

Applying Recent Enhancements in Reputation Based Mechanisms for Promoting Trust in Mobile Communities

Marcus Dobrinkat
Helsinki University of Technology
mdobrink@cc.hut.fi

Abstract

One inherent problem with virtual communities is how to assess the trustworthiness of unknown peers. Reputation based trust mechanisms are one means to help users decide whether to rely on a stranger or not. Using eBay's feedback forum as an example, we show how simple reputation based trust mechanisms work. We evaluate shortcomings of these mechanisms and give an overview of reasonable enhancements. Later we show how effectiveness is improved by various means including policy based trust measures and using more sophisticated models for calculating reputation values.

When applying the mechanisms to mobile communities, we realize that they have specific requirements of reputation based trust mechanisms such as decentralized reputation storage. With these results, we analyze which enhancements fulfill the given requirements. Next we describe a hypothetical reputation based trust mechanism for a virtual community.

KEYWORDS: Trust, Reputation, Mobile Community, Virtual Community, P2P

1 Introduction

The rise of the Internet has introduced many kinds of virtual worlds where we interact with identities that conceal the individuals behind. This limits our abilities to judge people in the same way as we do in real life, therefore assessing another's attitude and character has become more difficult. In addition, members in virtual communities can come from anywhere in the world, which considerably lowers the likelihood meeting somebody we already know. Thus, there is a genuine need for systems that help us determine the trustworthiness of a potential peer in a virtual community. Policy based systems have been used to manage trust by imposing a rule-based view using exact data [1]. This is often used when validating users' credentials in access control. Policy based trust management alone can not help in determining the trustworthiness of a peer, as its reputation is influenced by that peer's behavior in a specific context. This is easy to understand for humans, but very hard to measure and describe by exact data. However, as we will describe in section 2.3, policy based trust can help in evaluating trustworthiness.

A different approach is to use reputation mechanisms.

They rely mostly on community users' feedback and utilize rules and values of the specific community to maintain order based on these values. Individuals that do not act according to the community rules are punished with negative feedback and will be stigmatized with a low reputation value visible for the whole community, which effectively excludes that peer from being addressed.

One of the first prominent real world examples of reputation mechanism was eBay's feedback forum. Resnick and Zeckhauser have done a statistical analysis of the reputation and transaction data [2] and show that the system works although they reveal problems in the mechanism. They claim that eBay's users trust in the system makes the system work. Without that trust, the problems that the system has would appear: most users believe that dubious sellers can be identified by their low ratings. Few users know that very little negative feedback is given and that the rating system suppresses negative feedback to some extent. Probably many users would be more cautious given these details about the system. As people think that bad behavior is punished adequately, they will behave as they think they are supposed to. This makes the system work. We analyze the shortcomings in eBay feedback forum in section 2.2.

Another simple but time-consuming approach is to keep interactions under surveillance. This only works well for such public communication as discussion forums and can be performed by human moderators. One main requirement is a central information hub where all data is available. Many virtual communities today are based on peer-to-peer (P2P) systems. A large number of these systems do not have a central authority that could monitor user behavior, identify fraudulent peers and hence suspend them.

The outlined problems regarding trust are thus more significant in this form of virtual communities where no central trusted party exists. In section 3 we will explore whether the use of non-centralized P2P communities changes the requirements for reputation based trust mechanisms.

2 Reputation Systems

When we talk about a reputation system, we mean the virtual adoption of word of mouth recommendation. This system works well in real life, for example when a new garage does a bad job, a reputation can be destroyed by word of mouth more quickly than built. Reputation systems try to imitate this behavior: they allow individuals to assess others

and will then calculate reputation values that indicate the subject's trustworthiness. Calculating the trustworthiness can be very simple, such as using only the given feedback values, or based on a much more complex model taking many other factors into account (see 2.3). Reputation systems are based on the assumption that one receives feedback many times and that past behavior will be sustained.

When people understand that their actions influence their reputation, these systems foster proper behavior and discourage untruthful activity within the community.

Barber and Kim describe how peer behavior affects reputation [3]: Reputation is gained by interactions with a peer that fulfils the peer's expectations. The peer's expectations may be receiving valuable information, relying on responsible use of personal information, or supplying goods in accordance with their description. Reputation systems have three basic goals [4]: people should be motivated to give transaction feedback in the first place, be encouraged to behave properly and be discouraged to behave badly.

Dingledine et al. suggests the following measures for the effectiveness of reputation mechanism systems [5]:

1. *Accuracy for long term performance.* This measures how meaningful the rating is, for example how well the mechanism separates fresh users with little historical reputation information from users with low reputation ratings and long history.
2. *Weighting towards current behavior.* Measures how well latest quality variations are taken into account when calculating the reputation rating. Very old reputation data should have a lower weight.
3. *Robustness against attacks.* Measures the immunity against attempts to trick the reputation system in order to manipulate reputation ratings.
4. *Smoothness.* This metric demands the reputation ratings to change slowly without rapid changes due to one time outliers.

2.1 An Example: eBay's Feedback Forum

eBay's feedback forum is one prominent example of a simple, centralized reputation mechanism. It's data has been studied thoroughly and analysis suggests that it increases trust in the platform and leads to less frauds [2].

In eBay's feedback forum [6], both buyer as well as seller can rate their transaction partner with a positive rating (+1), a neutral rating (0) or a negative rating (-1) and add a short text comment. A user's reputation is calculated as the sum of all ratings, where a positive or negative rating of one and the same transaction partner only counts once each. For example, after doing eight transactions with Joe, we got three +1 ratings, which count as one +1 only, and we got five -1 ratings, which count as one -1 only. Thus the total rating would be $+1 - 1 = 0$. The reputation value is displayed along with the user's name and is visible to anybody. Displayed are also the number of positive, neutral and negative ratings for the last month, half a year and year so that trends might be seen.

Interested potential buyers can also see any text comments given to a seller in the past. The comments are displayed together with the buyer's name and rating. Although each user can choose to hide all his feedback comments from the public, this option is not applied often [2]. The text comments give each rating a context and may be used to categorize positive or negative feedback based on personal preferences (a positive rating for "quick delivery" does not tell anything about the actual state of the product).

2.2 Weaknesses

2.2.1 Difficulties to obtain ratings

It may be difficult to obtain feedback from users. One reason can be that the incentive is missing: there is no own profit when rating others. Few users rate others when the rating has negative effects on themselves. Usually, feedback is collected after some interaction or transaction between users and giving feedback usually costs a small amount of time and resources. When feedback is collected voluntarily, a situation known as free-riding can occur, meaning that some users do not provide feedback [2].

The reputation system only works if enough users create input for it by giving feedback, therefore too little feedback will adversely affect the system. Consequently it is important to elicit feedback in the first place. When the reputation system add no or only little value, users might be reluctant to give feedback. Users act similar if they would harm themselves: one might want to withhold good feedback about a sparse resource, which one makes regular use of [7]. This could happen if the user does not want to share the good resource with other people. As an example, take a clever family who would not announce the names of their babysitter everywhere.

This problem is not apparent in eBay's Feedback Forum [2]: over 50% of all transactions get some kind of feedback.

Another side of the issue is not related to the total amount of feedback, but to the low frequency of negative feedback with respect to positive feedback [2]. This is very drastically apparent in eBay's reputation system. In the analysis mentioned above, Resnick et al. found that less than 1% of all transactions are rated as neutral or negative. These results point to a problem: further analysis showed that ratings are strongly correlated when seller and buyer both give feedback. Resnick proposes two reason for this. First, some transactions fail due to third party problems such as delay in delivery where as a result buyer as well as seller get upset and give negative ratings. Second, when one of the two parties deserved a bad rating and the other did not, retaliation causes both to get negative ratings.

Resnick et al. suggest some potential reasons for these results: first, eBay encourages people to contact the other party and try to resolve problems before giving negative feedback. Second, people withhold negative feedback because they are scared of retaliation, meaning that users give positive feedback and expect the other side to do the same. Third, people are very courteous and rather do not give any feedback than negative one and fourthly, users are likely to behave in the

same way as the majority of the community demands, which is being polite [2].

The analysis also shows that high volume sellers could threaten low volume buyers with negative feedback. This works as negative feedback is boosted with lower total ratings [2] and thus has more impact on the low volume buyer.

2.2.2 Unfair Ratings

A basic challenge in reputation systems is how to deal with dishonest positive or negative ratings. There are two circumstances where this happens. The first is when a group of users agree to discriminate against another person and subsequently give this person intentional dishonest negative feedback. The second case, also called "ballot stuffing" means that a group of users agree to intentionally give unjustified positive feedback to somebody who will be promoted by this fraud [8].

In the example of eBay, this risk is reasonably contained by the following rules: feedback can only be given after a finished transaction; there is a charge for each transaction involved; the feedback by each person is aggregated and then limited to +1, 0 or -1 rating points.

2.2.3 Ephemeral Identities

The main goal of reputation mechanisms is to use previous behavior as a predictor for future behavior. This requires that previous behavior can be captured. Each user's identity is linked to the user's previous behavior. If the user abandons its old identity and acquires a new one, past reputation information is effectively removed.

The whole reputation system loses its value if it does not track negative behavior reliably. Therefore an important issue of a reputation mechanism is when past negative ratings can easily be removed by changing one's identity.

In many systems it is possible to change identities. One example is eBay, where changing identities has been made more difficult, but it is still possible. eBay has tried to prevent erasing the prior rating by using an email address as well as a credit card number for creating a new eBay account.

2.2.4 Improper Reputation Calculation and Display

This problem has primarily to do with the way how the reputation rating is calculated. Many different ways exist, ranging from simple sums to complex statistical models. An improper algorithm will result in low values for some of the efficiency measures of the mechanism (measures are described earlier in this section), particularly for the meaningfulness of the rating, the weighting towards current behavior and the smoothness measure.

This can be seen in the example of eBay: with higher ratings, negative feedback is hardly visible in the calculated total rating, therefore the meaningfulness is unsound. This is the case as eBay calculates the overall score simply by subtracting the negative from the positive feedback, given about 99% positive and under 1% negative feedback [2]; the later is more informative and should be highlighted. However, in the user comments negative comments can be displayed, but are

hard to find: all comments are sorted chronologically and as negative comments are rare and there is no sorting function available, it is time-consuming to find them.

2.3 Enhancements

We will now focus on finding solutions for some of the problems presented in the last section.

2.3.1 Eliciting Feedback

Feedback in general can be elicited by creating incentives. These can be monetary or of some other kind. A model which includes monetary remuneration was suggested by Miller et al. [7]. A similar approach is conducted by Jurca and Faltings [9] who present an incentive compatible reputation mechanism which introduces different payments for correct and incorrect reputation information.

The problem of a low rate of negative feedback (see 2.2.1) can be tackled for example by removing the threat of retaliation: in eBay's reputation system, the buyer could be allowed to rate the seller, but not the other way around [2]. Another approach is to anonymize the ratings in order to protect both partners from retaliation (there is no risk in giving a negative user rating anonymously). This has been suggested in a system proposed by Ismail et al. [10] and also by Singh [16]. However, this is not possible for all application areas. When anonymizing communication, often public key cryptography is used which is an example for combining policy based trust with a reputation system.

2.3.2 Unfair Ratings

For tackling the problem of unfair ratings, many approaches have been proposed (see [8] for a list of references).

Josang et al. [8] categorized these into two classes:

1. The ratings that are classified as unfair are neglected or damped. This assumes that a statistical analysis reveals the unfair ratings. As one example, Dellarocas [11] proposes a combined approach using controlled anonymity and cluster filtering. He limits the proposed system so that only sellers can be rated by buyers. The enhanced cluster filtering algorithm proves to work well when unfair ratings are there. As any other filtering method, also this algorithm comes at a cost: when no unfair ratings are present, the reputation is decremented falsely by a small amount.
2. The ratings are weighted according to the reputation of the person who gave the feedback. This assumes that ratings from users with low reputation are probably unfair [12].

2.3.3 Ephemeral Identities

The solutions to this problem can be divided in two classes. First, preventing identity changes by assigning a real world identity and second, making the cost of identity changes high enough so that it is not profitable to change identities [13].

Preventing identity changes could clearly avert ephemeral identities as they would no longer exist: every natural person had one virtual identity which carried historical reputation data which could not be cleared. The solution would be effective, but had several problems. It is not easily feasible as it requires a central authority that can identify people worldwide in an authentic way. Another problem is that in many cases, the loss of anonymity is not desired, for example in a forum where HIV patients discuss matters, users might prefer to stay anonymous to avoid discrimination in real life. A possible solution would be a system, where a central authority grants personal certificates for each different social area which hide the real identity. A user is guaranteed to get only one certificate per area [13]. One advantage is that newcomers could be trusted while they have no negative reputation which would increase cooperation. This is in contrast with the next approach.

Increasing the cost of identity changes generally suspects newcomers to be users that changed their identity in order to clear low reputation. The cost is to get the lowest possible reputation value, which also is the initial reputation value. Newcomers will have to make great effort initially to become trustworthy. The initial effort also is the aspect that makes this mechanism work. Fridman and Resnick have proven that this results in the best cooperation among users [13].

Some derivations of the mentioned solutions to the problem include high cost for identity changes by paying initial fees [13], monetary penalties for misbehavior made possible by being obliged to deposit valid credit card details [7] and making identity changes harder by checking a valid email address, credit card or similar verification as in eBay [6].

2.3.4 Proper Reputation Rating

Dellarocas has suggested a two state approach in [11] for robustness against attacks. He combines controlled anonymity to tackle discriminative ratings and then applies cluster filtering to remove unfair ratings.

Mui and Halberstadt [14] present a computational model for trust and reputation that attempts to take the time variational aspect of reputation into account, differentiates clearly between trust and reputation and promises to take sociological characteristics into account. Time variational aspects are important as new ratings are closer related to the current behaviour of the rated user. Older ratings should therefore have less weight or be ignored completely [12].

3 Reputation in Mobile Communities

The enhancements for reputation mechanisms that were presented so far focus on traditional applications. In this section we will evaluate which special requirements exist in the domain of mobile communities. Most problems that we pointed out for reputation mechanisms in general also apply to mobile communities.

3.1 Requirements

In this section we will evaluate which special requirements exist in the domain of mobile communities. We mainly look at the design principles for reputation management in virtual communities suggested by Sarkio [15] and evaluate them. Some of the requirements Sarkio suggests for virtual communities are:

1. Need for a decentralized reputation information storage. There exist mobile communities where no centralized management is available. This is especially true for decentralized P2P networks such as some file sharing systems. However, there also exist mobile communities with centralized systems. Depending on the assumptions, this point needs to be addressed and will be discussed in the next section 3.2.1.
2. Measures for personalized trustworthiness. Sarkio suggests, that "... a reputation system lacks of personalized trustworthiness estimations and thus, ignores the fact that the members of the community need the reputation information in different situations." [15]. Some kind of ratings are always subjective personal impressions, which depend much on the preferences of the users. This has lead to systems favoring personalized reputation estimates. The technique is called collaborative filtering (CF) and delivers a personal reputation measure based on ratings of people with similar likings. Although CF systems are related to reputation systems, Josang [8] clearly separates the two with the reasoning that the basic assumption in reputation systems is to "... judge the performance of a transaction partner or the quality of a product or service consistently" [8]. In this paper we will adopt the later view and not consider CF as requirement for mobile communities.
3. In virtual communities there is no trusted third party (TTP) available. As for (1), we think this depends on whether the system is decentralized or not. As decentralized systems clearly exist, we adopt this requirement. The lack of a TTP in particular effects how identities are managed and distributed. As with all reputation systems, long term identities are important as reputation information will be based on previous interactions.

3.2 Enhancements

3.2.1 Decentralized Reputation Information Storage

In a decentralized system, there is no one place which stores all reputation ratings. The ratings must be stored either on decentralized dedicated systems like distributed reputation repositories or stored on the user's nodes like in P2P systems where each node is client and server at the same time. One example is that every node stores the reputation information it locally creates and distributes this information to other nodes on request.

In a centralized system, all reputation information is stored in one or more known places available to all nodes. When distributing this information to user nodes, new problems appear. One is improper calculation of the reputation

value. When the ratings of a node are spread through the whole community, the calculation of a reputation value takes only the ratings into account that are located on currently available nodes. All nodes that are off-line and have ratings stored about the node in question, will be ignored. A small number of fraudulent nodes which are always on-line and deliberately send incorrect ratings can significantly bias reputation values [16]. One solution for this problem is to assign for each one peer another peer that acts as anonymous reputation store [16].

Another problem is that fraudulent nodes can try to prevent bad ratings by denial of service attacks or menace when they know who reports bad ratings or which node acts as reputation value storage for which other node. This and other problems can be prevented when anonymity is made a design principle for the reputation system as shown in TrustMe [16].

Eigenrep [17] is one more solution, which employs a DHT based system without special focus on security issues. One other solution was pointed out earlier: in a simple architecture ratings can also be stored locally where they are created. At the time a user requests reputation information about some subject, the node must find other nodes that possess information about the subject [18].

3.2.2 Identities

In decentralized systems without TTP, one major problem is how to manage user identities as no central place exists for storing and authenticating user information. From the solutions presented earlier in section 2.3.3, we favor increasing the initial cost of identities for the use in mobile communities. This approach provides anonymity, which is an important element in a reputation system in order to avoid retaliation for eliciting negative feedback, and protect privacy [16]. The alternative approach which was to prevent identity changes by the use of a central authority is not feasible in a decentralized system.

The path conducted by TrustMe [16] claims to present a decentralized secure reputation management system, however, also here a central node called bootstrap server acts as a certification authority. It is used to initially assign for each node another node that stores the reputation values for the first node. The bootstrap server is required to hide these kind of node relationships from normal nodes, so that nobody knows where to find reputation values for a certain node. The values will be queried by some kind of broadcast request and replied anonymously.

3.3 An Example: P2P File Sharing

We found enhancements for reputation mechanisms in general and for mobile communities in particular. Now we will demonstrate how these enhancements could be realized in a mobile community. We take P2P file sharing as an example for a mobile community. As P2P file sharing we understand a decentralized open environment where peers can join to share and download files. Files are considered unique by the

use of a hash key that is created using the file content. Different parts of a unique file can then be received from different peers. The primary motivator to enhance the P2P file sharing by use of a reputation mechanism is to get an indicator of trust for peers so that the file sharing application can automatically decide whether to trust a peer or not.

3.3.1 Eliciting Feedback

We have seen that by using incentives (see 2.3.1), it is possible to motivate users to give feedback. The motivation to use a file sharing application is to receive the genuine version of a file which is not corrupted. This motivation alone would probably already drive users to give feedback if a corrupted file was received. We suggest additionally that the feedback will be enforced after each finished download by asking the user if the file is genuine or not. Further incentives can be created in form of benefits regarding the user's community status, in example download limits.

The user feedback consists only of the answer to the question if the downloaded file was genuine or not. Behind the scene of this simple Boolean question, the application rates the peers which were contacted to download the file parts. If the file was corrupt, all peers that provided corrupt parts receive negative feedback [18].

3.3.2 Ensuring Fair Ratings

One possibility is to choose a combined approach as pointed out in section 2.3.2. Dellarocas describes in [11] an architecture that employs controlled anonymity in order to prevent peer discriminations as a first step. Then an enhanced cluster filtering algorithm is applied to find and filter out unfair ratings.

A mechanism suggested by TrustMe [16] could be used to provide a single node that stores aggregated reputation values on behalf of another node so that all ratings are available, not only the ones from users that are on-line at the time of the query.

3.3.3 Distributed Reputation Storage

We propose a distributed reputation storage as described by Selcuk [18]. Every peer stores many "trust vectors", one for every remote peer that it has been dealing with. The trust vector contains information about a fixed number of previous interactions with the remote peer. The result of each transaction, whether genuine or not is stored as one bit in this vector. Therefore the trust vector maintains a history of transaction quality for each known remote node.

When the reputation of a remote node is to be evaluated, a process similar to searching for files in the P2P network is started. However, this time it is querying the reputation of the node in question. If the local node has enough information about the remote node, then no query is issued, but instead the local trust vector is used for evaluation. Otherwise, a specified number of randomly selected peers are queried. Using another measure, a credibility rating for each peer, the results are then weighted and the reputation is calculated. The credibility of a peer is evaluated based on that

peers previous responses and if they turned out to be correct or not.

4 Conclusion

As we have seen, reputation mechanisms are a research area that spreads via several domains from networking, security, statistics and artificial intelligence to game theory and sociology, which indicates that it is a complex field. Taking reputation mechanism from standard eCommerce applications to mobile communities poses additional requirements, especially regarding decentralized management and security.

In this paper we have shown how trust can be promoted in mobile communities by the use of reputation mechanisms. We identified special requirements for the use with mobile communities and evaluated recent enhancements as solutions. The requirements of a reputation mechanism for a mobile community depend much on the type and structure of community, for example whether it is purely decentralized or not. Therefore no one single solution exists which provides best results. We used the example of a P2P file sharing application and found how selected enhancements could be combined into a reputation mechanism that has less weak points.

We have seen that eBay's feedback forum is an economically successful, however simple and primitive mechanism when compared with the features of advanced reputation mechanisms. Much research effort is made in the different areas to improve these mechanisms. One question arose during our work: is all that effort justified? In the area of reputation mechanisms where social behavior plays a major role, it is not only the technical brilliance but also what users think or believe about a system, which can have significant consequences.

Finally we have seen that managing trust via reputation mechanisms is a topic that is researched in many fields, which makes this topic not only interesting but also challenging. In order to produce meaningful results in this area, it is, thus, necessary to first understand the complexity of the area being investigated, and then be able to identify and focus on the most pending problems within this area.

References

- [1] P. Bonatti, C. Duma, D. Olmedilla, N. Shahmehri. An Integration of Reputation-based and Policy-based Trust Management. In *Proc. 2005 Semantic Web and Policy Workshop*, 2005.
- [2] P. Resnick, R. Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. *The Economics of the Internet and E-Commerce*. Volume 11 of *Advances in Applied Microeconomics*, Amsterdam, Elsevier Science, 2002.
- [3] K. S. Barber, J. Kim. Belief Revision Process Based on Trust: Agents Evaluating Reputation of Information Sources. In *Proceedings of the workshop on Deception, Fraud, and Trust in Agent Societies held during the Autonomous Agents Conference*, Lecture Notes In Computer Science, vol. 2246. Springer-Verlag, London, 73-82. 2001.
- [4] P. Resnick, R. Zeckhauser, E. Friedman, K. Kuwabara. Reputation Systems: Facilitating Trust in Internet Interactions. In *Communications of the ACM*, 43(12), 12/1/2000.
- [5] R. Dingleline, M.J. Freedman, D. Molnar. Accountability Measures for Peer-to-Peer Systems. In *Peer-to-Peer: Harnessing the Power of Disruptive Technologies*, O'Reilly Publishers, 2000.
- [6] eBay Feedback Forum, <http://pages.ebay.com/services/forum/feedback.html>. Accessed on 20.3.2006.
- [7] N. Miller, P. Resnick, R. Zeckhauser. Eliciting Honest Feedback in Electronic Markets. Working Paper, 2002.
- [8] A. Jøsang, R. Ismail, C. Boyd. A Survey of Trust and Reputation Systems for Online Service Provision, Decision Support Systems, 2005.
- [9] R. Jurca, B. Faltings. An Incentive-Compatible Reputation Mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.
- [10] R. Ismail, C. Boyd, A. Jøsang, S. Russel. Strong Privacy in Reputation Systems. In *Proceedings of the 4th International Workshop on Information Security Applications*, WISA 2003, ser. Lecture Notes in Computer Science, vol. 2908. Springer, 2004.
- [11] C. Dellarocas. Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior. In *Proceedings of the Second ACM Conference on Electronic Commerce*, Minneapolis, MN, October 17-20, 2000.
- [12] B. Yu, M.P. Singh. A Social Mechanism of Reputation Management in Electronic Communities. In *Lecture Notes in Computer Science*, Volume 1860, Jan 2000, Pages 154 - 165
- [13] E. Friedman, P. Resnick. The Social Cost of Cheap Pseudonyms. In *Journal of Economics and Management Strategy*, Volume 10(2), Pages 173 - 199, 2001.
- [14] L. Mui, M. Mohtashemi, A. Halberstadt. A computational model of trust and reputation. In *Proceedings of the 35th Hawaii International Conference on System Science*, volume 7. IEEE Computer Society Press, 2002.
- [15] K. Sarkio, Design Principles for Distributed Reputation Management in Member-Initiated Virtual Communities. *Seminar on Network Security, T-110.5290*, Helsinki University of Technology, 2005.
- [16] A. Singh, L. Liu, TrustMe: Anonymous Management of Trust Relationships in Decentralized P2P Systems. In *IEEE Intl. Conf. on Peer-to-Peer Computing*, 2003.

-
- [17] S. Kamvar, M. Schlosser, H. Garcia-Molina. Eigenrep: Reputation management in p2p networks. In *Twelfth International World Wide Web Conference*, 2003.
 - [18] A.A. Selcuk, E. Uzun, M.R. Pariente. A Reputation-Based Trust Management System for P2P Networks. In *Proceedings of the 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CC-GRID04)*, Chicago, Illinois, April 2004.