

Effective Use of Reputation in Peer-to-Peer Environments^{*}

Thanasis G. Papaioannou and George D. Stamoulis

Department of Informatics, Athens University of Economics and Business (AUEB),

76 Patision Str., 10434 Athens, Greece.

{pathan, gstamoul}@aueb.gr

Abstract

Peer-to-peer environments have become popular as a framework for exchange of services. In these environments, certain peers may fail to provide their services. Reputation can be a proper means of discovering low-performing peers, without affecting significantly inherent characteristics of Peer-to-Peer environments, such as anonymity and privacy. However, the accurate calculation of the reputation metrics may not be sufficient to provide the right incentives to peers. In this paper, we show that the straightforward approach for peers to exploit the reputation metrics (i.e. by just selecting as a providing peer the one with the highest reputation) may lead to unexpectedly low efficiency for high-performing peers. We argue and justify experimentally that the calculation of the reputation values has to be complemented by reputation-based policies that define the pairs of peers eligible to interact. We introduce two orthogonal dimensions constituting the reputation-based policies: “provider selection” and “contention resolution”. We argue and show by means of simulation experiments that both these dimensions have a significant impact to the achieved efficiency of the peers. We also investigate experimentally the achievable efficiency of specific reputation-based policies for the case of short-lived peers of two different fixed-strategy types. Finally, we deal with the efficient computation of the reputation value by means of aggregation of the ratings’ feedback provided by the peers. We propose that this can be accomplished by aggregating only a small randomly selected subset of this feedback. Simulation experiments indicate that this approach indeed leads to the fast and accurate calculation of the reputation values even if the peer-to-peer population is renewed with a high rate.

1. Introduction

Peer-to-peer environments gain increasing acceptance in the information society as an overlay framework for exchanging services. The value of such a service for the client peer depends on the performance of the peers providing it. In fact, a peer may provide services with a low performance level. The reason for this can be either the peer’s hidden type (i.e. strategy) or his hidden quality (i.e. performing ability); however, since both of these reasons have the same effect, we do not distinguish between peers falling to the former or the latter case. Reputation can be a proper means of revealing low-performing peers in electronic environments, if it is calculated accurately [1], [3]. However, the accurate calculation of the reputation value by itself may not be adequate as a mechanism to improve the achievable efficiency of high-performing peers and to provide the right incentives for peers to offer services of high quality. Indeed, the straightforward approach for peers to exploit the reputation metrics is to just select as a providing peer the one with the highest reputation value. In this paper, we show that this approach leads to unexpectedly low efficiency for high-performing peers when no other incentive mechanism than reputation is employed in the peer-to-peer system. We argue and justify experimentally that the calculation of the reputation values has to be complemented by reputation-based policies that define the peers eligible to interact with each other. Two orthogonal dimensions of reputation-based policies are then introduced: “provider selection” and “contention resolution”. We show by simulation experiments the impact of reputation-based policies to the peers’ efficiency and analyze the incentives provided to peers by each such policy. Our objective is to differentiate the quality of service received by the various peers from others, depending on how much each peer contributes to the

^{*} The present work was partly funded by the IST project MMAPPs (IST-2001-34201).

overall provision of services. Through reputation-based policies, a cycle including both reputation calculation and exploitation of reputation is formed. We argue and justify experimentally that this reputation cycle greatly affects the speed and the accuracy of convergence of the reputation values to the real hidden information. Further simulation experiments were conducted on the efficiency of various reputation-based policies in cases of short-lived peers that follow fixed strategies. These experiments show the respective impact of each reputation-based policy on the efficiency of the Peer-to-Peer system. We also deal with the communication overhead for aggregating the ratings' feedback in a distributed way in order to accurately calculate a reputation value. We propose a method to reduce this overhead by aggregating just a small randomly selected subset of the complete ratings' information. We show experimentally that this aggregation approach does not essentially lead to a degradation of the speed and the accuracy of the calculation of the reputation values. Moreover, the simulation experiments indicate that this applies, even if the population of the Peer-to-Peer system is renewed with a high rate. Finally, we discuss the implementability of the various reputation-based policies. The remainder of this paper is organized as follows: In Section 2, Bayesian and Beta aggregation functions are described. In Section 3, the straightforward approach of using reputation in a Peer-to-Peer system is introduced. In Section 4, some reputation-based policies are defined. In Section 5, the cycle of reputation information is presented. In Section 6, randomized aggregation of the ratings' feedback is proposed. In Section 7, we present results on the effectiveness of the various reputation-based policies and the randomized aggregation of the ratings' feedback introduced in the paper. In Section 8, we analyze some implementation issues related to the proposed use of reputation in a Peer-to-Peer system. Finally, in Section 9, we conclude our work.

2. Accurate Calculation of Reputation

Peers according to their type, their inherent capabilities and/or their strategy, succeed or fail in offering services to other peers. After observing the *outcome* of his transaction, a client peer rates the providing one for his performance. Throughout the paper, we assume that peers truthfully report their evaluations for the performance of other peers. A mechanism ensuring truthful reporting is proposed in [7]. Actually, the outcome of their transactions is only

of interest to peers, rather than the hidden cause for this outcome. It has been documented [3], [4] that binary rating (i.e. success vs. failure) is appropriate for calculating the reputation value that expresses the expected outcome of the transaction with a specific peer. The aggregation of all the history of a peer's outcomes into a single reputation value is important for performance reasons (e.g. storage overhead is reduced). According to [4], *Bayes' rule* is an efficient aggregation function if there is defined an initial belief on the success probability of each type of peers and the proportions of the population of peers that belong to each type. If peers follow dynamic strategies over time and change their probability of success, then the fraction of the number of successful service provisions over the total service provisions of a peer could be used with more weight being placed to the recent history. This approach is called *Beta aggregation*, and was introduced in [3]. Specifically, it is described by the formulas below.

$$R = \frac{s'}{t'}$$

where $s' = s \cdot d + \mathbf{1}(\text{success})$ and $t' = t \cdot d + 1$

R is the new reputation value of the peer; s is the sum of as many unit as the number of the previous successes, with each unit term being discounted exponentially in the time distance from the present; t is the discounted number of services he provided; s' , t' are the updated values of s , t after a new service provision by the peer; $\mathbf{1}(\cdot)$ is the indicator function. Finally, d is a discount factor denoting the relative importance of the past history of ratings over against the recent one. Both Bayesian and Beta aggregation functions are used for the calculation of the reputation values of peers in the simulation experiments of Section 7.

3. Straightforward Approach of Exploiting Reputation

Assume the existence of a reputation system in the Peer-to-Peer services environment that accurately calculates a reputation value for each peer; this value reflects his probability to successfully provide his service the next time requested. Throughout the paper, we assume that reputation values are safely stored by the reputation system. How would peers use these reputation values? Reputation limits the risk for a peer to fail in a service transaction. Consider now a peer requesting a service. We assume that each such peer pursues self-interest, i.e. aims to maximize his obtained utility from service transactions. If all successfully provisioned services have the same value for the peer

considered, then he would select among other peers that offer the requested service the one that has the maximum reputation value. On the other hand, he is indifferent in selecting among requesting peers whom to serve with his limited resources, if no other mechanism than reputation is in place. The aforementioned approach is towards selecting transacting peers, in our opinion, the straightforward approach for employing reputation.

However, that way it appears that a high-performing peer is punished in two ways: a) the higher reputation a peer has, the more users he attracts to consume his own resources, and b) such a peer receives equal benefit from the peer-to-peer system as other peers that have a low performance level. Clearly, this straightforward response of peers to exploit the reputation metrics provides wrong incentives to both high- and low-performing peers. A high-performing peer is motivated to lower his performance, while a low-performing peer is motivated to keep his performance to the same level and continue to free-ride. These incentives lead to a market of “lemons”, and possibly to the gradual decomposition of the peer-to-peer system.

Thus, an appropriate reputation-based policy that changes this default response of peers has to be introduced in the peer-to-peer services environment. A reputation-based policy that also assigns higher benefit to higher performing peers provides the incentives for peers to improve their performance. We classify the reputation-based policies into two dimensions: “provider selection” and “contention resolution”. The former concern the selection of the providing peer among those offering the same service, while the latter concern the selection among the peers requesting for a service of the one to be served by the providing peer, who has limited resources. The assumption of a peer’s limited resources is a realistic one and can be related with link capacity, CPU time, etc.

4. Reputation-Based Policies

We have already defined the two “orthogonal” dimensions of rewarding and punishing policies: provider selection and contention resolution. In this section, we present certain potential policies for each dimension.

4.1. Provider Selection Policies

Highest Reputation: If a reputation metric for performance is existent in the Peer-to-Peer system, the most straightforward policy that each peer reasonably

follows (in absence of other policies) is to select among peers that provide the requested service the one with the highest reputation value. This policy was used for experiments on reputation in [6].

Comparable Reputation: A policy named “Peer-Approved” was studied in [2]. According to that policy, peers can download files only from other ones with lower or equal rating. This policy increases the probability for a peer that improves his performance (and thus his reputation value) to find the services requested. However, his received quality is questionable, as he may select services from lower reputed-peers. We propose a different policy (referred to as “Comparable Reputation”), whereby peers are able to request services only from peers that have reputation values comparable to theirs, i.e. within a pre-specified distance. The underlying idea of this policy is the matching of the performance level offered by a peer with the performance level provided to him. Thus, this policy results in *layered communities*, that is, services of similar quality are exchanged among peers of the same layer. The quality of offered services is high in the top layer if there are high-performing peers in the peer-to-peer system, while in the bottom layer the services offered are in most cases useless or even harmful for other peers.

Black List: This policy extracts from the Peer-to-Peer system peers that have a low performance, or equivalently that have reputation values below a certain threshold. Thus, peers offering services of low quality consistently for a certain period are excluded from the set of eligible providing peers. Thus, this policy improves the quality offered to the remaining peers.

4.2. Contention Resolution Policies

Highest Reputation: According to this policy, the peer with the highest reputation value is the one selected to be served by a peer among those simultaneously requesting a service from the latter and thus contending for his resources. This policy assigns absolute priority to peers with the highest reputation values. (Ties are resolved by resolved by means of a randomized symmetric rule). Using this reputation-based policy, a high-reputed peer is very likely to be provided the service even when he is contending with others. However, the outcome of the service provision depends on the provider selection policy that is employed in the Peer-to-Peer system. If the Highest Reputation policy is in use and there is a lot of contention for resources, then peers with low reputation values will not be able to be offered any services.

Probabilistically Fair w.r.t. Reputation: According to this policy (to be referred to as “Probabilistically Fair”), the peer to be served is selected according to the following rule: among the peers j that request the same service from a particular peer, the probability for each one peer i to be selected equals $r_i / \sum_j r_j$, where r_j is the reputation of the peer j . Note that for a highly-reputed peer that contends for the resources of another peer with a low-reputed peer the probability to be selected under this contention resolution policy is close to 1, similarly to the highest reputation contention resolution policy. On the other hand, under the present policy, peers with low reputation values have a small yet positive probability of receiving some services regardless of the contention level. Also, note that, in the case where all peers contending for a certain resource are equally reputed, the two contention resolution policies coincide.

4.3. Discussion

The contention resolution policies determine the probability with which a peer is offered a service, in the presence of contention. Note that in the absence of contention, the contention resolution policies have no effect. On the other hand, provider selection policies determine the probability with which a peer is served successfully, i.e. he is offered services of a high performance. If provider selection and contention resolution policies are employed jointly in the Peer-to-Peer system, then the expected probability of success for a service request of a peer is obtained by the multiplication of his selection probability (i.e. probability to be served) with the success probability of the providing peer, each resulting from the reputation-based policies. Depending on the specific kind of services that are offered in the Peer-to-Peer environment and the relative importance between the quality and the quantity of services exchanged the proper *pair* of provider selection and contention resolution policies should be employed in the Peer-to-Peer system. The various pairs of policies are evaluated experimentally in Section 7.

5. Reputation cycle

Reputation-based policies determine the pairs of peers that are eligible to interact. Recall that client peers rate the providing ones regarding the performance of the latter in their transactions with the former. This feedback is sent to the reputation system after a number of completed transactions. The reputation values of the respective peers are updated

based on this feedback. However, these updated values determine the new pairs of peers that are eligible to interact. These interactions will result in additional ratings’ feedback, etc. Thus, a cycle of reputation information is formed, when reputation-based policies are employed in the Peer-to-Peer system. Each pair of reputation-based policies determines the evolution of reputation information to the hidden true information, and thus determines the specific transactions, the number of successfully obtained services and the ratings involved in these transactions.

6. Randomized Aggregation

In this section, we deal with the efficient aggregation of ratings in terms of communication overhead in a Peer-to-Peer system in the absence of a central authority. In several approaches [5], [6] some peers are responsible for holding and providing upon requests the reputation information of one or more other peers. In both approaches, peers that store reputation information for another peer are determined by a number of hash functions that map the identifier of that peer to the identifiers of his reputation holders in a Distributed Hash Table (DHT) space. Reputation values are calculated using proper functions for the aggregation of votes like those described in Section 2. For the proper update of the reputation information a number of messages containing ratings’ feedback is sent to each reputation holder over time. The votes sent by a peer are associated to his transacted peers. Thus, if λ is the mean rate according to which a peer is served, then the number of feedback messages per unit time that have to be sent to the reputation holders is proportional to λ .

The set of messages required to be sent for the proper update of reputation information induce a significant traffic overhead to the underlay network of a Peer-to-Peer system. This overhead could be reduced by aggregation of ratings prior to their submission to the corresponding reputation holder(s) for the same providing peer within a time period. However, the achieved reduction would in general not be significant, because transactions of a peer with the same provider peer within reasonably small time periods are rare. Thus, the number of different feedback messages required is in fact close to λ .

We propose that a peer submits only a small randomly selected subset p of his ratings. Thus, the number of feedback messages sent is reduced to $p\lambda$. In particular, a potential implementation would be, after a transaction, a vote to be sent to the corresponding reputation holder with probability p . Another potential

implementation would be to send the votes aggregated in a time period to the reputation system with probability p . However, it is reasonable to ask: how much this reduction of feedback affects the accuracy of the reputation values of peers and the effectiveness of the reputation mechanism through reputation-based policies? What are the values of p that induce small losses in the efficiency of the Peer-to-Peer system while achieving considerable reduction of the communication overhead? Experimental results (described in Subsection 7.4) indicate that the accuracy of reputation information remains high even for very small values of p (~10%), as depicted in Figure 10. Furthermore, the efficiency losses induced for peers by the improper operation of the reputation-based policies due to the missing feedback of ratings are limited even for these small values of p , as depicted in Figure 12.

7. Simulation Experiments

7.1. The Model

Consider a peer-to-peer system where services of a certain kind are exchanged among peers. In this Peer-to-Peer system, there are two types of peers: *altruistic* and *egotistic*. The type of each peer is private information, i.e. it is known only to the peer himself. Each peer follows a mixed strategy regarding his performance in his service provisions that depends on his type. Specifically, each altruistic peer provides his service successfully with a high probability α , while an egotistic one succeeds in each of his service provisions with a low probability β . Furthermore, the Peer-to-Peer system is assumed to be *renewed* according to a Poisson distribution with mean rate $\lambda=10$ peers/time slot; the total size of the population is kept constant. In particular, each peer is assumed to live in the Peer-to-Peer system for a period determined according to the exponential distribution with mean rate λ/N , where N (=1000) is the total number of peers in the Peer-to-Peer system. When a peer leaves the Peer-to-Peer system a new entrant of the same type takes his place.

A distributed reputation system is employed. The reputation value for a peer is associated to his pseudonym. Each peer sends feedback to the reputation system on the performance of other peers based on the outcome of the services provided by the latter to the former. Randomized aggregation is not employed, except for the experiments presented in Subsection 7.4. Also, the Peer-to-Peer system is considered noiseless in the sense that the outcome of a transaction depends only on the performance of the providing peer in this transaction. A peer is assigned

an initial reputation value h_0 that expresses the probability that he is of the altruistic type. This initial probability is considered small ($h_0=0.1$), in order to limit the incentive for *easy name changes*. That is, if h_0 were high, then each peer would have the incentive to drop his pseudonym and obtain a new one, thus clearing his potential low-performance history. The votes are converted into reputation values using Bayes' rule (see Section 2). It is important to note that the same experiments were also conducted using Beta function for aggregating the votes into reputation values with similar results.

Time is assumed to be slotted. At each slot, each peer requests a service with a certain probability r . Service availability is distributed in the experiments described in this section according to Zipf distribution (i.e. assuming that services are ranked w.r.t. their popularity, then a service with rank x is found at a peer with probability x^{-k} , $k \geq 1$) in the Peer-to-Peer system. It is important to note that the experiments presented in this section were also conducted having services Uniformly distributed with similar results. A peer can serve only one peer at each slot due to his limited resources. Denial of service due to limited resources in our model, corresponds to realistic situations of temporary unavailability of the service, e.g. due to congestion in the network or in a server. Reputation-based policies can be employed in the Peer-to-Peer system. In the absence of reputation-based policies peers select their transaction party according to the straightforward approach described in Section 3. The *efficiency* in this Peer-to-Peer system can be measured as the average ratio of successful transactions for an altruistic peer over either i) the average number of service requests or ii) the average number of initiated services. Only successful transactions provide value for client peers. The total value provided to a peer should be in accordance to his performance, in order the right incentives for performance to be provided to him. Thus, altruistic peers should be offered high value by the Peer-to-Peer system, in order to stay and keep on offering value to the system. On the other hand, egotistic peers offer low value to the Peer-to-Peer system, and thus the value provided to them is not an important measure for efficiency. Nevertheless, on a per peer basis, this value should be lower than that provided to altruistic peers. The efficiency is affected by the fast and accurate revelation of the hidden type of peers, which in turn depends on the reputation-based policies employed.

7.2. Effectiveness of Reputation-Based Policies

We experiment on the effect of the various reputation-based policies in the efficiency of this Peer-to-Peer system. In all experiments, high-performing (i.e. altruistic) peers are considered to constitute a small subset (namely, 10%) of the total population of peers. If no reputation system were employed in the peer-to-peer environment, then the type of the peer would remain unknown, peers would select their providing peers randomly and the contention would also be resolved randomly. In this case, the cumulative ratio of successfully obtained services over the total requested services for the two types of peers is depicted in Figure 1. Recall that services are Zipf distributed. Notice that, in Figure 1, the inefficiency for an altruistic peer is even greater than that for an egotistic one. This happens because, due to the relatively large number of egotistic peers in the Peer-to-Peer system, the contentions for a service at a peer are resolved in favor of egotistic peers with higher probability per peer than that of altruistic peers.




Figure 1. Efficiency obtained in the absence of a reputation system.

In our simulation model, there is a distributed reputation system employed that accurately calculates the reputation values of the peers. In the absence of reputation-based policies, each peer is assumed to follow the straightforward approach of using reputation (see Section 3). The curve denoted “max-random” in Figure 2 shows the average cumulative ratio of successfully obtained services over the total number of initiated transactions by an altruistic peer following this straightforward approach. (Henceforth, in all figures “max” stands for Highest Reputation policy.) Clearly, this success ratio is greatly improved for altruistic peers, when reputation-based contention resolution policy is employed, as depicted by the other two curves of Figure 2. On the other hand, Figure 3 shows that this success ratio is decreased for egotistic peers when a reputation-based contention resolution policy is employed in the Peer-to-Peer system. Notice also, in Figures 2 and 3, that the success ratio of an

altruistic peer is lower than that of an egotistic peer when no reputation-based contention resolution policies are employed. Thus, the straightforward approach of using reputation is not only inefficient, but also unfair.




Figure 2. Efficiency of the Highest Reputation provider selection policy, when it is jointly enforced with various contention resolution policies for altruistic peers.




Figure 3. Efficiency of the Highest Reputation provider selection policy, when it is jointly enforced with various contention resolution policies for egotistic peers.




Figure 4. Efficiency of the Comparable Reputation selection policy, when it is jointly

enforced with various contention resolution policies for altruistic peers.

The success ratio is improved (resp. deteriorated) for altruistic (resp. egotistic) peers in the case where the providing peers are selected according to the Comparable Reputation policy for all contention resolution policies, depicted from Figure 4 (resp. Figure 5) as compared with Figure 2 (resp. Figure 3). Also, notice that all presented contention resolution policies (including random selection) when this particular provider selection policy is employed achieve similar success ratios for peers of the same type. This is reasonably expected, since Comparable Reputation policy constrains the contention only among peers having similar reputation values. Thus, this provider selection policy can be employed efficiently without being combined with a reputation-based contention resolution policy, which in some sense is accomplished by this provider selection policy itself! The resulting success ratio for altruistic peers in the case where Black List selection policy is employed in the Peer-to-Peer system is depicted in Figure 6 for the various contention resolution policies. Again, the achieved success ratios of the Highest Reputation and the Probabilistic Fair contention resolution policies are close for both types of peers. Thus, this conclusion applies for all the proposed provider selection policies, indicating that Highest Reputation and Probabilistic Fair contention resolution policies are almost equally efficient.




Figure 5. Efficiency of the Comparable Reputation selection policy, when it is jointly enforced with various contention resolution policies for egotistic peers.




Figure 6. Efficiency of the Black List Reputation selection policy, when it is jointly enforced with various contention resolution policies for altruistic peers.




Figure 7. Success ratio of altruistic peers, when most efficient reputation-based policies are jointly employed.

Next, we investigate the achieved efficiency of the various provider selection policies, where each of them is employed jointly with the contention resolution policy that maximizes the achieved efficiency. In this perspective, the achieved success ratio of each provider selection policy for altruistic peers is depicted in Figure 7. Observe that the Comparable Reputation selection policy outperforms the other two policies in terms of success ratio. This can be explained as follows: The Comparable Reputation policy has the effect that a peer requests services only from peers of his type, i.e., an altruistic peer requests for services only from other altruistic ones.




Figure 8. Throughput of successes for peers, when most efficient reputation-based policies are jointly employed.

On the other hand, the provider selection policies perform differently w.r.t. the ratio of successfully acquired services over the total service requests (as opposed to the total initiated services), as depicted in Figure 8. Note that, the number of the total service requests is the same for all policies, which is not the case with the number of the total initiated services. As depicted in Figure 8, the Highest Reputation provider selection policy outperforms the other such policies. This was expected, since under the Highest Reputation policy, for each particular peer, the set of potential providers is the whole Peer-to-Peer population. This does not apply for either the Black List or the Comparable Reputation provider selection policies. In our experiments, the set of altruistic peers is small (100 peers) and thus the resulting service availability is limited for these two policies. Clearly, if the set of peers of each type is large, and thus the probability for the requested services to be offered by a provider within this set is high, then the only important parameter of efficiency is the success ratio. The Comparable Reputation provider selection policy achieves the highest efficiency in this case. On the other hand, in Peer-to-Peer systems where only a small number of peers belong to the high-performing type, Highest Reputation is the best alternative. In certain cases of services the most important efficiency parameter for peers is the ratio of successfully provided services over the total service requests (throughput of acquired services); in other cases the most important efficiency parameter is the ratio of successfully provided services over the total number of initiated services. An example of such a service is the case of sharing amusing content files among peers.

7.3. Short-lived Peers

In this subsection, we analyze the efficiency of reputation-based policies for very short-lived peers. Specifically, we consider that the Peer-to-Peer system is renewed with a rate $\lambda=30$ peers/time slot, i.e. 3 times as fast as in the previous experiments. Recall that, according to our simulations model, when a peer leaves the Peer-to-Peer system a new peer of the same type enters into the Peer-to-Peer system. This is similar to having a peer of a specific type drop his pseudonym (and the reputation value associated with it) and re-enter the Peer-to-Peer system under a new one and a clear record. Also, this particular model is similar in its effect on the peers' reputation values with cases where peers modify dynamically their mixed strategies. In case of short-lived peers, the cycle of reputation information is short and the effectiveness of each reputation policy is tested under more "tight" conditions.




Figure 9. Success ratio of short-lived altruistic peers, when most efficient reputation-based policies are jointly employed.

In Figure 9, depicted are the achieved success ratios of altruistic peers for each provider selection policy jointly applied with the contention resolution policy that maximizes its efficiency. The Comparable Reputation provider selection policy still achieves the highest success ratio. Furthermore, the success ratio of the Black List policy is lower than that of the Highest Reputation policy. This is reasonable, as in case of short-lived peers the Black List provider selection policy has limited effect, as many low-performing peers may not be revealed and be selected as providing ones.

7.4. Using only small subsets of feedback information

Next, we experiment on the accuracy and the speed of convergence of the reputation, when only randomly

selected subsets of the complete ratings' information are used for the calculation of the reputation values. The population is renewed according to a Poisson distribution with mean rate $\lambda=10$ peers/time slot. In Figure 10, depicted are the average reputation values of altruistic and egotistic peers after 1000 time slots, as a function of the fraction p of submitted ratings. Observe that the two types of peers are adequately differentiated w.r.t. their reputation values even for very small values of $p \sim 10\%$.

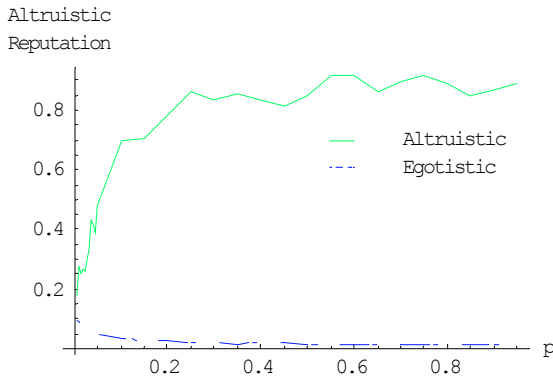


Figure 10. Evolution of convergence of reputation values for altruistic and egotistic peers as p increases.

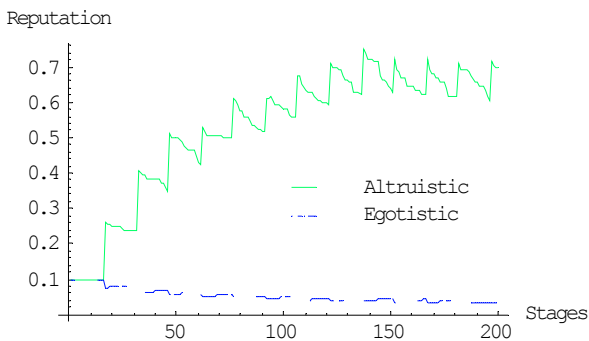


Figure 11. Reputation converges fast.
 Convergence is fast in time, as depicted in Figure 11 for $p=10\%$, where the selected ratings are first aggregated locally for 15 time slots and then submitted to the reputation system. In Figure 12, depicted are the achieved success ratios for altruistic and egotistic peers, if the Highest Reputation provider selection policy is jointly applied with the Probabilistically Fair contention resolution policy, as functions of the fraction p of submitted ratings. Again, the achieved success ratio converges for small values of p ($\sim 10\%$) to the value achieved using the complete ratings' information for the case of the same reputation-based pair of policies. Thus, the communication overhead for employing a reputation system in a Peer-to-Peer system can be decreased using this randomized

approach by an order of magnitude. The communication overhead can be further decreased, if each peer aggregates locally his ratings for a short period (e.g. 15 time slots) and then submits them to the distributed reputation system.




Figure 12. Evolution of the success ratio for altruistic peers as p increases.

8. Implementation Issues

The revelation of the peers' hidden information of interest (i.e. their type) can be achieved by means of reputation without sacrificing much of the desirable inherent characteristics of peer-to-peer systems, such as anonymity by attaching the reputation values to pseudonyms. However, ratings provide feedback for the complete history of transactions of a peer to reputation holders. Thus, privacy concerns are raised. Our proposed randomized aggregation approach limits these concerns, as feedback about a small random part of the total history of the transactions of a peer is provided.

Another potential concern is whether reputation-based policies are employable in a real Peer-to-Peer system, where there exist peers that can even "hack" their part of the Peer-to-Peer middleware (attempting to override the reputation-based policies), if they can gain in efficiency by doing so. Below, we discuss the employability of reputation-based policies in a smooth incentive compatible way. However, contention resolution policies have no direct impact to the providing peer himself if no external incentive mechanism (e.g. money-related) is in place for peer transactions. Thus, these policies can be pre-configured in the peer's part of the Peer-to-Peer middleware, as there is no incentive to change this configuration. Moreover, Comparable Reputation provider selection policy is effective regardless of the employment of a contention resolution policy (see Subsection 7.2). Regarding the provider selection policies, Highest Reputation is an incentive compatible

one for peers to follow, although it has to be combined with a contention resolution policy in order to be effective. Black list can also be easily applied, storing a warning flag in the reputation holder(s) of each peer that belongs to the black list, in order for other peers to avoid transactions with him. Comparable Reputation policy is trickier to apply. Below, we propose a related approach: We safely assume that peers will always tend to select among the providing peers the one that has the highest reputation. Consider that the Peer-to-Peer environment is divided into disjoint groups of peers, each constituting an independent system, in the sense that peers in a group cannot transact with peers of another group. (This should be enforced by the middleware). Each peer uses a unique pseudonym in the Peer-to-Peer environment. Each Peer-to-Peer group contains peers of different performance levels. New entrant peers in the Peer-to-Peer environment become members of the Peer-to-Peer group that has the lowest performance level. Peers are moved by the middleware across Peer-to-Peer groups w.r.t. their own performance level reflected by their reputation.

9. Concluding Remarks

In this paper, we have shown that the straightforward approach for peers to exploit the reputation metrics in peer-to-peer systems leads to unexpectedly low efficiency for high-performing peers. We have indicated and evaluated experimentally in terms of achieved efficiency specific reputation-based policies that define the peers eligible to interact based on reputation values. Comparable Reputation policy has the highest achievable efficiency among the reputation-based policies introduced in this paper. We also proposed a randomized approach for the aggregation of the ratings' information. A small randomly selected subset of the ratings' feedback is sufficient information for the fast and accurate calculation of the reputation values, even if the Peer-to-Peer population is renewed with a high rate. Throughout the paper we have assumed that peers report their ratings truthfully. How this can be enforced is the subject of [7]. We also intend to investigate applicability and efficiency of our

reputation-based policies in an e-commerce environment, where peer transactions also involve money.

10. Acknowledgements

We thank Costas Courcoubetis, Huw Oliver, Ben Strulo and the other members of the consortium of IST project MMAPPS for useful discussions on the subject of this paper.

11. References

- [1] C. Dellarocas. "Efficiency through feedback-contingent fees and rewards in auction marketplaces with adverse selection and moral hazard". In Proc. of the *3rd ACM Conference on Electronic Commerce (EC-03)*, June 9-12, 2003, San Diego, CA, USA.
- [2] K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. "'To Share or not to Share' An Analysis of Incentives to Contribute in Collaborative File Sharing Environments". In Proc. of the *Workshop on Economics of Peer-to-Peer Systems 2003*, Berkeley, CA, June 2003.
- [3] A. Jøsang, S. Hird, and E. Faccer. "Simulating the Effect of Reputation Systems on e-Markets". In Proc. of the *1st International Conference on Trust Management*, Crete, May 2003.
- [4] C. Dellarocas. "Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard". *MIT Sloan Working Paper No. 4297-03*, January 2003.
- [5] K. Aberer, and Z. Despotovic. "Managing Trust in a Peer-2-Peer Information System". In Proc. of the *10th International Conference on Information and Knowledge Management (CIKM01)*, New York, November 2001.
- [6] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. "EigenRep: Reputation Management in Peer-to-Peer Networks". In Proc. of the *WWW2003*, May 2003.
- [7] T. G. Papaioannou, and G. D. Stamoulis. "Enforcing Credible Reporting in Peer-to-Peer Environments". Working paper, AUEB, January 2004. Available at <http://nes.aueb.gr>