

Making multi-dimensional trust decisions on inter-enterprise collaborations

Sini Ruohomaa, Lea Kutvonen

Department of Computer Science, University of Helsinki, Finland

Email: {lea.kutvonen,sini.ruohomaa}@cs.helsinki.fi

Abstract—Enterprise computing is moving towards more open, collaborative systems. Joining a business network must be made efficient, despite the technical and semantic interoperability challenges involved in connecting different information and communication systems. Trust or lack thereof forms a pragmatic challenge: partners must continuously evaluate whether they trust each other enough to collaborate in the face of risk. Supporting technology is needed for making trust-based decisions on routine business transactions and observing the business peers for malicious or incorrect behaviour on interactions. We present a trust model for automating routine decision-making which considers both risk probabilities and tolerance valuations in the enterprise, and is dynamically updated based on new experience gathered both locally and from third parties.

I. INTRODUCTION

Enterprise computing is currently moving towards more open, collaborative systems, business networks [1]–[6].

Collaboration allows organizations to focus their resources on a few key fields of expertise, while continuing to provide broader services for customers. It also enables small and medium enterprises to compete in fields dominated by large corporations by joining together to gain more influence than they would have separately. The enterprises maintain their independence during the collaboration, and make local decisions based on the enterprise policy.

The value added by a collaborative business network must be balanced with its costs. It therefore becomes essential to make joining a business network efficient. Connecting to new information systems should be straightforward, and locating new partners should be made automatic when the business need is defined. Trust management between partners should be supported by automated trust decisions on routine business transactions, and observing business peers for malicious or incorrect behaviour on interactions. Contract negotiations should be based on machine-interpretable templates and routine negotiations largely automated.

There are technical and social challenges in the way of this development. Information systems that should be connected are incompatible both technically and semantically, and system integration is expensive and time consuming. Trust between new partners cannot be formed the same way as before when the entire process of setting up a collaboration is accelerated from several months, even years of negotiations to a few days or less. The actors are autonomous, and do not all use compatible information systems. The environment is also fully

distributed; there are no trusted third parties that would give consultation on whom to trust, or solve all contractual disputes.

A new kind of middleware is needed to help collaboration management. The tasks of the middleware include partner selection and negotiation, interoperability tests for technical and business aspects of services, collaboration lifecycle management with partner changes, and breach management [7], [8]. In developing this kind of middleware, a new approach to managing trust relationships between partners is required. The trust management system has two major tasks [9]: first, it should act as a guard for the service application, applying trust decisions to protect the enterprise from taking too high risks. Second, it should upkeep reputation information on other peers, to allow the system to adjust to how the peers have behaved in the past. Trust is required to balance risks, as an effective collaboration always contains an element of dependence and vulnerability. This paper discusses calculative trust that is used for making dynamic, situational trust decisions that balance estimated risks and the available incentives to trust. If the risk is found intolerable, either the incentives must be increased or the risks limited.

The paper presents our trust management system for inter-enterprise collaboration, TuBE (Trust Based on Evidence) [10]. The focus of the paper is in building the risk estimation for the trust decision. Section 2 presents the problem environment and outlines our solution. Section 3 presents the trust decision system. Section 4 presents related work.

II. TRUST DECISIONS ON INTER-ENTERPRISE COLLABORATION

Inter-enterprise collaboration builds larger or more complex services from pieces provided by autonomous organizations. Collaboration provides strategical benefits for agile enterprises, while it also poses new challenges for automating inter-enterprise business process management and supporting platforms. Collaborative business networks connect the smaller services together based on a business network model, which defines the roles and interactions between the smaller services. They are established dynamically in response to a certain business scenario or opportunity. The participants dynamically negotiate an electronic contract to govern the collaboration. Together the participants provide a composite service, where each peer's business service fulfils a given role in the network.

The Pilarcos approach to inter-enterprise collaboration is federated: all participants retain their autonomy within the

loosely-coupled business network, and make their own decisions on whether to join or continue in it. Partners are located in the open service market comprised of service providers' offers. Nothing can be said about the trustworthiness of a service provider based on their offers. For the offers to form a breeding environment, identity management is required: to be able to enter into binding contracts, enterprises must be able to identify each other. The details of the required identification process are determined by legislation, and will typically require an out-of-band registration with local authorities to form a legal entity. Self-generated identities natural to many peer-to-peer networks are too ambiguous for enterprise collaboration.

For collaboration management, the Pilarcos middleware provides several pervasive services. Tools and repositories support developing and publishing new business network models, and defining business services that match the needs of the roles defined in the models [11]. Service offer repositories allow enterprises to offer their service to the open market in a way that enables automated matching and interoperability testing [7]. Service matching and interoperability testing is done by a populator service based on information in the service offers [12]. The populator produces match proposals, and can be provided as a service by a third party, as it only relies on public information. Multilateral negotiations determine which proposal is accepted and whether some parameters of the service need to be further adjusted [8].

The TuBE trust management system forms a part of this middleware. It determines whether an acceptable level of mutual trust is present between the partners to support the collaboration, and upkeeps information needed for the analysis [10].

In TuBE and Pilarcos, trust is defined as the extent to which one party is willing to participate in a given action with a given partner in a given situation, considering the risks and incentives involved. A trust decision evaluates whether the willingness is sufficient; trust decisions are made on whether to join an inter-enterprise collaboration to begin with, and repeated during it whenever risk-relevant commitments are made. The TuBE system automates routine trust decisions, comparing the estimated risk and the tolerance for it.

Trust decisions are subjective evaluations made by the trustor, i.e. trusting party, targeting a given trustee and a given action. The trustee is a potential or existing partner enterprise or its representative. The action is represented as a group of messages exchanged in the context of a commitment of resources on the trustor's part: for example, reserving resources to complete a task or to deliver a set of goods.

Trust decisions are made to protect guarded assets from possible negative effects caused by defecting partners, disadvantageous contract clauses or other threatening elements. The guarded assets are resources that become vulnerable due to the uncertainty inherent in relying on a business partner. The potential for negative effect, the risks, can in some cases be mitigated through precautions, but the remaining uncertainty must simply be tolerated in order to collaborate.

We have chosen not to merge all assets into a single

resource, such as money, for the purpose of increased clarity: it is difficult to convert effects such as reputation loss or gain, human injury or a security breach to monetary terms, yet they are clearly important outcomes to consider in a trust decision.

We define a set of four standard assets shared between organizations: monetary, reputation, control and fulfilment. The monetary asset represents money and other artifacts in the enterprise that have a well-defined monetary value. The reputation asset represents the trustor's good reputation. The control asset is a joint representation for the trustor's security, privacy and general desired levels of self-protection. The fulfilment asset represents the fulfilment of the trustor's expectations of the trustee's participation in the action, such as the quality of the service the trustee provides or its efficiency in fulfilling its end of the agreement.

A trust decision is based on a comparison of the uncontrollable risks that allowing the action would cause, and the willingness to accept them—risk tolerance. Both are built by evaluating a lower-level factor, reputation and importance, respectively. The risk evaluation is expressed as probabilities of different outcomes, estimating how the partner will behave in the future. This estimate is based on earlier experience on the outcomes of earlier collaboration with the trustee. First-hand experiences and experiences shared by third parties form the trustee's reputation, which is the trustor's subjective perception of how trustworthy the trustee is, based on currently available information. Risk tolerance builds on the business importance of allowing the action: different kinds of benefits may be realized by a positive decision alone, such as building a partnership, helping the inter-enterprise collaboration towards realizing its goals, and not triggering compensation clauses in the contract.

The four factors of risk, risk tolerance, reputation and importance represent the valuations, expectations and policies across the enterprise and are not modified unless those change. However, the business world fits poorly into rigid frameworks: the risk evaluation of a particular partner should reflect whether they are forced to have an insurance which makes a considerable monetary loss impossible; another partner is a valued contact, and therefore is particularly important to collaborate with even in the face of some additional risk. A third partner's reputation may suddenly plummet due to what is suspected to be a misinterpretation.

Both the environment and the internal policies of an enterprise are in constant fluctuation. A trust management system for inter-enterprise collaborations must be prepared to handle frequent, often temporary adjustments to what is otherwise a clear set of policies and valuations. In order to retain clarity in modelling while catering for the messy reality, we have opted to add a fifth element to trust decisions: context filters.

Context filters are not a factor by themselves. Instead, they adjust the other four factors to temporary changes and special cases. When the constraints for triggering a context filter are met, a modification rule changes the values of the factor before passing it onwards in the decision-making.

III. MAKING MULTI-DIMENSIONAL TRUST DECISIONS

When a decision is taken to enter a collaboration or take part in an interaction, it is based on a set of measurable aspects that represent groups of assets that the enterprise wishes to protect. Trust decisions are formed from two tracks: on one track, calculating a risk estimate of the situation, and on the other, building a measure for risk tolerance. The risk estimate is built on a view of the reputation of the trustee, and tolerance builds on the importance of allowing the guarded action. Reputation represents earlier real-world experiences, and importance the business value of allowing the action; both are expressed through the effects on assets. We will first describe the assets, then present the risk and reputation and the risk tolerance and importance factor pairs. Lastly, we outline how the decision-making process is carried out in the middleware.

A. Assets guarded by trust decisions

A is the set of guarded assets, represented by integers $0..|A| - 1$. Assuming that the standard assets of monetary, reputation, control and fulfilment are used, $|A| = 4$.

The monetary asset forms the basis of decision-making in many situations, and it is also the simplest to measure. Even so, the monetary value of an item or a concrete service is not always straightforward to determine. The trustor can form a value for a target by deciding how much it would be ready to pay for it. For trust decisions, this measure is more valuable than the actual current market value of such a target, and simpler to determine. If, for example, the trustor is unaware or uninterested in a particular feature of a device it wishes to buy, it is irrelevant to the decision whether the device carries such a feature or not, or whether the feature works, even though the selling price of the device might depend greatly on it.

The reputation asset encompasses both the enterprise's reputation rating in any reputation systems, as well as the more abstract notion of its public relations, appearance in the media, and the attitudes of its partners and customers towards it. Since reputation is a broad phenomenon, a change for the reputation asset can never be measured as accurately as a monetary loss or gain, but it allows the enterprise to represent the risk of losing partners through a traditional drop in perceptions, even if for some reason the actual calculated reputation does not change. As partners are not solely dependent on information shared in a reputation network, other forms of reputation must be considered as well: for a well-informed decision, the enterprise should generate estimations of lost reputation for particular action outcomes, instead of simply monitoring a reputation network for changes which may not be directly connected to particular actions to begin with.

The control asset represents the general need for an enterprise to protect itself from outside influences: to maintain control over its security, privacy and other aspects of its autonomy. The security of an enterprise involves the physical safety of its people, equipment and goods, and less tangible aspects such as the continuity, availability and reliability of its services and the protection of its information and IT systems. The privacy of an enterprise encompasses its ability to control

information concerning it, beyond confidentiality alone: even if information becomes unconfidential when it is given to a partner, the partner can violate the enterprise's privacy by passing the information on without permission, or by releasing false information of its own. A reputation system is a privacy tradeoff in itself, and it can cause further privacy threats if a participant releases unfair or false experience information about its partners. Finally, the enterprise may feel its autonomy is threatened by some forms of collaboration, for example if an offered contract has severe enough compensation clauses to force local decisions.

The fulfilment asset is the one most tightly connected with a trustee. It encompasses whether the trustee does its part of what was agreed, leaves something relevant undone or does something it was not strictly expected to. Where the base for the monetary asset is the wealth of the organization, the base for fulfilment is the general trend of respected agreements, which reflects on the success of the organization. The asset has high value in evaluating the predictability of the trustee over a range of highly different actions. Like more traditional assets, it can be protected: leaving things undone can be avoided by putting effort into negotiating a more tightly binding contract, or by selecting trustees more carefully, with weight given on their earlier performance. No other asset can fully capture reliability, quality of service or competence: for example, if large deals with a particular trustee always end in less profit than was expected, but do not result in losses per se, the trustee has a spotless reputation money-wise, even if it is not as attractive a partner as a more reliable enterprise.

The first three assets represent the concrete effects to the business, which should be an important factor in a trust decision. The fulfilment asset is a new, more explicit formulation of the broadly-accepted main aspect of trustworthiness: doing what was agreed upon. The eBay reputation system [13], as well as any reputation system based on a simple model of cooperation or defection, measures only the trustor's subjective fulfilment of the deal. The fact that this outcome measure was the first one to be adopted to reputation systems speaks for its central nature. The fulfilment asset can help alleviate the common problem of trustees with the highest reputation being overwhelmed by requests [14]: once a provider either turns down proposals or responds to them sluggishly due to a high load, they will cause worse experiences related to the fulfilment asset, and a trust decision system can be adjusted to react to this kind of development quickly.

Outcomes of actions are represented through their effect on assets. J represents the set of possible outcomes as integers $0..5$: 0 for unknown effect, 1 for major negative effect, 2 for minor negative effect, 3 for no effect, 4 for minor positive effect, and 5 for major positive effect. $|J| = 6$. "No effect" differs from "unknown effect". For example, not losing or gaining money would represent a lack of effect, while an experience with a delayed payment on its way might be included as an unknown outcome in decision-making. The trustee defines how a series of events reflects on its assets.

The limits for minor and considerable change are subjective.

For the monetary asset, it is clear that a small company will focus on smaller amounts of money than a large one, and therefore the exact numbers to divide all real values to these five groups will be decided within each organization. This subjectivity is repeated over the other assets as well, albeit not quite as pointedly; differences between actors, their values and expectations complicate the semantics of all experience sharing, which must be taken into consideration in reputation systems design and when using information provided by them. For the control asset, it may be that none of the positive outcomes are ever used when storing experience, but they are retained for symmetry.

The limitation to a standard asset set has the benefit of increased interoperability: experiences based on these assets can be used across systems with less information lost due to unmatched or vaguely defined assets. An asset may have a clear and valuable role in one enterprise, but not be understood the same way in another.

B. Risk

The risk involved in a positive trust decision depends on the nature of the action and a prediction of how the trustee is likely to behave. A risk evaluation is an attempt to partially predict the costs and benefits resulting from different possible outcomes and their probabilities. The probabilities are based on the trustee's reputation, and are updated as more experience is gained on the trustee. The possible costs and benefits depend on the type of action and some of its parameters: for example, when arranging to buy a book costing ten euros, it is possible to lose the money if the book is not delivered at all, or the result is otherwise not considered worth the investment. However, losing more than the invested ten euros is unlikely. This means that even though a trustee's reputation may suggest they occasionally defect in a way that causes considerable monetary losses, it would only be possible for them to defect with a small loss within this particular action.

As risk builds on past information expressed as reputation, it must also include a measure for the quantity and quality of the information there exists about it. We measure the amount of information available, the amount of expressed uncertainty in it, and the credibility of its sources; these measures are discussed further in Section III-D.

C. Reputation

As risk estimates aim to define probabilities of different futures, they must examine the past, represented by reputation, to learn from it. This requires us to make the assumption typical for reputation systems: that trustees are sufficiently consistent in their behaviour for the past to indicate something about the future.

Reputation represents the current view of a trustor's trustworthiness formed from local experience and shared third-party experience. The reputation views building on these two very different sources are stored separately up until the moment of a trust decision.

Both external and local reputation views follow the same format. They store the number of experiences in different outcome types for each asset. A reputation view also stores the number of cases where no outcome type could be determined for the asset, which is a measure of the quality of the available information. To accommodate the credibility analysis of external information sources, reputation views also store the overall credibility value for the reputation view, a real number between 0 to 1 used in adjusting the overall uncertainty estimation.

The reputation of a trustee is independent of the action being considered. Experiences are expressed based on their effects, not on what kind of action caused them. This allows us to make decisions based on specific risks, while avoiding the problem of having information become too sparse. Competence in performing particular types of actions [15] is also not a part of our trust model: if an enterprise publishes a service offer that describes the interface to access its service and then repeatedly fails to provide it, it has either falsely claimed to be competent, or has refused to offer it. The end effect to our assets is not very different either way.

Local reputation consolidates single experiences gathered by the trustor. These, in turn, are formed by analyzing the output of the Pilarcos monitors [8]. The monitors are not aware of the particular assets being protected in the system, but only detect noteworthy events in the system connected to a business process; the trust management system needs a policy for translating the events into outcomes. An example event could be "product order", with the price or value of the product as a parameter. Further parameters would be the trustee's identifier and the identifier for the action whose business process the event connects to.

External reputation is based on information shared by third parties within a set of reputation networks. The trustor has representative agents to participate in each network and to pass information between it and the local reputation system. The agents report a trustee's reputation to the network in the native format of their represented reputation system, and feed back reputation information from the network into the local system, which is then transformed to experiences of the TuBE system format. The information is given a credibility rating based on a combination of the agent's view of its credibility, and the general credibility of the external reputation system. For example, a long-standing credit rating company might have a high credibility rating, while an eBay-style open reputation system would tend to have low credibility, depicting how vulnerable the system is to reputation attacks. If an external reputation system has a good representation of the source's own uncertainty and supports the evaluation of different sources' credibility well, the system's own credibility value will be higher.

The views of multiple actors must be combined to produce a unified reputation view. This can be done by several methods, such as preferring certain sources over others when they have a particular part of the information available, or using some kind of a combination formula such as weighted sums to produce the result.




Fig. 1. A local reputation view.

While a theoretical model of reputation can accommodate for infinite amounts of experience information, practical models must adapt to limited storage space and computational power: either information must be compressed, or old experiences must be purged after a while. We opt to compress experience items into outcome counters once it is relatively certain that there will be no further monitor information to change the experience; during the uncertain period, experiences can be kept open to modifications and consolidated on the fly for trust decisions. Figure 1 depicts a local reputation view.

We trade off timing information for space. The main value of discounting old information in favour of new [16, pp. 639]. is realized from reacting to changes in behaviour, i.e. by not allowing a good history to outweigh recent transgressions. For this goal, we find that time is not the optimal measure for determining the weight or value of a unit of experience, but rather whether the experience brings new information; something we did not already know.

We capture changes in behaviour by dividing reputation information into epochs, each a new leaf for experience gathering. While the latest turn of behaviour is most interesting, it is also typical that there is very little experience on it; hence information from older epochs must also be included. The weight given to the current epoch determines the speed in which the system reacts to changes in behaviour. The number of epochs also measures the consistency of the trustee.

Some central questions related to the implementation of reputation epochs remain outside the scope of this paper, such as what constitutes a strong enough change to trigger an epoch change, and how credible the information related to the change needs to be. These are issues of calibration, and must be explored further by simulating different models for reaction sensitivity.

D. Building a risk estimate from experience

The risk R of an action contains $|A|$ vectors \mathbf{r}_a , one for each asset. We omit the context adjustments here in favour of a clearer description of the central processes.

$$R = (\mathbf{r}_0, \mathbf{r}_1, \dots, \mathbf{r}_{|A|-1}), \text{ where } \mathbf{r}_a = (\mathbf{p}_a, |E|, c, q_a)$$

The vectors store the probabilities of different known outcomes for each asset, and three different measures of the amount and quality of the information used to produce the evaluation. The risk evaluation is specific to a given trustor, trustee and action; we omit references to these three parameters in the formalism for readability.

The first term \mathbf{p}_a is a vector that represents the probabilities of different outcomes: $\mathbf{p}_a = (p_{a,1}, \dots, p_{a,|J|-1})$, where each $p_{a,j}$ is the calculated probability of outcome j happening for asset a ; J is the set of outcome categories $\{0, 1, 2, 3, 4, 5\}$ presented in Section III-A. The latter indexes begin from 1; the probability of an unknown outcome is not considered, but information about unknown outcomes in the past is represented through another term, q_a . We require that the probabilities add up to 1 for all assets, i.e. for all $a \in A$, $\sum_{j=1}^{|J|-1} p_{a,j} = 1$.

E represents the group of all experiences the trustor has on the trustee in its reputation view, and the number of experiences, $|E|$, measures the amount of reputation information behind the risk estimation. The third term, c , is the combined credibility of local reputation and external, third-party reputation information at the time of evaluation. The last measure, q_a , is the number of experiences in E where the outcome was unknown for asset a : the higher this value is in relation to $|E|$, the lower the certainty of the risk analysis. We return to these quantity and quality measures in transforming reputation information into a risk analysis.

The reputation U of a trustee as viewed by a given trustor consists of two halves: a local reputation U^{local} , and a subjective evaluation of third-party reputation information U^{ext} .

The structure of both halves is the same: each contains $|A|$ vectors and a credibility score c^{local} or c^{ext} in the range $[0..1]$. The credibility value c^{local} of local reputation is 1, while the third-party reputation credibility value c^{ext} is set based on the trustor's analysis of the combined credibility of the reputation system the information comes from, and the credibility of the sources providing reputation information in that network. To avoid repetition, we present the two symmetric reputation structures as generic variables that are equal for both halves, denoting this with an asterisk (*). For example, U^* represents both U^{local} and U^{ext} ; the formulae are the same for both halves, while the values are different.

Both types of reputation consist of $|A|$ vectors \mathbf{u}_a^* , one for each asset:

$$U^* = (\mathbf{u}_0^*, \mathbf{u}_1^*, \dots, \mathbf{u}_{|A|-1}^*), \text{ where } \mathbf{u}_a^* = (u_{a,0}^*, u_{a,1}^*, \dots, u_{a,|J|-1}^*)$$

The counters $u_{a,j}^*$ express the number of experiences of outcome $j \in J$ for asset a , with $j = 0$ representing an unknown effect.

The set of experiences, E^* , is an abstract group

$$E^* = \{\mathbf{e}_k^* : e_{k,a}^* = \text{outcome value } j \in J, \forall a \in A\}.$$

That is, an experience consists of the effects of one action expressed for each asset. Each experience \mathbf{e}_k^* also belongs to one epoch. The index k ranges from 0 to $|E^*| - 1$.

Given the set of experiences, the reputation counters $u_{a,j}^*$ can be expressed as the size of the subgroup of E where the experiences had outcome j for asset a , that is: $u_{a,j}^* = |E_{a,j}^*|$, where $E_{a,j}^* = \{\mathbf{e}_k^* \in E^* : e_{k,a}^* = j\}$. The value $u_{a,0}^*$ is particularly interesting, as it expresses the number of experiences with unknown values for the asset: we name it q_a^* . When compared to the total number of experiences, $|E^*|$, it provides us a measure on the quality of the information.

Reputation is transformed into risk by 1) merging the local and external reputation views together with a weighted sum, 2) scaling the experience counters representing known effects to proportions in the range $[0..1]$ and 3) recalculating a joint credibility and information content score for the result, that is, the variables c , $|E|$ and q_a that appear in the risk vectors.

The local and external reputation views are merged based on the amount of information available in either, and the credibility attached to the views. Local reputation is more credible than external, but there is usually less local information available, and both should be reflected on the weight given to local reputation. We define two functions, μ^{local} and μ^{ext} , to determine the weights for both local and external reputation values. They use the corresponding credibility value c^* , amount of experience $|E^*|$ and a vector \mathbf{q}^* of the number of experiences where the effects are unknown for different assets: $\mathbf{q}^* = (q_0^*, q_1^*, \dots, q_{|A|-1}^*) = (u_{0,0}^*, u_{1,0}^*, \dots, u_{(|A|-1),0}^*)$. The multipliers produced by the μ^* functions add up to 1; the specific declaration of the μ^* functions depends on local calibration.

$$\mu^{local}(c^{local}, |E^{local}|, \mathbf{q}^{local}) + \mu^{ext}(c^{ext}, |E^{ext}|, \mathbf{q}^{ext}) = 1.$$

To support more than one epoch, the μ^* values can be divided beyond the two groups. Merging follows the same pattern, with each $*$ -marked variable appearing separately for each epoch.

Using the μ^* functions, we merge the experiences into a temporary $\mathbf{u}^{merged} = (\mathbf{u}_0^{merged}, \mathbf{u}_1^{merged}, \dots, \mathbf{u}_{|A|-1}^{merged})$, where each vector \mathbf{u}_a^{merged} contains six combined counters $u_{a,j}^{merged}$: the weighed sum of the local and external respective counters. Note that unlike the values in \mathbf{u}_a^{local} and \mathbf{u}_a^{ext} , these merged values are no longer integers, but real numbers. For all $a \in A, j \in J$, we have:

$$u_{a,j}^{merged} = \sum_{* \in \{local, ext\}} \mu^*(c^*, |E^*|, \mathbf{q}^*) * u_{a,j}^*$$

In the second phase, we scale the experience counters except the unknowns to the range $[0..1]$ to represent probability. To achieve this, we sum the values of known effects ($j \neq 0$) and divide each value by the sum. As a result, we get the $p_{a,j}$ values mentioned earlier in the risk representation.

$$\forall j \in J \setminus \{0\} \quad p_{a,j} = \frac{u_{a,j}^{merged}}{\sum_{j=1}^{|J|} u_{a,j}^{merged}}$$

In the third phase, we calculate combined measures of the quality of information: c , $|E|$, and the vector of $|A|$ different q_a values. The combined credibility c is determined by a μ -weighted average of the local and external credibilities. It depicts the weight given to each half in the probabilities as well.

$$c = \sum_{* \in \{local, ext\}} \mu^*(c^*, |E^*|, \mathbf{q}^*) * c^*$$

To calculate the total number of experiences, $|E|$, we sum the number of local and external experiences: $|E| = |E^{local}| + |E^{ext}|$. Although it is clear that not all experiences have been given equal weight in the evaluation, this measure gives an indication of how much information there is available on the actor overall. The combined number of experiences for each asset where the effect was unknown, q_a , is gotten by adding the values of the previously calculated \mathbf{q}^* vectors, for all $a \in A$:

$$q_a = q_a^{local} + q_a^{ext} = u_{a,0}^{local} + u_{a,0}^{ext}$$

Again, not all unknowns weigh equally in the probability calculations, so we could consider a μ -weighted average here similarly to the calculation of the credibility value c . On the other hand, the true total number of unknowns is a more useful value to use together with the amount of total experience, $|E|$, as $q_a/|E|$ gives the proportion of uncertain values.

E. Risk tolerance and importance

A trustor's risk tolerance is determined by the situation calling for a trust decision, independent of the trustees' behaviour. It depends on the business importance of the action, and local policy expressing the trustor's general risk attitude, encompassing the tolerance of both certain probabilities of various outcomes, and the uncertainty in the information. Tolerance is expressed as a set of constraints for the risk evaluation; if the constraints are met, a trust decision is positive. The constraints are asset-specific, and can give upper or lower bounds either to probabilities of particular outcomes, the sum of probabilities of a set of outcomes, or the measures of uncertainty. The bounds can be absolute or relative, containing comparisons between probabilities: for example the probability of monetary gain can be required to be larger than the probability of loss.

A trustor's risk attitude determines how risk-averse or risk-seeking the trustor is. A risk-averse trustor will require that an action have high importance to balance for the risk a positive decision would cause, while a risk-seeking trustor can accept a higher risk in relation to the baseline set by the action's importance.

Building a configuration system to help a trustor express their risk attitude through these formulae is an important item of future work. The aim is to bring the level of expression for configurations as close to the business processes and the language of the decision-makers as possible, and minimizing configuration work that requires expensive consultation.

The importance factor expresses the business value of the action, and the cost of a negative trust decision. The costs and benefits do not depend on the expected behaviour of

the trustee. For example, a negative trust decision blocking an action may result in compensation clauses being activated in the contract between the trustor and trustee. The required compensation may still be small enough that blocking the action is preferable to risking that the trustee causes greater losses by defection.

Importance information covers the investment required by the action and the guaranteed return of investment, when for example a certain group of actions are considered to be so valuable that requests for them get a high priority. To a bank, for example, a cheap loan may be a strategic way to attract customers to move all their banking services to it. Importance should also capture a lack of real choice, should it occur, and more generally the perceived cost of denying service to the trustee. The valuations considered may include the interests of the surrounding business network, adjusted based on how much weight the trustor decides to place on them.

Importance is expressed in the form of assets, and gains or losses to each asset caused by approving the action. For a particular action and its parameters, its importance is defined as the effect it has for each asset.

The risk evaluation must be compared to the risk tolerance to produce a trust decision. The risk tolerance \mathbb{T} of an action, given a particular trustor and trustee, is a vector of $|A|$ functions f_a , one for each asset.

$$\mathbb{T} = (f_0, f_1, \dots, f_{|A|-1})$$

The functions represent the acceptable limits for the risk values in the risk vectors \mathbf{r}_a : they evaluate whether the values are within bounds or not.

$$\forall a \in A, f_a(\mathbf{r}_a) = \begin{cases} 1 & \text{if the values of } \mathbf{r}_a \text{ are within bounds} \\ 0 & \text{otherwise} \end{cases}$$

Risk tolerance depends on the importance of an action. The importance factor \mathbb{I} contains $|A|$ values v_a , one for each asset.

$$\mathbb{I} = (v_0, v_1, \dots, v_{|A|-1})$$

The values express the known effects a positive trust decision has on different assets: for all $a \in A$, v_a = an effect value $j \in \mathbb{J} \setminus \{0\}$. There are no unknown effects ($j = 0$) for importance: it depicts only those assumed effects and valuations in the enterprise that affect decision-making.

Both importance and risk tolerance depend on the trustor, trustee and action. Risk tolerance is evaluated based on the importance value; each trustor determines the exact evaluation function $\Phi_{\mathbb{T}}(\mathbb{I})$ that produces the risk tolerance \mathbb{T} .

Once the reputation values have been transformed into a risk evaluation, and the risk tolerance functions generated from the importance values, the actual evaluation is straightforward: the evaluation result is positive (1) if all the risk tolerance functions evaluate to 1 with the risk vectors, i.e. $f_a(\mathbf{r}_a) = 1$ for all assets $a \in A$, and negative (0) otherwise.

F. Trust decisions in the middleware

The trust management system acts as a guard between a service application and the outside world. Service requests and

responses from the service are routed through the guard, which triggers trust decisions whenever risk-relevant commitments are being made, and blocks the messages when needed.

When a decision must be made, the guard calls the trust management system to evaluate the risk and risk tolerance [10]. The information the decision is based on is pre-processed and quick to access; reputation information is constantly gathered through Pilarcos monitors [8] and processed in the background, and importance information is readily defined.

IV. RELATED WORK

The TuBE system manages trust for enterprise collaboration in the open service market. Routine trust decisions are automated, evaluating the four factors of risk, reputation, risk tolerance and importance.

The TrustCoM project is developing a framework for trust, security and contract management for dynamic inter-enterprise collaborations. The framework includes a virtual breeding environment, a yellow pages service for finding partners, a reputation management service and a risk analysis tool to guide contract forming [17]. TrustCoM limits the group of potential partners beforehand by a pre-contract, and its contract negotiations are human-driven. TuBE operates in an open service market unbound by pre-contracts; the openness sets more strict requirements for trust management. The ECOLEAD project [18] aims to build infrastructure services such as billing, tools for human collaboration and federated executing of joint workflow descriptions.

The SECURE project has produced a trust management system directed towards private people [19]. It includes a model of trust which is updated through observations, and a single-asset risk evaluation based on cost-benefit probability distribution functions. The SECURE system learns from experience and models the situation the trust decision is made in, such as risks involved. The trust model is aware of the amount of information available, which is a necessary feature for inter-organizational trust management as well. It has no concept of business value, but defines fixed decision policies.

Risk tolerance is typically only expressed indirectly, with, for example, a fixed threshold set for minimum reputation [14], or policies which require the trustee to prove their trustworthiness by presenting a given set of certificates [20]. In the latter case, the required certification can vary based on what kind of access the trustee requires; there is a clear indication for a need for action-aware risk tolerance.

Reputation systems are prevalent in electronic commerce, and have been identified as a central means to support a healthy electronic marketplace [21]. Typically, decisions and processing the credibility and relevance of the provided reputation information is left to the human user, which allows systems such as eBay [13] to operate without a formal model of risk analysis and tolerance.

Importance and business value are relatively rarely seen to factor into automated trust decisions [15]; Poblano uses importance as a factor to extend evaluations based on reputation with the trustor's local valuations [22]. Importance is bound

to the objectives of the trustor, and in simple systems where all trustors are assumed to share a single objective, modelling importance is not a central issue.

V. CONCLUSION

We have presented a model for automating routine trust decisions in an inter-enterprise collaboration context. The model identifies and addresses the business-level goal of guarding identified assets through trust decisions. It supports evaluating risks to assets based on knowledge of the trustee's past behaviour, which is encoded in reputation. The trustor's risk tolerance is determined through evaluating the known costs and benefits of a positive decision, such as building a partnership or avoiding the activation of a compensation clause in an existing contract.

The TuBE trust management system acts as a guard between enterprises that wish to collaborate to realize a business opportunity, but cannot fully trust each other. In an open, market-driven environment, the autonomy of potential partners will always cause uncertainty which must be either accepted or avoided at the cost of missing a potential opportunity. A decision on commitment must consider the incentives and risks involved; measuring or simulating the actual emotional state of any human decision makers is not a primary goal. Trust covers the gap between risks and measures in place to reduce them, but risk tolerance policies determine exactly how trusting the enterprise is willing to be.

Future work involves refining the mechanisms involving reputation epochs further, and evaluating the model through its responses to threat scenarios. The communication between the TuBE reputation system and external reputation networks will also be further refined. Standards and a basic ontology are required to ensure that reputation and trust information for inter-enterprise collaboration are well-defined.

We find that trust management between partners must be supported by a system for automating routine trust decisions. No collaboration is feasible without mutual trust between partners.

ACKNOWLEDGEMENT

This work has been performed at the Department of Computer Science at the University of Helsinki, where the Collaborative and Interoperable Computing research group builds on work done in various projects funded by the national technology development center TEKES and industrial partners.

REFERENCES

- [1] European Commission, "EC FP7 ICT Work Programme," EC, Tech. Rep., June 2007. [Online]. Available: <http://cordis.europa.eu/fp7/ict/>
- [2] M.-S. Li, R. Cabral, G. Doumeingts, and K. Popplewell, "Enterprise interoperability. research roadmap," EC, Tech. Rep., July 2006. [Online]. Available: http://cordis.europa.eu/ist/ict-ent-net/ei-roadmap_en.htm
- [3] M. P. Papazoglou, P. Traverso, S. Dustdar, F. Leymann, and B. J. Krämer, "Service-oriented computing: A research roadmap," in *Service Oriented Computing (SOC)*, ser. Dagstuhl Seminar Proceedings, no. 05462. IBFI, Schloss Dagstuhl, Germany, 2006. [Online]. Available: <http://drops.dagstuhl.de/opus/volltexte/2006/524/>
- [4] B. Fitzgerald *et al.*, "The software and services challenge," NESSI, Tech. Rep., Jan. 2006. [Online]. Available: ftp://ftp.cordis.europa.eu/pub/ist/docs/directorate_d/st-ds/fp7-report_en.pdf
- [5] F. Nachira, P. Dini, A. Nicolai, M. Louarn, and L. Léon, *Digital Business Ecosystems*. European Commission, 2007. [Online]. Available: <http://www.digital-ecosystems.org/book/de-book2007.html>
- [6] M. N. Huhns, "A research agenda for agent-based service-oriented architectures," in *Cooperative Information Agents X*, ser. Lecture Notes in Computer Science, vol. 4149, 2006, pp. 8–22.
- [7] L. Kutvonen, T. Ruokolainen, and J. Metso, "Interoperability middleware for federated business services in web-Pilarcos," *International Journal of Enterprise Information Systems, Special issue on Interoperability of Enterprise Systems and Applications*, vol. 3, no. 1, pp. 1–21, Jan. 2007.
- [8] L. Kutvonen, J. Metso, and T. Ruokolainen, "Inter-enterprise collaboration management in dynamic business networks," in *On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE: OTM Confederated International Conferences, CoopIS, DOA, and ODBASE*, ser. Lecture Notes in Computer Science, vol. 3760. Agia Napa, Cyprus: Springer-Verlag, Nov. 2005, pp. 593–611. [Online]. Available: http://dx.doi.org/10.1007/11575771_37
- [9] S. Ruohomaa and L. Kutvonen, "Trust management survey," in *Proceedings of the iTrust 3rd International Conference on Trust Management, 23–26, May, 2005, Rocquencourt, France*, ser. Lecture Notes in Computer Science, vol. 3477. Springer-Verlag, May 2005, pp. 77–92. [Online]. Available: http://dx.doi.org/10.1007/11429760_6
- [10] S. Ruohomaa, L. Viljanen, and L. Kutvonen, "Guarding enterprise collaborations with trust decisions—the TuBE approach," in *Interoperability for Enterprise Software and Applications. Proceedings of the Workshops and the Doctoral Symposium of the Second IFAC/IFIP I-ESA International Conference: EI2N, WSI, IS-TSPQ 2006*. ISTE Ltd, Mar. 2006, pp. 237–248.
- [11] T. Ruokolainen and L. Kutvonen, "Service Typing in Collaborative Systems," in *Enterprise Interoperability: New Challenges and Approaches*, G. Doumeingts, J. Mller, G. Morel, and B. Vallespir, Eds. Springer, Apr. 2007, pp. 343–354.
- [12] L. Kutvonen, J. Metso, and S. Ruohomaa, "From trading to eCommunity management: Responding to social and contractual challenges," *Information Systems Frontiers (ISF) - Special Issue on Enterprise Services Computing: Evolution and Challenges*, vol. 9, no. 2–3, pp. 181–194, July 2007. [Online]. Available: <http://dx.doi.org/10.1007/s10796-007-9031-x>
- [13] eBay, an online marketplace, 2007, <http://www.ebay.com/> [19.12.2007].
- [14] S. Ruohomaa, L. Kutvonen, and E. Koutrouli, "Reputation management survey," in *Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES 2007)*. Vienna, Austria: IEEE Computer Society, Apr. 2007, pp. 103–111.
- [15] L. Viljanen, "Towards an ontology of trust," in *Trust, Privacy, and Security in Digital Business. Second International Conference, TrustBus 2005*, ser. Lecture Notes in Computer Science, vol. 3592. Springer-Verlag, 2005, pp. 175–184. [Online]. Available: http://dx.doi.org/10.1007/11537878_18
- [16] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decision Support Systems: Emerging Issues in Collaborative Commerce*, vol. 43, no. 2, pp. 618–644, 2007.
- [17] M. Wilson *et al.*, "The TrustCoM approach to enforcing agreements between interoperating enterprises," in *Interoperability for Enterprise Software and Applications Conference (I-ESA 2006)*. Bordeaux, France: Springer-Verlag, Mar. 2006. [Online]. Available: http://epubs.cclrc.ac.uk/bitstream/898/Trustcom_Interoperability_France.pdf
- [18] R. J. Rabelo, S. Gusmeroli, C. Arana, and T. Nagellen, "The ECOLEAD ICT infrastructure for collaborative networked organizations," in *Network-Centric Collaboration and Supporting Frameworks*, vol. 224. Springer, 2006, pp. 451–460.
- [19] V. Cahill *et al.*, "Using trust for secure collaboration in uncertain environments," *Pervasive Computing*, vol. 2, no. 3, pp. 52–61, 2003.
- [20] T. Grandison and M. Sloman, "A survey of trust in Internet applications," *IEEE Comm. Surveys and Tutorials*, vol. 3, no. 4, pp. 2–16, Dec. 2000.
- [21] P. Resnick, R. Zeckhauser, E. Friedman, and K. Kuwabara, "Reputation systems," *Comm. of the ACM*, vol. 43, no. 12, pp. 45–48, Dec. 2000.
- [22] R. Chen and W. Yeager, "Poblano—a distributed trust model for peer-to-peer networks," Sun Microsystems, Tech. Rep., 2001. [Online]. Available: <http://www.jxta.org/docs/trust.pdf>