
Tailored trustworthiness estimations in Peer-to-Peer networks¹

Katri Sarkio*

Helsinki Institute for Information Technology,
Advanced Research Unit,
P.O. Box 9800, FIN-02015 HUT, Finland
Fax: +358 9 694 9768 E-mail: katri.sarkio@hiit.fi
*Corresponding author

Silke Holtmanns

Nokia Research Center,
Software and Application Technologies Laboratory,
Itämerenkatu 11-13, 00180 Helsinki, Finland
Fax: +358 7180 36139 E-mail: silke.holtmanns@nokia.com

Abstract: In Peer-to-Peer (P2P) communities, users make personal trust evaluations of each other based on their experiences and observations. The available information of the peer's past behaviour, i.e., reputation, is often incomplete, so the credibility of evaluations is a concern and the relevance of the available information varies. In this paper, we propose functions for producing tailored trustworthiness estimations in P2P communities based on the peers' past behaviour. The presented mechanism provides some flexibility for applying it to different kinds of P2P networks.

Keywords: trust; reputation; reputation management; reputation management mechanism; peer-to-peer; P2P.

Reference to this paper should be made as follows: Sarkio, K. and Holtmanns, S. (2007) 'Tailored trustworthiness estimations in Peer-to-Peer networks', *Int. J. Internet Technology and Secured Transactions*, Vol. 1, Nos. 1/2, pp.95–107.

Biographical notes: Katri Sarkio is a Researcher at Helsinki Institute for Information Technology (HIIT). She received her MSc in Computer Science from the University of Helsinki in 2002. Currently, she is working on her PhD at Helsinki University of Technology in the area of Computer Science and Digital Economy. Her research interests include privacy, trust, reputation management and social networks in P2P networks and virtual communities.

Silke Holtmanns works as a Senior Research Engineer at Nokia Research Center Helsinki. She received her PhD in Mathematics from the University of Paderborn (Germany) in 2000. There she worked as a Scientific Assistant till she started in Ericsson Research in the Service Networks and Applications Lab as a Research Engineer. In 2004, she joined the Software and Application Technology Lab at Nokia Research. The main focus areas of her research activities are privacy, trust and identity management in the mobile environment.

1 Introduction

In virtual communities a group of users interacts with each other by digital means and has a mutual area of interest. Users make personal trust evaluations of each other based on their experiences and observations. An example of this type of a community is eBay, which has a centralised evaluation mechanism for managing the information about the users' past behaviour, i.e., reputation. This reputation information is meant to assist the users in their personal trust decisions.

Some virtual communities are based on P2P networks, such as file sharing, movie recommendation and free programming communities. In these interaction environments, i.e., contexts, the actual interactions happen directly between the users, i.e., peers – instead of a centralised trusted third party being in control. In a typical P2P system, the peers contribute in some means such as providing storage capacity to the community.

Typically, a virtual community has a centralised reputation system with one fixed mechanism processing the information and representing the users' trustworthiness. These systems usually have a centralised storage system that contains detailed transaction data, evaluations and other personal data of the users. However, this kind of an approach does not take into account the users' personal preferences and in addition the large centralised storage systems are tempting targets of attacks to the data.

Distributing the reputation management into the P2P network changes the situation somewhat. Processing locally the reputation information enables tailoring the trustworthiness estimations to meet the user's personal preferences. The reputation management mechanism can be personalised by giving weights to transactions that are considered important to the user's current situation. Nevertheless, distributing the reputation management without a trusted third party also produces challenges for managing the information. The system has to take into account that not all reputation information is available and also that its credibility is a concern due to spoofed transactions and ephemeral identities.

In this paper, we examine how to locally produce a tailored trustworthiness estimation of a user from the available reputation information in a P2P network. The main objective is to define a fine grained and context-dependent reputation evaluation mechanism that is generic for P2P applications and to present generic functions that may be instantiated in different problem areas. To achieve this goal, the mechanism focuses on capturing and processing the information that is linked to user level transactions conducted by an application.

2 Related work

A large part of the available distributed reputation management mechanisms focus on reliability in routing and message forwarding (Buechegger and Boudec, 2002). Moloney and Ginzboorg (2004) argue that, in principle, the idea of using the past behaviour of a peer in trustworthiness estimations maps well to pervasive networks utilising short-range wireless technologies. Supporting this, Moloney (2005) also presents simulation of distributed recommendation system.

Typically, mechanisms that capture and process users' evaluations are technology and security oriented such as the privacy-enhanced P2P reputation system that Kinateder and Pearson (2003) propose. Despite the slightly different approaches, the mechanism we present fits partly into the proposition of architecture and algorithms for a distributed reputation system by Kinateder and Rothermel (2003). The presented idea of weights in the different categories of the content in a context supports our approach in evaluating similarity of the content in the conducted transaction. Additionally, we consider it important to be able to evaluate the expertise of those providing the evaluations and to exploit the evaluations of peers that belong to the user's social network.

Compared to the amount of work done in the distributed reputation management, relatively few functions, e.g., Liu and Issarny (2004), and notes for mechanism implementation applicable to different contexts are proposed. Part of the research proposing the functions are related to protocols for exchanging the trust decision related information such as, e.g., works by Aberer and Despotovic (2001) and Selcuk et al. (2004). Some of the earlier works propose functions and factors for calculating peers' reputation from past ratings, others opinions and evaluations, e.g., the models that Chen and Singh (2001) and Xiong and Liu (2003) propose. Even though these models are distributed, the former relies on information available on different sites in the internet and the latter does not take into account the difference in credibility between peers' own experiences and others evaluations. The model that Abdul-Rahman and Hailes (2000) present is wider than these and takes into account agent's own experiences and recommendations in formulating the trust decision. Our approach fits into their model so that the trustworthiness is calculated in an ad hoc manner at any given time. Moreover, our approach extends the approach by defining and producing the weights for evaluating the available information as it is typical, e.g., in partial differential equation theory. In short, the work done in the area of distributed mechanisms having a user application centric approach is still comparably young.

3 Producing the estimated trustworthiness

In producing the trustworthiness estimations, we follow the Resnick et al. (2004) observations of the evaluation of the eBay reputation system. The most important finding from our angle is that the buyers do not often check the details of bad reputation but tend to rely on an overall figure and that the sellers with long-lived high reputation get premium prices. Hence, we add the reputation information as input values to the trust value functions. But, producing impartial trustworthiness estimation also requires evaluating the relevance and credibility of the available reputation information.

3.1 Terminology

Before going into the details of our mechanism, we shall first look at the terminology. Firstly, we consider trust a subjective expectation that a user has of its peer's behaviour, adapted from Baier (1986), Gambetta (1988), Ostrom (1998) and Mui et al. (2002). Here, trust is neither assumed to be transitive nor reciprocal, because the transactions can be nonrecurring and the users strangers to each other. Trustworthiness describes how users consider each other. Based on the trustworthiness, a user makes a trust decision, whether or not a transaction with another peer should be conducted.

The collected information about the past behaviour of a user is considered reputation – good behaviour indicates the peer to be more trustworthy and vice versa. A recommendation is a subjective, either positive or negative, evaluation that a user receives from its peers about another peer. Further, a reputation system captures this information, has a mechanism to process it and represents an estimation of the peer’s trustworthiness to provide the user support to make the trust decision. Our approach relies on the idea of context-specific roles, which are, e.g., in virtual marketplaces such as eBay’s: a buyer and a seller. Each role has a set of possible transactions that have pieces of descriptive information, e.g., value and time of the transaction. These pieces of information are utilised in producing the estimated trustworthiness. The outlined mechanism also examines the similarity of the content of the conducted transaction in evaluating the relevance of the received evaluations.

3.2 *Assumptions for the properties of the mechanism*

We make the following general assumptions for the reputation management mechanism. Firstly, the reputation is not straightforwardly transferable from one community to another and secondly, we assume a dynamic community or a system, where the number of peers is not fixed. Additionally, a user assuming a new identity has no reputation information and previous evaluations to provide.

We assume that N peers participate in the community, denoted by P_1, P_2, \dots, P_N . For clarity reasons, P_1 is the peer who needs to make a decision whether to trust the peer P_2 and peers P_3, \dots, P_N can provide evaluations of P_2 . Additionally, P_1 can have M friends in the community. Trust to these persons may differ from the trust the P_1 has to a stranger. The friends are denoted by F_1, F_2, \dots, F_M and are a subset of P_3, \dots, P_N .

The mechanism has the following properties:

- i is the transaction identifier. For the purpose of this paper, we assume that each peer has a sequence number for their transaction, $i \in \{1, 2, \dots, n\}$ with $n \geq 1$ is the number of all transactions of P_2 . However, defining the details of the transaction identifier for i is left to the mechanism designer.
- The functions related to the trustworthiness evaluation are scaled and in $[0, 1]$.
- Each peer has a unique identifier within the community or the system that is called the user’s identifier. The form of the identifier depends on the implementation and can be, e.g., a public key, a hash of a public key or a pseudonym.

For the implementation, we suggest that the number of transactions taken into account in calculating the trustworthiness is not too large to avoid time and resource consuming computations. This can be reached, for instance, by dropping old information from the system.

3.3 *The underlying architecture*

We assume that the P_1 receives the reputation information upon request and then calculates the estimated trustworthiness using this mechanism. The integrity of the received evaluations is not discussed in this paper. Hence, we assume that double data sets are identified, e.g., by the transaction identifier i , and removed before doing the

calculation. Some form of free textual feedback on the transaction experience is regarded as useful, but should not go into the computation itself.

We set no specific requirements for the P2P system and underlying architecture, thus we do not discuss or make assumptions about and in the distribution and exchange of the reputation information. Additionally, despite the relevance and importance of the identity management, authentication and security considerations, they are beyond the scope of this paper.

4 The reputation mechanism

The mechanism consists of five advanced functions that use three basic functions, namely, *success*, *value* and *time*. The first two advanced functions examine the *relevance* of the evaluations P_1 receives from his current interest point of view and the level of *experience* of the evaluation provider. Based on these and the basic functions, the latter three then produce trust values for the evaluations received from strangers, friends belonging to P_1 's social network and the P_1 's own experiences for the tailored trust value.

4.1 The basic functions

The first basic function is the (1) *success function* f , where s_i denotes the success of the performed transaction i . It is a subjective and numeric evaluation value of a predefined scale that the peers provide. The scale varies depending on the used context, for example, in eBay users evaluate each other on scale $\{-1, 0, +1\}$. If the scale consists of only positive integer values, then the scale is changed into positive and negative values that are balanced against 0. For example, if the scale is $\{1, 2, 3\}$ then the corresponding new scale would be $\{-1, 0, +1\}$.

This mechanism is designed so that earning a good reputation can take a longer time than loosing it: dishonest behaviour can collapse the good reputation fairly quickly, depending on the collapse factor $y \in [0, 1]$, which must be defined scale dependently. The main idea is that for good evaluations it is 1 or close to it and for bad evaluations it gets closer to 0. The closer to 0 the factor is, the quicker the reputation collapses. Nevertheless, intentional vilifying of others with untruthful evaluations should not immediately collapse the reputation of the well-behaving peers. Here, we approach this question partly in the experience function to reduce the effects of the problem.

The success function f is the normalised sum of all given evaluations. Max_s denotes the maximum positive value the evaluation can reach.

$$f(s_i, y; i \in \{1, \dots, n\}) = \frac{1}{n} \sum_{i=1}^n y \frac{s_i}{\text{Max}_s}. \quad (1)$$

The *value function* (2) is denoted by g and the v_i is a measurable value of the transaction i . v_i depends on the context and can be, e.g., the monetary value of purchased good at an online marketplace or the size of the file transferred. Most importantly, it is a measurable unit that is relevant to the particular context. Again, defining the type of v_i in the particular context is left to the person responsible for the mechanism implementation. For the purpose of this paper, $v_i \in \{1, \dots, u\}$ is defined, where $u \geq 1$ is a context-dependent upper limit for the transaction value. The upper limit is defined

for the purpose of scaling the value factor. Therefore, the mechanism designer should define a limit that is reasonable for the particular context. If the transaction value shall not be part of the reputation system, then v_i can be set to 1 for all i .

To reduce the significance of transactions with a high value, the value of function g increases logarithmically. This reduces the risk that one transaction can have too high a significance and overpower all other evaluations previously recorded.

$$g(v_i, i \in \{1, \dots, n\}) = \frac{1}{n} \sum_{i=1}^n \frac{\ln v_i}{\ln u}. \quad (2)$$

In the last basic function (3) *ageing function* h , the t_i denotes the time that the transaction i was performed and t_0 the time P_1 makes the trustworthiness evaluations of P_2 . For instance the numeric value of the internet time Universal Time Coordinate (UTC) and the RFC1305 (1992) and RFC1361 (1992) specifications can be used as the time format. Nevertheless, the means for obtaining a reliable time is out of the scope of this paper and is dependent on the used device and network.

To increase the importance of the latest transactions, and vice versa, the weight of the transaction decreases along the time. The value x ($0 \leq x \leq 100$) quantifies how rapidly the effect of old transactions on the estimated trustworthiness decrease. The mechanism designer must define x to be sufficient for the particular community or system utilising the reputation information.

$$h(t_i, x; i \in \{1, \dots, n\}) = \frac{1}{n} \sum_{i=1}^n \left(1 - \frac{x}{100}\right)^{(t_0 - t_i)}. \quad (3)$$

4.2 *Relevance function*

The examination of the trustworthiness of P_3 and its evaluations about other peers is actually a recursive problem. To be able to judge the recommendation of P_2 received from P_3 , information also about P_3 's past behaviour, i.e., its reputation is needed and so on. To avoid the issue, we present the relevance and experience functions, which aim to capture some information representing the credibility and relevance of P_3 's evaluations.

The *relevance function* d , takes into account roles and the type of the content. For the role, we consider the (dis)similarity of the P_1 's current role r_0 and the P_3 's role r_i in the transaction i that relates to the evaluation of P_2 . For example, r_0 might be 0 for being a buyer and $r_1 = 1$ for being a seller. If P_1 is considering whether or not to buy something from P_2 , the evaluation provided by P_3 is more relevant, if P_3 has evaluated P_2 as a buyer than as a seller. For calculating the similarity of the roles, we assign a value $\{0, 1, 2, \dots, l\}$ to each possible role r_i in a set of R of all roles, where $l + 1$ indicates how many roles can be assumed. It is possible that two or more roles have the same value assigned. However, the similarity of the roles is context dependent and therefore, no value assignment suggestion is given to the roles here. If and when there are more than two possible roles in the system it is necessary to determine the relationships between the roles when implementing the mechanism. For example, it is possible that an intermediary is selling the goods on behalf of the owner.

The second consideration is the (dis)similarity of the content c_i of the transaction i that the P_3 's evaluation relates to and the content c_0 of the transaction P_1 is considering. For calculating the similarity, the content is divided into categories C . Again, these

categories are context dependent and can be, e.g., product categories or themes in discussion forums. Ontologies, such as, e.g., that Liu and Issarny (2004) propose can be used assigning for the values for the content. Here we assign an integer value $\{0, 1, \dots, L\}$ to each category $c \in C$ in a system, where $L \geq 1$ indicates how many categories exist. Every c_i is an element of C for all $i \in \{1, \dots, n\}$.

$$d(r_i, c_i; i \in \{1, \dots, n\}) = \frac{1}{n} \sum_{i=1}^n \left(1 - \frac{1}{l+1} |r_0 - r_i| \right) \left(1 - \frac{1}{L+1} |c_0 - c_i| \right). \quad (4)$$

4.3 Experience function

The *experience* E_{P_3} is the other function that is used to examine the trustworthiness of P_3 and its evaluations. The experience depends on the three following factors. The first is $T \subset \{0, \dots, N\}$ with $N \geq 1$ is the set of former transactions of P_3 . The second is the success of these transactions, illustrated in equation (1). The third is the relevance of the transactions P_3 has performed, as illustrated in equation (4). Note that in the formulas for f and d the n has to be replaced by the size of T and in the sum, the n is $|T|$.

For the implementation it is important to try to avoid too small content categories leading to situations that there are only a few values within one category. Merging of the categories is then useful.

$$E_{P_3}(s_i, r_i, c_j; j \in T) = f(s_i; j \in T) d(r_j, c_j; j \in T). \quad (5)$$

4.4 Trust value for evaluations from strangers

For producing the tailored trustworthiness estimation, a weighted trust value T_{rec} is formed, based on evaluation of the P_2 provided by the P_3 and other peers that do not belong to P_1 's social network. The value is composed of the presented success function f , the value function g , the ageing function h , the relevance function d and these are then multiplied by the sum over all the experiences of all peers that gave the evaluations. With Q we denote the set of all these peers, hence $|Q| \leq n$. Only one experience value E_{P_q} per peer P_q is inserted for efficiency and also to reduce the significance of multiple evaluations given by one experienced peer.

$$T_{\text{rec}}(s_i, v_i, r_i, c_i, t_i, x; i \in \{1, \dots, n\}) = \frac{1}{|Q|} \sum_{q \in Q} [f(s_i; i \in \{1, \dots, n\}) \\ g(v_i; i \in \{1, \dots, n\}) h(t_i, x; i \in \{1, \dots, n\}) d(r_i, c_i; i \in \{1, \dots, n\}) \\ E_{P_q}(s_j, r_j, c_j; j \in T)]. \quad (6)$$

4.5 Trust value for evaluations from friends

To take a step towards the challenge of modelling a human notion of trust, we propose the function T_{friends} for calculating the trust value for evaluations the P_1 may receive from friends F_1, F_2, \dots, F_M belonging to his social network. The P_1 may know the friends in true life or they can be the P_1 's friends only in the virtual community. The P_1 may have the friends' identifiers, such as nicknames, for instance stored in his terminal device. Nevertheless, we omit the details of assigning and receiving the identifiers in this study.

We denote with Z the set of the friends who have conducted transactions with P_2 and provide the evaluations; hence $|Z| \leq n$. F denotes the previous transactions between a friend F_z , $z \in Z$ and P_2 . The F_z 's experience and the relevance of the evaluation are not examined because the F_z is trusted on the friendship basis. Instead, the friends' evaluations are weighted on percentage basis, $0 < w_z < 100$. The closer to 100 the w_z is the more importance the F_z 's evaluations have and vice versa. T_{friends} is a summary of the friends' recommendations of P_2 weighted by the P_1 's relationship to his friends. However, the actual weight of the friends' evaluations depends again on the mechanism implementator and is context specific.

$$F_z(s_b, v_b, t_b; b \in F) = f(s_b; b \in F)g(v_b; b \in F)h(t_b; b \in F) \quad (7)$$

$$T_{\text{friends}}(s_b, v_b, t_b, w; b \in F) = \frac{1}{|Z|} \sum_{z \in Z} \frac{w_z}{100} F_z(s_b, v_b, t_b; b \in F). \quad (8)$$

4.6 Trust value for own experiences

In the presented mechanism, the function T_{my} is used for calculating the trust value for the P_1 's own subjective personal experiences affecting to the trustworthiness evaluation of P_2 . $P \subset \{0, \dots, m\}$ denotes the set of previous transaction identifiers for transactions between P_1 and P_2 and m equals the number of transactions of P . The function is composed of the presented success function f , the value g and the ageing function h that take the data of the previously performed transactions $k \in P$ as input parameters. The relevance function d is not included in the calculation of T_{my} , because it would end up reducing the significance of P_1 's own experiences. Note that in the formulas for f , g and h the n , respectively, has to be replaced by the size of P .

$$T_{\text{my}}(s_k, v_k, t_k; k \in P) = f(s_k; k \in P)g(v_k; k \in P)h(t_k; k \in P). \quad (9)$$

4.7 A tailored trust value

The tailored trust value T is obtained by summing the $T_{\text{rec}}(>0)$, $T_{\text{my}}(>0)$ and $T_{\text{friends}}(>0)$ and dividing the sum by three to take into account all three types of evaluations. Basically, if only one or two values are >0 then they represent the tailored trust value. The combined trust value provides a single value to give an estimated trustworthiness to the user to guide him on the trust decision for the next transaction; the higher the trust value the more trustworthy the peer is.

$$T = \frac{1}{3}(T_{\text{rec}} + T_{\text{friends}} + T_{\text{my}}). \quad (10)$$

A remaining question is whether or not the P_1 should consider the P_2 trustworthy, if no reputation information is available, i.e., $T_{\text{my}} = 0$, $T_{\text{friends}} = 0$ and $T_{\text{rec}} = 0$. To be able to bootstrap the system in the first place and to deal with the situation of no evaluations available, the peers are assigned a very low trust value. In case of 'no evaluations available', the trust value can be 0, 001 or such as the T is in $[0, 1]$. But, the final decision is left to the implementator.

Based on the defined functions, it is possible to present different kinds of figures and values to the user representing the estimated trustworthiness of its peer (the P_2). For instance, in addition to showing the single combined trust value T as a single point, i.e., one number in $[0, 1]$, it is possible to present the experience value for the peer P_3 who provides evaluations of P_2 and the number of transactions P_2 has performed. Nevertheless, we omit further the details of presenting the trust value, which is left to the user interface designer. Also, making the final trust decision is left to the user (here the P_1).

5 Analysis

Next, we shall look at how this mechanism actually works. We examine how the type of the transactions and a bad evaluation affects the P_2 's reputation.

5.1 Settings for the case analysis

For analysing the case, we assume that the P_2 has conducted approximately one deal per month during the last year ($i = 12$) and prior to that the P_2 has not conducted any transactions, i.e., the P_2 has no earlier reputation. The statistical data of the example transactions are illustrated in Tables 1 and 2. It is worth noticing that, according to the figures in the eBay (2003), there were about 100 million users and 1 billion listed items. On this basis, on average each user had about ten listed items in a year. So, in this case, we can consider the P_2 an average active seller.

Table 1 P_2 's transactions: data set 1

i	1	2	3	4	5	6	7	8	9	10	11	12
s_i	1	1	1	0	0	0	1	1	1	-1	0	1
v_i	100	50	100	100	10	10	100	100	50	100	50	10

Table 2 P_2 's transactions: data set 2

i	1	2	3	4	5	6	7	8	9	10	11	12
s_i	1	1	1	0	0	0	1	1	1	-1	0	1
v_i	1	1	1	1	1	1	1	1	1	1000	1	1

The evaluation scale in s_i corresponds to the eBay's scale $\{-1, 0, 1\}$. The collapsing factor y is set to 1 for evaluations 0 and 1, and to 0.3 for -1 evaluations to collapse the reputation fairly quickly. Further, the P_1 that evaluates the P_2 's reputation has conducted one successful transactions ($s_8 = s_k = 1$) with the P_2 . P_1 receives two evaluations ($s_4 = 1$) and ($s_7 = 1$) from his friend F_1 whose opinions he appreciates, hence we set $w_2 = w_1 = 80$. P_3 provides the rest of the evaluations.


Additionally, we set random integer values $v_i \in \{10, 50, 100, 1000\}$ for the values of the transactions and $u = 1000$ as the upper limit. Here, the values represent the monetary value of the conducted transactions. For the simplicity, we consider all of the conducted transactions fresh enough and, therefore, we ignore the details of the ageing function.

For the roles, we assume that the P_1 and the evaluator P_3 were in the same role, i.e., ($r_0 = r_i$) and all of the purchased goods are from the same category ($c_0 = c_i$). Also, for the experience function E_{P_3} , we assume that we have available the information of one other P_3 's transaction, where $d(r_j, c_j; j \in T) = 1$, because the role and the category are again the same and assume that the transaction was successful, i.e., $s_j = 1$.

5.2 Discussion


The mechanism provides variation for the tailored trust value, even though the data set is quite limited, as we can see from Figure 1. The user can gain a good reputation with a reasonable amount of successful transactions, whereas the bad evaluation has an impact on lowering the trust value.

Figure 1 The P_2 's reputation: data set 1



Actually, the mechanism also deals well in separating endeavours to gain a good reputation with the transactions of very low value, such as selling lots of low-cost items (Figure 2). Here, the attempt to gain good reputation with the low-cost items and cheat in a big deal does not succeed as the reputation remains very low.

Figure 2 The P_2 's reputation: data set 2



In fact, the mechanism deals well with locally producing tailored trustworthiness estimation from the available reputation information. The defined functions also can easily be instantiated, e.g., in online marketplace type of transactions. One disadvantage here is that when the transactions involve money, wealthy users can gain reputation easier, which again reflects social schemes.

Another benefit of the mechanism is that the multi-dimensional factors it considers results in estimations that correspond to the character of the conducted transactions. Nevertheless, especially the factor γ needs application area, specific fine tuning to achieve the desired speed to collapse the reputation. In the example settings, this collapse factor in the bad evaluation has a greater impact on the P_2 's reputation in the first data set. Even if the reputation remains low in Figure 2, this collapsing speed gives partly false impression. This is because in Figure 1 the P_2 has earlier behaved in an expected manner, whereas in Figure 2 we can consider that the P_2 intentionally tried to cheat in the one big deal.

Admittedly, the mechanism solves only a subset of the problems related to reputation management. In a P2P community, the peers can create new and ephemeral identities, which enables spoofed transactions and further untruthful reputation generation. Revealing this kind of malicious behaviour is demanding if there exists no authority in control managing and verifying the identities. Currently, a peer with a bad reputation can assume a new identity and again start building up a new reputation.

This relates also to the issue of a new peer joining the system having little or no reputation data and that the evaluation can therefore, be based on a very small data set. One possible approach to distinguish the newcomers and the peer, e.g., selling the low-cost items, is to use the interval arithmetic. For the newcomers the interval is large and vice versa. Furthermore, the recursive problem of an increasing number of examined recommendations also leads to increased requirement for computational capacity. This again makes the trustworthiness evaluations more complex. One possible approach to distinguish a user systematically vilifying a peer's reputation is, for instance, to evaluate the similarity of the received recommendations against others' evaluations, thus examining these problems are again a topic for another paper.

6 Conclusions

Reputation systems can assist users in building desired level of trust, even though these systems are, realistically, not the only sources affecting these trust decisions and relationships. However, we find that the presented mechanism is a suitable tool in producing reputation information-based tailored trustworthiness estimations locally. It takes into account the multi-dimensional factors in the estimations, which correspond to the character of the conducted transactions.

The mechanism succeeds in providing variation to the tailored trust value and in separating the endeavours to gain good reputation with the transactions of very low value.

The user can gain reputation with a reasonable amount of successful transactions and the bad evaluations have an impact on lowering the trust value. Also, attempts to gain good reputation with low-cost items and to cheat in bigger deals do not succeed as the reputation remains very low. Still, many decisions about the details remain as the responsibility of the implementor.

Finally, defining the most relevant form of presenting the tailored trustworthiness estimations to the users must be defined in planning the implementation of the mechanism in the particular context. The factors related to the trust decision have different importance to different users.

Acknowledgements

The authors wish to thank Professor Martti Mäntylä, Seamus Moloney, Juha Päivärinta and Yki Kortnesniemi for feedback on this paper.

References

- Abdul-Rahman, A. and Hailes, S. (2000) 'Supporting trust in virtual communities', *Proc. the 33rd Annual Hawaii International Conference on System Sciences (HICSS-33)*, Hawaii, USA, pp.1769–1777.
- Aberer, K. and Despotovic, Z. (2001) 'Managing trust in a peer-2-peer information system', *Proc. the 10th International Conference on Information and Knowledge Management (ACM CIKM)*, ACM Press, Atlanta, Georgia, USA, pp.310–317.
- Baier, A. (1986) 'Trust and antitrust', *Ethics*, Vol. 96, No. 2, pp.231–260.
- Buchegger, S. and Boudec, J.L. (2002) 'Nodes bearing grudges: towards routing security fairness, and robustness in mobile ad hoc networks', *Proc. 10th Euromicro Workshop on Parallel, Distributed and Network-based Processing*, Canary Islands Spain, pp.403–410.
- Chen, M. and Singh, J.P. (2001) 'Computing and using reputations for internet ratings', *Proc. EC'01*, Florida, USA, pp.154–162.
- eBay (2003) *eBay, Annual Report 2003*, <http://investor.ebay.com/annual.cfm>.
- Gambetta, D. (1988) 'Can we trust?', in Gambetta, D. (Ed.): *Trust: Making and Breaking Cooperative Relations*, Electronic ed. (2000), Department of Sociology, University of Oxford, Oxford, UK, Vol. 13, pp. 213–237.
- Kinader, M. and Pearson, S. (2003) 'A privacy-enhanced peer-to-peer reputation system', *Proc. the 4th International Conference on Electronic Commerce and Web Technologies (EC-Web 2003)*, Springer-Verlag, LNCS 2738, Czech Republic, pp.206–211.
- Kinader, M. and Rothermel, K. (2003) 'Architecture and algorithms for a distributed reputation system', *Proc. the First International Conference on Trust Management (iTrust)*, Springer-Verlag, LNCS 2692, Crete, Greece, pp.1–16.
- Liu, J. and Issarny, V. (2004) 'Enhanced reputation mechanism for mobile ad hoc networks', *iTrust*, pp.48–62.
- Moloney, S. (2005) 'Simulation of a distributed recommendation system for pervasive networks', *Proc. 20th Annual ACM Symposium on Applied Computing – SAC'05*, Santa Fee, Mexico, pp.1577–1581.
- Moloney, S. and Ginzboorg, P. (2004) 'Security for interactions in pervasive networks: applicability of recommendation systems', *Proc. 1st European Workshop on Security in Ad-Hoc and Sensor Networks (ESAS 2004)*, Heidelberg, Germany, pp.95–106.
- Mui, L., Mohtashemi, M. and Halberstadt, A. (2002) 'A computational model of trust and reputation', *The 35th Annual Hawaii International Conference on System Sciences (HICSS-35)*, Hawaii, USA, Vol. 7, p.188.
- Ostrom, E. (1998) 'Behavioral approach to the rational choice theory of collective action: presidential address, American political science association, 1997', *The American Political Science Review*, Vol. 92, No. 1, pp.1–22.

- Resnick, P., Zeckenhauer, R., Swanson, J. and Lockwood, K. (2004) 'The value of reputation on eBay: a controlled experiment', *Working Paper Presented at the ESA Conference*, Boston, MA, USA.
- RFC1305 (1992) *Internet Engineering Task Force, RFC1305, Network Time Protocol (Version 3) Specification, Implementation and Analysis*, March, <http://www.ietf.org/rfc/rfc1305.txt?number=1305>.
- RFC1361 (1992) *Internet Engineering Task Force, RFC1361, Simple Network Time Protocol (SNTP)*, August, <http://www.ietf.org/rfc/rfc1361.txt?number=1361>.
- Selcuk, A.A., Uzun, E. and Pariente, M.R. (2004) 'A reputation-based trust management system for p2p networks', *Proc. 4th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid04)*, Chicago, Illinois, USA, pp.251–258.
- Xiong, L. and Liu, L. (2003) 'A reputation based trust model for peer-to-peer ecommerce communities', *The IEEE International Conference on E-Commerce (CEC)*, Newport Beach, California, USA, p.275.

Note

¹This is an extended and revised version of the paper originally prepared for the IEEE/Create-Net SecureComm, Value of Security Through Collaboration (SECOVAL) workshop, Athens Greek 9 September 2005, and that was published in the proceedings of the IEEE/Create-Net SecureComm workshops 2005.