Trusting Collaboration in Global Computing Systems

Colin English, Waleed Wagealla, Paddy Nixon, Sotirios Terzis,

Helen Lowe and Andrew McGettrick

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

Abstract. A significant characteristic of global 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 will function as expected. In this paper, we explore an architecture for interaction/collaboration in global computing systems. This architecture reflects the aspects of the trust lifecycle in three stages: trust formation, trust evolution and trust exploitation, forming a basis for risk assessment and interaction decisions.

1 Introduction

The future of distributed computing is likely to bring a massively networked world supporting a diverse population of hardware and software entities [1]. In this global computing environment, many of these will be mobile entities that stand to benefit from the ability to interact and collaborate in an ad-hoc manner with other (possibly unknown) entities and services to succeed in the tasks allocated to them. In such large systems, spanning multiple administrative domains, autonomous operation is an essential characteristic of entities that cannot rely on specific security infrastructures or central control to help in security related decisions. The composition and characteristics of these systems will be both highly dynamic and unpredictable. Entities will have to deal with unforeseen circumstances ranging from unexpected interactions to disconnected operation, often with incomplete information about other principals and the environment.

Freedom for collaboration between entities is an important benefit of such a dynamic environment. Collaboration can be defined as a joint interaction between two or more principals (P), which must perform one or more specific actions (A) on one of the principals' resources.

- E.g. Simple collaboration: (A, P_1, P_2)
  $\Rightarrow P_1$ (initiator) must decide if P_2 (executor or requesting entity) is authorized to carry out A.

To allow a secure collaboration to proceed, it is necessary to predict the behaviour of other principals. The current security focus on the protection of data in transit using cryptographic measures does not address the issues of undesirable behaviour of entities at either end of the communication channel. Entities must be able to make autonomous decisions about entering into a collaboration and configure themselves dynamically according to changes in expected behaviour.

The remaining sections of the paper are arranged as follows. Section 2 provides an insight into the human phenomenon of trust as a security mechanism for the type of system considered in this paper. Section 3 describes a computational model of trust, which, in conjunction with the trust information structure outlined in section 4, forms the basis for a collaboration model architecture. The collaboration architecture discussed in section 5 incorporates a trust box, which encapsulates the trust model and implements the trust information structure, to allow a principal's trustworthiness to be established based on the available information. Section 5 also describes a risk model, which utilizes the established trust to perform risk assessment upon which to base the decision to collaborate. Sections 6 and 7 consist of ongoing work and conclusions respectively.

2 Trust

In real life, humans use the mechanism of trust to cope with the inherent risks when dealing with only partial information about people and the environment. Accepting risk via this mechanism allows humans to interact on the basis of available evidence, assigning privileges or tasks to others accordingly. Similarly, computational interactions require an adequate level of trust between the principals, which is currently pre-configured by a system administrator. We assume that the administrator will not be present and measures must be in place to allow entities to form their own opinions of the trustworthiness of others. The pre-configured, coarse and static configuration of trust in traditional systems is not consistent with human intuitions of trust as a subjective and situation specific notion [2], being an individual's opinion of another entity. Trust is also dynamic, as an individual's opinion can evolve and develop based on the evidence available for subjective evaluation. Due to the complex subjective nature of trust, people have formed many different views of what exactly trust is. While this makes it difficult to form an exact definition, we assert that a model of trust can be developed in sufficient detail for use in a security model.

Within this paper distrust is not considered, as in the type of global systems considered here, it is possible for entities to change identity when they are distrusted in order to avoid negative evidence. Distrust is therefore not represented to reduce the incentive for change of identity.

It is proposed that the development of trust-risk based security architecture for collaboration incorporating a dynamic model of trust will provide devices with the abil-

ity to operate and make security-related decisions autonomously, on the basis of changing evidence. With the use of explicit representation trust, enhanced information is available on which to base decisions. It is proposed that the collaboration architecture can be used either to augment other security mechanisms or as a basis for unencrypted interactions.

2.1 Sources of Trust

There are three main sources of trust information about another entity. Personal observations of the entity's behaviour, through recording the outcome of an interaction, are essential for the subjective evaluation of trustworthiness. Recommendations from trusted third parties provide the possibility for trust regarding unknown entities to be propagated in a similar manner to the deferment of trust as seen in current trust models (e.g. [3]). Recommendations are based purely on the recommender's personal observations and as such it is possible to associate a measure of trust in the opinion of the recommender (this is not the same as trust in the recommender for other actions). The reputation of an entity can be consulted in the absence of experience or recommendation. Reputation is anonymous in the sense that it is an aggregation of trust information from different sources (including recommendations that are passed to us via intermediate parties) and as such we cannot associate a level of trust with the opinion expressed. Trust information relevant to specific action can be of more use than trust information about general activities, 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. Recommendations may be evaluated subjectively within a similar context to the recommendation evidence source. Clearly personal experience influences trust to a greater degree than recommendation, therefore it is important to weight the evidence dependant on the source of the information. The process of recommendation becomes more important in cases where we have no personal experience with the entity in question. Requesting recommendations allows us to consider interacting with unknown entities. The paper does not currently consider reputation for simplicity although this will be examined in future work by the authors.

2.2 Dynamic Aspects of Trust

The subjective nature of trust based on evaluated evidence has been introduced. The dynamic aspects of how trust is formed, how trust evolves over time due to available information and how trust can be exploited are equally important in striving for an intuitive representation of trust. These aspects of the model are collectively referred to as the trust lifecycle [4] and will provide an entity with the ability to reason about and make security-related decisions autonomously. 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.

Evaluating the trustworthiness of a principle is referred to as trust formation. An entity's trustworthiness can be synthesized from available evidence of past interactions, to be used when allocating privileges for specific tasks. 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. Recommendations may be used to establish collaboration between entities that have never met, but who trust a common third party.

The evolution process takes place, 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 trust information to maintain accuracy. Evidence from the outcome of interactions must be evaluated against the expected behaviour of the principal.

The essential problem in exploitation is to determine behaviour on the basis of trust, by determining the risk of interacting with a particular principal for a particular action. The calculated trust values enable a full assessment of risk to be carried out to allow a decision whether or not to collaborate to be made. The decision to collaborate will be determined by the security policy or the particular entity. Through the evolution process outlined above, there is feedback from this risk assessment process, demonstrating the cyclic nature of the relationship between trust and risk.

The next section introduces the basis of trust in the collaboration model that can represent trust in this dynamic manner.

3 The Trust Model

The basis for the collaboration architecture is a formal model of trust developed by Mogens Nielsen et al in [5]. Each principal in the system has a trust box, a component that processes evidence and principals to return trust values. The trust box has a state, represented by the trust information structure detailed later. Within the trust box, the model represents trust values as elements of a domain forming a complete lattice. This allows two new structures to be constructed based on intervals for all subsets of the lattice of trust values. The actual set of trust values used may be application independent. For example, if we use the simple set of integers from 0 to 100, the interval [10, 50] means that the appropriate trust value lies somewhere in the range of from 10 to 50, but we cannot be more precise given the information we have.

The first new structure is a lattice of intervals lifted from the lattice of trust values, providing an interval ordering (trust ordering) that allows the qualitative comparison of trust intervals. The second new structure is a complete partial ordering (trust information ordering) on intervals to represent quantity of trust evidence for a principal. For example, an unknown entity has the complete lattice of trust values as an interval such that its trust value could be anything. The narrower the interval, the less possible trust values we have. This serves as the basis for least fixed-point calculations to pinpoint an entity's trust value. The local trust policy of each entity determines how it computes trust. The collection of all local policies determines a global trust function, which serves as basis for least fixed-point calculations to determine trust in others.

The following section describes the structure developed to represent trust information, which in conjunction with this trust model facilitates the evaluation of dynamic trust values.

4 Trust Information Structure

The idea of a trust information structure is to provide information representing the state of the trust box in the trust based security model. The use of a layered structure (Figure 1) to store trust information provides a greater depth of information upon which to base any decision than merely storing the individual trust values relating to the entity in question. Each entity has one trust information structure, the basis of which is the store of all trust information available on every other entity with which it has been associated.

The structure has four layers: the collection of all known trust information in the first layer, the relevant separated evidence relating to personal experience and recommendations in the second layer, separate trust values (Tv_{OBS} and Tv_{REC}) derived from personal experience and recommendation in the third layer and a fourth layer containing general trust values. The goal behind these layers is to provide more fine-grained levels of trust information on entities for which we are unsure about the accuracy of the stored trust value. While the structure allows trust values to be used without further examination of available evidence in situations such as low risk assessment or high stored trust values, it may be necessary to re-evaluate trust afresh. If further information is required, the desired number of previous experiences can be examined, particularly those relevant to the currently requested action. If the final collaboration decision does not yield a positive result it is possible to seek further supportive evidence to re-evaluate a higher trust value.

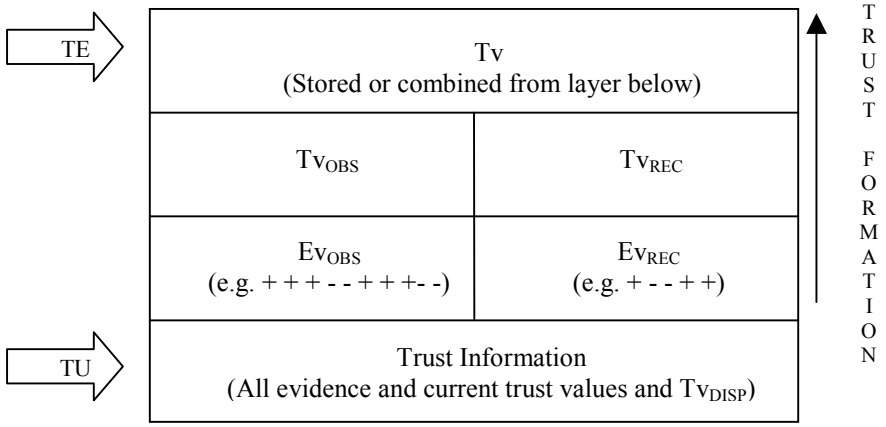

Fig. 1. The Trust Information Structure

The following points discuss these layers of trust information in more detail:

1. The base layer contains all of the trust information available to the entity, comprising of personal experience, recommendations and stored trust values. To prevent this stored information growing to an unmanageable quantity with new evidence, the information may be temporally limited to allow out of date information to be discarded. This layer also contains information on the entity's trusting disposition (Tv_{DISP}), i.e. whether or not it is generally trusting. This information can be used to initiate interaction in the absence of evidence, by selecting a node from the whole lattice of trust values. This dispositional trust may also be of use in the evaluation of evidence offered as a recommendation. After an interaction has terminated, the evaluation process will update this layer with evidence of the outcome, to keep the store of trust information up to date.
2. The information in the second layer is dynamically extracted from the trust information layer below upon request and contains evidence relevant to the requested action. In the absence of evidence related to the specific action, general evidence for the requesting entity may be extracted to provide some basis for the trust evaluation to proceed. The approach taken is that rather than have two separate structures for experiences and recommendations, one structure is used, and the recommendations are treated in a different manner to personal observations. This is necessary to ensure that the process of recommendation does not become merely delegation of trust values from other entities, and that the information passed can be evaluated subjectively, dependant on the trustworthiness of the recommender.
3. The third layer in the structure contains trust values specific to observation/experience evidence and recommendation evidence. From evaluation of each of the two sources of relevant trust information, trust values can be established to represent personal opinion/belief and the opinion/belief from other parties, as to the trustworthiness of the entity in question. Trust values will take the form of intervals on the trust lattice and can be compared both quantitatively (quantity of information supporting the values) and qualitatively (comparison of values directly) using the two orderings described in the trust model. This allows representation of how much recommendations and experience influence the final decision individually.
4. The fourth and top layer contains the final trust value upon which to base a decision in the risk model. This layer may contain the stored trust value for the requesting entity, extracted directly from the base layer of trust information on the fly or combine the separate trust values from the layer below in a suitable manner to derive a single value for trust. A consensus operator (described in the trust formation section below) combines the two trust values, taking into account the quantity of evidence that has contributed to the evaluation of each individual trust value. In any collaboration, when the stored trust value of the requester is high we can use that trust value without re-evaluating the other two lower layers, or seeking new evidence. Evaluating the trustworthiness from base trust information upon each request requires more processing to extract the necessary information relating to a particular entity than if the trust value itself were used on the fly. This overhead is acceptable only when further evidence has become available or must be sought to enable the interaction to proceed.

There are some additional points that should be considered, which are outlined here. Each piece of trust evidence in the trust information store must be linked to the action type from which the evidence originated, whether this is a personally observed outcome, or a recommendation from a trusted third party. This is merely to provide a simple notion of context, to represent the situational nature of trust by evaluating only the pieces of evidence of relevance to the currently requested action. It should also be possible to base the trust evaluation on all the pieces of evidence relating to an entity, to provide additional support for the decision if necessary.

The representation of evidence from experience and recommendations must be carefully considered. Each piece of evidence may take the form of a tuple containing the evaluated outcome or recommendation and parameters to represent the action. The exact representation of personal observations will depend upon the function used to evaluate the outcome of an interaction, such as comparison of the expected outcome with the actual outcome of the interaction (e.g. determination of cost incurred relative to the cost-PDF predicted by the Risk Model). It may be reasonable to consider any deviation from expected outcome should be deemed unsatisfactory and reduce trust because the outcome is not what was expected, regardless of whether the outcome is positive or negative. Recommendations may take the form of personal evidence offered for evaluation or trust values offered for evaluation. The use of separate trust values based on experience and recommendation allows an entity to offer the trust value based on its own experience as a recommendation, to ensure that the value passed is based purely on the personal experience of the recommender, rather than just hearsay (e.g. a police investigator, needs to know what you saw, not what someone else told you they saw). Also, if trust values based on recommendations from other entities down a chain are passed, we run the risk of double counting trust information and distorting the final trust decision. It is likely that the transfer of trust recommendations will take place through the use of certificates, to allow verification of the source. For example, it will be possible for each entity to offer the other collaborator a certificate containing its opinion of the outcome, which could be used as a recommendation in future interactions.

Evolution of trust values and update of evidence is an important part of the dynamic nature of trust, and the trust structure must facilitate both these functions. After any interaction, the store of trust evidence should be updated (TU in Figure 1) to contain the new evidence as a result of collaboration evaluation. The evolution function (TE in Figure 1) should provide the functionality to alter the stored trust value without the excess overhead of full re-evaluation. In situations where trust must be re-evaluated afresh due to the availability of new evidence, the trust formation process can be re-invoked. The following section will detail the collaboration model architecture, which uses the trust information structure and trust model outlined above to model the lifecycle of collaboration.

5 Collaboration Architecture

The collaboration architecture outlined in Figure 2 is an expanded version of work by David Ingram et al in [6], with the addition of the notion of the trust lifecycle to pro-

vide dynamic secure collaboration support. This architecture shows how we model these dynamic aspects of trust, in terms of the available trust evidence both from personal observations and recommendations from trusted entities. The architecture outlines how trust is exploited in the process of risk assessment [6], to allow security decisions to be made on the basis of probable cost of the outcome of an action. The lifecycle of the collaboration follows a series of steps described in the following subsections. These are entity recognition, trust formation, risk assessment, collaboration monitoring and collaboration evaluation. These steps incorporate the functionality of the trust lifecycle of formation, evolution and exploitation, through the use of a trust box, which encapsulates the trust model and the trust information structure defined above.

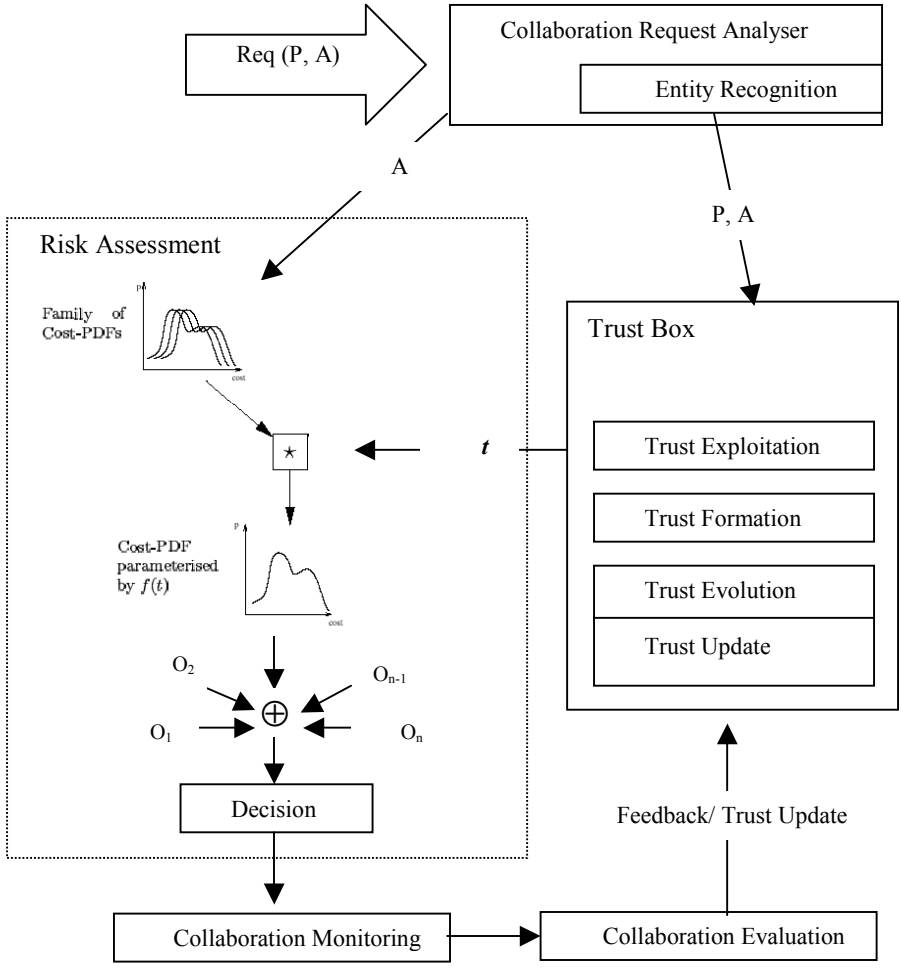

Fig. 2. The Collaboration Lifecycle (an extension of figure 1.1 in [6])

5.1 Collaboration Request Analyser

Upon receiving a collaboration request, we must analyze the contents of the request to determine whether we have the necessary resources to allow the action to take place. If the resources are not available, no further processing of the request is carried out, and a message is sent to notify the requesting entity. This message may also contain information recommending another entity, which can fulfill the request. Another important feature of the request analyzer is the process of entity recognition described below.

5.2 Recognition Mechanism

It may be impossible to establish the identity of unknown entities via intersecting certificate hierarchies, when entities roaming between administrative domains may be disconnected from their home network. Even when authenticated identity can be established in this manner (e.g. PKI [8]), as in most security mechanisms on the Internet, it conveys no a priori information about the likely behaviour of an entity. Work by J. M. Seigneur et al in [7] is relied upon to provide entity recognition in the collaboration architecture. It is proposed that all entities be assumed virtually anonymous, placing importance on recognition of entities rather than identity. In this way, the necessity for prior configuration of collaborative entities is removed, allowing unforeseen interactions to take place as the need arises. Recognition based on previous experiences allows the relevant evidence to be linked to the relevant entity. Auto configuration and dynamic enrolment measures must therefore be in place to remove the reliance on centralised certification authorities and allow the formation of an initial level of trust (dependant on an entity's trusting disposition) when entities meet for the first time, allowing enrolment in an initial low risk collaboration. This will provide the necessary evidence for future recognition.

A number of entity recognition mechanisms may be available, each trusted to different degree. One of these mechanisms must be selected to establish the level of trust in the recognition infrastructure. End-to-end trust in a particular collaboration then combines both trust in the underlying recognition infrastructure and trust between the principals determined by the trust model. If the final trust evaluation is trust insufficient for collaboration to take place, it is possible to look for a more trustworthy recognition mechanism to increase the end-to-end trust.

5.3 Trust Formation

From the set of all available trust evidence, we extract the set of experiences that represent personal observations and the set of experiences that represent recommendations. Upon receiving a collaboration request, we can dynamically filter the available trust evidence to retain only that relevant to the requested action. If there is no evidence for a principal regarding the specific action, it is possible use available trust evidence from other actions with the principal. This can be seen in the trust information structure section as the second layer of the diagram. Recommendation evidence is

treated separately from evidence based on personal experience as the latter has a greater influence on the trust value.

If the stored trust value (Tv) is high enough an entity may decide to use that value without performing any further investigation of the evidence. An extreme case of this is that absolute trust may lead to collaboration without assessment of the risks involved, and absolute distrust might lead to automatic rejection.

If no evidence is available for an entity from experience or recommendation we must establish an initial trust value to encourage low risk collaborations. This collaboration will provide further evidence upon which to base future trust formation. There are several schemes available for determining the initial trust value. These include:

- Selecting the minimum trust value for any entity in the trust information structure, or
- Exploit the trusting disposition of the entity to select a suitable interval from the trust model.

Trust Formation Function (TFF): For evidence relating to personal observations, there exists a set of possible experience values Ex, which consists of two subsets, Ex_{NEG} and Ex_{POS} to represent negative and positive outcomes of interactions respectively. Each element $ev \in Ev_{OBS}$ takes a value from the set Ex. To assist in the evaluation of a trust value based purely on the set of observation evidence (Ev_{OBS}) we will use the TFF.

$$TFF = (\#\{ev_i \mid ev_i \geq 0\} - \#\{ev_i \mid ev_i \leq 0\}) / i$$

The TFF gives a value to indicate the strength of positive evidence relative to negative evidence to support collaboration. This value (TFF), in conjunction with the total quantity of observation evidence ($\#Ev_{OBS}$), allows us to determine an interval on the lattice of trust value intervals ordered by the trust information ordering (Tv_{OBS}). $\#Ev_{OBS}$ determines the width of the interval, while the TFF allows the pinpointing of the exact interval on the lattice, dependant on how positive or negative it is.

Recommendation Trust Operator (RTO, ⊕): This operator is used to combine all of the evidence obtained via recommendation. Recommendations will only be considered based on first hand experiences from a trusted entity, to avoid the possibility of double counting trust evidence from entities further down the recommendation chain. This is referred to as recommendation independence, avoiding second hand recommendations. When we have insufficient evidence from personal observations or when the evidence we have is not relevant to the collaboration request, we may seek the relevant additional recommendations to encourage collaboration. The RTO seeks consensus between the Tv_{OBS} values from trusted entities, collected as elements of a set of recommendations Ev_{REC}. Assuming that we have two recommendations for principal $X \in P$, R_X^A and R_X^B (where $R_X^A = Tv_{OBS}$ for X, from principal $A \in P$), then the RTO reaches a consensus for Tv_{REC} as follows:

$$Tv_{REC} = R_X^A \oplus R_X^B$$

Trust Consensus Operator (TCO): The two trust values, Tv_{OBS} and Tv_{REC} are combined according to the TCO, in a manner inspired by Audun Josang's work on combining beliefs [9]. The use of Josang's consensus operator assumes the consistency of evidence underlying the opinions; therefore we make a similar assumption, that the requesting entity behaves in a uniform manner when interacting with all other principals. This consensus will strike the relevant balance between trust from experience and trust from recommendation and may also give more weight to the narrowest interval (which by definition must have been determined from more evidence). It is only necessary to use evidence in the form of recommendations if there is not enough evidence from personal observations. Thus, if Tv_{OBS} is a very wide interval, we will consider Tv_{REC} in order to narrow the interval to obtain a more accurate final Tv.

5.4 Trust Exploitation for Risk Assessment

Having established the relevant trust value for the requesting entity, this is passed to the risk model [6], in order to determine whether the risks are acceptable to enable collaboration to proceed. The evaluation of risk involves a combination of the probabilities and costs of the possible outcomes of action. An assumption is that all possible outcomes of an interaction are known and that cost or benefit associated with each can be determined. The range of possible costs for each outcome can be expressed as a cost-probability density function. For each possible outcome, the trust value is used to select one from a family of cost-pdfs, to represent possible costs or benefits, should this outcome occur. The appropriate cost-pdfs for all possible outcomes are combined and analysed according to security policy, to facilitate a decision on accepting the collaboration. The answer set can contain more than one response if necessary, or contain a moderator to express low confidence in the response.

5.5 Collaboration Monitoring

The goal of this stage is to monitor the progress of the individual actions of which the collaboration is composed. Defining policy for an interaction does not guarantee the secure execution of that interaction; therefore monitoring plays a crucial role. This is essential to ensure that the interaction is progressing towards the desired outcome rather than towards a negative outcome predicted during the risk assessment for that interaction. Moreover, we can measure the state and behaviour of principals during the interaction process, in order to terminate prematurely when a security infringement or some other form of incorrect behaviour occurs, resulting in a drastic reduction in trust. This protects resources immediately rather than waiting for the interaction to complete and then modifying trust levels. While modifying the trust level is important, it is also important to ensure that an action can be terminated without further damage being permitted. Monitoring offers further opportunity for the exploitation of trust values. If trust in a principal is very high, it may be your policy not to monitor the interaction to reduce processing overheads. The monitoring process is of increased importance for interactions established with unknown entities.

5.6 Collaboration Evaluation

After the interaction is finished, the outcome will be recorded and an evaluation will be carried out. The outcomes of the interaction will be evaluated with respect to the range of outcomes established during risk assessment, recording any deviation from the expected outcomes. Based on this evaluation, each action should be classified as a positive or negative experience according to the overall outcome. This information is recorded for each action, but may be grouped together in terms of the overall collaboration within which the action took place. This evaluation of the experience can be fed back into the trust lifecycle and used to evolve the trust value for the entity in question. Our work in this area contemplates a similar approach to the work of Catholijn Jonker and Jan Treur [10], but differs in that update affects only the trust information, not the trust value. The view of evolution here also differs from [10] concerning the narrowing and pinpointing of intervals on the trust information ordering. The evaluation of an interaction as a negative (Ex_{NEG}) or positive (Ex_{POS}) experience will update the stored trust evidence layer of the trust information structure for possible re-evaluation in future interactions.

Trust Update (tu). After each evaluation the outcome, in the form of experience evidence must be added to the store of trust information in the base layer of the trust information structure. The store contains the set of all trust evidence, Ev, to which we add the latest experience, an element (ex) from the set of all possible experiences, Ex, to produce an updated set of trust evidence, Ev′. This can be represented as follows:

$$\begin{aligned} tu &: Ev \times Ex \rightarrow Ev' \\ tu\ (Ev, ex) &= Ev' \end{aligned}$$

By updating the trust evidence in this manner, it is possible to re-evaluate trust values based on the most recent evidence, by following the procedure of trust formation again.

Trust Evolution (te). It is also important to allow the evolution of trust values over time to take place on the fly without incurring the processing overhead of trust formation. For this reason, we define the Trust Evolution function, which takes a new piece of evidence from experience (ex) and modifies the stored trust value (Tv) directly, producing a new trust value (Tv′). The function can be represented as follows:

$$\begin{aligned} te &: Ex \times Tv \rightarrow Tv' \\ te\ (ex, Tv) &= Tv' \end{aligned}$$

To recap, this collaboration architecture will facilitate the establishment of secure interactions between autonomous entities, showing how it is possible to base the decision to collaborate with another entity on evaluation of trustworthiness. The following section will introduce some of the open issues for the model and examples of ongoing work.

6 Open Issues and Ongoing Work

The work in this paper is still in progress and as such there remains many open issues to be addressed. Examples of such issues include the examination of the notion of reputation and how to represent this. It clearly should convey less reliable information than the other two sources of trust evidence highlighted in the paper. The idea of second hand recommendations based on the Tv_{REC} of a recommender may constitute one view reputation, but the aforementioned issues of double counting of trust evidence must be addressed.

The process of monitoring requires some notion of how to model a principal's behaviour over the range of possible incorrect and correct behaviours with respect to the expected outcomes of an action.

Work is in progress to examine further the use of a set of policies, which affect the manner in which a decision is taken, dependant on existing trust values. The dynamic selection from a set of policies affords greater flexibility in decisions than hard-coded behaviours. For example, a family of evolution and update functions may exist, parameterized by policy, further increasing the flexibility of the system.

A scenario to explore an example implementation of the collaboration lifecycle is currently in development. In section 3.4 of [11], a smart space scenario is outlined, a context-aware distributed system that gathers context information about individuals. In a smart university campus or department, smart applications allow the tracking of student and staff activities through effective use of the context information. For example, staff and students can use the system via PDAs or mobile phones, to check the availability and location of colleagues for a meeting. An issue of growing concern in these systems is security and privacy. It is crucial to provide secure access to context information in order to prevent its misuse and breach of users' privacy. This challenge motivates the consideration of smart spaces as an application scenario for the application of trust-based security mechanisms. This scenario has characteristics such as different methods of information sharing and a large number of possible principals, which will be important in addressing aspects of complicated collaborations. An example interaction that may occur in this scenario, involves a student (P_1) wishing to access the supervisor's calendar (P_2) in order to book an appointment (A), which involves a variety of risks to security and privacy. More complex collaborations composed of interdependent interactions may impact upon the functions outlined in the paper, and will be the focus of future work. The scenario will also provide useful information on the practicality of such complex lifecycle processes in the context of smart environments using small devices with limited processing capability.

Also in development is a simulation framework, where entities are represented by agents, for the investigation of trust and collaboration 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 suffer 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 use of a dropping window of evidence to limit the trust evidence considered in trust formation, the use of time limited evidence to represent out of date information and the issues involved in more complex interactions with multiple principals performing different actions, which may rely upon one another. 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 propagating trust information will become clearer through these investigations.

7 Conclusion

The new paradigm of global computing requires a new methodology for tackling the problem of security of interaction. Conventional hard coded security mechanisms lack the flexibility required for use in such systems, where only incomplete information is available on which to base security decisions. A more flexible mechanism is the application of trust based security models to cope with the risk inherent in interactions in this environment. Current security mechanisms with pre-configured representation of trust fail to capture the notion and its relation to risk in a manner suited to systems with no form of central control. This paper proposes an architecture with the characteristics necessary to provide a basis for reasoning about trust in security related decisions for these systems. Although it is clear there are 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 FET project **SECURE:** Secure Environments for Collaboration among Ubiquitous Roaming Entities (IST-2001-32486) funded by the Global Computing Initiative: [http://secure.dsg.cs.tcd.ie]. The scenario ideas originated with 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] S. Marsh: “Formalising Trust as a Computational Concept”. Ph.D. Thesis, University of Stirling, 1994.

[3] A. Abdul-Rahman: "The PGP Trust Model" Department of Computer Science, University College London, 1996.

[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, pp 39-44 October 2002.

[5] M. Carbone, O. Danvy, I. Damgaard, K. Krukow, A. Møller, J. B. Nielsen, M. Nielsen: "SECURE Deliverable 1.1: A Model For Trust", December 2002.

[6] J. Bacon, N. Dimmock, D. Ingram, K. Moody, B. Shand, A. Twigg: "SECURE Deliverable 3.1: Definition of Risk Model", December 2002

[7] J. M. Seigneur, S. Farrell, C. Jensen, E. Gray and C. Yong: "End-to-end trust in pervasive computing starts with recognition". To appear in the Proceedings of the First International Conference on Security in Pervasive Computing, 2003.

[8] 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.

[9] A. Jøsang: "The Consensus Operator for Combining Beliefs" Artificial Intelligence Journal, 142(1-2), p.157-170, Oct. 2002.

[10] 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.

[11] C. Bryce, V. Cahill, G. Di Marzo Serugendo, C. English, S. Farrell, E. Gray, C. Jensen, P. Nixon, J-M. Seignuer, S. Terzis, W. Wagealla, C. Yong: "SECURE Deliverable 5.1: Application Scenarios", September 2002.