Technische Universität Darmstadt

Department of Electrical Engineering and Information Technology
Department of Computer Science (Adjunct Professor)
Multimedia Communications Lab
Prof. Dr.-Ing. Ralf Steinmetz

Towards A Lightweight Incentive Scheme for Peer-to-Peer Systems

Technical Report SR08-2

By

Sebastian Kaune, Gareth Tyson, Andreas Mauthe,
Nicolas Liebau, Ralf Steinmetz

Contact: kaune@kom.tu-darmstadt.de
http://www.kom.tu-darmstadt.de

First published: May 4, 2008
Last revision: May 4, 2008

For the most recent version of this report see
ftp://ftp.kom.tu-darmstadt.de/pub/TR/KOM-TR-2008-02.pdf

Contents

Abstract

Peer-to-peer (p2p) systems rely on the contributions of peers to operate. In order to not only depend on the altruism of individual peers, a number of incentive systems have been proposed. However many of them are complex and suffer from high overhead. In this work we propose BioTrust, a lightweight incentive scheme. In contrast to other approaches, it does not require reputation histories and can operate with minimal overhead in terms of bandwidth, infrastructure, memory and computation. BioTrust is based on observations of cooperation as observed in biological systems, which show that individuals try to ally themselves with others that can best increase their own standing. Such observations are prevalent in evolutionary systems where individuals must weigh off their own individual needs against the need for cooperation to survive. Using these models, BioTrust has been developed to extract and modify some of the most useful properties of biological systems to work effectively in a p2p environment. It is shown through simulation that this scheme encourages honest peers, which leads to a close to optimal behavior in the system context as well as for single peers.

Chapter 1

Introduction

It has long been understood that the performance of peer-to-peer (p2p) systems rely solely on the cooperation of the member nodes. This realization creates a social dilemma for the users of such systems as the necessity to altruistically provide resources goes against the selfish desire to limit one's own personal sacrifice. This issue has led to the development of a number of incentive mechanisms that attempt to either force or incentivize users to contribute their own resources [HC06] [KhW03] [FLSC04] [GLBM01] [BAS03]. These solutions are not, however, without flaws as these intelligent strategies can lead to non-optimal results.

The main problem in this context is free-riders [AH00] who try to exploit others while not contributing themselves. In this work, we propose BioTrust, a new reputation based incentive scheme that aims to encourage honest users to participate in the system whilst successfully blocking free-riders. This system is inspired by many similarities observed between self-organizing systems and biological life. Both, the biological evolution of life and the distributed evolution of peer-to-peer networking face the constant conflict between self-optimization and the necessity to cooperate. This stems from the selfish biological desire to optimize one's own position while being aware of the need for external assistance to achieve this. In order to attain its goal an individual will attempt to seek association with those that best improve its own position [NS05]. These observations can easily transcend to the world of peer-to-peer systems through the shared desire of peers to optimize their own situation.

BioTrust utilizes these observations to build a light-weight mechanism derived from similar biological systems. In contrast to other approaches (e.g. [KSGM03] [SXL05] [NT04]) it does not require reputation histories and can operate with minimal overhead in terms of bandwidth, infrastructure, memory and computation. Further, an adaptive trust policy towards newcomers allows it to alter its own behavior based on experience rather than the rigid policies mandated in other systems [KSGM03] [SXL05].

Central to BioTrust's operation is a peer's desire to affiliate only with those of a respectable stature. This results in a society that ostracizes those that free-ride whilst supporting those of a respectable standing. The mechanism works on the lightweight binary classification of nodes as either *good* or *bad*. Intrinsic to this operation is that nodes attempt to interact only with those considered of good standing. The main goal of this work is to show that this binary policy offers high performance whilst requiring only a very limited overhead. It successfully ostracizes free-riders whilst, through an adaptive and intelligent forgiveness policy, ensures that strangers and erroneous classifications are allowed to easily redeem past assumptions. Within this work the different options are explored and it is shown what the best strategies are for individuals as well as the overall system.

The rest of this work is organized as follows. In Chapter II the background and problem context is discussed and a brief overview of other incentive schemes and the correlations to biological systems is given. Chapter III describes the design of BioTrust, and Chapter IV evaluates the properties of the proposed scheme. The work is concluded in Chapter V.

Chapter 2

Background and Problem Context

The performance of p2p systems depends highly on the amount of voluntary resource contributed from individual peers. However, providing *services* incurs costs (e.g. bandwidth and CPU resources). Since *rational* users act to maximize their own utility they are tempted to free ride [AH00] [SGG02]. Moreover, free-riders will often become *whitewashers* by leaving the system and then rejoining it as an anonymous entity, therefore avoiding any punishment related to their free-riding [FPCS06]. As a result, individual rationality counteracts social welfare and cooperation amongst peers becomes sparse unless a special incentive scheme encourages peers to contribute services to other participants.

When considering decentralized online environments without a central authority such as p2p systems, a well-designed incentive scheme has to meet several challenges in order to be robust, viz.:

- *Variety of user behaviors*: Users pursue different behavioral strategies which can vary in complexity. They can dynamically change their behavior, leave the system arbitrarily and join as newcomers.
- *Rationality of users*: User can be classified into two categories: *honest* and *dishonest*. The former are consistent with the system specifications and thus contribute to the system. The latter can be further divided into *selfish* and *malicious* users. Selfish users try to maximize their benefit at the expense of other user whereas malicious users try to cheat the system by breaking down system specifications.
- *Newcomers*: In general, it is impossible to distinguish whitewashers from so called legitimate newcomers and therefore a special strategy is required for how such *strangers* should be treated within a system.

2.1 Incentive Schemes for Cooperation in P2P

In the area of p2p, various incentive schemes have been proposed to tackle these challenges, e.g. [HC06] [KhW03] [FLSC04] [GLBM01] [BAS03]. They are based on methods and strategies such as inherent generosity, monetary payment or reciprocity. In the inherent generosity approach, nodes decide whether to share resources or not depending on the global generosity [FLSC04] of the participants, which is computed by a maxflow-based algorithm within the network. However, the computation overhead of this algorithm O(V^3) grows exponential with the system size, which makes the scheme not scalable. Monetary payment schemes are based on micro payment methods in which peers have simply to pay for the resources they consume [GLBM01] [JHB03]. However, many of these algorithms are impracticable since they require a centralized infrastructure for micro payments and accounting. In reciprocity-based schemes peers use historical information of past behavior of other peers to decide whether they want to share resources or not. These schemes can be further separated into *direct* reciprocity and *indirect*. Direct reciprocity assumes frequent repeated meeting between the same peers which might not be the case in large, diverse p2p environments. In contrast, indirect reciprocity [KhW03] [KSGM03] [SXL05] [NT04] allows peers to claim back their contributions from any peer but requires a secure accounting system to document the transactions.

2.2 Parallels between Biological Systems and P2P

Considering the tensions between individuality and communal social welfare in biological systems, many interesting parallels between p2p systems and these biological systems can be drawn.

Natural selection is assumed to favor the selfish that maximize their utility at the expense of others. However, many natural organisms and especially human societies are organized around cooperation and a certain level of altruistic behavior [NS05]. A number of studies try to explain this evolution of cooperation in human societies by indirect reciprocity models to answer the question why natural selection can promote unselfish behavior.

Nowak and Sigmund's groundbreaking survey uses a game theoretical model alongside computer simulations for [NS98]. They introduce a novel mechanism called Image Scoring, which implements a reputation based scheme. In this scheme, each individual has an image score as a global reputation value reflecting its cooperativeness among other individuals. This value increases on every occasion s/he provides a service and decreases when there is a possibility to help someone in need but no help is provided.

This model has been adopted by Feldman [FC05] for p2p systems by studying the effect of different patterns of defections arising in p2p environments (i.e. whitewashing and free-riding). Further, [NS98] has as a reference reputation scheme also inspired several other p2p incentive systems [NS05] [HC06] [FLSC04].

However, the drawback of Image Scoring mechanisms is that peers cannot represent their true strategic interests. Peers who do not contribute to the system are discriminated against. This discrimination will decrease their reputation and decreases the probability to get help in future. A rational peer in this setting should then use a strategy which only takes its own global trust value into account, because its gain can only be increased when it contributes to the system irrespective if a peer that is interested to receive a service has a high or low global trust value. Recent studies in biological systems [OI04] analyzing the Image Scoring mechanisms among several other alternatives have also confirmed this as a problem.

Chapter 3

Design of BioTrust

Inspired by the observations made in [OI04] and the main weakness of the Image Scoring mechanism, BioTrust is based on a reputation-based incentive scheme. In [OI04] the "keys to success" in reputation building schemes have been defined in the following properties: maintenance of cooperation among contributors, efficient detection of non contributors, punishment, justification of punishment, and apology and forgiveness. These are also taken into account in the design of BioTrust.

In an ideal reputation system it would be possible to access ubiquitous knowledge about a peer's history. However, due to the fact that p2p systems can consist of millions of peers, it is impossible to monitor and store all the actions taken by users and would certainly counteract the scalability properties of the p2p paradigm. To avoid such difficulties, a simple mechanism is needed that encodes as much information as necessary to express the past behavior of peers within the system. Thus, simple global trust values for peers are adopted in the proposed system, represented by a binary digit which can be either *0* or *1*. Further, this value is only based on the peer's last action. More precisely, a global reputation value of 0 indicates that a peer is in *bad standing* (B) whereas 1 indicates a *good standing* (G). Let P be the population of peers in our system. Then, the global reputation score of an individual is given by

$$r : P \rightarrow \{0, 1\} \tag{3.1}$$

3.1 Behavioral strategies

We define the way a peer uses the reputation scores as its behavioral strategy (denoted by $\vec{s}$). There are four possible situations in which a peer *i* wants to assess another peer *j* with respect to the reputation scores:

- both peers are in bad standing (s_{00})
- peer *i* is in bad standing whereas peer *j* is in good standing (s_{01})
- peer *i* is in good standing whereas peer *j* is in bad standing (s_{10})

- both peers are in good standing (s_{11}).

Thus, $\vec{s}$ consist of four components, for each component the behavioral strategy s describes whether to *cooperate* (C) or to *deny cooperation* (D).

$$\vec{s}: \{0,1\}^2 \rightarrow \{C, D\} \tag{3.2}$$

For example, altruistic peers would follow behavioral strategy $\vec{s}_{alt}= (s_{00}, s_{01}, s_{10}, s_{11}) = (C, C, C, C)$ whereas free-riders are described by free $\vec{s}_{free}= (D, D, D, D)$. Hence, there are $2^4 = 16$ different behavioral strategies in total.

3.2 Dynamics of global reputation values

BioTrust dynamically assigns reputation values to peers based on their last action within the system. More precisely, if a peer (acting as a service provider) takes an action A (either C or D) when there is the option of providing a service, BioTrust assesses the goodness of this action by using so called reputation transitions. In general, each reputation transition m depends on three factors:

- the reputation score of the service provider
- the reputation score of the service consumer
- the taken action A (either C or D) by the service provider.

This can be specified in the following formula:

$$m(r_p(i), r_c(j), A) \rightarrow \{0,1\} \tag{3.3}$$

Fig. 3.1 shows a state diagram of this transition process, highlighting the 8 possible steps between states leading a node to either a good or a bad standing.

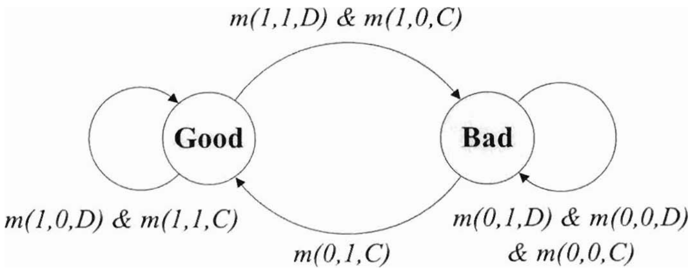

Figure 3.1: The 8 reputation transitions of BioTrust

3.3 Reputation Transitions: Basics

Although some of the basic assumptions behind reputation transitions seem to be obvious, the interdependency between them is crucial for the system performance, and the right selection is none trivial. Consider that there are generally $2^3 = 8$ possible reputation transitions due to the three influencing factors (cf. Formula 3.3) with each of them offering two alternatives to choose from. Each transition leads to either a good or bad standing. Accordingly, there are $2^8 = 256$ possible reputation transitions. If we combine them with all kinds of behavioral strategies, we have $256 \cdot 16 = 4096$ pairs in total to design our system.

The scheme adopted in BioTrust takes into account the key properties of successful reputation building schemes from evolutionary biology [OI04]. These schemes have been proven to be highly robust and stable in biological environments against different patterns of defection even in the presence of observation errors of individuals' reputation scores. In addition to this, we overcome the earlier stated weakness of the Image Scoring mechanism by using a special punishment policy. The basic insides behind the chosen reputation transitions will be explained in the following:

1. *Maintenance of cooperation: m(1,1,C)=Good.* It is important to ensure that nodes which have been proven to be cooperative and are contributing to the system maintain a high reputation score when meeting each other. Thus the following property seems to be intuitive and essential: If two nodes in good standing meet each other and cooperation in terms of providing a service takes place, the donor node has to keep its good standing.

2. *Identification of Dishonesty: m(0,1,D)=Bad and (1,1,D)=Bad.* A good incentive mechanism has to protect against nodes that only want to exploit the system and has to identify them quickly. This implies that a node not providing a service has to fall into bad standing, irrespective of the current reputation score. Thus, a potentially dishonest node will be excluded from the system.

3. *Forgiveness: m(0,1,C)=Good.* Sometimes it may happen that nodes which are cooperative in their nature fall mistakenly into bad standing due to observational errors of the underlying system (e.g. messages are getting lost, false reports). Thus, if a node recognizes this, there should be an opportunity to allow immediate forgiveness to regain a good standing again. Otherwise these nodes would experience refusal from future interactions and the social welfare of the system would decrease.

4. *Punishment and Justification of Punishment: m(1,0,D)=Good.* When a dishonest node is detected and identified, other nodes contributing to the system should refuse to provide services to it in the future. Hence a node in good standing encountering an opponent labeled as bad should refuse to provide services. In addition to this, a node in good standing who refuses to cooperate with a node in bad standing should not be punished for this. Note, this reputation transition is in contrast to the image-based incentive scheme proposed by [FC05] [NS98] in which the reputation score decreases after a peer denies granting a service regardless of the partner's status.

3.4 Rationality of User Behavior

The fundamental economic requirement to be achieved by the p2p system is that the benefits from participation exceed their costs. The benefits from participation lie in the received services. The costs of participation stem from providing a service and are predominantly caused by associated communication cost c. A rational node would save these costs by denying services. As long as there is no benefit associated with providing a service a peer's optimal strategy would be to free-ride causing the system to collapse. Thus, our incentive scheme has to ensure that the benefit b of peers having a higher global trust value must exceed the immediate costs of providing a service ($b > c$). We will show that our proposed scheme can enforce this if cooperative peers are using the behavioral strategy $\vec{s}_{BioTrust} = (s_{00}, s_{01}, s_{10}, s_{11}) = (D, C, D, C)$.

3.5 Adaptive Stranger Policy

When peers join the system they do not have a history and differentiation between whitewashers and legitimate newcomers is not possible; thus, they are all strangers. How they should be treated is therefore a problem since on the one hand always providing services to strangers encourages peers to free-ride and whitewash, whilst on the other hand, always denying services to strangers punishes potential altruistic and cooperative user behavior. This also aggravates the problem of bootstrapping. To tackle this issue, each peer therefore adopts an adaptive stranger policy in which peers treat strangers based on the experiences this peer had with strangers in the past. More precisely, peers provide services to strangers with probability q_s, which is given by:

$$q_s = q_{s-1} + \alpha \times [\sigma(X) - q_{s-1}] \tag{3.4}$$

Note that each peer encountering a stranger updates this probability each time this stranger has an occasion, X, to provide a service to him and is doing so (C) or not (D). Therefore, $\sigma(X)$ is either 1 or 0, respectively. Thus, using this policy we can highlight the cooperative trend of strangers and being generous to them when they are being generous. Short-term fluctuations can be smoothed out by a constant smoothing factor α. When the system bootstraps some peers initially have to cooperate, in order to generate the first positive trust values. Therefore, q_0 is set to 1.

Note that the adaptive stranger policy is solely based on local experience, does not build a history about individual peers, and therefore does not induce communication overhead.

Chapter 4

Evaluation

In the following section the effectiveness and robustness of BioTrust against different patterns of defection of dishonest users will be examined. We adopt a game theoretic model to evaluate which behavioral strategy is the dominant one among a set of chosen strategies by combining several aspects of the game theoretic models presented in [FC05][NS98][FLSC04].

4.1 Context of Evaluation

For the purposes of evaluating BioTrust it is assumed that peers have globally known binary reputation values, which are stored in the system and are accessible to all peers. Further, it is also assumed that users do not pollute the reputation values with collusions and false reports. This is a valid assumption since our goal is to initially explore the best strategy for individuals as well as the overall system.

In this system, peers act rationally and try to optimize their own *payoff*. Time is modeled in rounds. In each round peers are interacting as clients and servers. Accordingly, there exist the four following actions: client requests, client receives, server receives and server service provision. We will model the *benefits* and the *communication costs* of each interaction (the game) in our system. Each game consists of two peers randomly chosen from the population. One peer acts as a server having the opportunity to provide a service and one peer acts as a client interested in receiving this service. The outcome only depends on the decision of the server and thus on its *behavioral strategy*. Depending on how the server acts, both the client and the server will receive a payoff from the matrix depicted in Tab. 4.1. The values of this matrix are equivalent to [FLSC04] and are derived from a model for p2p file sharing applications. More precisely, by using this general form of an asymmetric payoff matrix, we satisfy the inequalities of the *Generalized Prisoner's Dilemma* (GPD) presented in [FLSC04] while still creating a social dilemma. Thereby, we are able to reflect the asymmetry of nodes interests in a p2p file sharing scenario where only the decision of the server is meaningful for the outcome of each interaction. Experiments with other asymmetric payoff matrices varying b and c ($b > c$) show similar results.

Payoff-Table Client/Server		**Server** Service Provided	Service Denied
Client	Request Service	7/-1	0/0
	Do not request service	0/0	0/0

Table 4.1: [FLSC04]: The payoff matrix for a p2p file sharing application

In order to concentrate on the basic principles we assume further that providing a service in each game incurs the same costs to all servers, and clients receive the same benefit, respectively.

To model how users tend to behave, peers can change their behavioral strategy after each round due to *learning effects* [FL98] or suffer *turnovers*.

- *Learning*: Each peer takes into account the moving average of the payoff of its current strategy and the moving average of the payoff of other peers it knows using the same strategy. If its current moving average achieves lower payoffs than the moving average of other peers, then the peer will switch to the most promising strategy. This assumption models that peers can exchange information on the success of their strategies.
- *Turnover*: Peers can leave or join the system arbitrary after each round. This is modeled by the substitution of a peer with a stranger who adopts the same behavioral strategy of the dropout.

4.2 Simulations

To study the robustness of BioTrust with a large number of peers, we have implemented a round-based simulator that corresponds to the above stated game theoretical model. We performed several simulation runs to reflect a particular situation in which two types of dishonest users are trying to exploit a p2p system dominated by honest nodes. We consider the first type as *traitors* as they acquire a good standing before turning into defectors. The second type is represented by *whitewashers* who are always free-riding, escaping from the consequences of punishment each round and joining in the next round as strangers again.

In our simulations, the population consists of 100.000 peers and is equally divided into whitewashers $\vec{s}_{white}$, altruistic peers $\vec{s}_{alt}$, traitors $\vec{s}_{traitor}$ and nodes following the behavioral strategy $\vec{s}_{BioTrust}$. Additionally, $\vec{s}_{BioTrust}$ -strategists are further divided into three equally sized groups differing in the way they treat strangers. Two groups follow a rigid stranger policy, either always (cooperate) or never provide services to strangers (defect), and one group is using the proposed adaptive stranger policy. Tab. 4.2 gives an overview about the four components of the behavioral strategies used in our experiments as well as the stranger policies they apply.

Further more, in each round, each node plays one game as a server and one game as a client. To calculate the payoffs our simulation framework uses the asymmetric payoff matrix presented earlier and we assume a fixed turnover rate of 0.1 which means that 10% of all peers leave and join in each round.

Name	s_{00}	s_{01}	s_{10}	s_{11}	Stranger
BioTrust	D	C	D	C	Adapt./Def./Cop.
Altruistic	C	C	C	C	Cooperate
Traitor	D	C	D	D	Cooperate
Whitewasher	D	D	D	D	Defect

Table 4.2: Overview about behavioral strategies.

4.3 Results

Fig. 4.1 and 4.2 show the results of our evaluation. Fig. 4.1 depicts peers of the different strategy groups switching to the strategy with the highest payoff over time as their rationality suggests them to do so by having the own payoff as optimization criteria (learning behavior).

Alongside this, Fig. 4.2 shows the achieved mean payoff of each strategy on average as well as the mean overall payoff of the all users of the whole system per round. Thus, the latter represents the degree of cooperation in the system. The highest level of cooperation would be 6 indicating that all peers are contributing to the system and everyone in the network is able to receive a service.

From these results it can be seen that both types of dishonest users cannot gain ground in BioTrust as they achieve low payoffs. Moreover, the $\vec{s}_{BioTrust}$ behavioral strategy seems to be especially effective against these patterns of defections as it achieves the highest payoffs over time and rational users very quickly switch to this strategy as it is the most promising one. Further, Fig. 4.2 indicates that this behavioral strategy drives the system to the highest level of cooperation (cf. mean average payoff of all users) by successfully eliminating dishonest users. Further, the effectiveness of the adaptive stranger policy can be observed as it achieves the highest payoffs over time compared to the rigid stranger policies applied by the two other $\vec{s}_{BioTrust}$-groups.

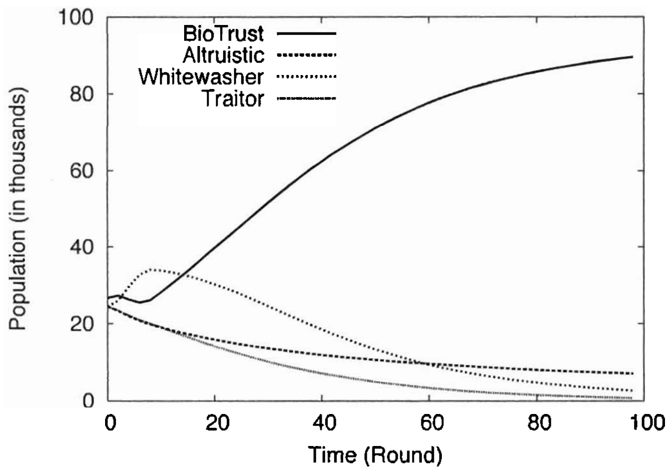

Figure 4.1: The dominance of BioTrust strategist. Time moves in rounds. Population indicates the number of peers pursuing a specific strategy.

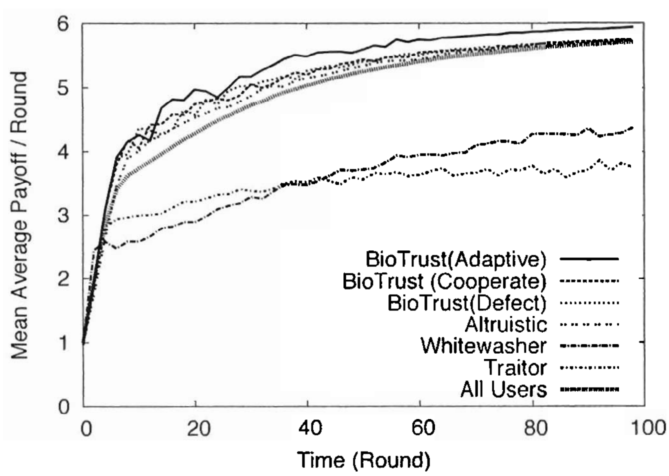

Figure 4.2: The mean average payoff per round of each strategy and of the whole system.

Chapter 5

Conclusion and Future Work

We have presented BioTrust, a robust p2p reputation-based scheme that extract key properties from an indirect reciprocity scheme which is proven to be robust in biological systems. BioTrust encourages honest users to participate in the system using an adaptive newcomer policy. Further, dishonest users who want to exploit the generosity of honest nodes can be efficiently eliminated. Apart from its robustness against these patterns of defections, BioTrust also comes with minimal overhead in terms infrastructural, computation and storage complexity since global reputation is represented as a binary value whilst still encoding the desired information.

The future work will now focus on two areas. Further studies related to the security constraints of the proposed mechanism are necessary to assess the impact of how collusion between dishonest peers or false reports of peers who lie about providing or receiving service can affect the system performance and stability. Moreover, in this work we have assumed that the reputation value of individual peers is globally known. This can, for instance, be achieved through adopting a system similar to EigenTrust [KSGM03]. In future, a gossiping based alternative will be explored that aggregates the subjective trust values of peers in a fully decentralized manner.

Bibliography

[AH00] Eytan Adar and Bernardo A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.

[BAS03] Chiranjeeb Buragohain, Divyakant Agrawal, and Subhash Suri. A game theoretic framework for incentives in P2P systems. *CoRR*, cs.GT/0310039, 2003. informal publication.

[FC05] Michal Feldman and John Chuang. The evolution of cooperation under cheap pseudonyms. In *CEC*, pages 284–291. IEEE Computer Society, 2005.

[FL98] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games.* Series on Economic Learning and Social Evolution. MIT Press, Cambridge, MA, 1998.

[FLSC04] Feldman, Lai, Stoica, and Chuang. Robust incentive techniques for peer-to-peer networks. In *CECOMM: ACM Conference on Electronic Commerce*, 2004.

[FPCS06] Feldman, Papadimitriou, Chuang, and Stoica. Free-riding and whitewashing in peer-to-peer systems. *IEEEJSAC: IEEE Journal on Selected Areas in Communications*, 24, 2006.

[GLBM01] Golle, Leyton-Brown, and Mironov. Incentives for sharing in peer-to-peer networks. In *CECOMM: ACM Conference on Electronic Commerce*, 2001.

[HC06] Ahsan Habib and John C.-I. Chuang. Service differentiated peer selection: an incentive mechanism for peer-to-peer media streaming. *IEEE Transactions on Multimedia*, 8(3):610–621, 2006.

[JHB03] Jakobsson, Hubaux, and Buttyan. A micro-payment scheme encouraging collaboration in multi-hop cellular networks. In *FC: International Conference on Financial Cryptography*. LNCS, Springer-Verlag, 2003.

[KhW03] H. T. Kung and Chun hsin Wu. Differentiated admission for peer-to-peer systems: Incentivizing peers to contribute their resources, June 09 2003.

[KSGM03] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in P2P networks. In *WWW*, pages 640–651, 2003.

[NS98] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity by image scoring. *Nature*, 393(6685):573–577, 1998.

[NS05] M. A. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, October 2005.

[NT04] Nikos Ntarmos and Peter Triantafillou. Seal: Managing accesses and data in peer-to-peer sharing networks. In Germano Caronni, Nathalie Weiler, and Nahid Shahmehri, editors, *Peer-to-Peer Computing*, pages 116–123. IEEE Computer Society, 2004.

[OI04] H. Othsuki and Y. Iwasa. How should we define goodness? - repuatation dynamics in indirect reciprocity. *Journal of Theoretical Biology*, 2004.

[SGG02] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. A measurement study of peer-to-peer file sharing systems. Technical report, July 23 2002.

[SXL05] Mudhakar Srivatsa, Li Xiong, and Ling Liu. Trustguard: countering vulnerabilities in reputation management for decentralized overlay networks. In Allan Ellis and Tatsuya Hagino, editors, *WWW*, pages 422–431. ACM, 2005.