Noname manuscript No.
(will be inserted by the editor)

A Mechanism that Provides Incentives for Truthful Feedback in Peer-to-Peer Systems

Thanasis G. Papaioannou ·
George D. Stamoulis

Received: date / Accepted: date

Abstract We propose a mechanism for providing the incentives for reporting truthful feedback in a peer-to-peer system for exchanging services (or content). This mechanism is to complement reputation mechanisms that employ ratings' feedback on the various transactions in order to provide incentives to peers for offering better services to others. Under our approach, each of the transacting peers (rather than just the client) submits a rating on the performance of their mutual transaction. If these are in disagreement, then both transacting peers are punished, since such an occasion is a sign that one of them is lying. The severity of each peer's punishment is determined by his corresponding non-credibility metric; this is maintained by the mechanism and evolves according to the peer's record. When under punishment, a peer does not transact with others. We model the punishment effect of the mechanism in a peer-to-peer system as a Markov chain that is experimentally proved to be very accurate. According to this model, the credibility mechanism leads the peer-to-peer system to a desirable steady state isolating liars. Then, we define a procedure for the optimization of the punishment parameters of the mechanism for peer-to-peer systems of various characteristics. We experimentally prove that this optimization procedure is effective and necessary for the successful employment of the mechanism in real peer-to-peer systems. Then, the optimized credibility mechanism is combined with reputation-based policies to provide a complete solution for high performance and truthful rating in peer-to-peer systems. The combined mechanism was experimentally proved to deal very effectively with large fractions of collaborated liar peers that follow static or dynamic rational lying strategies in peer-to-peer systems with dynamically renewed population, while the efficiency loss induced to sincere peers by the presence of liars is diminished. Finally, we describe the potential implementation of the mechanism in real peer-to-peer systems.

Thanasis G. Papaioannou · George D. Stamoulis
Department of Computer Science,
Athens University of Economics and Business,
Patision 76, 10434 Athens, Greece
Tel.: +30-210-8203549, Fax: +30-210-8203686
E-mail: {pathan, gstamoul}@aueb.gr

1 Introduction

Peer-to-peer systems have become very popular as environments for exchanging services, i.e. files, storage capacity, video streams, etc. Commercial exploitation of peer-to-peer systems is also under way due to their unpreceded scalability. For example, commercial peer-to-peer systems, such as BBC iPlayer[1] and Kontiki[2], are employed for content delivery. Another example is that, in many peer-to-peer video streaming systems such as P2PLive[3] and Sopcast[4], the initial encoder and uploader of a free video clip or channel often embeds advertisements in the video stream obtaining some value for his effort to encode and upload the original video stream. If there is no accounting of information about who is offering what to whom in such systems, then peers have the opportunity for free-riding, and for providing malicious services or services of unacceptably low quality. Due to this information asymmetry among transacting peers, the risk for a peer of placing some individual effort and receiving much less in return is high. Reputation on the basis of ratings can be a proper means for achieving accountability, since it reveals hidden information regarding the inherent quality and the behavior (i.e. performance) of peers [1], [2]. However, as we showed in [2], a reputation metric should be exploited by reputation-based policies that determine the pairs of peers eligible to transact. When such policies are employed, the total value generated within the system is shared to peers according to their performance, thus, providing the right incentives to peers for exerting effort and offering services of high quality. However, reputation mechanisms are vulnerable to false or strategic rating. For example, a particular peer may benefit by submitting unjustified positive ratings for his friends and/or negative ratings for his competitors. This problem is further augmented in case of pseudo-spoofing, i.e. use of multiple false identities, which may appear in a peer-to-peer system. In this paper, we deal with the issue of credibility, i.e. truthfulness of the submitted ratings' feedback. Many reputation systems deal with this issue together with that of promoting high performance [3], [4], [5]. Such an approach provides peers with the incentive for employing various malicious strategies; e.g. an adversary peer may obtain a high reputation by offering services of high performance and subsequently exploit it as a rater to demote his competitors or to promote his colleagues. Moreover, poor performance and lying are not necessarily related; e.g. poor performance may be inherent for a peer due to his limited resources. In the present work, we deal with credibility separately from performance. In particular, we propose a proper mechanism for promoting truthful reporting of feedback information that was first presented in [6,7]. This mechanism detects and penalizes peers that lie. A non-credibility value as well as a punishment state is maintained for each peer. The effect of our credibility mechanism in a peer-to-peer system is modeled as a Markov chain. We experimentally prove that this model is very accurate. Employing this model, we prove our mechanism leads the peer-to-peer system with very large fractions of collaborated liar peers to a desirable *steady state*, where almost all liars are almost always under punishment, while sincere peers are almost never under punishment. Using this Markov model, we also define a fixed-point procedure for optimization of the punishment parameters of the credibility mechanism for peer-to-peer systems of different characteristics based

[1] www.bbc.co.uk/iplayer
[2] www.kontiki.com
[3] www.pplive.com
[4] www.sopcast.com

on ergodic arguments. We experimentally prove that this procedure is Pareto optimal for sincere peers and necessary for maximizing the effectiveness of the mechanism in peer-to-peer systems of different characteristics. Moreover, we show that the credibility mechanism can be combined very effectively with reputation-based policies that promote high performance, thus providing a *complete* and practically implementable solution for accountability in peer-to-peer environments. We experimentally justify that the optimized credibility mechanism deals successfully with very large fractions of liars in peer-to-peer systems with dynamically renewed population. Even if liar peers follow various static and dynamic lying strategies and are collaborated in order to gain unfair advantage, our experiments reveal that the efficiency attained for sincere peers by the optimized credibility mechanism combined with reputation-based policies is comparable to that of the case where no liar peers are present in the system. The mechanism provides peers with the right incentives for truthful reporting of feedback information, as sincere peers always receive more benefit from the peer-to-peer system than liar peers, whose benefit is minimal. Thus, the credibility mechanism is *strategyproof.* Finally, we describe how our credibility mechanism can be implemented in a real peer-to-peer system without central trusted authorities. We also prove with simulation experiments the effectiveness of the proposed architecture for dealing with large fractions of liars.

The remainder of this paper is organized as follows: In Section 2, we overview the literature related to truthful ratings. In Section 3, we define the proposed credibility mechanism. In Section 4, we model the effect of the mechanism regarding the punishment states of the peers as a Markov chain and, in Section 5, we introduce the procedure for optimizing the parameters of the mechanism. In Section 6, we overview our approach for the assessment of the mechanism, of the Markovian model and of the optimization procedure, while, in Section 7, we describe the simulation model that we employ in the experiments of this paper. Then, in Section 8, we prove the accuracy of the Markovian model, the effectiveness of the optimized mechanism, and the applicability and the necessity of the optimization procedure for maximizing the effectiveness of the mechanism for any peer-to-peer system of different characteristics. Also, in Section 9, we combine the credibility mechanism with reputation-based policies and experimentally prove that the right incentives regarding both reporting and performance are provided to peers. Furthermore, in Section 10, we describe the potential implementation of the credibility mechanism in a peer-to-peer system in the absence of trusted entities and experimentally prove the effectiveness of the proposed approach. Finally, in Section 11, we conclude our work.

2 Related Work

Below, we overview a variety of articles dealing either explicitly or implicitly with the consequences of lying in electronic environments, and, in certain cases, with how to alleviate them. We emphasize on the differences of these works with our assumptions as well as with our credibility mechanism and its effectiveness, in order to clarify our contribution.

Dellarocas [8] addresses the problem of unfair high or low ratings to sellers ("ballot stuffing" or "bad-mouthing") and positive or negative discriminatory behavior against clients in on-line trading communities where collaborated liars constitute at most 10% of the entire population of buyers. Only ballot stuffing and positive discrimination are dealt in [8], by employing collaborative filtering techniques to weight ratings in trust estimation proportionally to the similarity of preferences between the estimator and the raters. Moreover, this approach is not directly amenable to peer-to-peer environments where consumers are also producers of services, and bad-mouthing and negative discrimination can also arise due to peers' personal interest. Also, finding buyers with common taste requires a global view of the transaction history and raises privacy issues. An approach for improving the effectiveness of collaborative filtering for smaller sets of "similar" raters (i.e. neighbors) selected for predicting ratings has been proposed in [9], where rating prediction errors on different items are found to be correlated to the similarity of these items and to the shared neighbors of the items. However, the approach in [9] does not consider untruthful recommendations. Chen et al. [10] deal with the credibility of raters based on the quality and the quantity of the ratings they provide. However, the method assigns high confidence to ratings that agree with a majority opinion. Therefore, lying adversaries can still improve their credibility by submitting a large amount of feedback and thus forming the majority opinion.

Schillo et al. [11] deal separately with strategic performance and credibility using the so-called disclosed prisoners' dilemma game with partner selection. Credibility and performance (due to strategic behavior) of other agents are updated by an agent's own observations. Testimonies of witness agents are used for partner selection. It is assumed in [11] that witnesses may hide positive feedback but not tell lies in order not to be discovered. The approach approximates hidden feedback of witnesses and calculates a transitive credibility metric over a path to an agent using Bayes' rule. However, an adversary may still strategically gain high credibility by being truthful in his claims about his high offered performance and then manipulate as a witness the partner selection of other agents. Furthermore, collaboration among lying agents is not considered in [11]. The need for discovering witnesses for an agent is also a drawback of applying this approach in large electronic communities where the same agents meet very rarely. Damiani et al, in [12], extend Gnutella protocol to calculate performance and credibility of other peers based on a peer's own experience and votes from witnesses. Credibility is calculated in a similar way in [13], where trustworthiness of a peer is based on five factors, namely the feedback it receives on its performance from other peers, the number of transactions, the credibility of the feedback source, the transaction context factor (i.e. size and kind of transaction) and the community context factor (e.g. common incentives or beliefs). [12,13] approaches for calculating credibility are similar in many aspects to that of [11] and hence they have the same limitations. The same idea with [11], yet for evaluating direct and indirect recommendations, also taking into account context similarity of raters is proposed in [14].

Credibility and performance (due to strategic behavior) are also addressed by Yu and Singh [3]. However, this approach has no explicit mechanism for assessing the credibility of the witnesses; this issue is dealt together with a trust metric regarding behavior, which is determined by direct observations or by asking witnesses. Therefore, it is possible for an adversary peer to maintain a good reputation by performing high quality services and send false feedback for its competitors or his colleagues. A similar approach that has the same limitations with [3] is followed by Malaga in [15], where each rating is weighted by a function of the reputation of the rater. Credibility is addressed jointly with performance by Kamvar et al. in [5]. Therein, a global reputation metric regarding performance of each peer is calculated in distributed way and each peer's local beliefs (based on observations) on the performance of other peers are weighted by the others' beliefs on his own performance. [16] improves the convergence speed of global reputation of peers as related to [5]. It employs a gossiping protocol according to which local reputation is preferentially sent to power peers (i.e. peers that attract most of the requests). [16] takes as credibility of the raters their global reputation values. This approach is argued to counter dissemination of false local reputation values by malicious peers. However, as simulation experiments in [16] reveal, it has low accuracy even if only 10% percentage of collusive peers clustered in small groups are present in the system.

Aberer et al. [4] present an approach to evaluate trustworthiness (i.e. the combination of credibility and performance) of peers based on the complaints posed for them by other peers following transactions. The approach also aims to provide incentives for truthful submission of complaints. The main idea is that a peer is considered less trustworthy the more complaints he receives or files. An agent trusts another if the latter is at least as trustworthy as the former. The experiments conducted showed that the approach does not succeed in identifying a significant part of liar peers if they constitute 25% of the population. Note that the effectiveness of the approach in the case of collaborated liars was not examined and the approach is not robust against various types of peers' misbehavior. Feldman et al. [17] address the problems of free-riding (i.e. poor performance) and misreporting of feedback on contributions (i.e. low credibility) by an indirect reciprocity scheme. Their objective is for each peer to offer to any other peer roughly equal benefit as indirectly offered by the latter to the former. However, their approach provides opportunity for peers to lie about the contribution of other peers in order the latter to be unfairly exploited or for another liar collaborated with the former to prevail in competition. Ngan et al. [18] have proposed another indirect reciprocative approach for avoiding free-riding and false claims in a peer-to-peer system for sharing storage capacity. This approach requires peers to publish auditable records of their capacity and their locally and remotely stored files. However, collaborated adversaries can exploit this mechanism by claiming to have stored huge files of one to another. It is important to note that, to the best of our knowledge, our credibility mechanism is able to effectively deal with the highest fractions of collaborated liar peers that follow various static or dynamic strategies in the literature, as explained in Section 9.

A side payment approach for eliciting honest feedback in electronic markets has been proposed by Miller et al. in [19]. In particular, a payment charged to a buyer is paid to a second buyer according to a scoring rule for his prediction of the rating of a later buyer for their common seller. In the environment considered, honest reporting proved to be Nash equilibrium. However, strategic voting was considered to generate no value for buyers, which is not the case in general, particularly in cases of strategic collaborations. This approach does not deal with collaborated liars, while it is not appropriate for peer-to-peer systems, as it involves the employment of a central bank that distributes payments to peers. Jurca and Faltings [20] have proposed a similar approach that also has similar limitations. Another budget-balanced rewarding mechanism is proposed in [21] for providing incentives to participants to truthful report their subjective distributions on their beliefs over a hidden variable, so that it is collectively revealed. The approach seems promising for a limited set of privately observed variables by a large set of agents. However, it is deemed as an adequate approach for revealing the hidden performance of peers due to the large amount of information that has to be exchanged among all raters for each peer and the necessary exchange of rewards.

An approach for providing incentives for truthful reporting of feedback in e-markets has been proposed by Jurca and Faltings in [22]. This approach, similarly to ours, employs disagreement in feedback messages for discovering potential lying. However, upon disagreement different fixed side-payments are fined to the transacting agents with the one fined to the seller being higher. This approach is not directly amenable to peer-to-peer systems since side payments require the existence of a bank for mediating the transactions, while sellers and buyers are not supposed to exchange roles. Also, in [22], strategic voting and collaborated lying agents are not considered.

In [23], Sybil attacks are encountered based on a PKI approach. Peers employ self-created certificates to sign their identities, which are split to groups and resigned by group certificates. Upon identity creation, the peers are assigned to groups based on user credentials that prove that the identity corresponds to a real person. However, lying on recommendations is still possible in [23]. This approach could be used complementary to our credibility mechanism to deal with Sybil attacks. Finally, in perfect pseudonymity settings, Resnick and Sami propose in [24] an approach for limiting the total trust that can be exploited by Sybil attacks; the total trust is kept bounded by its initial value after any transactions. We agree with [24] that, in this context, some social loss due to Sybil attacks is unavoidable. Although, our credibility mechanism diminishes social loss even for very large fractions of adversary identities, as experimentally shown in Section 9.

3 The Credibility Mechanism

Consider a peer-to-peer system for exchanging services that employs a distributed reputation system for performance. The client peer, after a transaction, sends feedback that rates his offered performance. For example, he may rate the transaction as "successful" (i.e. high offered performance) or as "unsuccessful" (i.e. low offered performance). Simple binary feedback mechanisms are not only sufficient to appropriately reveal the hidden performance and quality, but, as proved in [25], the most efficient cooperation equilibrium is the one where participants group arbitrary ratings into two disjoint sets: positive and negative. We assume that votes are aggregated into reputation values using the Beta aggregation rule [26]. That is, each peer's reputation equals the fraction of the "weighted number" of his successful service provisions over the "total weighted number" of his service provisions, with the weight of each service provision being a negative exponential function of the elapsed time. The feedback messages are useful only if their content is true. Unfortunately, peers actually have the incentive of strategic rating of others' performance, since they can thus hide their poor performance, improve their reputation, and possibly take advantage of others. Thus, a proper mechanism should make lying costly or at least unprofitable. "Punishing liars" has already been proposed in [27] and [17]. Nevertheless, two questions arise: How can lying peers be discovered? How can they be punished in a peer-to-peer system, where there is no central control? Under our approach peers submit ratings' feedback according to the following rules: i) after a transaction, both peers involved have to send one feedback message each, and ii) besides voting the transaction as successful or not, each feedback message also contains a quantifiable performance metric, e.g. the number of transferred bytes of useful content. We assume that the observed performance is with high probability the same with that actually offered. (The opposite may only occur due to unexpected events during a transaction like network congestion etc.) Thus, if feedback messages for a transaction are in disagreement (either in their performance metric or in their vote), then, with high probability, at least one of the transacted peers is lying and has to be somehow "punished", in order for the right incentives to be provided. However, the system cannot tell which of the peers does lie, and consequently whom to believe and whom to punish. Thus, according to our approach, both peers are punished in this case. This idea was initially introduced in [27]. However, by simply applying it, a sincere peer is often punished unfairly.

Therefore, we need a complete mechanism specifying how to punish peers in a system without central control and how to limit potential unfairness. To this end, we introduce for each peer: i) the non-credibility metric $ncr \in [0, +\infty)$, which corresponds to reputation for non-credibility, and ii) a binary *punishment state* variable, declaring whether the peer is "under punishment" (if the variable is "true") or not (if the variable is "false"). For each peer, both *ncr* and punishment state are public information, i.e. they are appropriately stored so that they are available to other peers (see Section 10 for practical implementation details). Upon entering the peer-to-peer system, each peer is assigned a moderately high initial non-credibility value ncr_0, while he is not under punishment. (Note that the lower *ncr* the better for the peer.) This choice of ncr_0 offers to peers limited gain from whitewashing their non-credible record and reenter the system under new pseudonyms. The flowchart of the credibility mechanism is depicted in Figure 1. In particular, after a transaction between two not punished peers i, j their feedback messages f_i, f_j are sent as input to the mechanism: Upon disagreement (i.e. if $f_i \neq f_j$), the non-credibility values of the transacted peers are both increased by 1 while both peers get punished. The duration of a peer's punishment equals b^{ncr}, i.e. is exponential in his non-credibility ncr, with a base $b > 1$. Upon agreement (i.e. if $f_i = f_j$), the non-credibility values of the transacted peers are decreased (i.e. improved) by d, where $0 < d < 1$, without ever allowing them to drop below 0. The common feedback is forwarded to the system computing reputation for performance. If the reputation mechanism employed more than two feedback levels or the ratings involved subjectivity, then the matching rules for determining agreement or disagreement should be properly adjusted. For example, feedback agreement could be observed by examining if the actual distance between the two ratings was within a certain threshold that depends on the subjectivity level in the system and the semantic proximity of the different feedback levels. Decrease of non-credibility in cases of agreement serves as a rehabilitation mechanism. This is crucial for the efficient operation of the credibility mechanism, because, as already mentioned, upon disagreement in reports, most probably one peer is unfairly punished. The ratio 1:d determines the speed of restoring a non-credible reporting behavior. We employ additive increase/decrease of the non-credibility values for simplicity. Other approaches such as multiplicative increase/additive decrease are also plausible.

Punishing peers is not an easy task to achieve in the absence of any control mechanism, particularly if peers have full control over their part of the peer-to-peer middleware. In our mechanism, a punishment amounts to losing the value offered by other peers for the period of punishment. That is, a peer under punishment should *not* transact with others during his punishment period, while, if this happens, his ratings for such transactions are *not* taken into account. The latter measure provides incentives for peers to abide with the former one! Indeed, first, note that sincere peers under punishment are not expected to be willing to offer services as they would be subject to strategic voting without being able to disagree. On the other hand, liar punished peers collaborated with other liar peers that strategically vote them (i.e. always positively) can raise their reputation without high service performance. Thus, they have no incentives to perform well during their punishment. Based on the above, no peer has any incentives to *ask* for services from a punished peer except for the purpose of

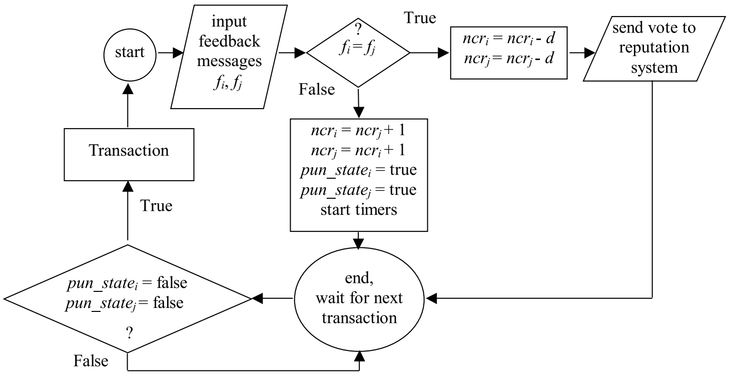

Fig. 1 The credibility mechanism.

strategic voting. Moreover, no peer has any incentive to perform well when offering services to a punished peer, because the corresponding feedback is not taken into account. To strengthen these incentives prohibiting transactions with punished peers we introduce a rule: If a peer transacts with a punished one, then both of the transacting peers are punished as if they were involved in a new disagreement. Note that, if no feedback is submitted for such a transaction, then this transaction is not traced by the mechanism. However, in such a case, the transaction would have taken place for the sole purpose of altruistic or collusive service exchange and it would not affect the accuracy of reputation information of the service performance incentives. Therefore, the non-credibility value of a peer remains unchanged during his punishment period unless he is discovered to transact with other peers; in such a case it is further increased.

4 The Markovian Model

In this section, we analytically study the *effectiveness* of the proposed mechanism in *equilibrium* for providing incentives to peers for truthful reporting. For this purpose, we define a discrete-time Markov-chain model of a peer-to-peer system where the credibility mechanism is employed. Then, we derive the *steady-state* distribution of the punishment state of sincere and liar peers of the modeled peer-to-peer system. Modeling of time is slightly different than that introduced in Section 3 for convenience. In particular, for the purpose of specifying and analyzing this Markov chain, we define as time step of our discrete-time model the interval between two successive service requests by any peer, henceforth referred to as *transaction unit*. We assume that in this interval at most *one* transaction takes place. Thus, transition from one state to another can *only* happen after a transaction between any two peers. This time modeling significantly facilitates the analysis of the Markov-chain model and the study of the performance of the original system defined in Section 3. Performance measures can be easily translated from the new "transaction units" to actual time slots; see Section 5. Note that at the beginning of each time step, a peer is randomly selected to be the client of the only transaction that takes place in this step.

We assume that there are two types of peers, namely sincere and liar ones. Sincere peers always report their feedback truthfully, while liar peers always disagree in their transactions, unless they transact with other liar peers collaborated with them. The total populations of sincere and liar peers in the peer-to-peer system modeled as a Markov chain are S_0 and L_0 respectively. The population of the peer-to-peer system can be dynamically renewed as long as S_0 and L_0 remain fixed. Consider that a state is a snapshot of the system where state variables are the number s of sincere peers not under punishment, the number l of liar peers not under punishment, and the number k of peers under punishment. Clearly, this Markov chain has $(S_0 + 1)(L_0 + 1)$ different states. Observe also that the state variable k can be computed by the formula $k = S_0 - s + L_0 - l$, but k is still used for readability reasons. Let q be the probability that a requested service is found at a certain peer and r to be the probability that a peer asks for a service. Recall that credibility values and punishment state are *public* information, and that not punished peers are not allowed to transact with punished peers. The probability y that a selected client peer finds a requested service is given by:

$$y = r(1 - (1 - q)^{s+l-1}) \tag{1}$$

A client sincere peer is punished if he finds his service at a liar peer. The probability P_S of this event is given by:

$$P_S = \frac{l}{s + l - 1} y \tag{2}$$

A client liar peer is punished if he interacts with a sincere peer or with another liar peer that is not collaborated with. Thus, the probability of punishment for a client liar peer is given by the formula below:

$$P_L = \frac{s}{s + l - 1} y + \frac{l - 1}{s + l - 1} y(1 - \theta) \tag{3}$$

θ is the fraction of liars that are collaborated to each other or alternatively the probability that two liar peers are collaborated. In the analysis that follows, we study the case where all liar peers are collaborated with each other, which is the hardest one for the mechanism to deal with.

Recall that at the beginning of each time step, a peer is randomly selected to be the client of the only transaction to take place. The probability P_T that the two peers of a transaction are punished, i.e. they disagree in their feedback messages is given by:

$$P_T = y \left(\frac{s}{s + l} P_S + \frac{l}{s + l} P_L \right) \tag{4}$$

For modeling purposes, we assume that during a time step, a sincere (resp. liar) peer that is under punishment can be "rehabilitated", i.e. stop being under punishment in the next step, with probability P_{RHS} (resp. P_{RHL}). Thus, when there are $k = S_0 - s + L_0 - l$ peers under punishment in the current state, the average number of rehabilitated peers in the next state is $(S_0 - s)P_{RHS} + (L_0 - l)P_{RHL}$. Next, we relate the Markovian model with the original mechanism of Section 3.

Suppose that the peer-to-peer system is currently in state (s, l, k), i.e. there are s not punished sincere and l not punished liar peers, while k peers are under punishment. Then, in the next time step (i.e. after a transaction), the system may move to various states with the transition probabilities given in the Table 1. Term A corresponds to the transition arising when the transacting peer are punished, while term B corresponds to the transition arising when they are not punished. Both terms also involve the probability of rehabilitation of the number of liar and sincere peers necessary for the transaction to happen.

Table 1 The transition probability from current state (s, l, k) to another.

Transition Probability
$Probability[(s,l,k) \rightarrow (s-1+i, l-1+j, k+2-i-j) = A + B$, where
$A = \begin{cases} P_T \binom{S_0 - s}{i} P_{RHS}{}^{i} (1 - P_{RHS})^{S_0 - s - i} \binom{L_0 - l}{j} P_{RHL}{}^{j} \cdot \\ (1 - P_{RHL})^{L_0 - l - j}, \text{ for } 0 \leq i \leq S_0 - s \text{ and } 0 \leq j \leq L_0 - l \\ \\ 0, \text{ otherwise} \end{cases}$
$B = \begin{cases} (1 - P_T) \binom{S_0 - s}{i - 1} P_{RHS}{}^{i-1} (1 - P_{RHS})^{S_0 - s - i + 1} \binom{L_0 - l}{j - 1} P_{RHL}{}^{j-1} \cdot \\ (1 - P_{RHL})^{L_0 - l - j + 1}, \text{ for } 1 \leq i \leq S_0 - s + 1 \text{ and } 1 \leq j \leq L_0 - l + 1 \\ \\ 0, \text{ otherwise} \end{cases}$

Under the Markovian model, the distribution of the punishment period is geometric; i.e. the duration of the punishment period is independent of the peer's past history. Clearly, this is only an approximation of our credibility mechanism that was described in Section 3, which is very complicated to model accurately and has a huge state-space. Indeed, recall that a peer upon disagreement is punished for a time period that is exponential to his non-credibility value, which should be maintained as part of the state for all peers! However, as the results of Section 6 reveal, this approximation is indicative of the performance of the actual mechanism provided that rehabilitation probabilities are successfully selected. Indeed, let us denote as c the period of conviction for a peer with a certain punishment record. Then, for a geometric-distribution approximation of this period, the probability of rehabilitation of this peer in the next state should be estimated as $1/c$. The probabilities P_{RHS} and P_{RHL} that lead to the same expected punishment time per type of peer (throughout a peer's lifetime) depend on the parameters b, ncr_0, and d of the credibility mechanism. All these parameters can be inter-related by means of the optimization procedure presented in Section 5. Thus, for given b, ncr_0, and d, appropriate values of P_{RHS} and P_{RHL} can be derived that render the Markov-chain model a good approximation of the evolution of the actual system. The steady state distribution of the model is depicted in Figure 2 for a certain peer-to-peer system with $S_0 = 30$, $L_0 = 20$, $r = 0.5$, $q = 0.1$ and rehabilitation probabilities $P_{RHS} = 0.1$ and $P_{RHL} = 0.0024$. As already discussed, these values of P_{RHS} and P_{RHL} result from the proper selection of the punishment parameters of the credibility mechanism according to the procedure described in Section 5. The z axis in Figure 2 is the equilibrium probability $\pi(s, l, k)$ that the system consists of s sincere and l liar peers not under punishment, while $k = S_0 + L_0 - s - l$ peers are under punishment. Clearly, in the peer-to-peer system of Figure 2, sincere peers are almost never under punishment during their lifetime, while liar peers are under punishment almost all of their lifetime. Thus, the credibility mechanism is very effective in expelling liar peers from the peer-to-peer system if its punishment parameters are properly selected.

Note that collaborated liar peers should be fewer than sincere ones in the system in order to dealt with effectively by the credibility mechanism. Otherwise, sincere peers

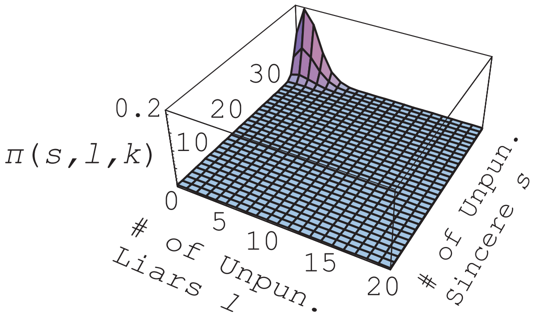

Fig. 2 Equilibrium probabilities $\pi(s,l,k)$ for the punishment states of peers in a system with $S_0 = 30$ sincere and $L_0 = 20$ liar peers when the credibility mechanism is employed.

would be involved into more disagreements than liar ones and consequently they would be under punishment most of their time; see equations (2) and (3). This is an inherent limitation of the credibility mechanism. However, if liar peers are not collaborated, then a fraction of liar peers higher than that of sincere ones can be tolerated.

5 The Procedure for Optimizing the Mechanism

As shown in Figure 2, the credibility mechanism is capable of providing the right incentives to peers for truthful reporting of feedback. However, this result applies for certain rehabilitation probabilities (essentially for certain expected punishment periods) that are determined by the punishment parameters of the mechanism (i.e. the initial non-credibility ncr_0, the base b of the exponential punishment, and the restoration factor d). These parameters have to be properly selected on the basis of the peer-to-peer system's, i.e. peers' lifetime, service availability, service request probability etc. in order lying to be effectively punished without inducing an unacceptable overhead for sincere peers. In this section, we propose a methodology for the calculation of the proper parameters of the mechanism for any peer-to-peer environment. We specify two ideal *objectives* on the achievable effectiveness when employing the credibility mechanism in a peer-to-peer system:

- Objective 1: Sincere peers must *not* be punished more than *once* during their lifetime.
- Objective 2: Liar peers must *always* be punished when they transact with other peers.

Specifically, consider the Markov-chain model of peer-to-peer system described in the previous section. Recall that we have defined as the time step of our discrete time Markov chain the time between any two successive transactions, i.e. the transaction unit. Furthermore, recall that, for the peer-to-peer system originally defined in Section 3, we assume that time is slotted, while the population of the peer-to-peer system is dynamically *renewed*, and S_0, L_0 are kept constant. Moreover, each time slot equals the minimum time interval between two successive service requests by the *same* peer. Next, we explain how we can inter-relate the two aforementioned systems. We denote as t_{life} the mean lifetime of a peer in time slots. We also denote as t_s (resp. t_l) the mean number of time slots that a sincere (resp. liar) peer is not under punishment during his lifetime, when our credibility mechanism is employed in the peer-to-peer system. Thus, $S_0(t_s/t_{life})$ [resp. $L_0(t_l/t_{life})$] is the mean number of sincere (resp. liar) peers not under punishment at a certain time slot. Recalling that y given by equation (1) is the probability to find a requested service, then the mean total number N_{trans} of transactions per time slot is given by the following equation:

$$N_{trans} = y\frac{t_s S_0 + t_l L_0}{t_{life}} \tag{5}$$

Furthermore, we denote as n_s and n_l the mean numbers of transaction periods that a sincere and a liar peer respectively are not under punishment during their lifetime that is denoted as n_{life} in transaction periods. Specifically, $n_s = N_{trans} \cdot t_s$, $n_l = N_{trans} \cdot t_l$ and $n_{life} = t_{life} \cdot N_{trans}$. Recall now that according to the Markov model, the distribution of each punishment period is geometric with expected value equal to $1/P_{RHS}$ for sincere peers and $1/P_{RHL}$ for liar peers. Using ergodic arguments, Objectives 1 and 2 lead to the following equations[5]:

$$\frac{1}{P_{RHS}} = n_{life} - n_s \tag{6}$$

$$\frac{1}{P_{RHL}} = \frac{n_{life} - n_l}{yt_l} \tag{7}$$

Indeed, Objective 1 amounts to equation (6), which implies that the expected punishment time for sincere peers equals the mean duration (in transaction periods) of the one and only punishment during their lifetime. Objective 2 amounts to equation (7), which implies that the expected punishment time for liar peers equals the mean punishment time of a liar peer in transaction periods divided by the mean number of time slots where: (i) he is not under punishment, and (ii) he transacts with another peer. Specifically in the denominator of equation (7), the term yt_l expresses the number of transactions of a liar peer. Note that equations (6) and (7) express the most conservative bounds arising from Objectives 1 and 2 for the mean punishment periods of a sincere and a liar peer respectively. Equation (6) [resp. equation (7)] involves n_s (resp. n_l), which determines the mean fraction of a sincere (resp. liar) peer's lifetime that he is not under punishment, namely n_s/n_{life} (resp. n_l/n_{life}). In equations (6) and (7), the values of these fractions are treated as inputs. However, these values actually arise as a result of the operation of the credibility mechanism. Thus, the input values in equations (6) and (7) have to be *consistent* with those resulting due to the mechanism. Therefore, in order to determine the values of n_s, n_l that render the objectives feasible a *fixed-point* approach is followed:

1. Initially, we take that $t_s = \max\{t_{life} - 1, 1\}$, $t_l = \max\{0.1 \cdot t_{life}, 1\}$ and calculate the corresponding n_s, n_l values. Note that, ideally, $t_s = t_{life} - 1$ and $t_l = 0$ should be used according to the Objectives 1 and 2; however, the chosen initial values for t_s, t_l have been experimentally verified to speed up the convergence of the fixed-point optimization approach.

[5] Equation (6) is tighter than the corresponding one in [7], as it is now expressed in transaction units instead of time slots. Also, equation (7) follows the Objective 2 closer than the corresponding one in [7], which was unnecessarily taking into account the fraction of sincere peers in the system for feasibility reasons. However, feasibility is satisfied by our fixed-point optimization approach that determines the values of n_s, n_l.

2. We calculate the mean fraction of a peer's lifetime that he is not under punishment, which equals n_s/n_{life} for sincere and n_l/n_{life} for liar peers.
3. From equations (6) and (7) we calculate P_{RHL} and P_{RHS}. These are employed in the Markov-chain model and the steady-state distribution of the punishment state is calculated.
4. Then, the mean fraction of a peer's lifetime that he is not under punishment is re-calculated for sincere and liar peers based on the steady state probabilities, i.e. n'_s/n_{life} for sincere and n'_l/n_{life} for liar peers.
5. If the convergence criteria are met, e.g. $|n'_s - n_s| < \epsilon$ and $|n'_l - n_l| < \epsilon$, with $\epsilon \leq 0.03$, then a fixed point has been reached, and the proper values of n_s, n_l have been found for this peer-to-peer system. Otherwise, we set $n_s = (1-\delta)n_s + \delta n'_s$ and $n_l = (1-\delta)n_l + \delta n'_l$, with $\delta \in (0.5, 1)$ as a relaxation parameter, and the control is transferred back to step 2.

Having determined the values of n_s and n_l that give rise to Objectives 1 and 2, the proper parameters of the credibility mechanism have also to be derived. The expected value of total punishment period in time slots for a liar peer that is punished in all of his transactions is at most $E[b^{ncr_0}(1 + b + b^2 + .. + b^v)]$, where v is the number of transactions. This is approximated as $b^{ncr_0}(1 + b + b^2 + .. + b^{y \cdot t_l})$, since $E[v] = y \cdot t_l$, which is henceforth treated as integer for simplicity. The total punishment period for a liar peer should be equal to the mean total punishment time for that peer $t_{life} - t_l$, see equation (8) below. (Note that this is a bound because the last punishment period may not be fulfilled until the end of the lifetime of the peer. However, again we take the equality, as it is the most conservative relation.) Similarly, the total expected punishment time of a sincere peer is taken as $b^{ncr_0 - d \cdot rh}$, see equation (9). rh, given by equation (10), is the expected number of time slots where transactions are conducted by a sincere peer until his one and only punishment and d is the restoration factor; thus $rh \cdot d$ is the expected decrease in the sincere peer's non-credibility value until his punishment. Specifically, in equation (10) the term $\frac{1}{P_T} - 1$ expresses the expected number of transactions of a sincere peer until punishment which are translated to time slots dividing by N_{trans}. Note that the relations for b and ncr_0 involve d as a parameter as well. Instead of setting one more objective and devise one more equation in order to determine d, we take $d = 0.5$ for illustrative purposes. This is a meaningful choice for the restoration of a disagreement to require two agreements. Therefore, b, ncr_0 (and rh) can be determined by the equations below:

$$t_{life} - t_l = b^{ncr_0} \frac{b^{yt_l+1} - 1}{b - 1} \tag{8}$$

$$b^{ncr_0 - d \cdot rh} = t_{life} - t_s \tag{9}$$

$$rh = \frac{y(\frac{1}{P_T} - 1)}{N_{trans}} \tag{10}$$

6 Methodology for the Evaluation of the Optimized Mechanism

In this section, we present the methodology for the assessment of the credibility mechanism that is followed in the subsequent sections: Initially, we describe the simulation model of the credibility mechanism in a peer-to-peer environment. Both in this model and in the Markovian model of Section 4, the pairing of peers that transact is random. Employing the simulation model, we prove the accuracy of the Markovian model on the prediction of the resulting punishment periods for sincere and liar peers in a system where the credibility mechanism is employed. Moreover, we perform sensitivity analysis of the credibility mechanism to the punishment parameters and prove that the the effectiveness of the credibility mechanism is Pareto optimal for sincere peers for the optimized punishment parameters. We also prove the applicability of the credibility mechanism with optimized punishment parameters for real peer-to-peer systems of different characteristics. Then, we combine the credibility mechanism with reputation-based policies, i.e. the pairing of peers that transact is done by means of reputation-based policies and we experimentally assess the effectiveness of the optimized credibility mechanism in isolating liar peers that submit ratings' feedback according to various fixed strategies or a rational dynamic strategy.

7 The Simulation Model

We consider a peer-to-peer system where services of a certain kind are exchanged among peers. Similarly, with other articles [2], [11], [17], we assume that there are two types of peers with different performance in this system: altruistic and egotistic. Each peer exhibits (either inherently or intentionally) a mixed strategy regarding his performance in his service provisions; this strategy depends on the peer's type. In particular, each altruistic (resp. egotistic) peer provides a service successfully with a high probability $\alpha = 0.9$ (resp. with a low probability $\gamma = 0.1$). Different service provisions by the same peer are taken as independent. At the same time, each peer exhibits a reporting strategy regarding the sincerity of his feedback: he is either (always) sincere or liar. The lying strategies considered are defined in Subsection 9.1. In each experiment, all liars follow the same such strategy. The performance and the reporting types of each peer are private information, i.e. only the peer himself knows them.

Furthermore, the population of peers is assumed to be renewed according to a Poisson process with mean rate $\lambda = 10$ peers/time slot, while the total size N of the population is kept constant, with $N = 1500$. That is, each peer is assumed to live in the peer-to-peer system for a period determined according to the exponential distribution with mean N/λ. When a peer leaves the system, a new entrant of the same type takes his place. To make matters worse, the vast majority of peers (90%) are taken to be egotistic. The percentage of liar peers in each experiment varies. In fact, for each lying strategy, we present the results for the maximum such percentage that can be dealt with effectively by our mechanism.

Time is assumed to be slotted. The duration of the time slot is of the same order of magnitude as the average interval between two successive service requests by the same peer. At each slot, every peer requests a service with a certain probability $r = 0.5$. The relative large value of this parameter is not important for the effectiveness of the mechanism and it just accelerates its convergence. Service availability is Zipf-distributed, i.e. assuming that services are ranked with respect to their popularity, a service with rank z is found at a certain peer with probability z^{-1}. For each service instance, popularity is randomly selected in the range $[1, 300]$. A peer can serve only one peer per slot due to his limited resources.

The credibility mechanism of Section 3 with optimized punishment parameters b, ncr_0 is employed in this system. Therefore, each peer is assigned an optimized initial non-credibility value ncr_0. The non-credibility values are increased upon disagreement with his transacted peer in their feedback by 1 and decreased upon agreement by $d = 0.5$. Upon disagreement, both peers i, j are exponentially punished for b^{ncr_i}, b^{ncr_j}, where ncr_i, ncr_j are their non-credibility values prior to disagreement.

8 Assessment of the Optimization Procedure

8.1 Accuracy of the Markovian Model and Effectiveness of the Credibility Mechanism with Optimized Punishment Parameters

First, we assess the accuracy of the optimization procedure of Section 5 regarding the expected punishment periods resulting by the credibility mechanism with optimized punishment parameters b and ncr_0 for sincere and liar peers. Due to the ergodicity arguments employed in the optimization procedure, it is expected that if the Markovian model is an accurate proxy of the employment of the mechanism in a peer-to-peer system, then the mean punishment periods for sincere and liar peers resulting by simulation experiments for a long period are expected to approximate the ones calculated by equilibrium analysis of the Markovian model. Recall that the mean punishment periods for sincere and liar peers depend on the b and ncr_0 parameters according to the optimization procedure of Section 5.

To this end, we denote as PM_S, PM_L the mean lifetime fractions where sincere and liar peers respectively are not under punishment, which are calculated by the Markovian model, and PS_S, PS_L the corresponding mean lifetime fractions for sincere and liar peers respectively that result after long simulation experiments. As depicted in Table 2, the absolute differences $|PS_S - PM_S|$ and $|PS_L - PM_L|$ are very small (i.e. ≤ 0.02 and ≤ 0.06 respectively) for all different peer-to-peer systems considered. Thus, the Markovian model indeed approximates very accurately the punishment effect of the credibility mechanism for both sincere and liar peers in a multitude of peer-to-peer systems of different characteristics.

Another result depicted in Table 2 is that the credibility mechanism with the punishment parameters b and ncr optimized for different peer-to-peer systems is very effective, as it always results to severe punishments for liar peers at equilibrium. Therefore, a liar peer is almost always under punishment during his lifetime as PS_L, PM_L are close to 0, while sincere peers are almost never punished as PM_S, PS_S are close to 1. Therefore, the optimization procedure of Section 5 is very effective. In Subsection 8.3, we perform sensitivity analysis of the effectiveness of the mechanism to the selection of the punishment parameters and prove that the optimization procedure is also necessary.

We now discuss how the optimized punishment parameters depend on the various parameters of the peer-to-peer system. We found that b increases and ncr_0 decreases the larger the system with the same other characteristics. This effect is the same with increasing the probability of requesting a service per time slot and with increasing the probability of finding a requested service at a peer, as they all result into an increased number of transactions per time slot. When the number of transactions per time slot increase, a larger number of punishments is expected and then the mechanism converges faster. Therefore, for achieving the objectives of Section 5, ncr_0 should be smaller, so that the credibility mechanism results to a smaller unfairness for sincere peers; on the contrary, b should be larger, so that the resulting punishment for liar peers remains long

Table 2 The accuracy of the Marovian model and the effectiveness of the credibility mechanism employed with punishment parameters b and ncr_0 optimized for various peer-to-peer systems.

lifetime	q	r	S_0	L_0	b	ncr_0	PM_S	PS_S	PM_L	PS_L
150	0.1	0.5	22	18	1.57	6.19	0.966	0.947	0	0.038
150	0.1	0.5	40	10	1.44	8.07	0.97	0.982	0	0.03
150	0.1	0.5	35	15	1.44	8.02	0.95	0.963	0	0.03
150	0.06	0.5	37	30	1.82	4.15	0.968	0.954	0	0.036
150	0.1	0.5	30	10	1.87	3.85	0.997	0.982	0	0.03
150	0.1	0.3	30	10	1.25	14.41	0.99	0.966	0	0.044
150	0.025	0.5	30	10	1.24	15.05	0.995	0.965	0	0.056
400	0.1	0.5	30	10	1.36	5.16	0.995	0.981	0	0.028
800	0.1	0.5	30	10	1.22	15.47	0.995	0.991	0	0.016

enough. Also, as depicted in Table 5, the same trend (although with small fluctuations) is observed for punishment parameters when increasing the percentage of liar peers in a system while keeping the other characteristics constant. Again, this is because the expected number of punishments per time slot is larger in a system with more liars. Finally, keeping the other system characteristics constant and increasing the average lifetime of peers, the reverse trend for optimized punishment parameters is observed, i.e. the base of punishment decreases and the initial exponent of punishment increases. This is because the larger the lifetime of a peer, the larger should be the exponent of punishment in order for liars to get punished for almost all their lifetime, while having sincere being punished the minimum possible.

8.2 Applicability of the Optimized Mechanism

In this subsection, we study the applicability of the optimization procedure for finding punishment parameters b and ncr_0 for arbitrary peer-to-peer systems and especially for those of large populations. This is very important as the Markovian model does not scale well with the number of peers due to its large transition matrix, i.e. $(S_0+1)(L_0+1) \times (S_0+1)(L_0+1)$. We consider a peer-to-peer system with fixed rate of service request $r = 0.5$ and fixed mean lifetime period of 150 time slots. The percentage of liars in the systems is taken to be 45%. We calculate the optimized b and ncr_0 parameters for a fixed population mix as population scales up. The probability q of finding a service at each peer is taken to be 0.1 for a system of 40 peers. However, q is properly adjusted for larger population sizes, so as the probability of finding the requested service in the system remains constant as the system scales up and equal to 0.983. As depicted in Table 3, the periods t_s, t_l for sincere and liar peers respectively of not being under punishment at equilibrium converge to very close values, which even become constant as the system scales up. This is very important, because knowing t_s, t_l, one can estimate the expected number of sincere and liar peers at equibrium. Therefore, one can calculate the probability y, given by equation (1), for a peer to conduct a transaction at a time slot and the probability P_T, given by equation (4), to be punished at equilibrium for a peer-to-peer system with arbitrary population and the same other characteristics. To this end, for calculating the optimal parameters for a peer-to-peer system with large population, one should run the optimization procedure for a very small system of the same other characteristics, namely population mix, average lifetime, and probability of finding a requested service in the system, and obtain t_s and t_l. Then, calculate y and P_T for the larger system and calculate the optimal punishment parameters b, ncr_0 using equations (8), (9), (10). Using this methodology, we calculate the optimal punishment parameters for larger peer-to-peer systems of 45% liars, with r=0.5, $life$=150 and constant probability of finding a service in the system equal to 0.983 (see Table 4). As shown in Table 4, the effectiveness of the optimized credibility mechanism is very high and remains almost constant for larger systems, i.e. sincere and liars are not under punishment ∼95.5% and ∼4.5% of their lifetime.

Table 3 Optimized punishment parameters b and ncr_0 for peer-to-peer systems with $r = 0.5$, lifetime of 150 time slots and $q = 0.1$, $L_0/(S_0 + L_0) = 0.45$ but with different total population sizes.

S_0	L_0	b	ncr_0	t_s	t_l
22	18	1.57	6.19	145.81	3.36
27	22	1.65	5.38	145.91	3.36
32	26	1.74	4.67	146.01	3.36
37	30	1.82	4.15	146.11	3.36
40	34	1.89	3.74	146.2	3.36
49	40	1.94	3.54	145.9	3.36
55	45	1.99	3.3	145.9	3.36
60	49	2.04	3.11	145.9	3.36
70	57	2.13	2.8	145.9	3.36

Table 4 Effectiveness of the optimized punishment parameters b and ncr_0 as the peer-to-peer system scales up.

$S_0 + L_0$	b	ncr_0	PS_S	PS_L
500	2.6	1.73	0.956	0.044
1000	2.73	1.53	0.954	0.044
2000	2.80	1.43	0.955	0.044
3000	2.82	1.4	0.955	0.043
4000	2.83	1.39	0.954	0.043
5000	2.84	1.38	0.954	0.043

Another important issue for the applicability of the optimized credibility mechanism is the sensitivity of its effectiveness to the inaccuracy of the characteristics of the peer-to-peer system for which the punishment parameters are optimized. Indeed, in reality some differences between the estimated system characteristics and the real ones are expected. One could rightfully argue that the population mix, i.e. estimating the percentage of liars, might be hard to predict. Fortunately, as we observe in Table 5, t_s and t_l do not significantly change for different population mixes. Therefore, one can select t_s and t_l for a value close to the estimated fraction of liars in the optimization procedure without introducing significant inaccuracy to the optimization procedure. Running the optimization algorithm for various different systems, we observed that t_s, t_l at equilibrium mostly depend on the average lifetime of the peers in the system and to a much smaller extend to probability r of requesting a service at a time slot and the probability q of finding services at a peer. We experimentally studied the sensitivity of the effectiveness of the credibility mechanism to the inaccuracy of the estimated mean lifetime and the mean service request ratio of the real peer-to-peer system. We found that the effectiveness of the credibility mechanism remains roughly the same for optimized punishment parameters calculated for a system with ±10% different mean lifetime, mean service request rate and probability q. Therefore, the optimized punishment parameters are very efficient despite small errors in the estimation of the parameters of the peer-to-peer system for which the credibility mechanism is optimized. Also, note that these parameters of the peer-to-peer system are easier to be accurately estimated, than the population mix.

Table 5 Optimized punishment parameters b and ncr_0 for peer-to-peer systems with $r = 0.5$, lifetime of 150 time slots and $q = 0.1$, $S_0 + L_0 = 50$ but with different fractions of sincere and liar peers.

S_0/L_0	b	ncr_0	t_s	t_l
45/5	1.35	10.27	149.2	3.36
40/10	1.44	8.07	146.4	3.36
35/15	1.39	9.17	144	3.36
30/20	1.66	5.26	146.5	3.36
27/23	1.64	5.44	145.6	3.36

8.3 Sensitivity Analysis of the Optimized Mechanism

We now investigate the sensitivity of the effectiveness of the credibility mechanism to the punishment parameters selected by the optimization procedure of Section 5. To this end, we consider a peer-to-peer system with $N = 1500$ participants of which 45% are liars, with service request rate per peer $r = 0.5$, with mean lifetime 150 time slots and with probability of finding a requested service in the system equal to 0.983. Employing the Markovian model and the optimization procedure of Section 5, we find that the optimized punishment parameters for this system are $b = 2.77$ and $ncr_0 = 1.46$. Using the simulation model of Section 7, we run experiments for a long time and observe the mean fractions of time PS_S and PS_L that sincere and liar peers are not under punishment. We modify b and ncr_0 parameters by ±50% and [-50%, 350%] respectively from their optimized values and we run experiments again to observe resulting PS_S and PS_L. The resulting percentage differences in the effectiveness of the credibility mechanism for sincere and liar peers are depicted in Figures 3(a) and 3(b) respectively. Note that the sensitivity of the effectiveness of the credibility mechanism to the punishment parameters b and ncr_0 may be different for peer-to-peer systems of different characteristics. As observed by Figure 3(a), the effectiveness of the optimized credibility mechanism is Pareto optimal for sincere peers, while it can be significantly reduced (i.e. up to 20% in this experiment) for different punishment parameters. The effectiveness of the credibility mechanism for liars is not Pareto optimal, but the mechanism with optimized parameters achieves that liars are under punishment ∼95% of their lifetime. However, this result may significantly deteriorate (i.e. over 50% in this experiment) for punishment parameters other than optimized ones, as depicted in Figure 3(b). Therefore, the effectiveness of the credibility mechanism is guaranteed to be high only for the optimized punishment parameters.

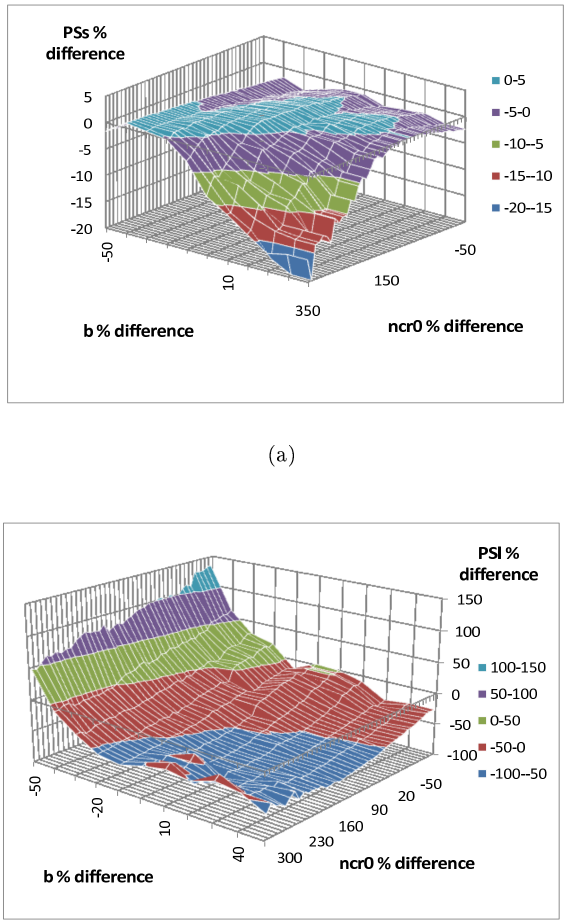

(a)

(b)

Fig. 3 The percentage difference of PS_S (a) and PS_L (b) that result by simulation experiments when the credibility mechanism is employed with punishment parameters b and ncr_0 that have a percentage difference than their optimized values.

9 Integration with Reputation-based Policies

In this section, we experimentally investigate the effectiveness of the credibility mechanism when combined with reputation based policies, such as those that we described in [2]. This is very important considering that the introduction of reputation in every peer-to-peer system affects the eligibility of peers to be selected for transaction by other peers, as explained there. Furthermore, it was established in [2] that reputation provides the right incentives for high performance only when proper reputation-based policies are employed in the peer-to-peer system.

In the experiments of this section, we employ the simulation model of Section 7 along with the Max-Max reputation-based policy pair: each client select to transact with the provider that has the highest reputation among those that offer the requested service and providers select to transact with the client that has the highest reputation among those that compete for the same service at the same time slot. We have also run experiments where we employed other policies of [2] and we noticed similar effects. Note that the Markovian model of Section 4 is not adequate to measure the effectiveness of the credibility mechanism when combined with reputation-based policies, as the random selection of the transacted parties is no longer valid: here peers are selected to transact based on their reputation values. If we assumed high inaccuracy of reputation in the peer-to-peer system, then peers could be assumed to be selected ad hoc initially. However, as reputation values become more accurate due to the presence of the credibility mechanism in the system, a cycle of reputation information is formed: two peers are selected to transact based on their reputation values, both provide feedback on the provider's performance after transaction, upon agreement the reputation of the providing peer is updated affecting its probability to be selected in the future while upon disagreement both transacted peers are punished and then the transacted peers are again selected based on their reputation and so forth. This reputation bias could only be described by a Markovian model that would have a much larger number of states, e.g. comprising all the discrete reputation levels that a peer could be assumed to be categorized into. However, we would thus end up with a computationally non-tractable Markovian model. On the other hand, we determine the punishment parameters of the credibility mechanism according to the optimization procedure of Section 5, so as to optimize the effectiveness of the mechanism in the worst case scenario of high inaccuracy of reputation values in the system.

After a transaction each of the peers involved sends feedback to the reputation system as explained in Section 3. Votes are converted into reputation values using the Beta aggregation rule discussed in Section 3. The reputation value for a peer is associated to his pseudonym, and expresses his probability of offering high performance given his past record. 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 a low initial reputation h_0 (i.e., $h_0 = 0.1$), in order to limit the incentive for name changes. In the experiments conducted, we assess the efficiency attained in this peer-to-peer system when the optimized credibility mechanism is employed, which is measured as the mean number of successfully provided services to each peer type. Particular emphasis is placed on the efficiency of sincere altruistic peers, as such peers offer the most of the value of the peer-to-peer system. We also assess the incentives offered per type of peers for truthful reporting. First, liar peers are assumed to employ static strategies, while next liar peers are assumed to be rational employing a dynamic strategy.

9.1 Static Lying Strategies

Recall that peers belong to two fixed reporting types namely sinceres and liars. Liars may follow various strategies for manipulating their ratings depending on their objectives. We considered four possible lying strategies, some of which are similar with those in other related works [4], [8], [18]:

- *Destructive*, in which liar peers reverse the feedback on the outcome of their transactions.
- *Opportunistic*, in which liar peers claim that they always succeed in their transactions and that all other peers not collaborated with them fail.
- *Mixed*, in which a liar peer randomly selects which of the above two lying strategies to employ. The selection probability may vary with time.
- *Discriminating*, in which a liar peer apart from being opportunistic, only serves peers collaborated with him, thus bypassing the Max-Max policy.

Liar peers may be collaborated to each other in the sense that they always rate positively each other.

In all the experiments of this subsection, we assume that liars are collaborated to each other. We also omit an initial "bootstrapping" period of operation of the peer-to-peer system in the beginning of which all peers are newcomers. (This period lasts for 250 slots; in general its duration depends on various parameters, but mainly on the service request probability.) Thus, we assess the efficiency of peers during the normal operation of the peer-to-peer system with dynamically renewed population. First, liar peers are taken to follow the *destructive* lying strategy and to constitute the 45% of the population of the peer-to-peer system. Using the optimization procedure of Section 5, we find that the optimized punishment parameters for this system are b=2.77 and ncr_0=1.46. In Figure 4(a), depicted are the mean reputation values of sincere peers, which are very accurate when the credibility mechanism is employed, as shown by arrow 1 for altruistic and by arrow 2 for egotistic peers. Indeed, the mean reputation values for sincere altruistic and egotistic peers are very close to their corresponding a priori probability for successful service provision $\alpha = 0.9$ (resp. $\gamma = 0.1$). On the contrary, if the credibility mechanism is not employed, then the two performance types cannot be distinguished by means of reputation. Also, note that altruistic liar peers benefit from the absence of the credibility mechanism as opposed to altruistic sincere ones! Therefore, peers have *wrong* incentives in the absence of our mechanism. On the other hand, the mean reputation values for liar peers are higher when the credibility mechanism is employed, as depicted in Figure 4(b). This is because liar peers agree on their feedback only in their transactions with liars and as a result due to the credibility mechanism they receive only positive ratings.

Next, we deal with efficiency issues for the same set of experiments. The number of total successfully obtained services per peer (i.e. efficiency) increases for both altruistic and egotistic sincere peers when the credibility mechanism is employed, as depicted by arrow 1 in Figure 5(a) and 5(b) respectively. On the contrary, when the credibility mechanism is employed, the efficiency of liar peers (which was greater than that of sincere ones) becomes almost zero as depicted by arrow 2. Also, when the credibility mechanism is employed, the efficiency achieved by sincere peers in the presence of liars is very close (i.e. up to 10% relative difference) to that achieved in the ideal case where no liar peers are present in the peer-to-peer system. The same conclusion also applies for egotistic sincere peers, whose efficiency is much lower than that of

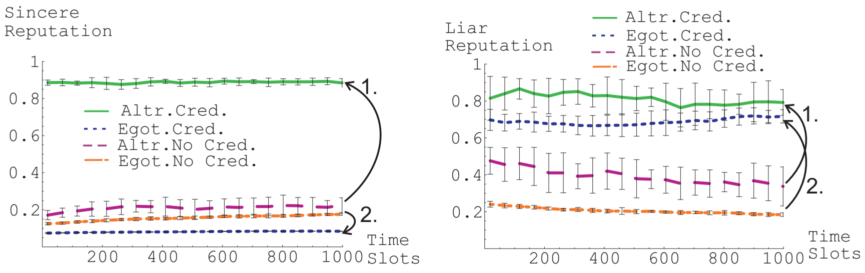

(a)

(b)

Fig. 4 The mean reputation values of sincere (a) and liar (b) altruistic and egotistic peers when the credibility mechanism is employed ("credibility") or not ("no credibility") in a peer-to-peer system with 45% collaborated liar peers that follow the destructive strategy.

altruistic sincere ones, as expected. That is, the credibility mechanism enables the proper operation of the reputation for performance when reputation-based policies are employed. Thus, when our credibility mechanism is employed, the disturbance of sincere peers by presence of liars is minimal. The introduction of the mechanism is very beneficial for sincere peers and very harmful for liar ones, who receive a much lower efficiency than sincere peers. Therefore, the strategy of collaborative destructive lying strategy is dominated by the "always be sincere" strategy. Our mechanism renders truthful reporting *incentive-compatible* for peers, as liars spend most of their lifetimes being under punishment. On the contrary, sincere peers recover very soon both from the initially high non-credibility value and from their possible unfair punishments imposed by the credibility mechanism.

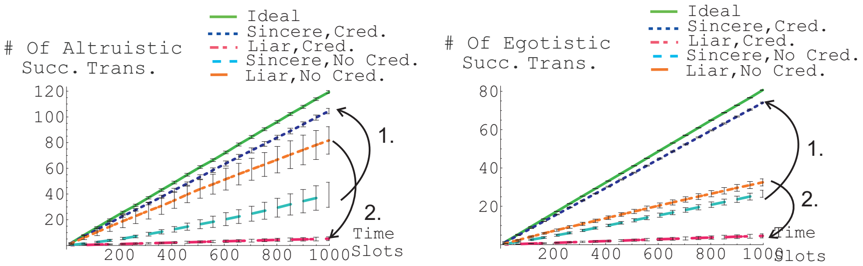

(a)

(b)

Fig. 5 The cumulative number of successfully offered services per altruistic (a) and per egotistic (b) peer when the credibility mechanism is employed ("credibility") or not ("no credibility") in a peer-to-peer system with 45% liar peer that follow the destructive strategy.

Next, we consider the case where liar peers are collaborated and follow the *opportunistic* lying strategy. In order for the optimized punishment parameters to properly calculated for this strategy, we replace the formulas of equations (2) and (3) of the probabilities P_S and P_L that a sincere or a liar peer respectively is punished at each time slot by the following ones:

$$\begin{aligned} P_S &= \frac{l}{s+l-1} y[l_{\mathrm{a}}(1-a) + l_{\mathrm{e}}(1-\gamma)] \\ P_L &= \frac{s}{s+l-1} y(1-l_{\mathrm{a}})a + (1-l_{\mathrm{e}})\gamma) \ , \end{aligned} \tag{11}$$

where s, l are the expected number of sincere and liar peers at equilibrium, estimated in Section 5, and l_{a}, l_{e} are the fractions of altruistic and egotistic liar peers in the system respectively. These formulas are derived by the definition of the opportunistic lying strategy. Note that these changes in the Markovian model do not change the optimization procedure defined in Section 5. Then, the calculated optimized punishment parameters are $b = 2.5$ and $ncr_0 = 1.9$. Indeed, ncr_0 should be higher in order for liar peers to be punished for long enough despite the fewer disagreements that they are involved into. When the credibility mechanism with these punishment parameters is employed in a peer-to-peer system with 40% collaborated liar peers that follow the opportunistic strategy, then the number of total successfully offered services per peer (i.e. efficiency) increases for both altruistic and egotistic sincere peers, as depicted by arrow 1 in Figures 6(a) and 6(b) respectively. On the contrary, the efficiency of liar peers (which was greater than that of the ideal case) becomes lower than that of the sincere ones, as depicted by arrow 2 in Figures 6(a) and 6(b).

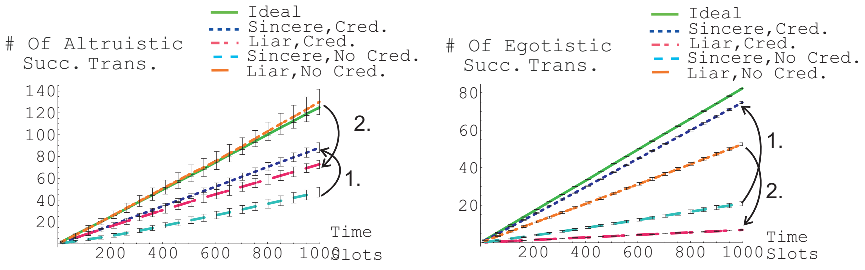

(a)

(b)

Fig. 6 The cumulative number of successfully offered services per altruistic (a) and per egotistic (b) peer when the credibility mechanism is employed ("credibility") or not ("no credibility") in a peer-to-peer system with 40% liar peer that follow the opportunistic strategy.

Liar peers are next supposed to follow the *discriminating* strategy. In this case, the optimized punishment parameters for the credibility mechanism are calculated by the optimization procedure of Section 5, but using the formulas in equations (12) for calculating the probability P_S (resp. P_L) that a sincere (resp. liar) peer is punished at each time slot. These probabilities are directly derived by the definition of the discriminating lying strategy described in the beginning of this subsection. The punishment parameters calculated by the optimization procedure of Section 5 are $b = 1.15$ and $ncr_0 = 23.2$. The same argument for ncr_0 in the case of the opportunistic strategy also applies for this lying strategy. However, b has to be low so as the unfair punishment for sincere peers to be minimized. We experimentally found that the optimized credibility mechanism effectively deals with up to 12% of peers that follow this lying strategy, as depicted in Figures 7(a) and 7(b).

$$
\begin{aligned}
P_S &= 0 \\
P_L &= \frac{s}{s+l-1} y\left(\frac{s_{\mathrm{a}}}{s} a + \frac{s_{\mathrm{e}}}{s} \gamma\right)
\end{aligned}
\tag{12}
$$

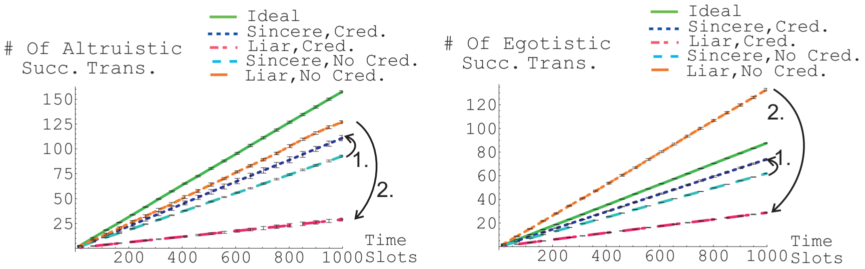

(a)

(b)

Fig. 7 The cumulative number of successfully offered services per altruistic (a) and per egotistic (b) peer when the credibility mechanism is employed ("credibility") or not ("no credibility") in a peer-to-peer system with 12% liar peers that follow the discriminating strategy.

We also consider the case that liar peers follow the *mixed* strategy. Employing the above methodology for optimizing the punishment parameters, we have experimentally found that the numbers of successfully offered services to sincere altruistic and egotistic peers respectively are always equal or greater than the corresponding ones offered to liar altruistic and egotistic peers respectively, when the optimized credibility mechanism is employed in a peer-to-peer system with up to 33% liar peers. Therefore, truthful reporting dominates the mixed lying strategy under the credibility mechanism and truthful reporting is incentive-compatible for peers.

Finally, we have experimentally observed the effect of subjectivity to the effectiveness of the credibility mechanism. We only describe the interesting case of severe subjectivity, i.e. when the subjectivity is so high that creates feedback disagreements between sincere peers and thus they get unfairly punished. We found that the credibility mechanism is still effective to isolate liars, as long as the sum of the total fraction of transactions that are subject to severe subjectivity and the fraction of collaborated liars is lower than the corresponding upper bounds that are presented earlier in this section for each lying strategy. However, the efficiency of the peer-to-peer system for sincere peers may be significantly affected. For example, if 2%, (resp. 5%) of the transactions are subject to severe subjectivity, ∼ 10% (resp. ∼ 20%) of the efficiency for sincere peers is lost.

Note that the above results were also experimentally verified for Uniform service distribution. Also, note that the effectiveness of the credibility mechanism increases as the renewal rate of the population of the peer-to-peer system decreases, as expected.

This is because, a number of time slots is needed for the non-credibility and the reputation values of the peers to converge to their proper values. Moreover, the effectiveness of the mechanism increases with the percentage of sincere peers. The upper bound on the percentage of liar peers that follow the destructive strategy and are not collaborated to each other and they can effectively dealt with by the credibility mechanism was experimentally found to be 55%. In particular, we found that the efficiency for sincere peers is near that of the ideal case, while the efficiency of liar peers is diminished when the credibility mechanism is employed. This was expected, as in this case liar peers disagree even when they transact with each other. Such fractions of liar peers that are not collaborated correspond to liar peers that do not belong to the same real entities. On the other hand, large fractions of collaborated liar peers correspond to artificially generated identities on behalf of a certain real entity. Such collectives of so high fractions of liar peers are difficult to emerge in large actual peer-to-peer systems where a proper membership mechanism that requires some real-life credentials for is employed.

9.2 Rational Dynamic Lying Strategy

So far, we have experimentally proved that the credibility mechanism, employed jointly with reputation-based policies, assigns very high efficiency to sincere peers despite the presence of high fractions of liars that all follow one of the fixed lying strategies of Subsection 9.1. Next, we assume that peers choose their lying probability according to a dynamic rational strategy so as to maximize their long term expected payoff by means of a *learning* algorithm explained below. We employ the simulation model defined in Section 7. The lifetime of each peer is exponentially distributed with mean value 150 time slots. After this lifetime period, the peer rejoins the system under a new pseudonym with clean transaction record and with the initial values of reputation and credibility. In this section, we assume that a peer belongs to the same authority throughout its consecutive lifetime periods during the operation of the system. Thus, a peer retains its probability to lie between two consecutive lifetime periods. This setting corresponds to a peer that periodically changes its pseudonym to clean its record of low performance and lying, but retains its probability to lie that has been learned in order to maximize its expected payoff. Under this model, each peer i follows a rational dynamic strategy according to which it selects its probability s_i to lie in feedback reporting according to a learning algorithm explained below. Upon lying, a peer follows the destructive lying strategy, while liar peers are supposed to be collaborated. According to the learning algorithm, after the expected lifetime period of a peer, its payoff (i.e. the number of time slots that the peer was not under punishment) is compared to the payoff of the peer at the end of its previous lifetime period and its probability s_i to lie is adjusted accordingly:

- If the probability s_i was previously increased during the last lifetime period and this was not beneficial, then decrease s_i in the next period. Otherwise, further increase s_i in the next period.
- If the probability s_i was previously decreased during the last lifetime period and this was not beneficial, then increase s_i in the next period. Otherwise, further decrease s_i in the next period.

The above scenario corresponds to a repeated game among peers that have two pure strategies: *Tell the truth* and *Lie*. Each peer tries to optimally choose its mixed strategy by learning. Consider that the initial fraction of peers that lie is very important for the evolution of the system; the credibility mechanism by nature demands a lower fraction of collaborated liar peers in the system than that of sincere ones in order to be effective. Therefore, we expect that the credibility mechanism provides the right incentives to peers only when fewer than 50% of peers of the system initially lie. Indeed, as depicted in Figure 8(a), if only 45% of the population initially lie with probability 1, while the other 55% of the population has zero lying probability (or equivalently if all peers have an initial lying probability 45%), then the peer-to-peer system asymptotically evolves to a stable equilibrium where all peers report their feedback truthfully. Thus, the credibility mechanism provides incentives to rational peers to be truthful. Otherwise, if 55% of the population initially lie, then the peer-to-peer system evolves to a stable equilibrium where all peers constantly lie in their feedback reports, as depicted in Figure 8(b).

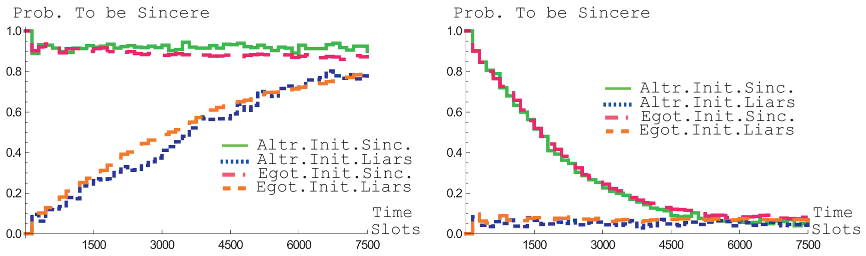

Fig. 8 a) The evolution of the mean probability of peers to rate truthfully, when the credibility mechanism is employed in a peer-to-peer system where initially 45% of peers lie. b) The evolution of the mean probability of peers to rate truthfully, when the credibility mechanism is employed in a peer-to-peer system where initially 55% of peers lie.

10 Implementation Issues

10.1 The Architecture

We have already demonstrated the effectiveness of our proposed mechanism for promoting credible reporting of feedback in a peer-to-peer system, as well as the right incentives provided thereby. Next, we discuss how this mechanism can be implemented in a completely insecure, anonymous and distributed peer-to-peer environment. The credibility information for each peer has to be efficiently stored and traceable. Authentication, integrity and non-repudiation of the credibility information and the feedback messages are also required. The security issues can be dealt with by means of the public-key infrastructure (PKI). Upon registering in the peer-to-peer system, each peer chooses a public-private key pair and creates his own certificate, which is signed by the system; that is, it is signed by a certain number of randomly selected peers, similarly to Pretty Good Privacy (PGP) [28].

Throughout the paper we have assumed that no peers are pre-trusted. Thus, we propose an implementation that does not rely on such a requirement. Peers are assumed

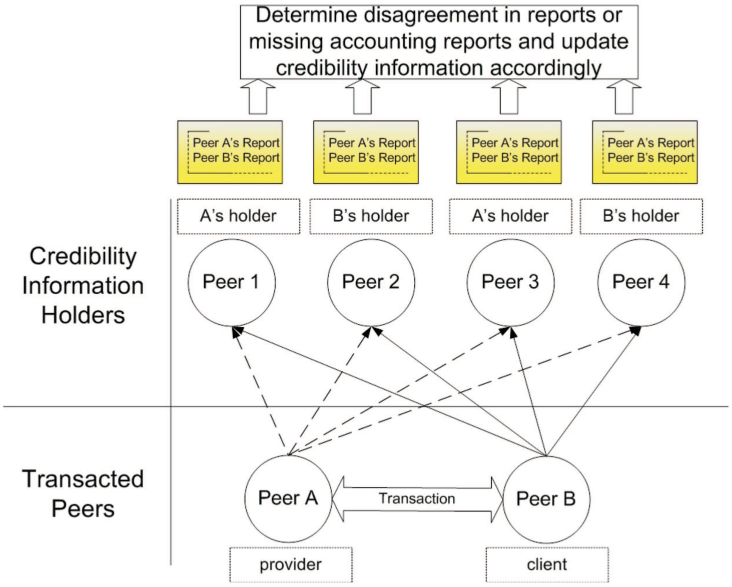

Fig. 9 Determining disagreement in feedback messages in a peer-to-peer environment.

to be organized in a hash-indexed structure enabling search of data. Such a structure is already available in systems such as Chord, P-Grid; see [4] and references therein. Peers are required to submit their feedback messages to other peers (referred to as credibility holders) based on their node identifier in the hash-indexed structure and on a number of hash functions employed for this purpose. Each peer is responsible for storing non-credibility values and punishment states of multiple other peers. Thus, multiple peers are responsible for holding credibility information of each fixed peer. After a transaction, each peer sends his feedback message (provider identifier, client identifier, rating and performance metric) and its digest signed by his private key to all peers that store credibility information of both transacted peers, as depicted in Figure 9. Peers that receive feedback messages verify the sender and the integrity of messages. Then, they detect agreement or disagreement of the feedback messages and compute non-credibility values and update the punishment states of the transacted peers as necessary. If only one feedback message is received, then this is also regarded as a disagreement and both transacted peers are punished. The credibility information is vulnerable to strategic modification by malicious peers. To avoid this, the credibility information provided by the majority of holders can be taken as valid. If there is enough redundancy in storing credibility information, then any malicious modification thereof can be observed by the peer himself. Indeed, the peer can monitor the credibility information about him periodically, by asking the corresponding information holders and comparing their responses. Thus, if a peer detects significant inconsistency in these responses, then the minority of holders should be punished for misreporting. The credibility holders of the misreporting peers should be informed for this inconsistency, which should be observable by these holders too. If there are fewer collaborated liars in the peer-to-peer system than sincere peers, then the inconsistency will be revealed and corrected, and the corresponding credibility information will be updated accordingly.

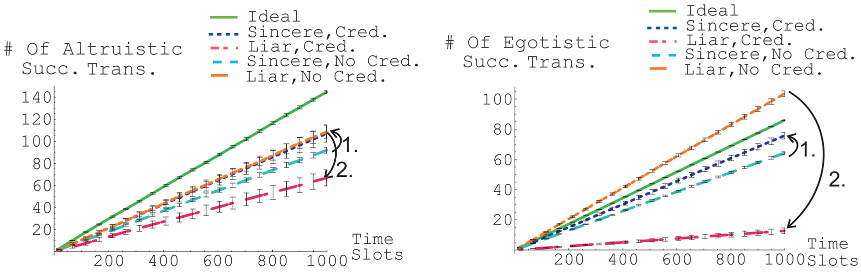

(a)

(b)

Fig. 10 The average number of successfully offered services to (a) altruistic and (b) egotistic peers when the credibility mechanism is employed or not in a system where peers store credibility information and 25% collaborative liars.

10.2 Collusive Liar Holders

In order to prove that our approach would still be effective in case that the holders of credibility information were collusive liars we have experimentally investigated such a scenario. In this scenario, holders of credibility information are normal peers, which have a fixed reporting type being either sincere or liar as before. Holder peers store and report credibility information (i.e. punishment state and non-credibility value) for specific peers selected by a hash function of the unique system identifier of the latter ones. The strategy that liar holders follow is *collaboratively opportunistic* in the sense that they always store and report negative credibility information for sincere peers and positive credibility information for their collusive partners. Specifically, liar holders:

- Always report agreement and discount non-credibility value for a liar peer that is involved as a client in a transaction.
- Always report agreement and record a positive vote for a liar peer that is involved as a provider in a transaction.
- Report disagreement and record a new punishment for a subject sincere peer that succeeded in a service provision.
- Report agreement and record a negative vote for a subject sincere peer that failed in a service provision.

Prior to a potential transaction between two peers, each of them asks the credibility holders of each transacted party for the punishment state of the latter. According to the credibility mechanism, only if neither of them is under punishment, the transaction should take place. However, the sincerity of the credibility information reported depends on the reporting type of the *majority* of the holders for each transacted party. After a transaction, the transacted parties submit their feedback to all the holder peers of both transacted parties. Credibility holders independently observe agreement or disagreement and store credibility information according to their reporting type. The credibility information stored by the majority will be taken into account in the future transactions of these peers.

There are two approaches regarding reporting minorities: i) ignore them, ii) punish them. The punishment of the minority holder peers after obtaining credibility information could be employed by reporting a disagreement to the respective holders of credibility information of each of them. We implement the aforementioned scenario with 10 credibility holders for each peer and ignoring minorities in the simulation model of Section 7 with parameters $N = 1500$, λ=10 peers/time slot, $r = 0.5$ and Uniform service availability with $q = 0.00275$. The punishment parameters are taken to be b=1.77 and ncr_0=1.46. Also, the Max-Max reputation based policy is employed in the system. However, for the clarity of results, reputation information is assumed to be accurately and centrally stored, as opposed to credibility information. The average number of successfully offered services for altruistic (resp. egotistic) peers in the presence of 25% liar peers when the credibility mechanism is employed or not is depicted in Figures 10(a) and 10(b) respectively. Therefore, if the percentage of collaborated liar holders is less than or equal to 25%, then there is no point in punishing reporting minorities, as in such a case, liar holders have a minor impact to the achievable efficiency of sincere peers. On the other hand, it has been observed in other experiments that if more than 25% liar holders are present in the peer-to-peer system, then punishing minorities is both necessary and effective for providing the right incentives to holders. Note that the punishment state of credibility holders (i.e. second-order credibility information) could also be employed for determining the credibility information of peers. However, then, the number of required messages for acquiring the credibility information of transacting peers would be multiplied by the number of holders.

11 Conclusion

In this paper, we have defined, analyzed and optimized a credibility mechanism that we first presented in [6, 7]. This mechanism provides incentives for truthful reporting of ratings' information in peer-to-peer systems by discovering and punishing liar peers. Based on a Markov-chain model, we found that the mechanism leads the system in beneficial steady states where almost all liar peers are under punishment, while almost all sincere ones are not. The punishment parameters of the mechanism were optimized for peer-to-peer systems of arbitrary characteristics by a fixed-point procedure. The optimization procedure was proved to be Pareto optimal and necessary for the effectiveness of the mechanism. Moreover, we experimentally proved that the optimized credibility mechanism combined with reputation-based policies assigns to sincere peers almost ideal benefit from the peer-to-peer system and diminishes the benefit of liar peers even when very high fractions of collaborated liars follow various static and dynamic rational strategies. Therefore, truthful reporting is *individually rational* and *incentive compatible* for peers under the optimized mechanism, which is thus *strategyproof.* Furthermore, the fractions of collaborated liars successfully dealt with by the credibility mechanism are the highest in the literature. Also, we have discussed the implementation of the mechanism in a real peer-to-peer system without central authorities for storing credibility information, and experimentally proved that the mechanism could effectively deal with up to 25% collaborated liars that follow opportunistic strategies. Overall, in this paper, we provided a *complete solution* against free-riding and low-performance in peer-to-peer systems.

As already explained, the optimized credibility mechanism is very effective in isolating liar peers. In further experiments omitted for brevity reasons, we found that different punishment approaches in case of a feedback disagreement do not improve this effectiveness. In particular, attempting to limit the unfairness introduced for sincere peers upon disagreement, an alternative approach could be to probabilistically punish the transacted peers according to their relative non-credibility values. Another approach could be, instead of probabilistically punish peers upon disagreement, to keep a counter of the potential punishment for each peer and increase it by the respective probability of punishment of each peer upon disagreement. Then, when the counter of potential punishment reaches 1 for a particular peer, then this peer is deterministically punished and his counter is set to 0. Both these alternative approaches fail to deal with large fractions of liar peers that were successfully dealt with by the original mechanism. We have also considered a potential improvement of the integration of the credibility mechanism with the reputation-based policies by weighting the importance of a vote by the non-credibility value of the client peer. Again, this approach was experimentally found to have almost no improvement to the efficiency of sincere peers.

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*, San Diego, CA, USA, June 2003.
2. T. G. Papaioannou and G. D. Stamoulis. Reputation-based Policies that Provide the Right Incentives in Peer-to-Peer Environments. *Computer Networks (Special Issue on Management in Peer-to-Peer Systems: Trust, Reputation and Security)*, 50(4):563–578, 2006.
3. B. Yu and M. P. Singh. Distributed Reputation Management for Electronic Commerce. *Computational Intelligence*, 18(4):535–549, 2002.
4. K. Aberer and Z. Despotovic. Managing Trust in a Peer-to-Peer Information System. In *Proc. of the 10th International Conference on Information and Knowledge Management*, New York, NY, USA, November 2001.
5. S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. EigenRep: Reputation Management in Peer-to-Peer Networks. In *Proc. of the Twelfth International World Wide Web Conference*, Budapest, Hungary, May 2003.
6. T. G. Papaioannou and G. D. Stamoulis. An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems. In *Proc. of the 5th IEEE/ACM International Symposium in Cluster Computing and the Grid*, Cardiff, UK, May 2005.
7. T. G. Papaioannou and G. D. Stamoulis. Optimizing an Incentives' Mechanism for Truthful Feedback in Virtual Communities. In *Proc. of the 4th International Conference on Autonomous Agents and MultiAgent Systems*, Utrecht, The Netherlands, July 2005.
8. C. Dellarocas. Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior. In *Proc. of the 2nd ACM Conference on Electronic Commerce*, Minneapolis, MN, USA, October 2000.
9. S. Ding, S. Zhao, Q. Yuan, X. Zhang, R. Fu, and L. Bergman. Boosting collaborative filtering based on statistical prediction errors. In *RecSys '08: Proceedings of the 2008 ACM conference on Recommender systems*, pages 3–10, New York, NY, USA, 2008. ACM.
10. M. Chen and J. P. Singh. Computing and Using Reputations for Internet Ratings. In *Proc. of the 3rd ACM Conference on Electronic Commerce*, New York, NY, USA, October 2001.
11. M. Schillo, P. Funk, and M. Rovatsos. Using Trust for Detecting Deceitful Agents in Artificial Societies. *Applied Artificial Intelligence*, 14(8):825–848, 2000.
12. E. Damiani, S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Managing and Sharing Servents' Reputations in P2P Systems. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):840–854, 2003.
13. L. Xiong and L. Liu. Peertrust: supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(7):843–857, July 2004.
14. M. G. Uddin, M. Zulkernine, and S. I. Ahamed. Cat: a context-aware trust model for open and dynamic systems. In *SAC '08: Proceedings of the 2008 ACM symposium on Applied computing*, pages 2024–2029, New York, NY, USA, 2008. ACM.
15. R. A. Malaga. Web-based reputation management systems: Problems and suggested solutions. *Electronic Commerce Research*, 1(4):403–417, October 2001.

16. R. Zhou, K. Hwang, and M. Cai. Gossiptrust for fast reputation aggregation in peer-to-peer networks. *IEEE Transactions on Knowledge and Data Engineering*, 20(9):1282–1295, September 2008.
17. M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-Riding and White-Washing in Peer-to-Peer Systems. In *Proc. of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, Portland, Oregon, USA, September 2004.
18. T.-W. J. Ngan, D. S. Wallach, and P. Druschel. Enforcing Fair Sharing of Peer-to-Peer Resources. In *Proc. of the 2nd International Workshop on Peer-to-Peer Systems*, Berkeley, CA, USA, February 2003.
19. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Honest Feedback: The Peer Prediction Method. *Management Science*, 51(9):1359–1373, September 2002.
20. R. Jurca and B. Faltings. An Incentive Compatible Reputation Mechanism. In *Proc. of IEEE Conference on Electronic Commerce*, Newport Beach, CA, USA, June 2004.
21. S. Goel, D. M. Reeves, and D. M. Pennock. Collective revelation: a mechanism for self-verified, weighted, and truthful predictions. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 265–274, New York, NY, USA, 2009. ACM.
22. R. Jurca and B. Faltings. Eliciting Truthful Feedback for Binary Reputation Mechanisms. In *Proc. of IEEE/WIC/ACM International Conference on Web Intelligence*, Beijing, China, September 2004.
23. P. Dewan and P. Dasgupta. P2p reputation management using distributed identities and decentralized recommendation chains. *IEEE Transactions on Knowledge and Data Engineering*, 99(1), 2009.
24. P. Resnick and R. Sami. Sybilproof transitive trust protocols. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, pages 345–354, New York, NY, USA, 2009. ACM.
25. C. Dellarocas. Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard. *Information Systems Research*, 16(2):209–230, 2005.
26. 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, Greece, May 2003.
27. P. Antoniadis, C. Courcoubetis, R. Mason, T. G. Papaioannou, G. D. Stamoulis, and R. Weber. Results of peer-to-peer market models., September 2004. Project IST MMAPPS: Deliverable 8. Available at: http://www.mmapps.info.
28. Pretty good privacy. http://www.pgp.com/.