper proposes trust and risk models to help automate interactions of this sort, making the computations as unobtrusive as possible while still respecting participants' trust beliefs.

2 Trust infrastructure

Our trust framework uses a homogeneous recommendation system, to allow users to share and exchange privileged information. This information can include conventional data such as personal contacts and calendar entries, and also trust beliefs about principals.



(a) Alice sends Bob her work phone number.

(b) Some time later, Bob forwards Alice's number to Charlie.

Figure 1. Recommendations in action.

For example, Alice gives a telephone number to Bob, together with recommendations that it is her telephone number and that it be considered privileged business information. This is illustrated in figure 1(a). Alice signs i to certify that she is the origin of the information and also signs her recommendations to allow the recipient to evaluate their relevance using Alice's trust-rating. t represents Alice's trust in her recommendation, that is, how much confidence she has in it. Later, Bob forwards Alice's number to Charlie (figure 1(b)), along with her recommendation that it is her number (R) and Bob's recommendation that it is her "work" number (R'') in which he has trust t'' . He could also forward

Alice's original recommendation R' that it her "business" number if he wished to, but he has chosen not to in this case as the transmission link is expensive and he thinks Charlie will find his recommendation more useful.

Existing trust models for pervasive computing typically represent trust using a security policy which explicitly permits or prohibits actions [2]. These policies are not well suited to dynamic environments, in which participants have only partial trustworthiness, and trust assessments must constantly change. To avoid this, Abdul-Rahman and others [1] have also proposed explicit recommendation systems, but with only very simple trust values. In our work, we use recommendations to control the flow of information, as well as for access control; we are also able to combine our more complex recommendations consistently, by formally ordering recommendations according to information content. This gives us a well-founded approach to trust management decisions, which is suitable for distributed computing applications.

2.1 Phone book example

A phone book exchange service illustrates the need for and advantages of trust-based information exchange for ubiquitous computing. Users of handheld computers currently exchange contact details laboriously on a one-to-one basis. Furthermore, there is no associated trust information, so users cannot recommend to whom the information should be redistributed — for example, private and business numbers are usually redistributed together. Our trust and risk framework makes this service more transparent for users, while preserving the privacy of personal information.

The phone book database consists of many items, each with associated recommendations. These may be signed to prove their authenticity, using a public key infrastructure.

Figure 2 illustrates how Charlie uses the trust model to display Alice's phone book information in the example above. When Charlie searches his phone book for "Alice", he finds the entry for Alice's name, ranked according to the strength of its recommendations.

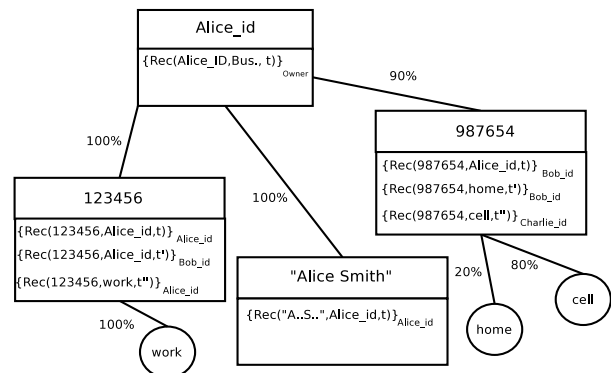


Figure 2. Each data item is stored separately and links between them determined by the trust model.

*This is an extended abstract of a paper with the same name that appeared at the *First IEEE Conference on Pervasive Computing and Communications*, March 2003, Ft. Worth, Texas, USA.

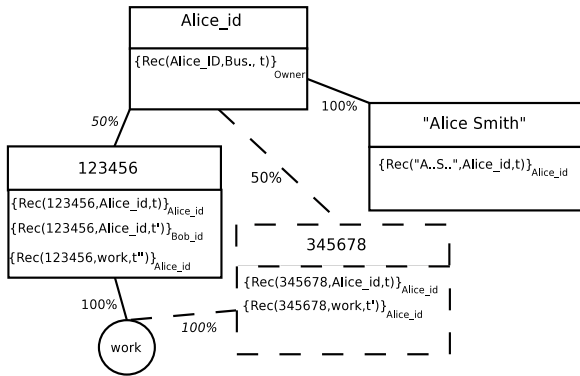


Figure 3. David receives a different view of the data, plus information from his own database (the dashed lines).

In contrast, consider David, Charlie’s colleague who is allowed to view business information in Charlie’s phone book, but nothing personal. If David views Alice’s information there, he is presented with the restricted view shown in figure 3. Furthermore, the importance of links might be different, if David had other knowledge of Alice, such as an old work number that is now out of date, as illustrated here.

In the next section we show how rôles and categories can be structured to preserve the *meaning* of recommendations. This ensures that user privacy is better protected in automated information transfers, by unifying trust assessments with access control.

3 Categories and rôles

Information exchange is restricted with the help of categories, arranged in a partial order, which are analogous to rôles in Rôle-Based Access Control. We extend RBAC rôles by associating a trust assessment with each category assignment. Users of the system can then combine a risk assessment, together with their trust in the information, to decide whether or not it should be used or displayed.

Each category has a list of privileges associated with it; these are action and category pairs which can be used by principals associated with the category. The overall trust assessment of an entity is thus a mapping from action and category pairs to primitive trust values.

Categories are arranged in a natural privilege hierarchy: when category c_0 extends the privileges of another c_1 , we write $c_0 \supset c_1$. Figure 4 illustrates a typical hierarchy, where the top category \top contains the owner of the PDA, c_n represent user defined categories such as immediate family, business colleagues, business contacts, friends and relatives. A are acquaintances, people known to the owner, but not cat-

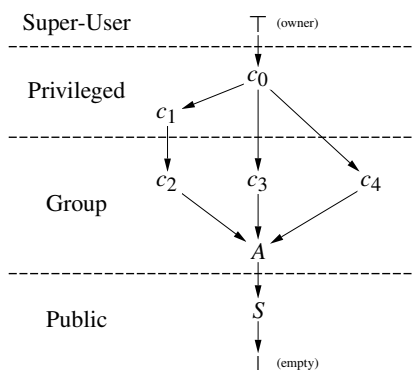


Figure 4. Example of a category hierarchy.

egorised, and S are strangers.

4 Trust and Risk Computations

Participants compute their trust in information, by combining their own trust assumptions with others’ recommendations. Recommendations associate one permission with another in our trust framework. By treating the identities of actors, categories and data entries as permissions, we can use a homogeneous recommendation structure for privilege assignment and for restricting the flow of information. Each trust value consists of a (*belief*, *disbelief*) pair, representing the weight of evidence for and against a particular trust assignment, with $belief + disbelief \leq 1$. No information is represented by (0,0), while (1,0) and (0,1) represent certain belief and disbelief respectively.

Users must combine their own recommendations with others’ to assess trust. This is achieved by forming a *policy function* for each principal’s recommendations (analogous to Weeks’ formalisation of access control system policies [4]). These policy functions are then combined to reach the appropriate trust conclusions [3].

4.1 Risk Assessment and Decision Making

We believe security measures must be proportional and appropriate for the risk involved: a user may happily distribute a business card to strangers to advertise their business, but may be quite careful as to whom they give their mobile phone number.

We use an outcome-based risk analysis to yield cost-benefit functions that are then used to make access control decisions. In line with existing literature on risk management, risk is defined as a function of the likelihood and impact of all the possible outcomes that can result from a decision. These cost-benefit functions (derived in [3]) not only incorporate the value of the data but also explicitly take into account the value of the PDA owner’s time if they must be interrupted to ask for guidance in making the decision.

5 Conclusions and Future Work

We have described how a PDA address book application could use our framework of trust-evaluated recommendations, combined with an explicit risk assessment, as an unobtrusive security model that permits transparent collaboration while still protecting the user’s privacy. We continue to work towards generalising the model so that it may be used with all PDA applications that allow collaboration and interaction, for example appointment diaries. We are also investigating more powerful and dynamic models of trust and risk that are suitable for use in this area.

References

- [1] A. Abdul-Rahman and S. Hailes. Supporting trust in virtual communities. In *Hawaii International Conference on System Sciences* 33, pp. 1769–1777, 2000.
- [2] T. Finin, A. Joshi, L. Kagal, O. Ratsimore, V. Korolev, and H. Chen. Information agents for mobile and embedded devices. *LNCS*, 2182:264–286, 2001.
- [3] B. Shand, N. Dimmock, and J. Bacon. Trust for transparent, ubiquitous collaboration. In *First IEEE Annual Conference on Pervasive Computing and Communications (PerCom 2003)*, pp. 153–160, Dallas-Ft. Worth, TX, USA, Mar. 2003.
- [4] S. Weeks. Understanding trust management systems. In *IEEE Symposium on Security and Privacy*, pp. 94–105, 2001.