Emerging Properties of Knowledge Sharing Referral Networks: Considerations of Effectiveness and Fairness

Priyadarshini Manavalan and Munindar P. Singh

Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206, USA
{ppaul2,singh}@ncsu.edu

Abstract. Referral-based peer-to-peer networks have a wide range of applications. They provide a natural framework in which agents can help each other. This paper studies the trade-off between social welfare and fairness in referral networks. The traditional, naive mechanism yields high social welfare but at the cost of some agents—in particular, the "best" ones—being exploited. Autonomous agents would obviously not participate in such networks. An obvious mechanism such as reciprocity improves fairness but substantially lowers welfare. A more general incentive mechanism yields high fairness with only a small loss in welfare. This paper considers substructures of the network that emerge and cause the above outcomes.

1 Introduction

Referral networks are a form of agent-based peer-to-peer systems [1]. Agents in such networks extensively use referrals to find other agents that can provide desired services. In knowledge-based referral networks, the focus of this paper, these services are primarily *knowledge services* [2]. For example, an agent seeking information on a subject searches for experts on the subject. Each agent maintains a set of neighbors, whom it contacts to initiate a search for experts. Unlike some conventional peer-to-peer approaches, we model the neighborhood relation as fundamentally asymmetric: Alice may not be Bob's neighbor even when Bob is Alice's neighbor. As a result, each agent can add or remove its neighbors unilaterally. Further, the in-degree of an agent may be larger or smaller than its out-degree, thus leading to interesting structures in the referral network.

Over time, as each agent finds or fails to find experts who can provide the knowledge services it requires, it may adjust its set of neighbors. The *local* adaptations of each agent cause the structure of the network to evolve. In many cases of interest, the agents evolves to form a stable network structure where most or all agents are able to obtain information more efficiently from the network. When the efficiency of individual agents in the network increases, so does the overall social welfare of the network.

In order to understand motivation behind an agent's interaction, we consider two key properties: performance and fairness. The *performance* of an agent at a specific time measures the usefulness of the surrounding network to the agent and indicates how capable the agents in the surrounding network are at providing information or referrals. The *fairness* experienced by an agent in a network measures how much the agent benefits from the network relative to how much work it performs.

Previous studies on referral networks focus on the properties of the network as a whole [2]. By contrast, we study the characteristics of agent interaction and have shown that in a typical referral network, performance and fairness are inversely related. This results in a structure with high agent exploitation or low performance. Autonomous selfish agents are not motivated to participate in such a setting. In addition, if we assume that most autonomous agents are selfish, their wellbeing usually takes precedence over the welfare of the network. In our study, we attempt to overcome this problem by experimenting with settings that create a network with both high performance and fairness.

Specifically, we model and consider three settings: Philanthropy, Reciprocity, and Incentives. Under *philanthropy*, our default typical network [2], agents help each other whenever they can. Under *reciprocity*, agents only help those who have helped them or whom they expect will help them. Under *incentives*, agents help others based on the incentives they receive from helping others; they can trade such incentives for their own searches, thus improving the value they obtain from the network.

Contribution. Through simulation, we find that Philanthropy is naive where although agents are successful and show high performance, the fairness of the network suffers. Some agents are heavily exploited. Reciprocity creates a fair network but the agents achieve low performance and are often unable to find the experts in the network. Incentives gets the best of both worlds: it yields fairness along with high performance.

Organization. Section 2 describes the specifics of our study, including the experimental setup and the key metrics. Section 3 describes the results of the experiments and the discussion. Section 4 concludes with a discussion of the literature and some future directions.

2 Technical Framework and Definitions

We can model a referral network as a directed graph each of whose nodes represents an agent and each of whose edges represents an agent (at the origin) having another agent (at the target) as a neighbor [3]. Each agent's *expertise* describes what knowledge it possesses and its *interest* determines what knowledge it seeks. Each agent generates outgoing queries based on its interest. Each agent may respond to an incoming query by giving an *answer* based on its expertise or a *referral* to one of its neighbors. An agent who sends out a query and receives a referral may, at its discretion, follow that referral by sending the same query to the target of the referral.

The performance of an agent reflects the good answers it can receive to its queries. Clearly, an agent's performance depends on its neighbors (modulo the setting, as we explain below). To explore the structure of the networks, we restrict each agent to have a small number of neighbors. Thus agents adapt to select neighbors that would yield them improved performance, in the process causing the network structure to evolve.

An agent's *acquaintances* are the agents with whom it has interacted. Each neighbor is also an acquaintance. Each agent maintains models that characterize the inferred expertise and *sociability* of each of its acquaintances [4]. The inferred expertise generally would not equal the actual expertise of the acquaintance. The sociability of an acquaintance corresponds to the presumed usefulness of the acquaintance in leading to a good answer to a prospective query.

Each agent evaluates the answers (if any) that it ultimately receives to its query. It upgrades the expertise of an agent that produces a good answer and simultaneously upgrades the sociability of the agents on the referral chain leading to that agent. For bad or no answers, it downgrades the expertise and sociability, respectively. Based on updates to its acquaintance models, an agent may modify its set of neighbors, in essence promoting some acquaintances to be its neighbors and demoting some neighbors to be mere acquaintances.

We use following metrics in our analysis.

- X: Set of agents
- E: The neighborhood relation
- $N_i = \{j : (i,j) \in E\}$: Set of neighbors of i
- $H_j = \{x : (x,j) \in E\}$: Set of agents of whom j is a neighbor
- $path(i,j)$: The path length of the shortest path from i to j
- I_i: The interest of agent i, modeled as a vector of dimension n
- E_i: The expertise of agent i, modeled as a vector of dimension n
- $\sigma_{j,i}$: Agent j's sociability of agent i

The similarity between two vectors of dimension n is given by

$$I \otimes E = \frac{\sum_{t=1}^{n}(i_t e_t)}{\sqrt{n \sum_{t=1}^{n}(i_t^2)}} \tag{1}$$

The Euclidean distance between two vectors of dimension n is given by

$$U \oplus V = \frac{e^{-\|U-V\|} - e^{-n}}{1 - e^{-n}} \tag{2}$$

The performance experienced by agent i is the summation of the contributions made by agents in the surrounding network. We define this as agents within a path of length $\log(|X|)$ from agent i. In most cases, these are the agents that provide responses to agent i. Agent j's contribution to agent i's performance is [2]:

$$\frac{I_i \otimes E_j}{path(i,j)} \tag{3}$$

The above metric reflects how similar the expertise of the agents in the surrounding network is to the agent's interest. The more similar the nearby agents are the better it is for an agent. For instance, if Agent A is interested in music and obtains high performance, this indicates that agent A's surrounding network contains experts in music.

The sociability of an agent i with respect to agent j measures i's usefulness to j. Agents that provide useful referrals tend to be rated at high sociability values and vice versa.

$$S(i) = \sum_{(j,i)\in E} (\sigma_{j,i}) \tag{4}$$

Interest Clustering measures whether the cliques formed by the agents reflect common interests among them [2]. Below $\gamma(i)$ compares i's interest with agents who are i's neighbors and have i as a neighbor. Informally, $\gamma(i)$ is high if the neighbors of i are neighbors with each other and have similar interests. Below $V_i = N_i \cup H_i$ is the set of agents that are either neighbors of i or of whom i is a neighbor.

$$\gamma(i) = \frac{\sum\limits_{(u,v)\in E} (I_u \oplus I_v)}{|V_i|(|V_i| - 1)} \tag{5}$$

PageRank measures the authority of an agent in the network [2]. An agent's PageRank depends on the PageRank of the agents of whom it is neighbor. The PageRank of each agent is divided equally among its neighbors, which makes the definition recursive. The following simplified definition of PageRank is adequate for our purposes and used to mesure authority under reciprocity.

$$P(i) = \sum_{j:(j,i)\in E} \frac{P(j)}{|H_j|} \tag{6}$$

The Relative Performance measures the benefit an agent receives as from others relative to the benefit it provides others. Below t_i and g_i are help taken and given, and equal the number of good responses received and sent, respectively.

$$R(i) = t_i - g_i \tag{7}$$

3 Experimental Results

We conducted a simulation study based on the above framework. Every agent is modeled with an interest and an expertise which remains constant over the course of the simulation. The network is seeded with each agent having some initial neighbors. Constrained only by the setting in effect, as described below, the agents generate queries in each round and exercise the referral process for each query. Therefore, we can reasonably compare the results across the three settings described below.

Philanthropy places no restrictions on an agent's interactions. Each agent always helps other agents whenever possible irrespective of how useful the other agents are to it.

Reciprocity is a variation of Philanthropy. The key difference is that, with Reciprocity, each agent helps only those agents in the network that have been helpful to it in the past or have high PageRank (which we use as a surrogate for reputation). Reciprocity ensures that agents who do not contribute to others eventually ceases to benefit from others.

If reciprocity is applied myopically, it has the risk of leading to agents not helping each other [5], because one failure by one agent to help a second agent is enough reason for the second agent to stop helping the first. To prevent this, we have each agent maintain the prospective value of each of its acquaintances. This value is adjusted upward based on good responses and downward based on bad responses. Each agent classifies its acquaintances into three primary groups and interacts with each group differently.

- High value. The agent responds to queries from high-value agents with direct responses if possible or referrals.
- Medium value. New acquaintances often fall into this category. The agent provides referrals but not answers.
- Low value. The agent disregards their queries unless they provide a referral from one of the agents' neighbors, in which it responds as usual.

Incentives is based on the idea—thinking of the incentives in monetary terms—that each agent pays for each response it receives. Each agent begins with a fixed endowment. But since each agent needs money to conduct a search, agents who help others continue to have funds to search, whereas agents who are not helpful eventually exhaust their endowments.

For a referral, an agent pays based on the quality of response received from the referral as well as the position of the referral in the referral chain. For a direct answer, the payment is predetermined. If two agents provide the same response, the response with the shorter referral chain is chosen. If this is not possible, the agent computes the similarity between its interest and the responding agent's expertise and chooses the response from agents with higher similarity. This is the same method adapted when agents do not have sufficient money to purchase all the responses received.

3.1 Agent Performance

We analyzed the outcomes of performance and fairness in a referral network based on three settings introduced above.

We expect that, as agents interact more in the network, their local performance increases and they locate the experts in the network. Moreover, the performance of an agent directly affects the manner in which the agent interacts and determines how its surrounding network evolves.

Figure 1(a) compares the performance of agents in the three different environments. Philanthropy yields higher local performance for most agents interacting in the network. Under Philanthropy, agents respond to each other freely. Thus each agent receives the best responses that it can from its surrounding network. Additionally, the number of interactions in the network is high.

Figure 1(a) also shows that under Reciprocity, the performance of each agent is significantly lower because fewer interactions take place. This is especially so for agents who do not contribute to the network.

Under Incentives, the number of responses an agent receives is proportional to the number of responses it gives. This is as in reciprocity. However, the agents don't need to have reciprocally matching interests with another agent in order to help them. The incentives can be traded in for responses from any agent. The agents, therefore have higher performance than when they are in a network defined with Reciprocity.

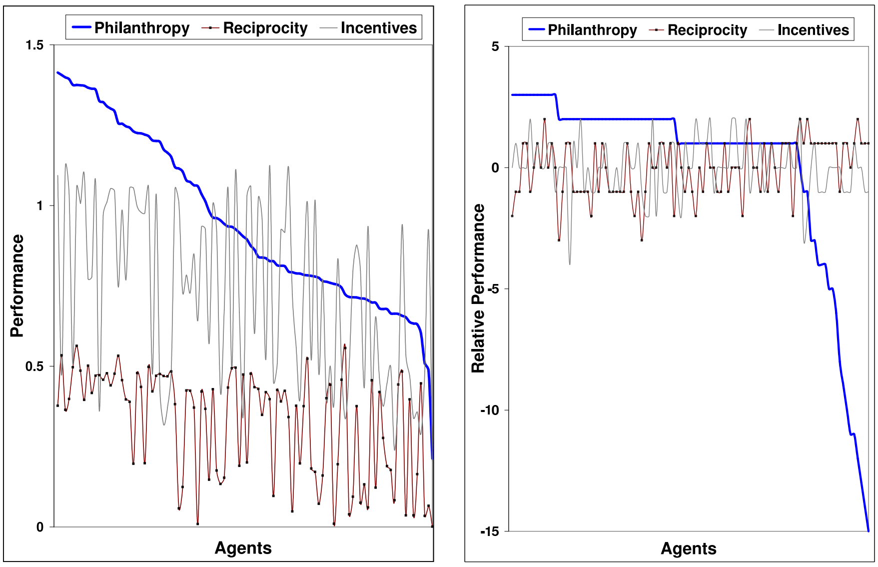

(a) Performance (b) Relative Performance

Fig. 1. Performance and fairness for a network of 100 agents; the X axes are agents sorted best to worst with respect to the Y axis for Philanthropy

3.2 Agent Performance and Clustering

The agents that have low performance are those that show high interest clustering and high cliquishness.

Figure 2(a) shows a common network structure for a low performing agent, called A. Network structures like these evolve over the course of the simulation if agents B, C, and D have similar interests as agent A. The interest clustering for agent A is high because the neighbors of agent A also have A as a neighbor. Moreover, A has a small surrounding network. The number of agents that are a distance of one from agent A is the same as the number of agents that are a distance of two from it. Therefore, the number of agents that A can reach is small and does not increase much over the course

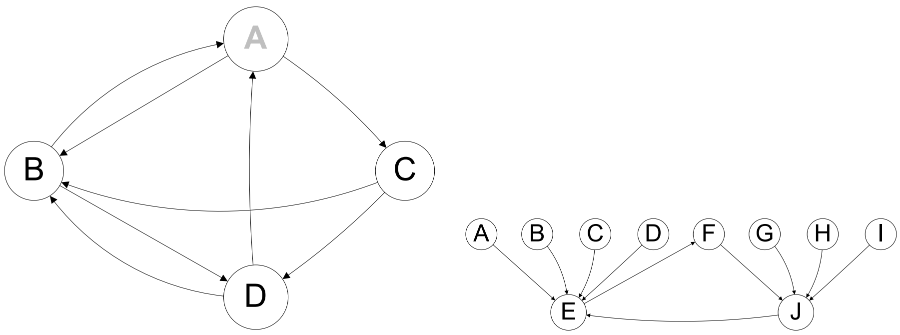

(a) Clustering (b) Exploitation

Fig. 2. Understanding the effects of Philanthropy: (a) Low performing agents show high interest clustering and cliquishness; (b) Situations where agents are exploited

of the simulation. The diminished size of the surrounding network causes a drastic reduction in A's local performance.

3.3 Fairness

The net benefit perceived by an agent is measured as relative performance: the number of responses the agent receives minus the number of responses it produces. A fair network is one in which all agents are treated fairly. That is, their relative performances are not widely distributed, which means each agent obtains a relative performance that is close to zero. An unfair network means that some agents are being exploited—they are the ones who do more work than they receive.

Under Philanthropy, agents may not receive sufficient help from the others. An agent may receive few or no responses, or responses of poor quality. Figure 1(b) depicts that, under philanthropy, over 20 percent of the agents in the network obtain low relative performance compared to other agents. This is depicted by the spread of the data points. In our simulation, relative performance ranges from -15 to +5. High negative values point to agents who are performing more work than they receive in the network. These agents are primarily those with high expertise values or high sociability values. Other agents gravitate toward them. As the exploitation of the agents increases, it leads to an additional problem in the network—the formation of bipartite graphs similar to the one shown in Figure 2(b).

Figure 1(b) shows that Reciprocity and Incentives result in a fair network. The range of the relative performance of the agents in the network is closer together on the vertical axis. The difference between the fairness values of both settings are very small. This is because both Reciprocity and Incentives control agent interaction and this enforces fairness. With Reciprocity, an agent only responds to those agents that have been helpful to it in the past and the responses given are usually good. Any agent that is not helping others receives limited help from others. No agent is exploited excessively. Under incentives, each agent that does not answer questions cannot ask any in return.

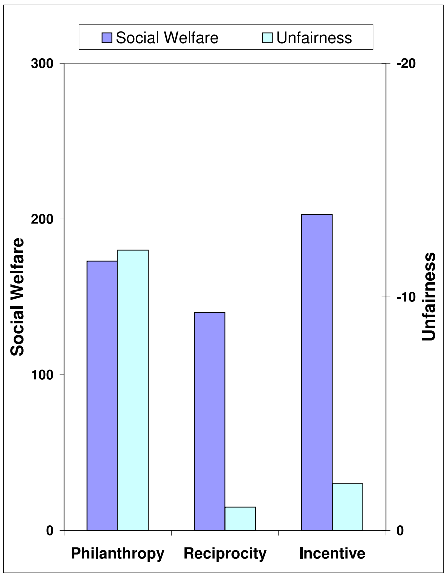

Fig. 3. Social welfare and Unfairness: for Philanthropy, Reciprocity and Incentives respectively

While Reciprocity and Incentives both result in a fair network, the number of interactions in the two models are significantly different. Reciprocity has the effect of reducing interactions. This results in formation of disjoint groups of agents. Over time, the cliquishness of the network can become much more pronounced than under Philanthropy. We observe that often the network splits into disconnected components. This is because agents choose to help only a select number of agents in the network.

The *social welfare* of a network is the summation of good responses received by all agents in it. Figure 3 compares the three settings in terms of Social Welfare. Under Philanthropy, every agent answers queries even from agents that have never helped it and therefore, Social Welfare is high. Under Reciprocity, Social Welfare is significantly lower than the other two settings. This is caused by the reduced interactions. Under Incentives, agents may ask questions as long as the they have money. As the agents find the experts in the system, they can obtain responses for a cheaper rate since the referral chains for the responses are shorter in length. Therefore, they pay less for referrals and this increases the number of questions that can be asked and therefore, the Social Welfare increases too.

When social welfare is compared with the degree of unfairness as in Figure 3 in the network, the incentive model emerges as superior in referral networks.

3.4 Performance and Fairness of Expert Agents

Fairness and Performance is specially important for experts in the network. These are the most valuable agents in our network and without their participation, the efficiency of the entire network would fall. The performance of the experts varies under the three settings. Figure 4(a) shows the performance of the top twenty experts under each setting. Most experts perform best under Incentives, in the middle under Philanthropy, and worst under Reciprocity. Under Incentives, as experts provide answers, they earn money and can ask more questions. Consequently, experts interact more and are able to receive a larger number of responses than the nonexperts. However, it is interesting to note that there are a small number of experts who are different. These experts show higher performance under a Philanthropic setting than an Incentive based one. In both Philanthropy and Incentives, the contribution of these agents to the network does not alter. However, they have an added advantage in the Philanthropic network since they benefit from their neighbors being exploiters and therefore indirectly exploit other experts themselves.

Additionally, figure 4(b) shows the relative performance of these agents. Since the dispersion of relative performance of the expert agents is low, we can conclude that they perform much better in terms of fairness with Reciprocity and Incentives, than Philanthropy. Agent exploitation has been significantly reduced. Therefore, autonomous agents who are classified as experts would prefer to participate in a Incentive setting as opposed to a Philanthropic one.

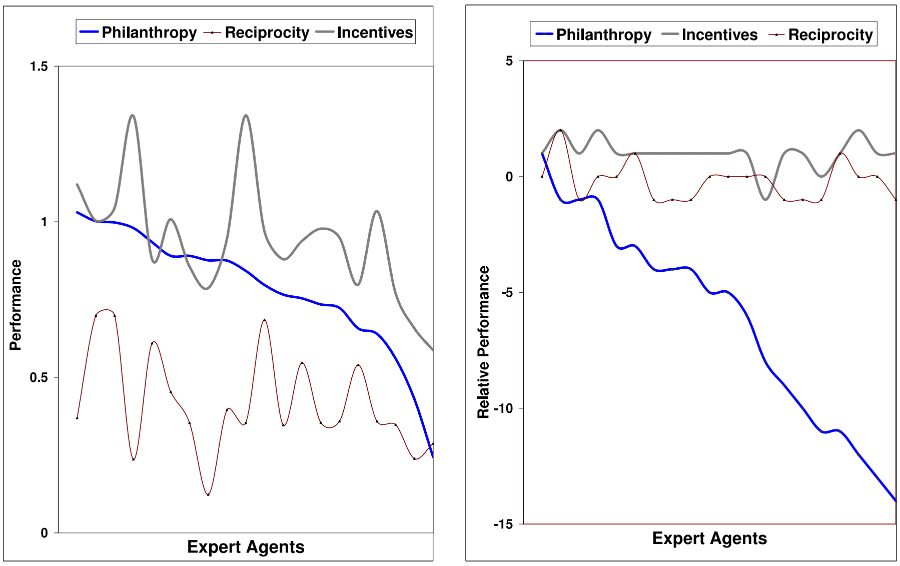

(a) Performance (b) Relative Performance

Fig. 4. Performance and Relative Performance for a network of 100 agents: (a) Performance of top 20 experts sorted best to worst for Philanthropy; (b) Relative Performance of top 20 experts sorted best to worst for Philanthropy

4 Discussion

Philanthropy results in networks where on average the agents perform well. However, the fairness of the network is reduced and, in particular, the experts are exploited. This often results in bipartite communities.

Reciprocity results in a network that shows high fairness but low social welfare. A fundamental shortcoming of Reciprocity is that it deals only with two-party interactions. For instance, if agent A fails to help agent B, the interactions between the two would mostly not proceed (unless there is a referral from another party). In other words, Reciprocity works best when two agents are such that each can help the other. Since such pairs of agents may be rare, a lot of potential social value is lost.

By contrast, Incentives naturally supports “trade” between multiple parties. This is why Incentives yields the best of both worlds. Under Incentives, we obtain networks with high agent performance. In particular, we find that experts perform well without being exploited.

4.1 Literature Review

Yolum and Singh [2] studied the emergent properties of referral networks with respect to the policies of agents for giving referrals or answers. Here we focus on the two properties, fairness and performance. We consider how this evolves in different settings and network structures. We focused on creating an environment in which the welfare of the network is not sacrificed for the wellbeing of the agent and vice versa.

In previous studies, researchers have tried to adopt policies that enforce agent cooperation in a network to decrease exploitation of individual agents. Hales and Edmonds apply the concept of *social rationality* to multiagent systems [6]. Agents in their study use tags to form socially rational groups and enforce cooperation among agents in the groups in the network. Hales and Edmonds extended this method to study cooperation among agents in peer-to-peer networks. However, by enforcing the tag system we limit agent interaction for the most part and agents are confined to social groups. This would lead to a structure with poor agent performance because the cliquishness of the network increases and the interaction decreases. Therefore, this model does not provide a solution to our problem. Additionally, in our simulation agents are unaware of the properties of other agents in the network and this creates an entirely different peer-to-peer network.

In other studies, the concepts of reciprocity and incentives have been applied to address the problem of *free riding*, i.e., the exploitation of some agents by others. Sen [5] compared the behavior of Philanthropic agents to Reciprocative agents and studied a probabilistic model of Reciprocity to increase cooperation among agents. Once again Sen’s work is limited to focus on the social aspect of agent behavior. He groups his agent into Philanthropic, Reciprocative or Selfish and compares the evolving structure. In comparison, we assume that agents are selfish and their wellbeing is more important that that of the network. We do not try to create a system of social cooperation but instead create a system where efficient agent interaction will lead to social welfare as a by product.

Yu and Singh [7] studied a dynamic pricing mechanism with the focus of studying the properties of incentive based models. In our simulation, we keep our incentive mechanism as basic as possible with fixed pricing policies. In addition, we only focused on how this mechanism affects the performance of agents with respect to fairness.

As a variety of policies based on Reciprocity and Incentives have been successfully applied to the problem of agent exploitation [5, 7–9] in previous studies, we chose to adopt similar mechanisms in our simulation too. We adopted a simple asymmetrical referral network setting based on these policies and focused on creating, not only a fair network but also an effective one.

4.2 Future Work

This paper has opened up several interesting problems. In future work, we will study different types of incentives, including a credit-based system that reveals further characteristics of agent interactions. We will consider mixed settings where different agents may follow different settings. The results of this paper indicates that results would improve if the agents reasoned based on the incentives to provide more referrals and increased the number of their interactions. Accordingly, we will study settings where we can incorporate strategic reasoning by the agents to maximize the incentives they obtain.

References

1. Foner, L.: Yenta: A multi-agent, referral-based matchmaking system. In: Proceedings of the 1st International Conference on Autonomous Agents. (1997) 301–307
2. Yolum, P., Singh, M.P.: Emergent properties of referral systems. In: Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS), New York, ACM Press (July 2003) 592–599
3. Singh, M.P., Yu, B., Venkatraman, M.: Community-based service location. Communications of the ACM **44**(4) (April 2001) 49–54
4. Yu, B., Singh, M.P.: Searching social networks. In: Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS), New York, ACM Press (July 2003) 65–72
5. Sen, S.: Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In: Proceedings of the 2nd International Conference on Multiagent Systems, AAAI Press, Menlo Park (1996) 322–329
6. Hales, D., Edmonds, B.: Evolving social rationality for MAS using "tags". In: Proceedings of the 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS), ACM Press (July 2003)
7. Yu, B., Singh, M.P.: Incentive mechanisms for peer-to-peer systems. In: Proceedings of the 2nd International Workshop on Agents and Peer-to-Peer Computing. (2003)
8. Krishnan, R., Smith, M.D., Telang, R.: The economics of peer-to-peer networks (2002) Carnegie Mellon University.
9. Golle, P., Leyton-Brown, K., Mironov, I.: Incentives for sharing in peer-to-peer networks. In: Proceedings of the 3rd International Conference on Electronic Commerce (EC). (2001) 264–267