

Comparing and Evaluating Metrics for Reputation Systems by Simulation

Andreas Schlosser, Marco Voss, Lars Brückner
Department of Computer Science
Darmstadt University of Technology
Hochschulstraße 10
D-64289 Darmstadt, Germany
{schlosser,voss,brueckner}@ito.tu-darmstadt.de

Abstract

Reputation systems evolve as a mechanism to build trust in virtual communities. In this paper we evaluate different metrics for computing reputation in multi-agent systems. We present a formal model for describing metrics in reputation systems and show how different well-known global reputation metrics are expressed by it. Based on the model a generic simulation framework for reputation metrics was implemented. We used our simulation framework to compare different global reputation systems to find their strengths and weaknesses. The strength of a metric is measured by its resistance against different threat-models, i.e. different types of hostile agents. Based on our results we propose a new metric for reputation systems.

1. Introduction

Reputation systems are an important building block for achieving trust within large distributed communities, especially when mutually unknown agents engage in ad-hoc transactions. Based on the knowledge of past behavior it is up to each agent to form his opinion on his potential transaction partner.

This paper is focused on comparing the effectiveness of different methods to compute reputation based on past behavior. A system is effective if it perceives and penalizes agents who try to undermine the community by cheating on their transaction partners. The effectiveness of the reputation systems is evaluated by simulation. Different threat types of misbehaving agents are modeled and for each system it is checked to which degree the system can stand the attack. Other important implementational aspects of reputation systems, including performance, quality of service, protection against hacker attacks, and privacy of transaction data are out of the scope of our analysis and are not discussed here.

The contributions of this paper are to present a formal

model of reputation and an overview of metrics used in different global reputation systems. Based on the formal model a simulation framework was developed to test and compare different reputation metrics. We summarize our simulation results and present a combined metric, called *BlurredSquared*, as an optimization.

The remainder of the paper is organized as follows. In section 2 we give an introduction into the field of reputation systems and present a formal model for metrics in reputation systems. Section 3 describes how our model is applied to different types of metrics. Section 4 deals with the simulation. There we present our simulation framework, introduce the agent models which are used for simulation, and evaluate the simulation results. Related work is discussed in section 5. We conclude the paper and give an outlook on future work in section 6. Appendix A presents some technical details of our simulation setup, appendix B provides supplementary simulation results.

2. Reputation Systems

Reputation is a subject of research in a lot of disciplines. There are many different definitions of the terms *reputation* and *trust* and no accepted common model of reputation exists. Mui et al. [15] have provided an overview of the different notions these terms have in various disciplines and have developed a typology of kinds of reputation. We do not repeat this discussion here but use an intuitive definition of reputation and trust:

Reputation is the collected and processed information about one entity's former behavior as experienced by others. "Trust is a measure of willingness to proceed with an action (decision) which places parties (entities) at risk of harm and is based on an assessment of the risks, rewards and reputations associated with all the parties involved in a given situation." [12]

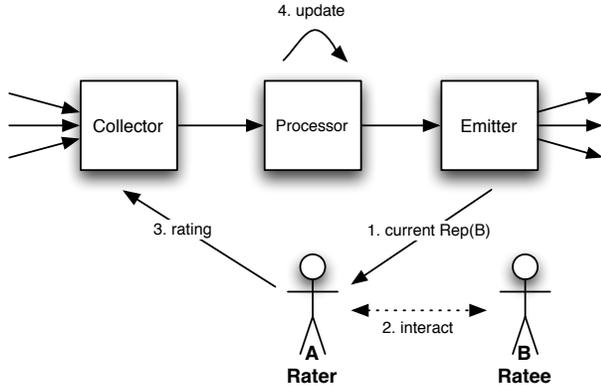


Figure 1. Architecture of a reputation system

A *rating* is a single opinion about the outcome of a transaction. Reputation systems [17] monitor an agent's behavior by collecting, aggregating and distributing such feedback. Conceptually, a reputation system consists of the following components and actors (see figure 1): The target of a rating is called *ratee*. The *collector* gathers ratings from agents called *raters*. This information is processed and aggregated by the *processor*. The algorithm used by the processor to calculate an aggregated representation of an agent's reputation is the *metric* of the reputation system. The *emitter* makes the results available to other requesting agents.

A reputation system is *centralized* if only a single or a small number of processing entities exist. It is *distributed* when every agent calculates reputation values about his partners or neighbors by aggregating all available information. Within a *global* reputation system, there is only a single reputation value per agent. *Local* reputation systems provide different reputation values between different sets of agents. An agent can have different reputation values depending on the requesting agent that issues the query. Mui et al. [15] call this *personalized* reputation. Local can also mean that not all rating information is accessible everywhere. For instance in a distributed system with propagation there is some delay before a submitted rating is available at all nodes. The range of propagation may be restricted to a certain number of hops like in [2]. Another example of local reputation is the approach to take the underlying social network into account [7, 20, 21, 26]. However, in this paper we focus on global reputation systems.

Reputation systems have to deal with a number of threats. Dellarocas [3] describes the problems of unfair ratings and discrimination. Important is the 'ballot stuffing' attack where a group of agents collude to give one member high ratings. In some systems there is an incentive for a bad behaving agent to switch its identity after some time and to start over again [6]. This happens if the reputation of new

members is better than the minimum possible reputation in the system. On the other hand it must not be too difficult for new members to enter the system. Consequently, the initial reputation of an agent has to be chosen carefully.

2.1. Context

Reputation is context dependent. In general reputation earned in one context cannot be applied in another context. For instance, a reputation for being an expert in computer science says nothing about being a trustworthy customer. As a consequence reputation cannot be easily aggregated into a single number. Additionally, a rating may contain different aspects. For customers of an online shop not only the quality and price of a product are important but also the delivery time and after sales services. Reputation is more a vector than a scalar value where each entry represents the reputation with respect to a well specified context.

Most existing reputation systems ignore this fact or are restricted to a single and simple context. eBay's feedback system [5, 18] – as a well studied example – does not distinguish between ratings for buyers or sellers, although these roles are very asymmetric. The seller of a product has much more information about its quality than the buyer and will receive payment before shipment in most cases. Consequently, the risk of a seller is limited. Additionally, the value of the traded goods is ignored by eBay's feedback system. This allows to build up a high reputation by selling goods with low value.

The context of a transaction between two agents contains information about its circumstances and the environment. Examples are topic, time, quality, value or role of the participants. Mui et al. [14] define the context based on a number of attributes which are present or not.

If there are two contexts which are compatible to some degree, a mapping between the reputation information should be possible. However, it is not clear when and how this transfer is possible. Approaches that make use of ontologies [13] may provide a solution to this problem.

In this paper we use an abstract scenario of a space where agents provide homogeneous symmetric services to each other. Thus we do not try to present a detailed model of context here. However, context and mapping between different contexts should be a topic of future research.

2.2. Formal Model

We now present a formal model for reputation that is used throughout this paper. It is not the aim of this model to include every aspect of a reputation system. Especially, the flow of information, the location of processing and storage, the query mechanism and incentives to give feedback are not part of the model. The model provides an abstract view

of a reputation system that allows the comparison of the core metrics of different reputation systems.

According to our definition of reputation a transaction between two peers is the basis of a rating. An agent cannot rate another one without having had a transaction with him.

A is the set of agents. C is the context of a transaction. In the following we assume a simple uniform context and set $C = T \times V$ where $T = \{0, 1, \dots, t_{now}\}$ is the set of times and V is the set of transaction values. We define E as the set of all encounters between different agents that have happened until now. An encounter contains information about the participating peers and the context:

$$E = \{(a, b, c) \in A \times A \times C \mid a \neq b\}$$

A rating is a mapping between a target agent $a \in A$ and an encounter $e \in E$ to the set of all possible ratings Q :

$$\rho : A \times E \rightarrow Q \cup \{\infty\}$$

where ∞ means undefined. Depending on the system Q can have different shapes. In the simple case Q is a small set of possible values: $Q_{ebay} = \{-1, 0, 1\}$ or an interval $Q_i = [0, 1]$. Complex schemes like in [23] are possible, too: $Q_r = \mathbb{R}_0^+ \times \mathbb{R}_0^+$.

E_a represents the subset of all encounters in which a has participated and received a rating:

$$E_a := \{e \in E \mid (e = (a, \cdot, \cdot) \vee e = (\cdot, a, \cdot)) \wedge \rho(a, e) \neq \infty\}$$

All encounters between a and b with a valid rating for a are:

$$E_{a,b} := \{e \in E_a \mid e = (a, b, \cdot) \vee e = (b, a, \cdot)\}$$

Furthermore we define

$$E_a^* := \bigcup_{b \in A} \{e \in E_{a,b} \mid \rho(a, e) \neq \infty \wedge \tau(e) = \max\}$$

as the subset of all most recent encounters between a and other agents. $\tau(e)$ gives the time of encounter e . $v(e)$ gives the value of the transaction e .

We define \overline{E}_a and \overline{E}_a^* to be time-sorted lists of the sets E_a and E_a^* . The Operator $\#$ gives the size of a set or a list. We get the elements of a list L through $L[i]$ where $i \in \{1, \dots, \#(L)\}$. An encounter between a and b at the specific time t is $e_{a,b}^t \in E_{a,b}$.

The reputation of an agent $a \in A$ is defined by the function $r : A \times T \rightarrow R$. The properties of R have already been discussed in the previous section. In most cases it is a subset of \mathbb{R} . We use $r(a) := r(a, t_{now})$ for short. r_0 describes the initial reputation $r(a, 0)$ of a new agent. A complete metric \mathcal{M} is defined as $\mathcal{M} = (\rho, r, Q, R, r_0)$, or for short by the pair $\mathcal{M} = (\rho, r)$.

3. Metrics in Reputation Systems

Within the model given above, a reputation system is described by its specific metric. This allows us to compare different systems by measuring the computed reputation values within certain communities. The next sections give an overview on the metrics that have been subject of our simulations. Most of the systems have been analyzed with and without taking the value of a transaction into account, and with or without multiple ratings per agent pair. The figures show a characteristic graph for each system. The horizontal axis is the time axis, and the vertical axis represents the reputation computed by the specific metric, based on the collected ratings so far. The depicted reputation is distinguished according to several simulated agent types, as described later in section 4. More details about how these figures are generated will be explained in this section, too.

3.1. Accumulative Systems

If a system accumulates all given ratings to get the overall reputation of an agent we call it an accumulative system. The well known feedback system of eBay [5] is an example. We have implemented simulations of this kind of systems with and without considering the transaction values and multiple ratings from the same agent. The possible ratings are $\rho : A \times E \rightarrow \{-1, 0, 1\}$. The basic idea of these metrics is, that the more often an agent behaves in a good way the more sure can the others be, that this agent is an honest one. It is accepted, that an agent can iron out some bad ratings he received just by further good transactions.

However, these systems also allows an agent to behave bad in a certain fraction of transactions and still to improve its overall reputation if this fraction is small enough. Recently, eBay has updated its feedback system to also include the percentage of positive transactions in the detailed member profile.

In the *eBay*-system itself, no transaction values and multiple ratings are considered. The reputation of an agent $a \in A$ computes with:

$$r(a) = \sum_{e \in E_a^*} \rho(a, e) \quad (\text{eBay})$$

With consideration of transaction values, the reputation in the *Value*-system computes with:

$$r(a) = \sum_{e \in E_a^*} \rho(a, e) \cdot v(e) \quad (\text{Value})$$

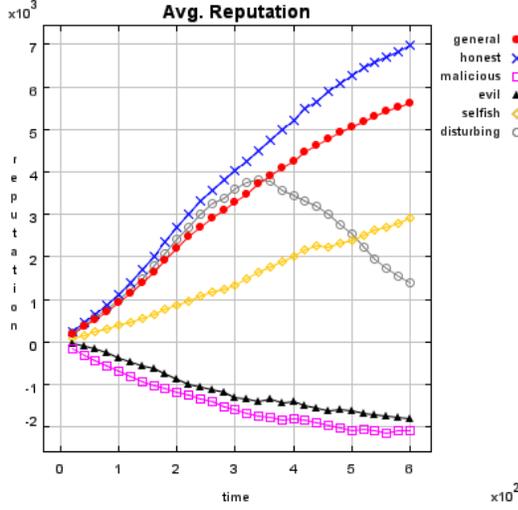


Figure 2. The *Value*-system

Adding multiple ratings, we get the *SimpleValue*-system with the following reputation computation:

$$r(a) = \sum_{e \in E_a} \rho(a, e) \cdot v(e) \quad (\text{SimpleValue})$$

The *Simple*-system considers multiple ratings, but no transaction values. Thus the reputation for an agent $a \in A$ computes with:

$$r(a) = \sum_{e \in E_a} \rho(a, e) \quad (\text{Simple})$$

Figure 2 shows the development of the reputation in the *Value*-system.

Dellarocas has done a theoretical analysis of accumulative systems in [4]. His setup is $Q = \{-1, 0\}$ and

$$r(a, t) = \sum_{e \in D_t} \rho(a, e)$$

where $D_t \subset E_a$ has n elements. This means that only negative ratings are recognized. In every round a random element $e_r \in D_t$ is replaced by the newest encounter: $D_{t+1} := D_t \setminus \{e_r\} \cup \{e_{t+1}\}$. If $n = 1$ we have a system as described in 3.4. Some of Dellarocas' results are reproduced by our simulations as we will describe in the evaluation.

3.2. Average Systems

This kind of reputation system computes the reputation for an agent as the average of all ratings the agent has received. This average value is the global reputation of this agent. An example can be found in [10]. The idea of this metric is, that agents behave the same way most of

their lifetime. Unusual ratings have only little weight in the computation of the final reputation. This could be used to place some bad transactions intentionally by bad agents. The simulated systems use $\rho : A \times E \rightarrow \{-1, 0, 1\}$. The computation is done several ways, with or without considering the transaction value and multiple ratings.

The reputation of an agent $a \in A$ in the *Average*-system without considering multiple ratings and transaction values is:

$$r(a) = \frac{\sum_{e \in E_a^*} \rho(a, e)}{\#(E_a^*)} \quad (\text{Average})$$

The reputation of an agent $a \in A$ in the *AverageSimple*-system without considering the transaction value but including multiple ratings is:

$$r(a) = \frac{\sum_{e \in E_a} \rho(a, e)}{\#(E_a)} \quad (\text{AverageSimple})$$

The reputation of an agent $a \in A$ in the *AverageSimpleValue*-system including multiple ratings and the transaction value is:

$$r(a) = \frac{\sum_{e \in E_a} \rho(a, e) \cdot v(e)}{\#(E_a)} \quad (\text{AverageSimpleValue})$$

The reputation of an agent $a \in A$ in the *AverageValue*-system including the transaction value but not considering multiple ratings is:

$$r(a) = \frac{\sum_{e \in E_a^*} \rho(a, e) \cdot v(e)}{\#(E_a^*)} \quad (\text{AverageValue})$$

Figure 3 shows the development of the reputation in the *AverageValue*-system.

3.3. Blurred Systems

These reputation systems compute a weighted sum of all ratings. The older a rating is, the less it influences the current reputation. An approach in a peer-2-peer environment is described in [22], a metric with an unspecified time-dependent weight-function is used in [7]. A blurred system is implemented with and without considering transaction values for the simulation. Multiple ratings are always considered in this system. This is no problem, because older ratings have a lower weight. Groups cannot manipulate their reputation by giving each other high ratings. Possible ratings are $\rho : A \times E \rightarrow \{-1, 0, 1\}$. This metric is based on the observation that agents do change their behavior during their lifetime. The assumption is, that they will behave more probably like they did in their most recent transactions than they did in transactions long ago in the past.

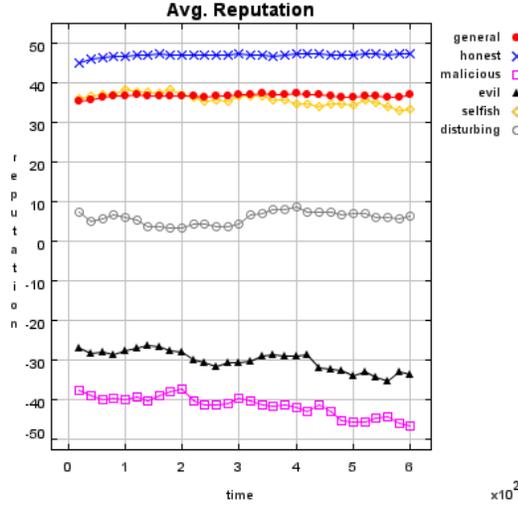


Figure 3. The *AverageValue*-system

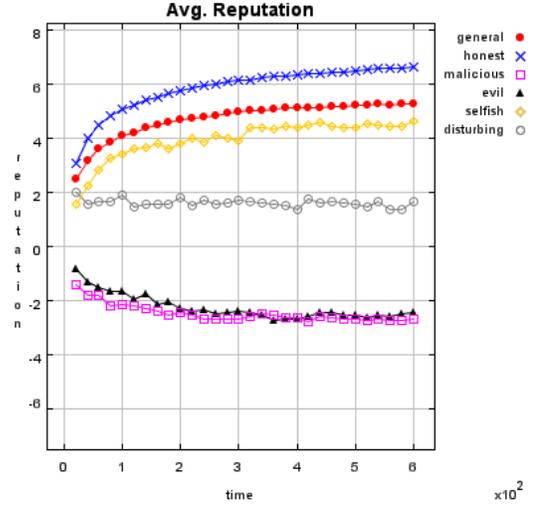


Figure 4. The *Blurred*-system

The reputation of an agent $a \in A$ without considering transaction values is:

$$r(a) = \sum_{i=1}^{\#(\overline{E}_a)} \frac{\rho(a, \overline{E}_a[i])}{\#(\overline{E}_a) - i + 1} \quad (\text{Blurred})$$

and with consideration of transaction values:

$$r(a) = \sum_{i=1}^{\#(\overline{E}_a)} \frac{\rho(a, \overline{E}_a[i]) \cdot v(\overline{E}_a[i])}{\#(\overline{E}_a) - i + 1} \quad (\text{BlurredValue})$$

Figure 4 shows the development of the reputation in the *Blurred*-system.

3.4. OnlyLast Systems

Even if only the most recent rating posted for an agent is regarded, the resulting reputation system is working. Delarocas has formulated the same result in his theoretical work [4]. This system is implemented in the simulation with and without consideration of transaction values, ratings are $\rho : A \times E \rightarrow \{-1, 0, 1\}$. This is an extreme variation of the *Blurred* system. Here we expect an agent to behave like he did last time, no matter what he did before.

Without considering transaction values in the *OnlyLast*-system the reputation of an agent $a \in A$ is:

$$r(a) = \rho(a, \overline{E}_a[\#(\overline{E}_a)]) \quad (\text{OnlyLast})$$

With consideration of the transaction value in the *OnlyLastValue*-system the reputation of an agent $a \in A$ is:

$$r(a) = \rho(a, \overline{E}_a[\#(\overline{E}_a)]) \cdot v(\overline{E}_a[\#(\overline{E}_a)]) \quad (\text{OnlyLastValue})$$

Figure 5 shows the development of the reputation in the *OnlyLast*-system.

3.5. EigenTrust System

This system combines the local reputation values of each agent iteratively to a global reputation [11]. This is done by modifying a target agent's reputation values stored locally at one agent a by the opinions of surrounding agents. These opinions are weighed according to the local reputation values a has about its neighbors. During this process the individual reputations are iteratively accumulated to one single global reputation for each agent. This system is a special instance of the metric described in [24]. In this metric the computed reputation depends on the ratings, the reputation of the raters, the transaction context (e.g. transaction value), and some community properties (e.g. the total amount of given ratings). Another example is the *Sporas*-system in the next section 3.6, where each rating is weighted by the reputation of the rater.

The algorithm for the *EigenTrust*-system is described below. Legal ratings are $\rho : A \times E \rightarrow \{-1, 1\}$. First we have to build a reputation matrix M , where (m_{ij}) contains the standardized sum of ratings from Agent i for Agent j :

$$(m_{ij}) = \frac{\max(\sum_{e \in E_{i,j}} \rho(j, e), 0)}{\sum_j \max(\sum_{e \in E_{i,j}} \rho(j, e), 0)}$$

$$\vec{r}^{(0)} = \begin{pmatrix} 1/\#(A) \\ \vdots \\ 1/\#(A) \end{pmatrix}$$

$$\vec{r}^{(k+1)} = (M)^T \vec{r}^{(k)}$$

$$r(a_i) = \lim_{k \rightarrow \infty} \vec{r}_i^{(k)} \cdot \#(A) \quad (\text{EigenTrust})$$

After about 10 iterations this value is sufficiently approxi-

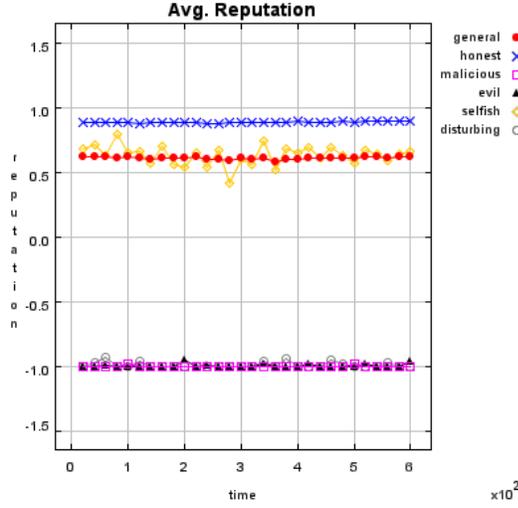


Figure 5. The *OnlyLast*-system

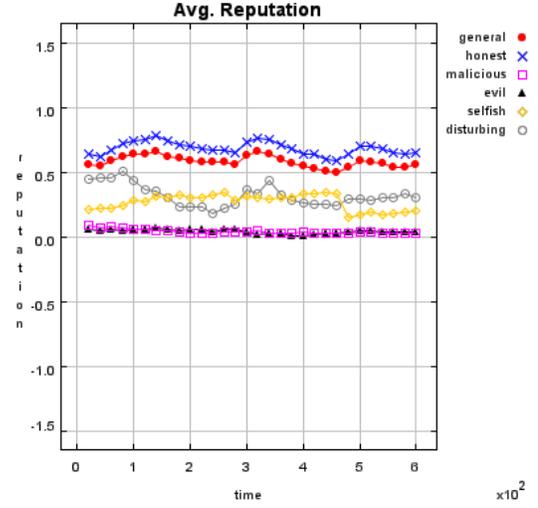


Figure 6. The *EigenTrust*-system

mated. Figure 6 shows the development of the reputation in the *EigenTrust*-system.

3.6. Adaptive Systems

Within adaptive systems the reputation of an agent changes differently when receiving a new rating, depending on his current reputation. I.e. the reputation of the ratee increases more, if he receives a positive rating and his reputation is low, and changes less if his reputation is already high. Reputation can also decrease in the same manner.

The *Sporas*-system [26] grants only positive ratings: $\rho : A \times E \rightarrow \{0.1, \dots, 1\}$. This causes that reentering the community under a new identity is useless, as the reputation will be lower than the current reputation. The resulting value of a given rating depends on the reputation of the rater, too. The higher the rater's reputation is, the larger is the value of his rating. The reputation of an agent $a \in A$ having received a rating from b at time $t \in T$ computes with

$$N_i(a) = r(a, i-1)/D$$

$$\Phi(r(a, i-1)) = 1 - \frac{1}{1 + \exp\left(\frac{-R(a, i-1) - D}{\sigma}\right)}$$

$$r(a, 0) = 1$$

$$r(a, t) = r(a, t-1) + \frac{1}{\Theta} \cdot \Phi(r(a, t-1)) \cdot r(b, t) \cdot (\rho(a, e_{a,b}^t) - N_t(a)) \quad (\text{Sporas})$$

at which D is the maximum reachable reputation and Θ and σ are constants for the time- and reputation-dependable weight. In the simulation we used $D = 30$, $\Theta = 10$, and $\sigma = 0.8$.

A similar system providing positive and negative ratings and reputations is suggested by Yu and Singh [25]. The possible ratings are $\rho : A \times E \rightarrow \{\alpha, \beta\}$, where $\alpha > 0, \beta < 0, |\alpha| < |\beta|$, thus it is easier to drop a good reputation than to build it up. We adapted their proposal to a global approach. The reputation of an agent $a \in A$ at time $t \in T$ computes with

$$\Phi(r, q) = \begin{cases} \alpha & \text{if } r = 0 \wedge q > 0, \\ \beta & \text{if } r = 0 \wedge q < 0, \\ r + \alpha(1-r) & \text{if } r > 0 \wedge q > 0, \\ \frac{r+\beta}{1-\min\{|r|, |\beta|\}} & \text{if } r > 0 \wedge q < 0, \\ \frac{r+\alpha}{1-\min\{|r|, |\alpha|\}} & \text{if } r < 0 \wedge q > 0, \\ r + \beta(1+r) & \text{if } r < 0 \wedge q < 0. \end{cases}$$

$$r(a, 0) = 0$$

$$r(a, t) = \Phi(r(a, t-1), \rho(a, \bar{E}_a[i])) \quad (\text{YuSingh})$$

In the simulation we set $\alpha = 0.05$ and $\beta = -0.3$. Figure 7 shows the development of the reputation in the *YuSingh*-system.

3.7. Beta Systems

The *Beta*-System [9] tries to predict statistically an agent's behavior in his next transaction. Therefore the data about earlier transactions is evaluated, and the probability with which an agent behaves good or bad is derived. The share of good (r) and bad (s) transactions an agent made in the past (e.g. neutral behavior means $r = 0.5$ and $s = 0.5$) is determined. These two variables are parameters for the beta-distribution, whose expectation predicts the future behavior of the agent.

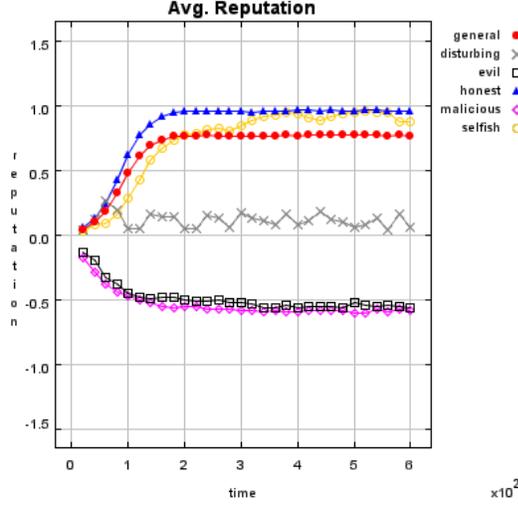


Figure 7. The *YuSingh*-system

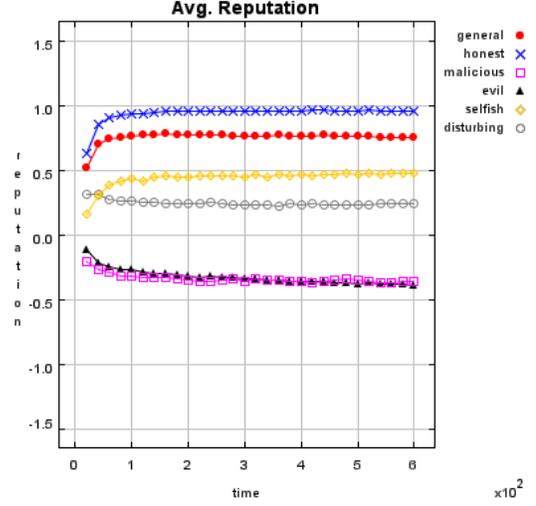


Figure 8. The *Beta*-system

The beta-distribution is a continuous distribution $\text{Beta}(a, b)$ between 0 and 1, with two parameters $a, b > 0$ and its expectation \mathbb{E} is $a/(a + b)$. If $a = b = 1$ the beta-distribution is identical to the uniform distribution. If $a > b$ the expectation is > 0.5 , if $a < b$ the expectation is < 0.5 . The larger a and b are, the smaller is the variance $\sigma^2 = ab/(a + b + 1)(a + b)^2$.

In this system the ratings $\rho : A \times E \rightarrow \{-1, \dots, 1\}$ are available. The reputation for an agent a in the

$$\begin{aligned}
 r^a &= \sum_{i=1}^{\#(\overline{E}_a)} \lambda^{\#(\overline{E}_a)-i} \cdot (1 + \rho(a, \overline{E}_a[i]))/2 \\
 s^a &= \sum_{i=1}^{\#(\overline{E}_a)} \lambda^{\#(\overline{E}_a)-i} \cdot (1 - \rho(a, \overline{E}_a[i]))/2 \\
 r(a) &= \frac{r^a - s^a}{r^a + s^a + 2} \quad (\text{Beta})
 \end{aligned}$$

where $0 \leq \lambda \leq 1$. When computing the reputation with consideration of the transaction value, the computation for r^a and s^a changes:

$$\begin{aligned}
 r^a &= \sum_{i=1}^{\#(\overline{E}_a)} \lambda^{\#(\overline{E}_a)-i} \cdot (1 + \rho(a, \overline{E}_a[i]))/2 \cdot v(\overline{E}_a[i]) \\
 s^a &= \sum_{i=1}^{\#(\overline{E}_a)} \lambda^{\#(\overline{E}_a)-i} \cdot (1 - \rho(a, \overline{E}_a[i]))/2 \cdot v(\overline{E}_a[i])
 \end{aligned}$$

We simulated this system with $\lambda = 0.99$. Figure 8 shows the development of the reputation in the *Beta*-system.

4. Simulation of Reputation Systems

There is no accepted test-bed scenario for reputation systems that allows to compare different metrics under the same conditions. Often the *Prisoner's Dilemma* or a custom scenario is used for experimental evaluation of a single system. Sabater [19] proposes a scenario inspired by a supply chain with a number of different markets that allows quite complex settings. In our work the simple scenario of a space where agents provide homogeneous symmetric services to each other is preferred. We have evaluated the efficiency of the different metrics by simulation of this scenario. As a simulation framework we used RePast [16]. RePast is an agent-based simulation toolkit which offers easy to use methods for visualization. The simulation is based on discrete time ticks. At each tick every agent is supposed to do something, in our case to trade with another agent and rate him. After the agents finished their actions the data is collected and visualized.

After a short description of the framework's capabilities, we present the different agent models we implemented. Then we evaluate how the different metrics given in section 3 performed against several attacks by these agents.

4.1. Framework

Our simulation framework is highly automated. The handling of the agents, the initiation of transaction, and the storage of the ratings are part of the framework. The only thing that must be implemented for simulating a new metric \mathcal{M} are the functions $\tilde{\rho}$ and r .

The steps of a transaction are depicted in figure 9. When the simulation engine selects an agent to initiate a transac-

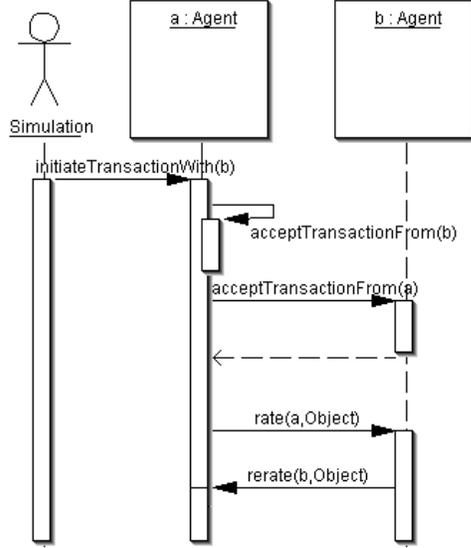


Figure 9. The rating process

tion with another agent, he first tests if the transaction is acceptable. It is also checked if the other agent is willing to transact with him. The acceptance function is described in section 4.2.1. Then the initiating agent determines the outcome of the transaction and both agents submit their ratings. These ratings depend on the agent type and may not match the real outcome. The different types of agents we used in our simulation framework are described in section 4.2.2. New agent types can be easily integrated by implementing a template class.

For the correct interpretation of the outcome of encounters, the framework needs a mapping of the possible outcomes of transactions to the set of possible ratings: $\tilde{\rho} : [-1, \dots, 1] \rightarrow Q$. This function is used by the rating function ρ such that the agent can rate his partner corresponding to his intention. The worst rating is $\tilde{\rho}(-1)$, the best one is $\tilde{\rho}(1)$. For example the *eBay*-system we implemented uses this function:

$$\tilde{\rho}(o) = \begin{cases} -1 & \text{if } o < -0.2, \\ 1 & \text{if } o > 0.2, \\ 0 & \text{else.} \end{cases}$$

The metric function r is evaluated by the framework after each step and can be used by the agents to retrieve information about other peers. It computes the current reputation of $a \in A$. The storage object and the agent a of interest are passed as arguments. Methods for retrieving all necessary information, i.e. the received ratings $\rho(a, e)$ and the context information $\tau(e)$ and $v(e)$ for each $e \in E_a$, are provided by the storage object.

To influence the environment of the simulation, our

framework provides many user-definable parameters. Appendix A gives a detailed introduction in the parameters of the simulation and their values during our simulation runs.

4.2. Agent Models

The effectiveness of a reputation system and its metric depends on its resistance against several types of attacks. This section describes the different simulated agent behaviors. The success of non-honest agents is the measurement for the quality of the metric. The different agent types implemented for the simulation are called *honest*, *malicious*, *evil*, *selfish*, and *disturbing*. These types differ in their behavior when transacting and rating. But they have in common, that their decision whether they are willing to transact with another peer or not, is based on the reputation of this peer. Thus we first present our model for acceptance behavior and after that the differences between the agent types.

4.2.1. Acceptance Behavior We assume that no agent will transact voluntarily with a non-trustworthy agent. Thus the peer's reputation has high influence on an agent's decision if he should transact or not. The higher the reputation of the transaction partner is the more likely he wants to trade with him. Another aspect which influences the agent's decision is the value of the transaction. When trading lower values an agent should not expect such a high reputation as he would when trading high value goods. The method which is described below models the behavior we expected the agents to have. The usability of this method is proofed empirically by our simulation runs. This method works better than other methods with simple thresholds like in [1] without considering transaction values. However, this mechanism is still far away from modeling a realistic human behavior.

We use a scale-based approach. For every reputation system, a reputation $s > 0$ must be known, at which an agent is expected to be 100% trustworthy. The other way around we expect an agent with the reputation $-s$ to be totally untrustworthy. Based on this we divide the reputation space in scale segments á $\frac{1}{10}s$. Based on an empirical study we define the values from table 1 for the value of s . The scale segment

$$\left[\frac{r_s - 1}{10}s, \frac{r_s}{10}s \right), r_s \in \mathbb{N}$$

has the number r_s , and for negative reputations the segment

$$\left[\frac{r_s}{10}s, \frac{r_s + 1}{10}s \right), -r_s \in \mathbb{N}$$

has the number r_s . During the simulation the agents trade goods with the value $v \in [1, \dots, 100]$. These values are divided in five equal segments, too. They are numbered

like the reputation scale. The value v lies in the segment $v_s = \lceil 1/20v \rceil$. For instance the value $v = 25$ falls in the segment 2.

Table 1. Acceptance reputation

System	Reputation s
eBay	120
Simple	100
SimpleValue	2000
Value	2000
Average	1
AverageSimple	1
AverageSimpleValue	40
AverageValue	40
Blurred	6
BlurredValue	250
OnlyLast	0.8
OnlyLastValue	40
EigenTrust	2
Sporas	30
YuSingh	1
Beta	0.8
BetaValue	0.8

Every agent $a \in A$ is classified by his reputation r and the transaction value v . The classification c_s is computed with $r_s(a) - v_s(a)$, where $r_s(a)$ computes the segment on the reputation scale for the agent a and $v_s(a)$ the segment on the value scale. Figure 10 shows the computation of the classification. The classification c_s is used as a parameter of the beta-distribution. We chose this distribution, because it is easily configurable by the two parameters to the desired distribution. If $c_s > 0$ the parameters for the beta-distribution are $a = c_s$ and $b = 1$, if $c_s < 0$ it is $a = 1$ and $b = -c_s$. Now every agent decides with a threshold $t \in [0, \dots, 1]$ if he wants to accept the transaction, $t < \text{Beta}(a, b)$, or deny it, $t > \text{Beta}(a, b)$. The higher the threshold t is, the less risk an agent is willing to agree to.

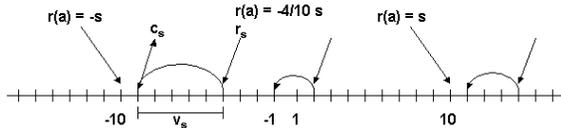


Figure 10. Computation of the classification

Reputation systems which have only positive reputations (cf. section 3.5 *EigenTrust*), can also be handled by this model. The weakness of such systems is, that a new-

bie and an agent with bad reputation do not differ. Thus it is acceptable, that malicious agents are treated like agents with neutral reputation by this model. In the simulation of *EigenTrust* in [11] agents accepted transactions from other peers with the (worst) reputation 0 with the fixed probability of 10%.

4.2.2. Agent Types The agents types we used for simulation differ in their rating-behavior and the sort of transactions they initiate. It was not our aim to model a realistic behavior or to include elaborate strategies, but to develop some simple test cases for our framework. The different types are:

Honest Agent This agent initiates only good transactions (i.e. he serves his trade partner what he expected he would get). His ratings are always correct (good transactions are rated good, and bad transactions are rated bad).

Malicious Agent This agent initiates good, neutral and bad transactions by chance. He tries to undermine the system with his rating behavior and rates every transaction negative.

Evil Agent or Conspirative Agent These agents try to gather a high reputation by building a group in which they know each other. If an evil agent finds another evil agent to trade with, they always give each other a good rating. If an evil agent does not find another evil agent, after seeking for a while, he transact neutral and rates neutral.

Selfish Agent This agent is a so called *freerider*. He blocks all inquiries by other agents and refuses to rate his transaction partners. He just initiates neutral to good transactions by himself.

Disturbing Agent This agent tries to build a high reputation, such that the other agents trust him, with making good transactions and correct rating. Then he switches to a malicious behavior until his reputation gets too bad and then starts from the beginning.

Example If an honest agent h initiates a transaction with a malicious agent m , the outcome of this transaction is positive, thus h tells m that he should be rated with $\tilde{p}(1)$. But m rates h with $\tilde{p}(-1)$, that is why we call him malicious.

The naming of these agent types is quite arbitrary and should just provide an intuitive meaning of their behavior. They may not be confound with different agent types used in other papers, as there is no consistent naming policy for different agent types in the context of reputation systems known to us.

4.3. Evaluation Criteria

First we have to define the criteria which determine if a metric works well. In the simulation we measure the average reputation, the reputation bias, the transaction rate, and the profit separately for each agent type.

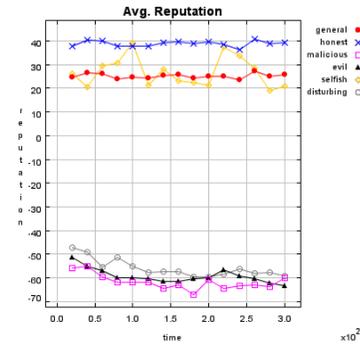
The *average reputation* is the mean value of all reputations of agents from a specific type. We store the correct rating an agent should receive for each transaction. The difference between the average reputation which is calculated with these stored ratings and the actual ratings is the *reputation bias*. The *transaction rate* is the portion of transactions completed successfully compared with all initiated transactions. The *profit* p of a transaction calculates by $p = \frac{1}{o+2}v$ where v is the transaction value and $o \in [-1, \dots, 1]$ is the real outcome of this transaction. The assumption is, that the profit is higher, if an agent cheats on another agent. Figure 11 illustrates an example of the measured values. The horizontal axis is the time axis in all cases, as usual. In figures 11(a) and 11(b) the vertical axis represents the computed reputation respectively the bias of this reputation. Figure 11(c) shows the amount of successfully finished transactions from the last 10 initiated ones, figure 11(d) has an abstract scale on the vertical axis, this values have to be interpreted relatively to each other.

The only agent types with concrete aims are the *evil* and *disturbing* agents. The evil agents want to raise their reputation and the disturbing agents try to gain high profit. The *malicious* and *selfish* agents do not receive any advantages from their behavior. Thus for these attacks it is only important, if the *honest* agents retain their good reputations.

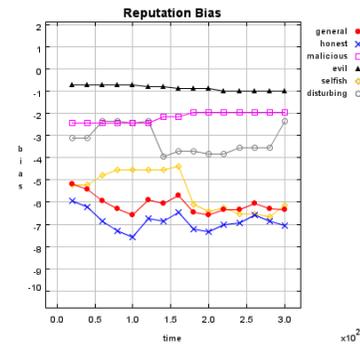
4.4. Simulation Results

We simulated the different metrics with an increasing amount of hostile agents to find the critical point, where the metric fails to maintain dependable reputation. It was not the aim to simulate a real world behavior but to use our framework under extreme conditions. This should give an example how it can be used to compare different metrics. Table 2 summarizes the results of the simulation. The *AverageSimple*-system is not listed there, because it turned out to be totally impractical. Even honest agents could not reach high reputation, after a while all agents had neutral reputation. The reason for this may be that the negative ratings from the malicious agents are weighed too much.

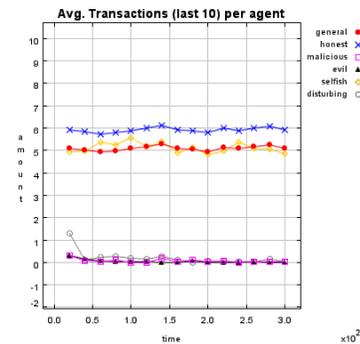
A first surprising result is the strength of the *OnlyLast*-system. This is the only system which can resist an attack from disturbing agents, and can also stand evil and selfish agents. This corroborates the theoretically derived results from [4]. Dellarocas proved that the efficiency of an accumulative system (he uses the term binary feedback system) is independent of the number of ratings n summarized by the mechanism. According to his results an *OnlyLast*-system (which is a binary system with $n = 1$) is more efficient in environments where agents can change their identities without cost. On the other hand it is the weakest systems against malicious agents. The reason for this effect is, that it is very hard for an honest agent, to trade with



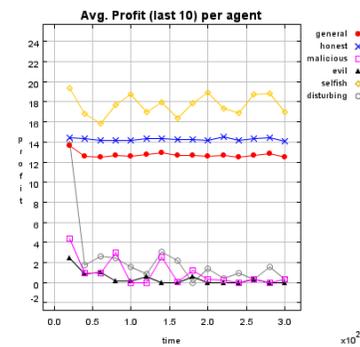
(a) average reputation



(b) reputation bias



(c) transactions



(d) profit

Figure 11. Evaluation Criteria

Table 2. Strengths & weaknesses of the metrics

System	Dist.	Evil	Mal.	Self.
eBay	–	50	60	+
Simple	–	30	70	+
SimpleValue	–	30	70	+
Value	–	50	60	+
Average	–	60	70	+
AverageSimpleValue	–	60	70	+
AverageValue	–	60	60	+
Blurred	–	50	70	+
Blurred-Value	–	50	70	+
OnlyLast	+	+	50	+
OnlyLastValue	+	+	30	+
EigenTrust	–	+	60	30
Sporas	–	30	70	60
YuSingh	–	50	50	60
Beta	–	50	70	+
BetaValue	–	60	70	+

- + : resistant up to 80% of this type
- x : at an amount of x% of this type, the system fails
- : the system does not protect against this attack

others and receive a good rating, if he already has a (false) negative reputation, because a malicious agent rated him negative.

The *EigenTrust*-system is the only system except from the *OnlyLast*-systems which is resistant to attacks from evil agents. But this system is very susceptible to attacks from selfish agents. The reason is that this system depends on groups of agents who trust each other. If only few agents rate other agents, the system fails to recognize groups and cannot compute dependable reputations.

Against attacks from malicious agents as well as selfish agents the most systems are equally good. Just the *OnlyLast* and *YuSingh*-systems cannot handle malicious agents, and the *EigenTrust*-, *Sporas*, and *YuSingh*-systems fail on large amounts of selfish agents. The weakest system in this simulation was the *Sporas*-system which supports only positive ratings and reputations. A possible reason for this is our model for the acceptance behavior. Additionally, it is hard to determine reasonably if the same positive reputation is good or bad without information e.g. about the amount of encounters an agent had.

With this results we tried to combine different metrics to compensate for the individual weaknesses. After some experiments with modifications of the *BetaValue*-system we

focused on the *Average*- and the *OnlyLast*-system. Both systems can be understood as summing up the previous ratings of an agent using different weights. In the *Average* system all ratings have the same weight, in the *OnlyLast*-system the recent rating has the weight 1 and all other ratings have the weight 0. The *Blurred*-system is somewhere in between, but could not handle the disturbing agents. Thus we decided to interpolate between the *Average* and the *OnlyLast*-system by weighting the ratings not linear, like we did in the *Blurred*-system, but quadratic, so that the recent ratings have more influence on the reputation. The resulting metric $\mathcal{M} = (\rho, r)$ is:

$$\rho : A \times E \rightarrow \{-1, 0, 1\}$$

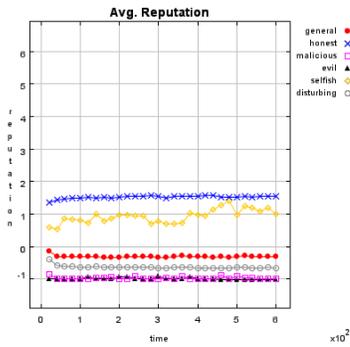
$$r(a) = \sum_{i=1}^{\#(\bar{E}_a)} \frac{\rho(a, \bar{E}_a[i])}{(\#(\bar{E}_a) - i + 1)^2}$$

We call this system *BlurredSquared*. This system is invulnerable to disturbing, evil, and selfish agents. It resists malicious agent up to an amount of 60%. The average reputation at the maximum amount of hostile agents is illustrated in figure 12. In figure 12(a) 80% of all agents are disturbing agents. Similarly, in figures 12(b), 12(c), and 12(d) we have the maximum amount of evil, malicious, and selfish agents. To us the resistance against disturbing agents is far more important than the resistance against an unnatural high amount of malicious agents. Thus with this new system we attenuated the weakness of the *OnlyLast* system against malicious agents a little.

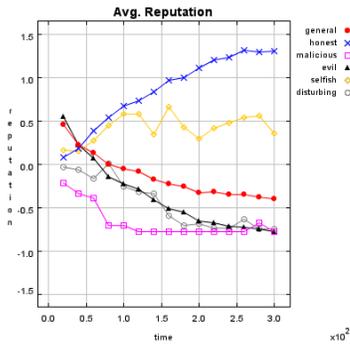
The amounts of hostile agents we are looking at in our simulation seem unnaturally high on first sight. But in the simulation these hostile agents spread their behavior equally over the whole community, while in reality they will focus on certain agents. Thus the amounts we found here would be in the same relation in reality, but much lower, if the aim of hostile agents were to harm individual agents and not the whole community.

5. Related Work

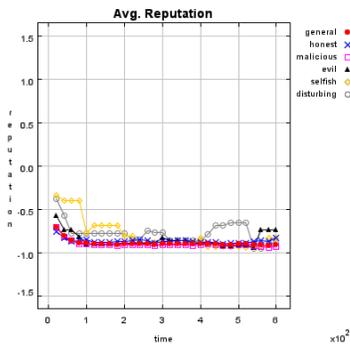
Most papers that propose a new reputation system evaluate its efficiency by simulation, but compare it only with one single system or a scenario without any reputation system. There is no work known to the authors that tries to provide an overall comparison of a number of reputation systems. Dellarocas [4] has provided some theoretical results that are applicable to accumulative systems. Mui has developed a typology of reputation system and developed a model of trust and reputation. However, his understanding of the term reputation is different to ours. Sabater [19] proposes a framework called *SuppWorld* as a test-bed scenario for reputation systems. It is not known if this framework



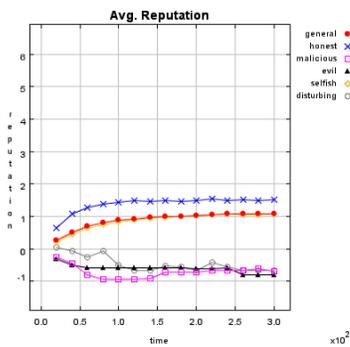
(a) disturbing



(b) evil



(c) malicious



(d) selfish

Figure 12. The *BlurredSquared*-system

has been used for other purposes than to test the ReGret system.

6. Conclusion

We presented a formal model which can be used to describe reputation systems, especially their metric. We gave an overview of the different kinds of metrics in global reputation systems and used our model to describe them. The different metrics were accumulative, averaging, iterative, adaptive, and probabilistic. Furthermore we gave an approach how a trust decision can be computed based on reputation, the scale-based acceptance behavior, and introduced models for the major types of agents related to reputation systems. Based on the model and the agent types we implemented a generic framework for reputation metrics that allowed us to compare arbitrary metrics. We hope that the framework is useful for others to provide a well-defined base for the test of new metrics. By simulation we found the strengths and weaknesses of the different metrics and proposed the *BlurredSquared* metric.

In the future, we plan to extend our model and the simulation framework to include support for local reputation systems and more complex contexts. We also see the need for more sophisticated types of agents and a better understanding of the acceptance function. Finally, a set of standard agent communities that model real-world application domains (e.g. auctions) would be required to enable more specific experiments. We will make the code, a description and additional simulation results available from our web page [8].

References

- [1] S. Buchegger and J.-Y. Le Boudec. A Robust Reputation System for Mobile Ad-hoc Networks. *Technical Report, EPFL, Switzerland*, 2003.
- [2] E. Damiani, S. D. C. di Vimercati, S. Paraboschi, P. Samarati, and F. Violante. A Reputation-based Approach for Choosing Reliable Resources in Peer-to-Peer Networks. In *Proc. of the 9th ACM Conference on Computer and Communications Security*, Washington, DC, USA, November 2002.
- [3] C. Dellarocas. Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior. In *ACM Conference on Electronic Commerce*, pages 150–157, 2000.
- [4] C. Dellarocas. Efficiency and Robustness of eBay-like Online Feedback Mechanisms in Environments with Moral Hazard. *Working paper, Sloan School of Management, MIT, Cambridge, MA*, January 2003.
- [5] eBay Homepage. <http://www.ebay.com>.
- [6] E. Friedmann and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

- [7] T. D. Huynh, N. R. Jennings, and N. Shadbolt. Developing an Integrated Trust and Reputation Model for Open Multi-Agent Systems. In *Autonomous Agents and Multi Agent Systems (AAMAS 04), Workshop on Trust in Agent Societies*, July 2004.
- [8] ITO. Project Reputation Homepage. <http://www.ito.tu-darmstadt.de/projects/reputation>.
- [9] A. Jøsang and R. Ismail. The Beta Reputation System. In *Proceedings of the 15th Bled Conference on Electronic Commerce, Bled, Slovenia, 17-19 June 2002*, 2002.
- [10] R. Jurca and B. Faltings. Towards Incentive-Compatible Reputation Management. In *Trust, Reputation and Security, AAMAS 2002 International Workshop*, volume 2631 of *Lecture Notes in Computer Science*, pages 138 – 147. Springer, 2003.
- [11] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The Eigentrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the twelfth international conference on World Wide Web*, pages 640–651. ACM Press, 2003.
- [12] G. Mahoney. Trust, Distributed Systems, and the Sybil Attack, 2002.
- [13] E. M. Maximilien and M. P. Singh. An Ontology for Web Service Ratings and Reputations. In S. Cranefield, T. W. Finin, V. A. M. Tamma, and S. Willmott, editors, *Proceedings of the Workshop on Ontologies in Agent Systems (OAS 2003) at the 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems, Melbourne, Australia, July 15, 2003*, volume 73 of *CEUR Workshop Proceedings*, 2003.
- [14] L. Mui, M. Mohtashemi, C. Ang, P. Szolovits, and A. Halberstadt. Ratings in Distributed Systems: A Bayesian Approach. In *Proceedings of the Workshop on Information Technologies and Systems (WITS'2001)*, 2001.
- [15] L. Mui, M. Mohtashemi, and A. Halberstadt. Notions of Reputation in Multi-Agents Systems: a Review. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 280–287. ACM Press, 2002.
- [16] RePast Homepage. <http://repast.sourceforge.net>.
- [17] P. Resnick and R. Zeckhauser. Reputation Systems. *Communications of the ACM*, 43:45–48, 2000.
- [18] P. Resnick and R. Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of ebay’s Reputation System. The Economics of the Internet and E-Commerce. *Advances in Applied Microeconomics*, 11, 2002.
- [19] J. Sabater. *Trust and reputation for agent societies*. PhD thesis, Universitat Autònoma de Barcelona, 2002.
- [20] J. Sabater and C. Sierra. REGRET: A Reputation Model for Gregarious Societies. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001.
- [21] J. Sabater and C. Sierra. Reputation and Social Network Analysis in Multi-Agent Systems. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, 2002.
- [22] A. A. Selçuk, E. Uzun, and 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 (CCGrid04)*, 2004.
- [23] A. Whitby, A. Jøsang, and J. Indulska. Filtering out Unfair Ratings in Bayesian Reputation Systems. In *Autonomous Agents and Multi Agent Systems (AAMAS 04), Workshop on Trust in Agent Societies*, July 2004.
- [24] L. Xiong and L. Liu. PeerTrust: Supporting Reputation-Based Trust in Peer-to-Peer Communities. In *IEEE Transactions on Knowledge and Data Engineering (TKDE), Special Issue on Peer-to-Peer Based Data Management*, 2004.
- [25] B. Yu and M. P. Singh. A Social Mechanism of Reputation Management in Electronic Communities. In *Proceedings of the 4th International Workshop on Cooperative Information Agents IV, The Future of Information Agents in Cyberspace*, pages 154–165, 2000.
- [26] G. Zacharia and P. Maes. Trust Management through Reputation Mechanisms. *Applied Artificial Intelligence*, 14, 2000.

A. Simulation Parameters

In this section we present the parameters we used for simulation. Table 3 gives an overview. Varying parameters are marked with *.

Table 3. Simulation Parameters

Parameter	Value
Num Start Agents	*
Add/Remove Agent Rate	1
Add/Remove Probability	20
Percentage*Agents	*
Transaction Probability	75
Reject	*
Search No for Other Evils	10
Graph Refresh Rate	20
Parallel	5

The size of the community at the beginning of every simulation run is set with the parameter *Num Start Agents*. Since the *EigenTrust*-metric needs very much computation-resources, we simulated it with only 50 agents, all other metrics are simulated with a starting size of 200 agents. The natural fluctuation in communities is modeled with the parameters *Add/Remove Agent Rate* and *Add/Remove Probability*. The latter fixes the probability at which an agent enters or leaves the community, the first how many agents may enter and leave at one time. The distribution of the different agent types is set up with the parameter *Percentage*Agents* for each type separately. Table 4 shows how the amounts are varied throughout the simulation to find the critical amounts. The *Transaction Probability* is the probability at which an agent initiates a transaction at a

Table 4. Amounts of different agent types during the simulation in %

vary	honest	others (each)
5	80	5.0
10	76	4.7
30	59	3.7
50	41	2.7
60	33	2.0
70	25	1.4
80	16	1.0

time tick in %. With *Reject* the treshold for the acceptance routine is set. These tresholds vary from system to system and are summarized in table 5. If an evil agent initiates a

Table 5. Tresholds for acceptance

System Type	Reject
Accumulative	0.5
Average	0.7
Blurred	0.5
OnlyLast	0.5
EigenTrust	0.5
Adaptive	0.7
Beta	0.7

transaction he first seeks for other evil agents. With *Search No for Other Evils* it can be set how often an evil agents tries until he makes a transaction with every agent type. The last two parameters are important for the evaluation of the simulation. With *Graph Refresh Rate* it is set how often the graphs shall be refreshed, *Parallel* sets the number of parallel simulation runs with the same parameter setting to interpolate between.

B. Simulation Graphs

In this section we have a closer look at some further graphs from the simulation. For each kind of metric we produce a figure like figure 12 with the respective critical amount of hostile agents. Due to spatial reasons we omit figures for the average and the blurred metric. The figures for the average metric are quite similar to those of the beta metric, the blurred metric produces similar graphs like the blurred squared metric. More graphs than we can present here will be available on the web [8].

B.1. Accumulative System

Figure 13 shows the reputation for the critical amount of hostile agents for the *eBay*-system. Since this metric cannot resist disturbing agents, the amount of them does not matter. Even at 5% they can build up their reputation and cheat the others. A high amount of selfish agents just causes a slower development of the reputation value, but does not influence the relation between the reputation of the different types, thus selfish agents are not dangerous to this metric. The evil agents have a slightly lower reputation at 50%, but in fact their reputation is quite good. At 60% they corrupt the system, the honest agents have only low reputation and the evil agents are the only with high reputation. The same border is valid for the malicious agents. At 50% the reputation develops quite normally, but at 60% the honest agents can hardly get a positiv reputation.

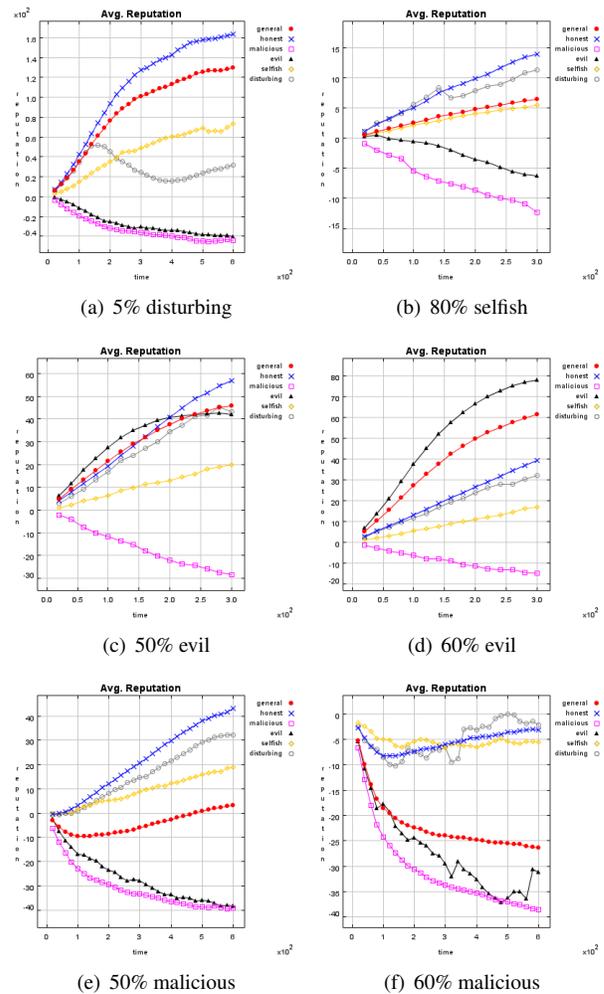


Figure 13. The *eBay*-system

B.2. OnlyLast System

Figure 14 shows one of the best systems in the simulation. Even at the high amount of 80% disturbing agents they are not able to reach a high reputation. They have the same bad reputation as evil and malicious agents and thus are not able to make any profit. A high amount of evil agent causes that the reputation for honest agents develops slower, but all the time the honest agents have a positive and the evil agents a malicious-like reputation. Many self-ish agents do not influence the effectiveness of the system at all. The big weakness of this metric is the susceptibility against malicious agent. Even at an amount of 30% the honest agents can just reach a neutral reputation. If there are more malicious agents, all agent in the community have the worst reputation after some time.

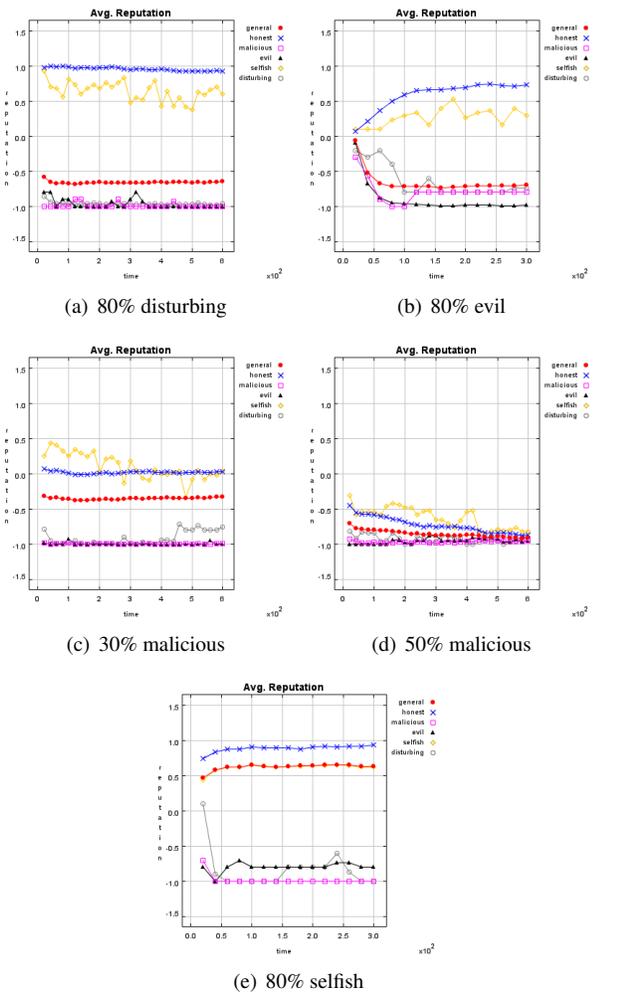


Figure 14. The *OnlyLast*-system

B.3. EigenTrust System

The behavior of the *EigenTrust*-system is depicted in figure 15. As many other metrics this metric cannot prevent attacks from disturbing agents. But this system is absolutely resistant against evil agents. Even at an amount of 80% they are not able to reach a nearly good reputation like the honest agents. If 50% of the community are malicious agents, it takes quite a long while till this metric recognizes honest agents. If the amount increases, the system fails completely. All agents have a reputation like the malicious agents. The biggest weakness of this metric are attacks from selfish agents. This system depends highly on the feedback of the participating agents. An amount of 30% selfish agents is enough to make this system absolutely useless.

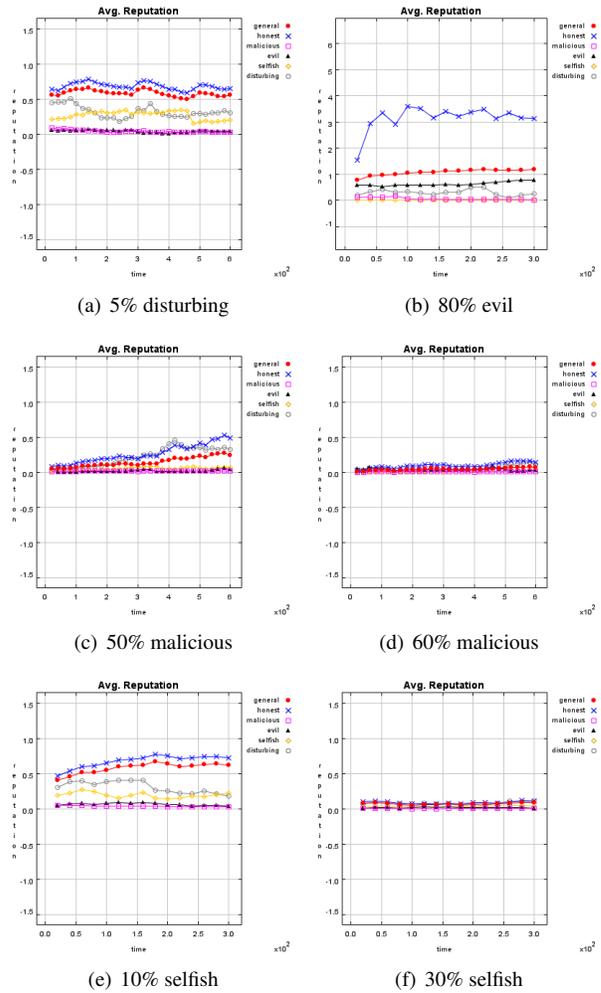


Figure 15. The *EigenTrust*-system

B.4. Sporas System

Figure 16 illustrates the result of attacks on the *Sporas*-system. This system cannot prevent attacks from disturbing agents. They can reach such a high reputation that the other agents trust them. A high amount of selfish agents does not make the system completely useless, but it takes too long to get differentiateable reputation values. The same argument counts for a high amount of malicious agents. If there are 70% malicious agents in the community they cannot pull down the reputation of the honest agents, but they can prevent the honest agents from rising their reputation quite fast. An absolute weakness of this metric is its susceptibility against evil agents. It can only handle an amount of 10%. Even 30% evil agents are able to push their reputation to the same the honest agents have.

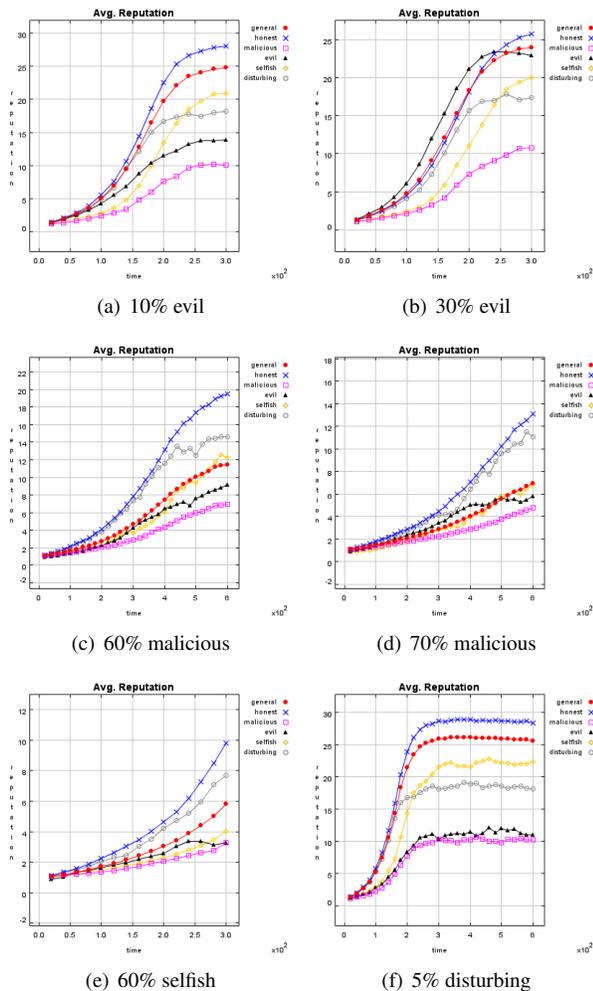


Figure 16. The *Sporas*-system

B.5. BetaValue System

Figure 17 shows the effects of high amounts of hostile agents on the *BetaValue*-system. Disturbing agents can always achieve that the other agents trust them and have high reputation. If there are many selfish agents in the community it just takes a little longer to compute adequate reputations, but this does not influence the system much. At an amount of 60% evil agents, evil and honest agent have the same middle reputation. Thus the reputation values get unreliable. At 60% malicious agents the honest agents can get a positive reputation after a while although they start with a negative one. If the amount of malicious agents increases, all agents have negative reputation yet the reputation of the honest agents is rising. But this happens unacceptable slowly.

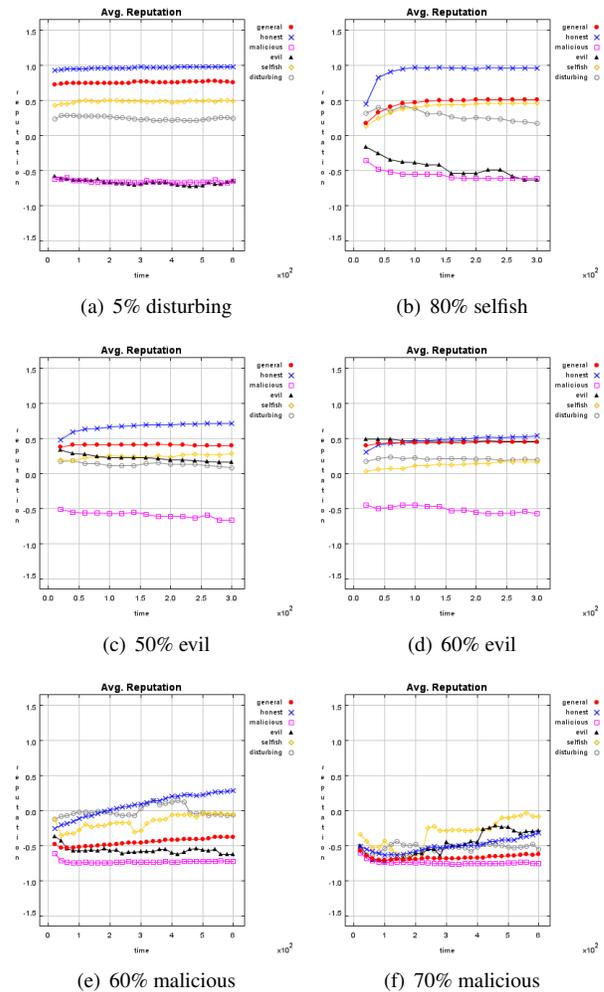


Figure 17. The *BetaValue*-system