Trust Dynamics for Collaborative Global Computing

Colin English, Sotirios Terzis, Waleed Wagealla, Helen Lowe, Paddy Nixon and Andrew McGettrick
University of Strathclyde in Glasgow
Colin.English@cis.strath.ac.uk

Abstract

Recent advances in networking technology have increased the potential for dynamic enterprise collaborations between an open set of entities on a global scale. The security of these collaborations is a major concern, and requires novel approaches suited to this new environment to be developed. Trust management appears to be a promising approach. Due to the dynamic nature of these collaborations, dynamism in the formation, evolution and exploitation of trust is essential. In this paper we explore the properties of trust dynamics in this context. Trust is formed and evolves according to personal experience and recommendations. The properties of trust dynamics are expressed through a formal model of trust. Specific examples, based on an e-purse application scenario are used to demonstrate these properties.

1. Introduction

In the future, there is likely to be increased use of the Internet and other wide area networks to provide a basis for collaboration between large numbers of diverse enterprise systems. To take advantage of the whole range of possibilities such a global computing environment creates for enterprise systems, it is essential to provide support for autonomous decision-making by its constituent entities. These entities must operate without the benefit of central control, with partial knowledge of the whole system and without reliance on a central security infrastructure inside unfamiliar administrative domains.

In this paper, the view is taken that the process of decision-making involves the estimation of likely behaviour; something not explicitly catered for by traditional security measures which focus on the identification/authentication of principals involved in an interaction. In addition to this unsuitable (for our purposes) focus on identity, the hard coded approach to centrally managed security domains is inflexible for environments with unpredictable composition; thus the responsibility falls to the entities themselves to make security related decisions to protect their resources from misuse by others and ensure payment is received for resource use.

The prediction of behaviour is very complex, particularly when consideration is given to the human mechanism for achieving this, specifically trust. This complex notion [6] has been studied extensively in sociological fields and some attempts have been made to establish its computational representation [8, 9], although most of these have fallen short of the various definitions intuitive to humans.

The aim of the SECURE [10] project is to define a computational model for trust-based decision making for the environments described above, which adheres to the commonly accepted characteristics of human trust [4]. Within this paper the approach adopted by SECURE is outlined, focusing on the use of trust information to provide dynamic trust evaluation, examining in particular a scenario from the project, the e-purse [2].

1.1. The E-purse Scenario

The scenario involves the use of an e-purse when interacting with a bus company. The purpose of the electronic purse is to hold a relatively small amount of electronic cash (in this scenario the e-purse is limited to 100 euro) that the owner can use as if it were real cash for buying bus tickets. The user can refill the e-purse when it is empty by contacting his/her bank. There are three different parties (principals) involved in the e-purse scenario: the user (owner) of the e-purse, the bus company, and the bank. For the purpose of this paper we are interested in modelling the relationship between the bus company and the user, taking the example interaction when users want to purchase a ticket using their e-purse. We focus on this aspect of the relationship as it highlights certain advantages for modelling the dynamics of trust as will become apparent later on in the paper.

E-cash is based on a protocol that, although protects user anonymity during normal transactions, enables identification of guilty parties in fraudulent transactions. Every time the bus company accepts e-cash in a transaction it takes the risk of losing money due to fraud. Therefore, for the bus

company to decide how to respond to a purchasing request, it needs to determine the trustworthiness of the user. Principals can assign different levels of trust to different entities based on the available information so as to evaluate the level of risk a transaction involving the user entails.

The rest of this paper is organized as follows: section 2 describes the trust-risk model of SECURE focusing mainly on the trust model. Section 3 analyzes principles for how trust is formed, evolved, and exploited, giving examples from the e-purse scenario. Section 4 concludes with discussion and examination of open issues.

2. The SECURE Model

The approach taken is based around the premise that trust and risk are inexorably linked and must both be considered when taking a decision about an ambiguous path, the outcome of which depends on the actions of another entity. At the heart of our approach lies the formal trust model [3] and risk model [1].

2.1. The SECURE Trust Model

Most existing systems follow the approach of binary trust values, which restricts the expression of trust to a certain degree (trusted or non-trusted). A finer granularity of trust values enables a wider variety of trust evolution strategies to be considered. For this reason, and since our aim is to follow human intuition on trust, we consider here more comprehensive and informative trust values. In the trust model, trust is expressed as a mapping (*m*) from pairs of principals (*P*) to trust information (*T*):

$$m : P \rightarrow P \rightarrow T$$

The trust model represents trust values as elements of a domain forming a complete lattice (i.e. a partial ordering where every subset has a greatest lower and least upper bound). Using such a lattice of values, it is possible to associate more than just one exact value with a principal. We can associate a non-empty subset of trust values in the form of a closed interval on the lattice, whereby our trust in a principal could be any value within the associated interval, although we cannot be more precise given the information available.

As intervals are used, methods for comparing them are required. For this purpose, two orderings on the set of possible intervals for a lattice of trust values are defined to allow us to compare trustworthiness (trust ordering) and precision of the information (trust information ordering) carried in the trust interval. Intuitively, the trust ordering allows comparison of intervals by comparing the interval bounds on the original trust value lattice. The trust information ordering can be seen as allowing higher precision to be represented as a subset relationship between intervals.

The trust ordering allows the definition of a new lattice of the intervals from the original trust value lattice. Similarly, the trust information ordering allows the definition of a complete partial ordering (i.e. a partial ordering with a least element) of these intervals, with the interval consisting of the full trust value lattice as the bottom element, representing the unknown trust in an unknown principal. The narrower the interval, the fewer possible trust values we have; thus an interval consisting of exactly one value implies a very precise opinion of trust.

Local trust policies are defined for each principal, which determine which interval to associate with other known principals. These local policies support trust based on collected evidence, in the form of personal observations of past interactions and recommendations from other entities. Another important feature of these policies is the ability to delegate decisions about the trustworthiness of entities to other principals. The difference between recommendation and delegation is that we delegate to principals similar to ourselves (e.g. the bus company might delegate to some other bus companies); however, in recommendation we gather trust information from all the principals in the application domain (e.g. recommendations about users from banks, bus companies and other users). Furthermore, recommendations are weighted according to the trust in the source as a recommender, whereas delegation is merely accepted, as the delegated trust interval comes from an implicitly trusted source. Note that the entities we delegate to may further delegate to other entities, forming long delegation chains.

The local policies (denoted π) are expressed in a formal policy language, and although a brief example is given here, it is not important for an understanding of this paper. For further details, the interested reader is referred to [3]. Let us assume that we have two bus companies C_1 and C_2, and each bus company has a list of registered users (users that have weekly or monthly tickets), then an example trust policy of C_1 could be that "if user u is one of our registered users, evaluate his/her trustworthiness using the relevant stored evidence, otherwise delegate to C_2". This can be expressed in the policy language as follows:

$$\pi = (registeredTo(u) = C_1) \Rightarrow [e_0, e_1]; \lceil C_2 \rceil$$

where registeredTo() is a function that returns the bus company that a user is registered to, $[e_0, e_1]$ is the trust value interval for user u according to collected evidence and $\lceil\ \rceil$ represents delegation.

2.2. The SECURE Risk Model

After obtaining a trust interval for the principal, it can be exploited to make a decision about the request for in-

teraction. This is achieved through the process of risk assessment, outlined briefly here, with further detail available in [1].

In general, the lower the trust in an entity, the higher the risk in the interaction is perceived to be. Each application incorporating the SECURE model will identify a number of actions, each with a number of independent possible outcomes. Each of these outcomes will have a range of possible costs and benefits, which can be represented as a cost probability density function (cost-PDF). The trust interval for the specific principal determines the probabilities within each cost-PDF. The cost-PDFs for each outcome are combined to represent the likelihood of an interaction with a specific principal incurring a specific cost or benefit. This information can then be used, in conjunction with some notion of utility, to provide information for the decision making process.

2.3. The Trust Lifecycle

The trust and risk models are combined within an architecture that allows for trust lifecycle management. The trust lifecycle is the collective term for the dynamic aspects of how trust is formed, how trust evolves over time due to available information and how trust can be exploited. These aspects are important in striving for an intuitive, flexible model of trust. The architecture and issues relating to this are described in previous work by the authors [5].

For the purpose of this paper, the important feature of the architecture is the repository used to collect trust information (evidence). There are two main sources of this information, which allow us to dynamically form an opinion about another entity. Personal observations of the entity's behaviour, through recording of outcomes of interactions, are essential for the subjective evaluation of trustworthiness. Recommendations from trusted third parties provide the possibility for trust regarding unknown entities to be propagated to provide supporting evidence for decisions. The process of recommendation becomes more important in cases where we have no personal experience with the entity in question, allowing us to consider interacting with unknown entities. In the case of recommendations, the recommenders give their trust interval regarding some other principal.

Evaluating the trustworthiness of a principal includes both trust formation and evolution. Formation differs from evolution in that additional evidence might be actively sought for formation, while evolution refers to the process of changing one's estimation of trust based on stored evidence. In human terms, people tend to treat evidence from personal experience in a different manner to evidence from recommendations. For this reason it is important that these two types of evidence are treated separately.

The model introduces dynamism because trust evolves as evidence becomes available, dynamically adapting our trust in principals. In this paper we focus on the dynamics of trust information, presenting a set of principles upon which new evidence can be taken into account, to update trust intervals as additional information becomes available either through recommendation or observation. We model the effect new evidence has on the current stored trust interval through the introduction of the notion of a *pulling force*.

The following section describes the principles, on which the processing of trust information depends, which in conjunction with the model outlined above supports the dynamic evolution of trust.

3. Trust Dynamics for the SECURE Model

The focus of this section of the paper is how to estimate trustworthiness of an entity from evidence. The general principles are first examined before giving a concrete example of how these principles can be applied in the e-purse scenario.

3.1. Principles of Trust Formation and Evolution

Trust formation and evolution, although differing in some respects, both involve altering the current estimation of trustworthiness (trust interval) with respect to some new piece of evidence. As the current trust interval represents all trust evidence so far, this is akin to updating the summary of all known trust information about the principal. Initially entities have no evidence of past behaviour of unknowns to establish a base for interaction, thus allocating the full lattice of trust values as a trust interval.

Given the current trust interval based on observations (T_{obs}) and a new piece of evidence we dynamically evaluate a new trust interval that reflects T_{obs} and the effect of the piece of evidence. In all cases, the effect of new evidence is to move the current interval closer to it (or at least no further away from it), giving a new interval that better reflects both. This is the intuition behind the concept of *pulling force* exerted by a piece of evidence. We shall examine the two kinds of evidence, observations and recommendations, which have different characteristics and thus differing principles.

Observations of the outcomes of interactions are evaluated against the expected behaviour of the principal. The expected behaviour can be established as a result of the risk assessment process that led to the decision to carry out the interaction in question. The evaluated observations are referred to as experiences [7]. There exists a range of experience values that reflects the effect of an observed outcome on T_{obs}. These values are ordered and are classified into two

sets, a trust positive and a trust negative one. Every new experience exerts a *pulling force* on T_{obs}, which will move it lower or higher, depending on the class of the experience. The actual choice of experience values will be application dependent and may be complex since it is dependent on the type of cost/benefit involved.

To relate the experience value to the expected behaviour of the user, a mapping of experience values to the range of trust values is introduced to allow comparison with trust intervals. The assumption is that the loss or gain during a transaction, represented by the experience value, provides further information about the user's actual trustworthiness. As a result, the position of the new mapped experience value relative to T_{obs} determines how the *pulling force* affects the bounds as follows: In all cases the force will pull the bounds of the interval closer to the value.

Note that, if the new experience value is inside the interval, then the result of the force will be a more precise interval, while if the new experience value is outside the interval, then the result of the force will be to pull the interval towards the value potentially resulting in reduced precision. The force determines both the direction of the bounds movement as well as the level of movement, i.e. a stronger force causes greater movement, although the exact distance of movement is application specific. The movement could be done in steps, the granularity of which is determined by the individual entity. Additionally, each entity may require a certain number of experiences, an experience count, to be kept before a movement is allowed. This feature incorporates some notion of the entity's disposition, similar to the principles of slow and fast trust dynamics described in [7].

As mentioned above, recommendations may be used to provide further information particularly when the precision of T_{obs} is considered insufficient. The interval encodes some measure of precision of the estimation of trustworthiness, incorporating the notion of weighting recommendation according to the recommender's precision of estimation. This relies on the assumption that an entity behaves in the same way towards all entities (uniform behaviour assumption). This assumption affects the way recommendations are taken into account. Since the main reason for using recommendations is to improve the precision of our estimation, imprecise recommendations could be discarded as they provide little additional information. The decision, though, on this issue is left to the individual entity.

An important property of any function that combines recommendations is to provide the same results regardless of the order in which the recommendations are considered. It is possible that the recommendations may have a shelf life to enable those that are out of date to be discarded. This might be the reason for incorporating a notion of time into our model, but does not affect the requirement of order independence. In any case, it is unlikely we will be able to determine the exact time at which the recommender calculated this estimate.

Recommendations are in general treated in a similar manner to experiences, by evaluating the *pulling force* that each one of them exerts on T_{obs}. The *pulling force* of each recommendation depends on how it compares to T_{obs}:

- Conflicting recommendation (i.e. a recommended interval with no overlap with T_{obs}). Because of the uniform behaviour assumption a conflicting recommendation is taken as an indication that our estimation might be incorrect. As a result, it will lead to a reduction in the precision of T_{obs} and a movement towards the recommendation.

- Partially conflicting recommendation (i.e. a recommendation interval with partial overlap with T_{obs}). In this case the recommendation will move T_{obs} towards the direction of the overlap.

- Non-conflicting recommendation (i.e. a recommendation interval fully overlapping with T_{obs}). In this case there are two sub-cases. If the recommendation includes T_{obs}, then it does not offer any additional information and it is discarded. If T_{obs} includes the recommendation then the precision of T_{obs} is improved.

The *pulling force* of a recommendation is inversely proportionate to the precision of T_{obs} and proportionate to the precision of the recommendation itself. At the same time, because of the uniform behaviour assumption, the *pulling force* of a recommendation is also proportionate to its distance from T_{obs}, i.e. the further away the recommendation the more doubtful we are of our current estimation. The individual *pulling forces* from the recommendations are combined, resulting in an overall *pulling force*, which, in practice, will have a separate effect on each bound of T_{obs}. As with observations, it is possible to incorporate the trust dynamics disposition of an entity in the application specific formulae, to incorporate a notion of how easily influenced it is by recommendations in general.

This section has outlined the principles according to which trust intervals are formed and evolved with new experiences or recommendations. The following section gives an example of how these principles might be applied to the e-purse scenario.

3.2. Applying the Principles: the E-Purse Scenario

3.2.1 Formation and Evolution. In the e-purse example the range of possible trust values is [0, 100] reflecting the amount of e-cash that will be accepted from the requesting user, leaving the remainder of the ticket price to be paid in cash. The use of such a range of trust values simplifies the decision making process for this scenario.

For observations in the e-purse scenario, the experience value range is [-100, 100] denoting the maximum gain or loss in a transaction due to the limit on the e-purse. In this case it is convenient to represent experience values as the direct measure of cost or benefit derived from the interaction; obviously a complex function could be substituted. The trust interval is mapped to the range of experience values to represent the user's trustworthiness as expected loss or gain during a transaction involving him/her. The size of movement is determined by the individual entity's trust dynamics disposition.

Given a set of recommended trust intervals of the form $[R_i, R_i']$, on the same range of possible trust values $[D, D']$ (0 to 100 as mentioned above) and the current observation based trust interval $T_{obs} = [O, O']$, the combined *pulling force* is the sum of the individual *pulling force* from each of the recommendations. The *pulling force* on each bound of T_{obs} is different depending on the position of the recommendation. If we denote *L* as the combined *pulling force* on the lower bound and *U* as the combined *pulling force* on the upper bound, then:

$$L = \sum_{i=1}^{N} \left(\frac{(O' - O)}{(R_i' - R_i)} (R_i - O) \right),$$

$$U = \sum_{i=1}^{N} \left(\frac{(O' - O)}{(R_i' - R_i)} (R_i' - O') \right)$$

where *N* is the number of recommendations.

The sign of these combined *pulling forces* determines the direction of movement of each bound and their value could be used to determine their relative effect. So assuming the lower bound *O* moves by some amount *m*, the movement of the upper bound would be *(U/L)m*.

It is necessary to limit the value of *m*, so that the resulting bounds still determine an interval and remain within the range of trust values. As *L* and *U* are signed, there are four cases (aside from the case of forces being zero where there is no movement caused by recommendation), which will each result in different conditions (see Table 1) for the limits of m in terms of lower bound movement (C_1), upper bound movement (C_2) and convergence of bounds (C_3).

The formulae in table 1 are used to determine the maximum value *m* can take, to determine how much the bounds can move with respect to the *pulling forces*. A measure of how susceptible the deciding entity is to recommendation (α) is necessary. This α is dependant on the entity's policy and represents its disposition in terms of trust dynamics. The direction of movement is already known from *U* and *L*, and the limit value for m is used in conjunction with α to determine the movement of the lower bound (*LM*) and upper bound (*UM*):

$$LM = \alpha m, UM = \alpha \left| \frac{U}{L} \right| m$$

The discussion in this section focused on how to determine the range of possible trust values to allocate to the requesting entity. This range must then be used to determine whether or not to grant the request, which is the process of trust exploitation.

3.2.2 Exploitation. The essential problem in exploitation is to determine expected behaviour on the basis of trust intervals. In the SECURE project, this is achieved using the trust interval to determine the risk of interacting with a particular principal for a particular action, as described in [1]. The calculated risk allows entities to decide whether or not to allow the action. In the e-purse scenario, a simplified view is taken, whereby the trust value directly determines the amount of e-cash a bus company is willing to accept. For example, if the price of the requested ticket is lower than the lower bound of the trust interval then the whole transaction amount could be paid using e-cash. If the price of the ticket is higher than the upper bound of the trust interval then e-cash is not accepted and the full amount has to be paid in cash. If the price of the ticket is within the trust interval then a percentage of the transaction could be paid in e-cash requiring cash for the rest.

4. Discussion and Open Issues

We have explored the dynamic aspects of trust using as an example the interaction between a bus company and a user from the e-purse scenario. This example shows how to produce estimates of trustworthiness (intervals of trust values) from experience and how recommendations influence these estimates. It also highlights the practical ramifications of the properties of experience and recommendation based trust formation, evolution and exploitation in large-scale global enterprise systems. The use of a trust-based approach to autonomous decision making in such systems, seems to provide a more flexible and dynamic solution compared to current security approaches, and leads to interesting and challenging research issues.

One such issue is that the exact format of the trust values. Trust values are in general application dependent, thus the format of the trust values/intervals differs from one application to another. The values used in the example were fairly simple, just intervals of numbers on a numerical range. In general, this may not be the case. The trust values might not be intervals of numerical values, but can be intervals on any lattice of values. Moreover they could be multidimensional structures where each dimension refers to a particular aspect of competence or benevolence. In addition, there may be cases where the trust values differ even within the same application. For example, in the e-purse scenario, the trust intervals from the user perspective might be different from the bus company ones. In this case, functions would have

Condition	L and U both +	L and U both -	L+, U-	L-, U+
C_1	$m \leq D' - O$	$m \leq O - D$	$m \leq D' - O$	$m \leq O - D$
C_2	$m \leq (D' - O')\frac{L}{U}$	$m \leq (O' - D)\frac{L}{U}$	$m \leq (D - O')\frac{L}{U}$	$m \leq (O' - D')\frac{L}{U}$
C_3	$m \leq (O' - O)\frac{L}{L-U}$ where $L - U > 0$	$m \leq (O' - O)\frac{L}{U-L}$ where $L - U > 0$	$m \leq (O' - O)\frac{L}{L-U}$ where $L - U > 0$	No Convergence
Overall	If $L - U > 0$ then $\min(C_1, C_2, C_3)$ If $L - U < 0$ then $\min(C_1, C_2)$	If $U - L < 0$ then $\min(C_1, C_2, C_3)$ If $U - L > 0$ then $\min(C_1, C_2)$	If $L - U > 0$ then $\min(C_1, C_2, C_3)$ If $L - U < 0$ then $\min(C_1, C_2)$	$\min(C_1, C_2)$

Table 1. Conditions for m.

to be developed to allow the translation of trust values from one format to another. A similar situation arises if we want to use trust values between applications. This whole issue leads to the problem of trust value contextualization. This is a difficult problem that requires further investigation.

Another challenging issue relating to the evolution of trust values is that the outcome might not be measurable in monetary terms. This makes the definition of the experience value range quite tricky. So far, we assume that a monetary value can be assigned to any sort of outcome, which is in accordance with practice in industries such as insurance. From our point of view, the only requirements on the experience value range are that its values are ordered and that they are classified into trust positive and trust negative ones. Of course, if the experience value range does not follow a monetary value then the mapping of the experience values to the trust values might be more difficult. This is another issue that requires further investigation.

An additional issue also remains in terms of the weighting of recommendations by trust in the recommender. This is important in order to allow subjective evaluation of the recommendations. For example, the proximity of some entity's recommendations relative to T_{obs} over the set of principals, may give an indication of how reliable his/her estimation is. This allows us to adjust the *pulling force* of a recommendation according to the reliability of the recommender. Adding this sort of weighting on the example formulae for the e-purse application scenario is straightforward.

Our immediate plans for the future include an investigation of more complex application scenarios, following though the general principles outlined in this paper to determine alternative functions and to evaluate the approach further. It would also be beneficial to consider more complex interactions between more than one principal. We are also developing a simulation framework that will allow us to evaluate the effect of the specific policies and functions on enterprise collaboration.

Acknowledgements

The work in this paper is supported by SECURE (Secure Environments for Collaboration among Ubiquitous Roaming Entities, IST-2001-32486) funded by the EU FET Programme under the Global Computing Initiative.

References

[1] J. Bacon, N. Dimmock, D. Ingram, K. Moody, B. Shand, and A. Twigg. Secure deliverable 3.1: Definition of risk model, December 2002. Available on SECURE website.

[2] C. Bryce, V. Cahill, G. D. M. Serugendo, C. English, S. Farrell, E. Gray, C. Jensen, P. Nixon, J.-M. Seignuer, S. Terzis, W. Wagealla, and C. Yong. Secure deliverable 5.1: Application scenarios, September 2002. Available on SECURE website.

[3] M. Carbone and M. Nielsen. Towards a formal model for trust. Technical report, BRICS, University of Aarhus, January 2003. http://www.brics.dk/RS/03/Ref/BRICS-RS-03-Ref/BRICS-RS-03-Ref.html.

[4] C. English, P. Nixon, S.Terzis, A. McGettrick, and H. Lowe. Security models for trusting network appliances. In *Proceedings of the 5th IEEE International Workshop on Networked Appliances*, pages 39–44, October 2002.

[5] C. English, W. Wagealla, P. Nixon, S.Terzis, A. McGettrick, and H. Lowe. Trusting collaboration in global computing. In *Proceedings of the First International Conference on trust management*, volume 2692 of *LNCS*, May 2003.

[6] D. Gambetta. Can we trust trust? In D. Gambetta, editor, *Trust: Making and Breaking Cooperative Relations*, pages 213–237, Oxford, 1990. Basil Blackwell.

[7] C. Jonker and J. Treur. Formal analysis of models for the dynamics of trust based on experiences. In *Proceedings of the European Workshop on Multi-Agent Systems: Modelling Autonomous Agents in a Multi-Agent World*, pages 221–231, 1999.

[8] S. Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, University of Stirling, 1994.

[9] S. Weeks. Understanding trust management systems. In *Proceedings of IEEE Symposium on Security and Privacy*, Oakland, 2001.

[10] Secure website: http://secure.dsg.cs.tcd.ie.