

PathTrust: A Trust-Based Reputation Service for Virtual Organization Formation

Florian Kerschbaum, Jochen Haller, Yücel Karabulut, and Philip Robinson

SAP Research, CEC Karlsruhe, Germany
{florian.kerschbaum, jochen.haller, yuecel.karabulut, philip.robinson}@sap.com

Abstract. Virtual Organizations enable new forms of collaboration for businesses in a networked society. During their formation business partners are selected on an as-needed basis. We consider the problem of using a reputation system to enhance the member selection in Virtual Organizations. The paper identifies the requirements for and the benefits of using a reputation system for this task. We identify attacks and analyze their impact and threat to using reputation systems. Based on these findings we propose the use of a specific model of reputation different from the prevalent models of reputation. The major contribution of this paper is an algorithm (called PathTrust) in this model that exploits the graph of relationships among the participants. It strongly emphasizes the transitive model of trust in a web of trust. We evaluate its performance, especially under attack, and show that it provides a clear advantage in the design of a Virtual Organization infrastructure.

Introduction

We consider the problem of member selection in Virtual Organizations (VO). A VO is understood as a temporary coalition of geographically dispersed individuals, groups, enterprise units or entire organizations that pool resources, facilities, and information to achieve common business objectives. The partners in a VO enjoy equal status and are dependent upon electronic connections (ICT infrastructure) for the coordination of their activities [1]. This concept of VOs is advocated as a promising model for e-activities and it is strongly supported by the European Union Sixth Framework Program. Each VO has an initiator who is responsible for creating and managing the VO. The VO management function can be performed by a group of persons delegated by the VO initiator. A person becomes a VO initiator when he notifies the system of his intention to create a VO.

A VO has a lifecycle which is a state model, which we have adopted from [15] and extended [12]:

1. *Identification*: the preparatory phase of a VO, where the initiator specifies required business roles (e.g. storage provider or data analyzer in a Collaborative Engineering scenario), high-level work units and interactions in what is referred to as the “Collaboration Definition” (CD) and defines control requirements [4].
2. *Formation*: the phase of the VO where members are discovered, selected and assigned to fulfilling the identified service requirements derived from the CD.

2 Florian Kerschbaum, Jochen Haller, Yücel Karabulut, and Philip Robinson

There is therefore a period of negotiation between the initiator and members that concludes with these agreements being signed by the relevant interacting parties.

3. *Operation*: the Collaboration Definition is enacted and the end-points of the role assignments (i.e. which are selected VO members) interact and exchange messages, documents and production information. Operation also has the implicit sub-state “dormant” when all members are inactive due to some technical or contractual exception that needs to be handled.
4. *Dissolution*: the final phase of the VO, where the business objective specified in the CD is met, or some technical or contractual violation occurs that invalidates the existence of the VO.

The focus of the paper is on the Formation phase and, especially, member selection during this phase. In the Formation phase the initiator has to perform following actions:

1. *Query*: the initiator sends a query containing keywords derived from the roles in the CD to a public registry and receives a list of candidates that have previously registered.
2. *Invitation*: the initiator contacts the candidates, informs them of his intention to form a VO and invites them to play a specific role. He sends them the partner profile detailing the expectations derived from the CD.
3. *Negotiation*: the initiator engages in negotiation about contractual terms with the candidates that have expressed interest in joining the VO. The initiator can negotiate with multiple candidates in parallel and pause or resume a negotiation to achieve the best result possible.
4. *Selection*: The initiator chooses the best-suited candidate and assigns him a role in the VO. The chosen candidate now becomes a member of the VO and other candidates are finally rejected.

If we expect a VO initiator to use a reputation system for member selection, it has to provide a benefit for him. There is likely no direct monetary benefit in using the reputation system rather he is likely to receive better service (or in general performance) by using a high reputable provider. So, one expectation would be that the overall number of positive transactions increases when using a reputation service for member selection.

The second benefit of using a reputation system arises when the initiator has to deal with unknown parties. Their business record may be unknown to the initiator and a reputation system can help establish trust. In our model one would then expect that a certain percentage of reputation values is based on the transitive trust evaluation.

We consider the question what kind of reputation system can support the selection of members for the VO. First, the initiator can invite only candidates whose reputation is above a certain threshold. The threshold can be fixed or adaptive to the candidates found (e.g. the ten best reputation values). Second, the initiator can choose from the set of candidates based on reputation. There are many other differentiating factors for candidates, such as price or delivery time, which must be considered in this decision, but reputation can be used as another weighted component in this mix or it can be used to make the final decision among a group of equally well-suited candidates.

The degree of influence reputation has varies with the selection method, but in any case: the higher the reputation the more likely a candidate is to be selected. And a higher selection ratio means more business and more profit. This implies that there is

an incentive for attacking the reputation system, such that the attacker's reputation increases. A reputation system suitable for member selection needs to be resistant against this kind of attack.

We present a model of reputation that is derived from the way business partners are currently selected and differs from most other models of reputation, as discussed in the related work section. It is particularly well-suited to withstand attacks from participants trying to increase their reputation. We present an algorithm that implements this model in our framework of VOs and evaluate its performance.

The paper is organized as follows: the next section reviews related work, section 3 presents our model of reputation and system architecture. Section 4 analyzes the attacks and outlines the design requirements on a reputation system for member selection. We then present in section 5 the design of our algorithm and show its evaluation in section 6. The last section concludes the paper.

Related Work

From a productive use perspective, reputation systems already play a role in several online businesses, such as eBay or Amazon. As in the work presented in this paper, business in those communities exhibits a transactional behavior and the partner selection for transactions is supported by reputation systems. Since the transactions are real business transactions involving money transfer, their reputation systems were subject of several published vulnerability and attack analysis [1][5][11]. Especially Resnick et al. in [11] classified the most common forms of attacks on reputation systems like badmouthing, liars or collusion attacks. They also put an emphasis on initial values, what kind of reputation value is initially assigned to a newly arriving entity without available prior knowledge or history. Josang et al. in [5] provide a quite exhaustive survey of reputation systems in industry and academic research. They also address the previously mentioned attacks for particular reputation systems. Bolton's analysis in [1] revealed that most productive reputation systems are susceptible to fraudulent behavior, for instance cheaters and liars in an eMarketplace. Addressing this particular issue of liars, Padovan et al. [9] and Sen et al. [14] present reputation systems which try to counteract fraudulent behavior or provide an augmented decision process. The work we present in this paper is rather changing the internal reputation mechanism/algorithm than working around a vulnerable system. An experimental evaluation of reputation systems was done in [13]. Many attacks, including fake transactions have been considered, but their main draw-back was the model of reputation that only considered global reputation values. We have designed a reputation algorithm that uses personalized reputation ratings and can show that it significantly performs better against this very important attack.

For Peer-to-Peer (P2P) networks several reputation algorithms have been proposed. These algorithms are related to ours, but usually need to consider different kind of attacks as they occur in real P2P networks. The algorithm proposed in [8] implicitly uses a personalized model of reputation, but is simpler than ours due to the restriction that it needs to be computed in a distributed fashion. The EigenTrust algorithm suggested in [7] is a global reputation system, but explicitly builds on the notion of a

web of trust by computing the global reputation from the entire matrix of ratings. It has been evaluated against attacks in P2P networks and furthermore in [13]. It is based on Google's PageRank algorithm and therefore has a well established basis. We use it as our reference to compare against. A personalization of EigenTrust has been attempted in [3] by applying an extension for PageRank to EigenTrust, but the personalization is very limited. A related approach has been followed in [10], but the reputation is not feed-back based.

A reputation algorithm for eCommerce P2P networks has been suggested in [18]. It is a global reputation algorithm, but uses the reputation of the rater in a restricted fashion without explicit reference to a web of trust. It has been evaluated against some attacks, but not fake transactions. In [19] a reputation algorithm also for electronic marketplaces is described that exploits the graph for properties of the ratings to compute a personalized reputation. The algorithm itself uses all paths instead of our maximum-weight path which deteriorates in cyclic, fully connected graphs like ours and, most importantly, is not evaluated (or has any new design properties) to resist attacks from fake transactions. The Beta reputation system by Josang et al. [5] tries to predict future performance based on a statistical approach. It follows the global model of reputation and is suggested for eCommerce applications. In [17] it has very successfully been made resistant to the related attack of unfair ratings, but a brute-force attack of fake transactions has not yet been evaluated.

Voss suggests the use of reputation for VOs [16], but does not detail its suggestions, nor evaluates the threats that are derived from the suggested uses. The main contribution of the paper, an algorithm to privately leave feed-back ratings, is unrelated to our contribution.

Model of Reputation

In the non-electronic business world business partners are selected based on personal relationships. A business owner has experience of interacting with his partners and therefore bases his trust in them performing business transactions as expected on this experience. The more (positive) experience he has with a partner, the more trust he usually places in that partner. In a highly dynamic, electronic, geographically dispersed environment such as VOs it is difficult to form such personal relationships. Often one is confronted to make choices among candidates with which one has no previous experience. The reputation system can help form trust in such candidates. We view such relationships as the combination of previous performance and recommendation trust, since we believe that an established positive relationship will foster honesty in future recommendations and vice-versa.

In most reputation systems [7][13][18] reputation is scalar value $R(A)$ for each participant A that is a global ranking of the participants. Our reputation model views the system as a web of trust relationships, such as the personal relationships formed by the business owners. Reputation is the relation of a participant A wishing to engage in business with participant B : $R(A, B)$. It is a two-variable function of the two participants, i.e. two participants A and C may have very different views $R(A, B)$ and $R(C, B)$ of B 's reputation.

The idea of using a web of trust is not new and many other reputation system involve the relationships of the participant in the computation of the reputation []. Our algorithm operates directly on the trust relationships and combines transitive trust (as in e.g. certificates or PGP keys) with a reputation rating: If a participant A trusts participant B (with a certain rating) and participant B trusts participant C (with a certain rating), then participant A trusts participant C (with a rating as a function of the other two ratings). One can also compare our algorithm to a recommender system. In some sense, B recommends C to A in the example above. It is also necessary to evaluate this model of reputation and its specific algorithm under the attacks important in its intended area of use (here selection of members for VOs).

This model of reputation (using the trust relationships amongst the participants) particularly lends itself to resistance against the attack of faking positive feedback. A group of attackers collaborate in order to boost their reputation rating by leaving false, positive feed-back for each other. In our model of reputation this will only strengthen the trust relationship among themselves, but not necessarily strengthen the path from an honest inquirer to the attacker, such that the reputation from the honest inquirer's point of view should remain unaffected. We test this hypothesis in the evaluation section of our algorithm.

The trust relationship between two participants is formed based on the past experience they had with each other. A participant leaves a feed-back rating after each transaction and these ratings are accumulated to a relationship value. The reputation $R(A, B)$ can therefore also be seen as a function of all ratings left in the system, i.e. the ratings are the only input to form the pairs of reputations.

Another benefit of exploiting established relationships in member selection is the formation of long-term relationships. By relying on positive past experience well-performing members are likely to be selected again and business networks of participants can form. Such networks have the benefit that they can exploit further long-term optimization of business processes by investing in infrastructure and business process adaptation technology rather than just the short-term satisfaction of a common (temporary) business objective.

System Architecture

Underlying each VO there is an Enterprise Network Infrastructure (EN). This infrastructure provides basic services, such as registration and notification. It also provides the reputation service.

Each participant of a VO must first register with the EN in order to be eligible for membership status in a VO. He must present some credential (e.g. an entry in local administration's business registry) in order to obtain membership status in the EN. Each VO in turn is registered with the EN, as well. The set of registered participants is queried for candidates for a role in a VO during Formation phase. This service is also provided by the EN.

The reputation service is a centralized service offered by the EN. We anticipate there being one reputation service for all VOs, but different EN providers might choose to allow competing reputation services to cater for different needs and

preferences. In the dissolution phase of each VO all members leave feed-back ratings with the reputation server for the other members with whom they have completed transactions. Each such rating can be authenticated to be associated with a specific VO and one cannot leave unsubstantiated feed-back. Nevertheless it is difficult for the EN provider to verify that a business transaction has taken place and an attacker can create fake VOs and leave feed-back for these with the reputation server, i.e. it is still possible for an attacker to create fake transactions.

Since each EN participant needs to register with some real-world credential in order to obtain EN member status, the multiple identities attack on the reputation system, where a participant always starts with a new identity once he has ruined his reputation of the old one, is sufficiently deterred, if not impossible.

Since the reputation service is central, it has access to all ratings and can do its computation locally instead of distributed, preventing difficulties in the reliability of the computation and the overhead of communication cost. Each query just sends the two parties (*A* and *B*) to the reputation service, which does a local trusted computation and returns the result.

Analysis of Attacks and Design Requirements

Analysis of Attacks

As described in the introduction the use of reputation in member selection can provide substantial gains to participants with high reputation, it is therefore necessary to prevent attacks that raise reputation.

The first attack we consider on the reputation system is the creation of fake transactions with positive feed-back. In most reputation systems this clearly raises the expected reputation of the participants the positive feed-back was left for. It therefore has the potential to increase profit when reputation is used for member selection and the attack is very critical. We evaluate the performance of our algorithm under this attack in section 6. A potential mitigation of this attack is to collect fees for every transaction that are supposed to capture the additional profit gained by the fake transactions, but the more vulnerable a reputation system is to this attack, the higher the fees have to be. A built-in resistance to this attack allows the fees to be lower covering the costs of the transaction rather than being used as a deterrence to create fake transactions. We don't consider using the value of the transaction in the reputation a useful deterrence of this kind of attack as suggested in [18], since the value of the transaction can be faked as well. Even if combined with fees, the attacker then can just replace several small fake transactions with one big one or vice-versa. Also the value of the transaction might be confidential in several business cases.

An attack on the overall system rather than on the reputation system itself is to consistently deliver bad performances. This attack is commonly considered for reputation systems in P2P networks, since it is actively being pursued in many real P2P networks. We do not consider this attack here, since we do not believe that any successful business model can be built on consistently performing badly. Differently

from P2P networks, we do not see a motivation for this attack in our scenario and therefore ignore it in our evaluation. We consider however subtle differences in well performing participants, which are supposed to be highlighted by the reputation system.

A third attack is to leave false or no feed-back at all. First, currently methods are being researched in the TrustCoM¹ project that leave feed-back automatically and, second, leaving no or false feed-back has an immediate negative impact on the participant's own feed-back left by the partner. If a participant is allowed to change his feed-back he is capable of reacting to such actions by the business partner, even after he has left feed-back. Since in our setting it is in the attacker's best interest to raise (and not lower) his reputation this attack seems unlikely and we do not evaluate its impact.

There have been "successful" fake business attacks where the attacker offered some services, engaged in many business transactions, collected payments, but never delivered the goods or services. One could imagine the attacker exploiting the reputation service to lure customers to his business. This corresponds to the erratic or changing behavior attack considered in other reputation systems. Luckily there are some economic deterring factors to using reputation for this kind of attack, besides the "legal deterrence" of prosecution. First, building reputation can be a slow process and requires real (successful) transactions. Therefore the attacker would be required to at least set up a minimal real business which is, of course, associated with the initial investments. Second, there are many other differentiating factors, such as prices or advertisement, which can attract customers to a business that work much faster than building a good reputation. We opt for the "legal deterrence" and leave this attack as a whole to the authorities.

The last attack on the reputation system is to create new identities every time one's reputation drops below a certain threshold. This attack is prevented in our system by requiring a real-world credential (such as an entry into local administration's business registry) to enter the system. Furthermore, the attacker always starts out with an initial reputation that is lower than the one of established successful businesses leaving him at a competitive disadvantage.

Design Requirements

Besides attacks on the system and the reputation system there are other scenarios that a reputation system might have to deal with. A business' reputation might be subject to a rapid decline, e.g. if it has entered an insolvency process. Such participants should not be selected as members in a VO, but it is very difficult to represent this scenario using a reputation system, since reaction would need to be immediate and harsh (upon the first indication of such circumstances). Such harsh action often invites another kind of attack where the attacker leaves false feed-back in order to eliminate a competitor (similar to spreading false rumors). Although, one can design for such cases, e.g. using authorization for very negative feed-back, we didn't

¹ www.eu-trustcom.com

and would like to see such cases handled outside the scope of the reputation system, since they only provide means for an “emergency” case.

Another important aspect for a B2B system, such as the VOs, is to support growth. The system will need to start slowly and continuously attract more and more participants. New participants need to be able to enter the market. We believe that VO system offers sufficient differentiating factors for business to be able to enter established markets and build good reputation. Furthermore, new services are offered all the time and allow business to build a good reputation that can be transferred to markets of established services in order to enter those markets as well.

Algorithm

Based on our model of reputation, the requirements and attacks, we designed an algorithm for a reputation system used for member selection, called *PathTrust*. As described earlier, the input to *PathTrust* is the set of all ratings. For each transaction in the system, the user of a service can leave feed-back for the provider. A feed-back rating r is a binary value, either positive or negative. Let $pos[i, j]$ be the number of positive feed-back ratings left by participants i for participant j and $neg[i, j]$ be the negative ones.

PathTrust sees the system as fully connected graph with edges between all participants registered with the EN. Each edge c_{ij} is a function of $pos[i, j]$ and $neg[i, j]$:

$$c_{ij} = \max \left(0.001, \frac{pos[i, j] - \max \left(1, \frac{\sum_{k=0}^n pos[i, k]}{\sum_{k=0}^n neg[i, k]} \right) \cdot neg[i, j]}{\sum_{k=0}^n (pos[i, k] + neg[i, k])} \right)$$

We lower-bound the system to the interval by 0.001 and normalize each edge by the number of total transaction a participant has performed, thereby limiting the weight to the interval $[0.001, 1]$. This provides a relative measure of trust for the participant in another participant (compared to his overall experience), but prevents comparison between edges from different participants. It allows us nevertheless to interpret the weight in our path-searching algorithm as a probability value. The lower bound allows our selection algorithm to choose edges with no experience, even if there are edges with experience from that participant. We weight negative feed-backs by the ratio between positive and negative feed-backs a participant has given to allow the algorithm to react even to fine-grained performance differences. This normalizes average performances to the lowest possible rating. If a ratio is not defined, because the denominator is zero, we default to the other option of the max operation.

We define the weight of a path $\langle i, j, k \rangle$ from participant i to participant k via participant j as: $w_{\langle i, j, k \rangle} = c_{ij} \cdot c_{jk}$. Upon receiving a query $R(A, B)$ for reputation of B

from A PathTrust computes the path with the maximum weight from A to B . Since $0 < c_{ij} \leq 1$ (and therefore each path weight is constantly decreasing), we can do this simply using Dijkstra's shortest path algorithm. The maximum path weight is returned as the reputation for $R(A, B)$.

The algorithm fully exploits the graph properties of the system, and therefore should provide the required resistance against fake transactions. An attacker generating fake (positive) transactions just increases the weight of the edges with his colluders, but no trust relationship is formed with the other participants. Therefore the path between the honest participants and attackers is only strengthened if they engage in real (positive) transaction with each other. We evaluate the algorithm's performance against this attack in the next section. Nevertheless the algorithm can form indirect paths based on transitive trust between participants allowing successfully querying the reputation of participants with whom there is no prior experience.

The algorithm supports the growth of the system for providers as described in the previous section, but the first query of an initiator will return equal reputations for every other participant. This applies to the very first query only, and therefore an initiator entering the system should be offered to choose a small set (one is actually enough) of trusted business partners. The value $pos[; \cdot]$ will be initialized to 1 for these participants simulating one positive transaction. The first query will then return the trust of those trusted partners. Over time as the initiator engages in more and more transactions the influence of the initial choice will be marginal. If an initiator is entering the system for the purpose of engaging in a specific transaction, this step can be replaced by the first transaction.

Evaluation

We ran several simulations to evaluate the performance of our proposed algorithm for VO member selection. The design of the experiments and their results are described in this section.

First, we need to describe how we intend the reputation system to be used for member selection. The service registry returns to the initiator a list of candidates from which the initiator chooses one (after negotiation). In our experiments we do not model negotiation or other differentiating factor between candidates, such as price. We assume that all candidates offer similar conditions and propose the weighted reputation selection algorithm: Let Φ be the set of candidates and let I be the initiator, then for each candidate $C \in \Phi$ the probability that she is chosen is

$$p(C) = \frac{R(I, C)}{\sum_{A \in \Phi} R(I, A)}$$

This approach supports our notion that reputation is a soft criterion for choosing candidates, since it is probabilistic and allows lower ranked candidates to be selected

as well, e.g. they could have differentiated using additional services, such as payment options or price.

Besides the actual algorithm we proposed a specific model of reputation that views reputation as a function of the inquirer and the queried. We argued that this model provides inherent benefits in attack resistance compared to models that see reputation as a function of the queried only. We therefore compare our algorithm to the EigenTrust [7] algorithm. The EigenTrust algorithm also works on the web of trust, since it uses the rater's reputation in computing the final reputation. Nevertheless it still adheres to the model that reputation is a global function (i.e. equal for every inquirer). Furthermore it has performed well in studies of such algorithms [13].

We used 1000 participants (nodes in the graph) in our test bed. 30 services were available to initiators and each participant offers 3 services. The providers for a service were uniformly chosen from the set of participants, i.e. there are 100 providers for each service.

We then simulated the formation of a VO. Each VO has an initiator which has the need for a specific service. The initiator was uniformly chosen among all participants and the requested service uniformly among all services. The initiator then queries the registry for all available providers of that service and chooses a business partner using the weighted selection algorithm explained above. Each such transaction has a value associated with it. The value was chosen uniformly from the domain $[1, 100]$ and given to the initiator. It represents the profit the service provider makes when being chosen for that VO. We did not simulate the profit of the initiator since the inception of the VO is random. The goal of each participant is to maximize its profit and since so far all choices are random, the means to achieve that is to boost reputation which has a direct impact on the probability being chosen for a VO. This models the situation and risk we have been discussing for choosing VO members using the reputation service.

We divided the simulation into rounds. During each round 100 VOs were formed in parallel and there were 100 rounds, i.e. we simulated 10000 transactions per test run. The reported numbers are averages of 3 test runs.

Resistance against fake transaction

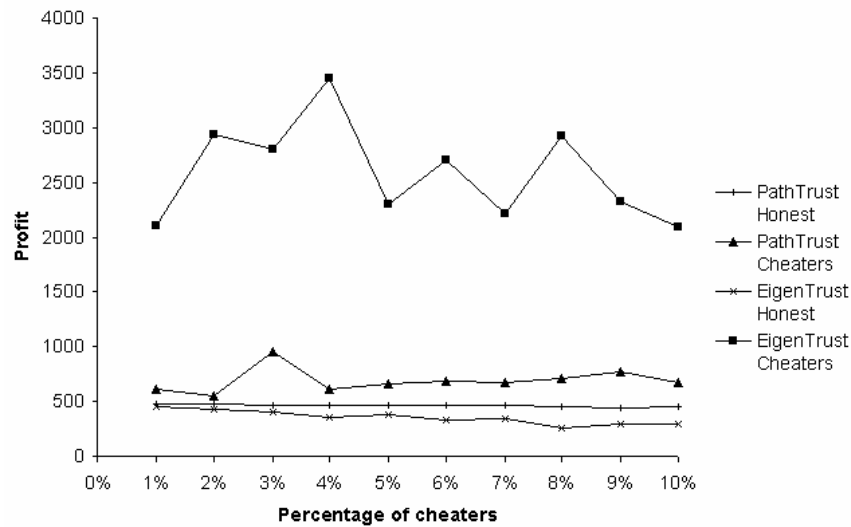


Figure 1 Resistance to cheating

Our first test was to create fake transactions and see if the profit of cheaters increases compared to honest participants. Each cheater created one false transaction per round, i.e. about 10 fake transactions per 1 real transaction. He always chose an assigned collaborator for the fake transaction and its value does not count towards the accumulated profit. In general, each transaction was positively rated, whether it was real or fake.

The results are summarized in Figure 1. We increased the percentage of cheaters from 1% to 10% and depicted the average profit a cheater and an honest participant makes. From the graph we can see that EigenTrust is clearly more vulnerable to this kind of attack than PathTrust, since the average profit of a cheater in EigenTrust exceeds the one in PathTrust up to a factor of 5.6. From these numbers we can conclude that transaction fees that consume the additional profit of a cheater would need to be 10 times higher in EigenTrust consuming 47% of the profit of an honest participant compared to 4.8% using PathTrust.

Percentage of positive transactions




Figure 2 System performance

Our second test is supposed to measure the impact of the reputation system on overall system performance. A reputation value is supposed to predict the performance of a participant. It therefore should help choose the best provider for a given service. Besides acquiring trust in unknown candidates this is a further benefit for the initiator. We divided the set of participants into two: good performers which provide good service in 99% of the transactions and not-so-good performers which provide good service in 95% of the transactions. This reflects our view that all businesses need to achieve a reasonable level of performance to be successful and it makes it difficult for the reputation system to operate on those small differences. There were 100 bad performers, i.e. 10% of the participants. This implies that the expected average percentage of good performances of all transactions is 98.6% when using a random choice of VO members (i.e. no reputation system at all). An improvement over this number indicates an advantage of using this reputation system, i.e. the higher this number the better the reputation system. Even if the reputation system managed to separate the two groups completely and only chose good performers, the percentage of good performances would be 99%. So the possible improvement from using a reputation system in this scenario is small and even small improvements are difficult to achieve.

The results of this experiment are summarized in Figure 2. We increased the percentage of bad performers that cheated using fake transactions attack as above from 0% to 100%. We thought that bad performers might be particularly inclined to conceal their disadvantage by resorting to cheating. No additional (good performing) cheaters were introduced. The graph depicts the percentage of good transactions given the percentage of cheaters. We can see that the EigenTrust algorithm loses its advantage over random choice once we introduce cheating. Furthermore, we see that this loss is much lower in the PathTrust, but still it loses its advantage to random

choice suggesting that cheating annihilates one of the benefits of using a reputation system. We therefore suggest using transaction fees to deter cheating (which can be much lower in the PathTrust algorithm than in the EigenTrust algorithm as discussed in the previous section) and then both systems provide nearly the same benefit in performance gain to the initiator.

Conclusion

We evaluated the requirements for a reputation system to be used for VO member selection. We identified threats and attacks that can be used against the whole system and the use of the reputation system in particular. Based on these findings we developed a model to be used for reputation system for VO member selection that seems particularly well suited to resist the major threats. Then we built a new reputation algorithm in this model and evaluated its performance in a simulation of VO formation against a chosen candidate from the prevalent model of reputation. The evaluation shows that our algorithm provides clear benefits in the presence of attacks. It is therefore beneficial to the operators of a VO infrastructure while preserving the advantages of using a reputation system to the users of that system, the VO initiators.

Currently, a VO infrastructure is being developed by the TrustCoM project that is supposed to be made available for use by business. It would be a great enhancement to this work to study the use and impact of a reputation system and the PathTrust reputation algorithm in particular in a real-world system

References

- [1] G. Bolton, E. Katok, and A. Ockenfels. How Effective are Online Reputation Mechanisms? Technical Report 2002-25, Max Planck Institute of Economics, Strategic Interaction Group, 2002. Available at <http://ideas.repec.org/p/esi/discus/2002-25.html>.
- [2] R. Bultje, and J. van Wijk. Taxonomy of Virtual Organizations, Based on Definitions, Characteristics and Typology. VOnet Newsletter 2(3), 1998.
- [3] P. Chirita, W. Nejdl, M. Schlosser, and O. Scurtu. Personalized Reputation Management in P2P Networks. Proceedings of the Trust, Security and Reputation Workshop, 2004.
- [4] J. Haller, Y. Karabulut, and P. Robinson. Security Controls in Collaborative Business Processes. Proceedings of the 6th IFIP Working Conference on Virtual Enterprises, 2005.
- [5] A. Josang, and R. Ismail. The Beta Reputation System. Proceedings of the 15th Bled Conference on Electronic Commerce, 2002.
- [6] A. Josang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. (to appear) Decision Support Systems, 2005. Available at <http://security.dstc.edu.au/papers/JIB2005-DSS.pdf>.

- [7] S. Kamvar, M. Schlosser, and H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. Proceedings of the Twelfth International World Wide Web Conference, 2003.
- [8] S. Marti, and H. Garcia-Molina. Limited reputation sharing in P2P systems. Proceedings of the 5th ACM conference on Electronic commerce, 2004.
- [9] B. Padovan, S. Sackmann, T. Eymann, and I. Pippow. Prototype for an Agent-based Secure Electronic Marketplace including Reputation Tracking Mechanisms. Technical Report 0204002, Economics Working Paper Archive, 2002. Available at <http://ideas.repec.org/p/wpa/wuwpco/0204002.html>.
- [10] J. Pujol, R. Sangüesa, and J. Delgado. Extracting Reputation in Multi Agent Systems by Means of Social Network Topology. Proceedings of the first international joint conference on Autonomous agents and multiagent systems, 2002.
- [11] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation Systems. Communications of the ACM 43(12), 2000.
- [12] P. Robinson, Y. Karabulut, and J. Haller. Dynamic Virtual Organization Management for Service Oriented Enterprise Applications. Proceedings of the 1st International Conference on Collaborative Computing, 2005.
- [13] A. Schlosser, M. Voss, and L. Brückner. Comparing and Evaluating Metrics for Reputation Systems by Simulation. Proceedings of the IEEE Workshop on Reputation in Agent Societies, 2004.
- [14] S. Sen, and N. Sajja. Robustness of reputation-based trust: boolean case. Proceedings of the first international joint conference on Autonomous agents and multiagent systems, 2002.
- [15] T. Strader, F. Lin, and M. Shaw. Information Structure for electronic virtual organization management, Decision Support Systems 23, 1998.
- [16] M. Voss, and W. Wiesemann. Using Reputation Systems to Cope with Trust Problems in Virtual Organizations. Proceedings of the 3rd International Workshop on Security in Information Systems, 2005.
- [17] A. Whitby, A. Josang, and J. Indulska. Filtering Out Unfair Ratings in Bayesian Reputation Systems. The Icfain Journal of Management Research, 4(2), 2005.
- [18] L. Xiong, and L. Liu. A Reputation-Based Trust Model for Peer-to-Peer eCommerce Communities. Proceedings of the IEEE Conference on E-Commerce, 2003.
- [19] G. Zacharia, A. Moukas, and P. Maes. Collaborative Reputation Mechanisms in Electronic Marketplaces. Proceedings of the 32nd Hawaii International Conference on System Sciences, 1999.