

Game Theory As A Tool To Strategize As Well As Predict Nodes' Behavior In Peer-to-Peer Networks

Rohit Gupta and Arun K. Somani

Dependable Computing and Networking Laboratory

Department of Electrical and Computer Engineering

Iowa State University

Ames, IA 50011

E-mail: {rohit, arun}@iastate.edu

Abstract

In this paper we use game theory to study nodes' behavior in peer-to-peer networks when nodes receive service based on their reputation. Reputation is used as a mechanism to incentivize nodes to share resources and provide services to others. The probability of a node obtaining service is directly proportional to its current reputation, and the only way to enhance reputation is by serving others. Thus, the problem of free-riding is minimized. Game theory can be used by individual selfish nodes to determine their optimal strategy for participation level in such a system. Moreover, game theory gives us interesting insight into the overall nature of nodes' interactions and system efficiency, and how system efficiency can be improved.

Index Terms

Peer-to-peer, economics, incentives, reputation, game theory.

I. INTRODUCTION

Peer-to-peer (P2P) networks ([1], [2], [3], [4], [5]) are flexible distributed systems that allow nodes (called peers) to act as both clients and servers, and to access and provide services to each other. P2P is a powerful emerging networking paradigm and permits sharing of virtually unlimited data and computational resources in a completely distributed, fault-tolerant, scalable, and flexible manner.

Free-riding is widely acknowledged to be plaguing the current growth and widespread deployment of P2P systems. In [6] it is mentioned that almost 70 percent of the nodes in a Gnutella system never share their resources. Several mechanisms have been proposed to address this problem of free-riding: (1) monetary payments (service providers get suitably compensated by service receivers), and (2) reputation schemes (nodes with higher reputation get better service from others). The monetary payment scheme involves a fictitious currency, and requires an accounting infrastructure to track various resource transactions, and charges for them using micropayments. While the monetary scheme provides a clean economic model, it is difficult to implement such schemes in practice. The reputation based incentive model seems more promising.

In this paper we study the behavior of nodes in peer-to-peer networks when reputation is used as a mechanism to incentivize nodes to share resources and provide services. The probability of a node obtaining service is directly proportional to its current reputation, and the only way to enhance reputation is by serving others. This minimizes the problem of free-riding without relying on any centralized entity and/or coordination among peers. Coordination among peers is not required since peers decide their

The research reported in this paper is funded in part by Jerry R. Junkins Chair position at Iowa State University.

optimal strategies independently. The strategies used are symmetric, i.e. the same strategy is used by all the nodes (as opposed to protocols, which require different nodes to behave differently - this is difficult as all nodes in a P2P system assume the same role, thereby complicating different strategy assignment to different nodes).

Currently, P2P systems are used primarily for file sharing, such as audio, video etc. However, it is widely acknowledged that other resources, such as compute power, bandwidth, storage etc., can also be potentially shared using a P2P paradigm. Moreover, it is predicted that in future large ad-hoc grids would be organized in a P2P configuration to execute large-scale complex applications (see [8]). In such systems, stakes for individual peers would be higher as the cost of providing (and receiving) services would be very high (as opposed to almost zero cost associated with an audio file sharing, for example). It is thus reasonable to assume that all peers would behave selfishly in order to maximize the utility that they derive from the system.

Game theory [7] is an ideal tool to model a system with selfish nodes. We model the interaction of peers in a P2P system as an infinitely repeated game and compute the Nash equilibrium strategies (i.e. the participation level) of nodes in such a game. Peers use game theory principles to determine when they should or should not serve others. We treat each peer as a rational, strategic player, who wants to maximize its utility by participating in the P2P system. Peers gain utility by obtaining services (resources) and loose utility (i.e. incur cost) while serving others. Since probability of obtaining service is dependent on one's reputation (which is gained only by serving others), peers strategize their actions such that their overall utility is maximized, i.e. they serve at a minimum level that maximizes their probability of obtaining service in future. Moreover, a system designer can use the game theoretic notion of *Nash Equilibrium* to analyze the strategic choices made by different peers, and study the overall efficiency of the system (including how it can improved).

When nodes also derive utility out of altruism, the probability of providing service would be higher than that given by the model presented in the paper. Therefore, in some sense the model presented here provides a lower bound on the participation level for nodes below which they should not provide service.

The paper is organized as follows. Section II develops the model of interaction among selfish peers as a game, which we call as a *service game*. In this game, reputation is used as a basis on which peers receive service. In Section III we derive the pure- and mixed-strategy Nash equilibrium of the service game. Section IV examine several interesting and practical properties of the Nash equilibrium derived in Section III. Section V is on related work, especially on some other game-theoretic approaches that have been proposed to address the free-riding problem in P2P networks. We conclude the paper in Section VI.

II. SERVICE GAME

We assume the network lifetime to be infinitely long and divide the total time into individual time periods, represented by t for $t = 0, 1, \dots, \infty$. In every time period each node gets a request for service. The other activity of each node during a time period is to obtain service for itself. Every service request maps to one or more service providers, which can be requested in parallel or sequentially. A request is assumed to be fulfilled when any of the requested service provider agrees to serve. For simplicity, we assume that a node always request exactly one service and is also requested exactly once for service in every time period. (We assume a node's utility to be zero if it receives service more than once in a time period. Moreover, in actual implementation a node might receive multiple requests, however, some of those requests might be from low reputation clients and thus might be ignored. So the action {Serve} (as described below) corresponds to a situation when any one of the received request during a time period is served.)

We model the interaction among peers as an infinitely repeated game. A game is played during each time period. In a game, denoted by G , nodes request service for themselves and decide whether to serve others or not. Precisely, the game G is defined as follows:

Players: All the peers.

Actions: Each player's set of actions is $\{\text{Serve}, \text{Don't serve}\}$.¹

Preferences: Each player's preferences are represented by the expected value of a payoff function that assigns value U when service is received and cost C when service is provided.

The *service game*, which is an infinitely repeated version of game G (i.e. when G is played over and over again in successive time periods) is represented by G^∞ .

As stated earlier, the probability with which a player receives service is dependent on its current reputation. Reputation of player i in some time period t is denoted by R_t^i . Reputation of a node depends on its performance in the current time period as well as in prior time periods. Formally, we define R_t^i as follows:

$$R_t^i = R_{t-1}^i(1 - \alpha) + \omega * \alpha, \quad 0 \leq \alpha \leq 1 \quad t \geq 2 \quad (1)$$

In the above equation, ω is 1 when service is provided by player i in time period t , and 0 otherwise. Therefore, we have $0 \leq R_t^i \leq 1$, i.e. reputation of a player in every time period is always a value between 0 and 1 (including). Moreover, at $t = 0$ reputation of all players is 0 and at $t = 1$ reputation is given simply by ω . The parameter α is a constant and captures the importance assigned to the current performance of a player as opposed to its past performance for estimating its reputation. A high value of α means that more importance is assigned to a player's service in the current time period than its previous service record, and vice versa. Thus, when α is high, a node with even low reputation value can significantly improve its reputation by providing service in the current time period.

We assume that service information is readily propagated in the network and is available to all the players, i.e. each player when it serves another player propagates that information to as many other players as possible. The received service information is then recursively forwarded to other players. Therefore, it can be assumed that service information is propagated using a mechanism, such as flooding, and is available to all the players.

III. NASH EQUILIBRIUM OF THE SERVICE GAME (G^∞)

Now we evaluate the possible Nash equilibria of the service game (G^∞). By *Nash Folk Theorem* [7], if a^* is a Nash equilibrium action profile for some game G' then it is also the Nash equilibrium action profile when G' is played repeatedly infinite number of times.²

Therefore, finding the Nash equilibria of G^∞ reduces to finding the Nash equilibria of game G . We evaluate both pure- and mixed-strategy Nash equilibria of G . Since the proposed incentive mechanism based on a player's reputation links the benefit that a player draws from the system to its contribution - the benefit is a monotonically increasing function of a player's contribution, Thus this is a non-cooperative game among the players, where each player wants to maximize its utility. The classical concept of Nash equilibrium points a way out of the endless cycle of speculation and counter-speculation as to what strategies the players should use, and is defined formally below (from [7]).

¹The action corresponding to obtaining a service in a time period is not explicitly included in the model, as it is always assumed to take place.

²The notation a^* denotes the Nash equilibrium action profile of all the players, such that the Nash equilibrium action of player i in this equilibrium is denoted by a_i^* .

A Nash equilibrium is an action profile a^* with the property that no player i can do better by choosing an action different from a_i^* , given that every other player j adheres to a_j^* .

An equilibrium point is a locally optimum set of strategies (service probabilities, i.e. how much to serve others), where no player can improve its utility by deviating from the strategy.

A. Pure Strategy Equilibria

The action profile where all the players select the action {Don't serve} is a Nash equilibrium. This is because if any player i decides to serve, by selecting the action {Serve} instead, its payoff is $-C$, which is less than a payoff of 0 that it gets when it provides no service. The payoff of player i is $-C$ when it decides to serve because all the other players choose the action {Don't serve}, and therefore player i is unable to utilize its increased reputation to obtain service from others (and derive utility U in return).³

However, the action profile wherein all players choose the action {Don't serve} is an undesirable Nash equilibrium, since it means that no service is provided in the network. As a result the whole P2P system breaks down, and therefore, this equilibrium is an undesirable one. In light of this we argue that this action profile is an unstable equilibrium and is not likely to be reached (especially when there is also altruism among network nodes to some extent).

The action profile where all the players select the action {Serve} is not a Nash equilibrium. This is easy to see because if everyone else is serving requests than the best strategy for any player is to deviate by switching its strategy to {Don't serve}. By doing so a node gets a payoff of U instead of $U - C$ when it also serves.

Thus, we conclude that the only pure strategy Nash equilibrium of G is when players select the action {Don't serve} (such that action {Don't serve} is selected in each time period in G^∞), which, however, does not appear to be a likely convergence state for any useful P2P system.

B. Mixed Strategy Equilibria

We now consider the possibility of a mixed strategy Nash equilibrium of G , wherein players instead of deterministically selecting their actions randomize among their available set of actions. In other words, players select the action {Serve} in some time periods and the action {Don't serve} in others in G^∞ .

We want to find a symmetric Nash equilibrium because all the players belong to the same population (i.e. assume the same role) and it is therefore easier (i.e. requiring no coordination among players) to achieve such an equilibrium. (A *symmetric Nash equilibrium* is an action profile a^* , which is a Nash equilibrium and $a_i^* = a_j^*$ for any two players i and j . Stated simply, in a symmetric Nash equilibrium all the players take the same action (deterministically or probabilistically).) If the players in a game either do not differ significantly or are not aware of any differences among themselves, i.e. if they are drawn from a single homogeneous population, then it is difficult for them to coordinate, and a symmetric equilibrium, in which every player uses the same strategy, is more compelling. The argument of a single homogeneous population implies that all the peers in a P2P system have equivalent responsibilities and capabilities as everybody else.

Let there be such a mixed strategy Nash equilibrium in G , such that a player selects the action {Serve} with probability p and the action {Don't serve} with probability $1-p$. Here p is a non-zero value, i.e. both the actions are assigned positive probability by the mixed strategy of the player. Since the discussion here applies to all the players, therefore, for convenience we omit reference to a particular player and drop superscript i when defining reputation as given in Equation 1.

³Similar result involving non-cooperation among players is predicted by the classical Prisoner's Dilemma game.

The expected payoff to a player in time period t when it selects the action {Serve} is $p(-C + R_t^{serve} * U)$. This payoff value we denote by $Payoff_{serve}$. Likewise, the expected payoff to a player in time period t when it selects the action {Don't serve} is $(1 - p) * (R_t^{don't} * U)$. This payoff value we denote by $Payoff_{don't}$. A player's payoff when it provides service in a certain time period is $-C + R_t^{serve} * U$, which is the cost of providing service plus the utility it derives upon obtaining service. The term $R_t^{serve} * U$ captures the notion that the probability of obtaining service is directly proportional to one's reputation. Therefore, a player's expected payoff in a mixed strategy that selects the action {Serve} with probability p is given by $p(-C + R_t^{serve} * U)$. Likewise, one can obtain a player's expected payoff in a mixed strategy that selects action {Don't serve} with probability $1 - p$. (We have made an assumption that players first have an opportunity to serve others (and hence increase their reputation) before requesting service for themselves.)

R_t^{serve} is a player's reputation when it provides service in time period t and $R_t^{don't}$ is a player's reputation when it does not provide service in time period t . From Equation 1, we obtain

$$R_t^{serve} = R_{t-1}(1 - \alpha) + \alpha$$

and

$$R_t^{don't} = R_{t-1}(1 - \alpha)$$

An important characterization of mixed-strategy Nash equilibrium of finite games (one where action set of players is finite) is the following (from [7]).

Each player's expected payoff in an equilibrium is its expected payoff to any of its actions that it uses with positive probability.

The above useful characterization of mixed-strategy Nash equilibrium yields,

$$Payoff_{serve} = Payoff_{don't}$$

$$\Rightarrow p(-C + (R_{t-1}(1 - \alpha) + \alpha) * U) = (1 - p) * ((R_{t-1}(1 - \alpha)) * U) \quad (2)$$

Solving Equation 2, we get

$$p = \frac{R_{t-1}U(1 - \alpha)}{-C + 2R_{t-1}U(1 - \alpha) + U\alpha} \quad (3)$$

It must be noted that the value p obtained above is not a constant, but varies in each time interval depending upon a node's reputation at the end of the previous time interval. The mixed-strategy $(p, 1 - p)$ for actions {Serve, Don't serve}, respectively, is a mixed-strategy Nash equilibrium for the players. Assuming no collusion among nodes, if all the other nodes follow the above strategy, then the best strategy for any node is to also follow the above strategy (from the definition of Nash equilibrium). This is a symmetric mixed-strategy Nash equilibrium for G as well as G^∞ . We argue that it is a more stable equilibrium than the one in which no service is ever provided by the nodes. This is because of the following two reasons. First, when no service is provided the network is not useful to any user. Second, in practice users, which derive finite utility from altruism, would always provide some service irrespective of how much they obtain in return. Therefore, it is unlikely to have a scenario in which no cooperation among nodes take place.

IV. PROPERTIES OF THE NASH EQUILIBRIUM

In this section we study some interesting properties of the mixed-strategy Nash equilibrium derived above.

A. Simplicity of Calculating the Equilibrium Strategy

In the previous section we calculated the probability based on which nodes decide whether it is optimal for them to serve or not to serve. In each play of the game (or time period), players based on their reputation at the end of the prior time period decide whether they should provide service in the current time period or not. This probability as one can see does not remain constant from one period to another, and depends on a player's reputation at the end of the last time period. Players can calculate their reputation using Equation 1, since they know precisely their actions at each play of the game. Thus, determining the Nash equilibrium strategy is fairly straightforward for a player. It must be noted that there is an inherent assumption that peers get serviced based on their current reputation. The exact mechanism as to how that gets achieved (or is enforced) is outside the scope of the paper.

Figure 1 gives an example of how a peer's reputation might change over time by following the equilibrium strategy proposed above. In the figure, an increase in the reputation value corresponds to time intervals when service is provided by the peer, and vice versa. As can be seen, the equilibrium strategy of players guarantees that over time their behavior, in terms of providing service to others, is similar to each other (i.e. independent of the initial state or reputation value).

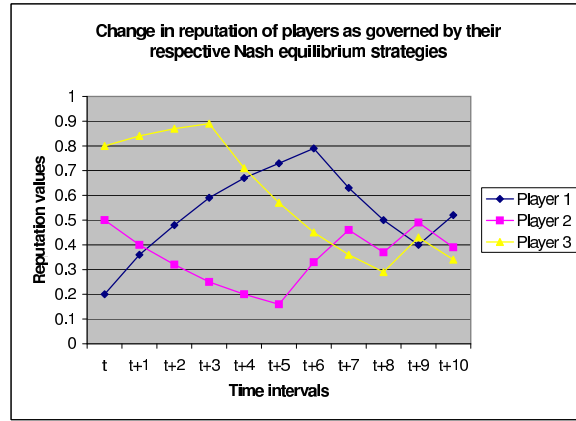


Fig. 1. The figure shows the change in reputation values of players over time starting at time period t . The three players are assigned arbitrary reputation values to begin with and the results are shown for α equal to 0.5. Also, we set $U/C = 100$.

B. Addressing the Problem of Free-Riding

Since the mixed-strategy Nash equilibrium assigns positive probability to the action {Serve} (i.e. players serve others with a positive probability in each time period), the problem of free-riding is minimized in the network. The simple game theoretic model presented here, wherein reputation is used as a basis for providing service, predicts that it is in every peer's (including free-riders') best interest to serve others. Our simulations support this behavior as we found that the total service received by a node is balanced by the total service that it has to offer to others, as shown in Figure 2.

C. 50 Percent Rule

An important property of the equilibrium emerges from Equation 3 that predicts the probability with which one should serve others. If we set $C \ll U$ (i.e. C can be ignored in Equation 3), then we have $p < 0.5$. In other words, if cost of providing service is negligible, Nash equilibrium of the service game predicts that players should serve each other less than 50 percent of the times when requested for service.

		alpha			
reputation		0.2	0.4	0.6	0.8
	0.2	14, 14	4, 4	2, 3	2, 2
	0.4	16, 17	3, 4	1, 1	0, 1
	0.6	12, 12	4, 5	2, 2	1, 1
	0.8	17, 18	4, 5	2, 2	1, 1

Fig. 2. The figure shows the total instances of service provided and received by a node over a period of 20 time intervals (first and second of the pair of numbers in each box represent the total instances of service provided and received, respectively). Irrespective of the initial reputation of a node and the value of α , the service received by a node is almost completely balanced by the service that it has to provide to others. Again, we set $U/C = 100$.

This, although it appears to be very restrictive, is a consequence of the fact that all peers are selfish and are better off free-riding than serving others. Intuitively, if a peer knows that everyone else behaves selfishly, i.e. provide as little service as possible, then the best strategy for the peer cannot be to serve others most of the time (i.e. with probability greater than 0.5).

In terms of Nash equilibrium, the above result can easily be understood by the following simple counter-example. If all players are known to service requests most of the time then any player can easily increase its payoff by switching to a strategy where it free-ride on others (i.e. serve with probability less than 0.5). Thus, an action profile where peers generously serve each other cannot constitute a Nash equilibrium.

We believe that the above result is an important outcome of our game theoretic model of nodes' interaction in a P2P system, where all participants behave selfishly. Although, the above result is intuitively appealing, our model provides a proof based on game theory that explains such a behavior of peers.

D. Fairness - Equal Sharing of Cost of System Inefficiency

From our discussion in the previous subsection, where we concluded that serving with probability less than 0.5 is an optimal strategy (when $C \ll U$), one can see that the overall system efficiency is severely reduced. This is because at least half of the service requests in the system are not fulfilled.

However, on the positive side, the equilibrium strategy provides fairness in the sense that the cost of system inefficiency is not borne by a single peer, but is shared (in inverse proportion to one's reputation) among all peers. This is because each peer's request is likely to be turned down by the serving peer. (We assume that if a peer's request at one peer is turned down it tries at some other candidate peer capable of serving the request. On average, the probability that a peer's request is successfully served in a time period is proportional to its current reputation).

E. Decreasing α for Higher Contribution

As can be seen from Figure 3, a lower value of α shifts the service probability curve upwards. In other words, when α is low peers serve each other with higher probability. This is to be expected, since

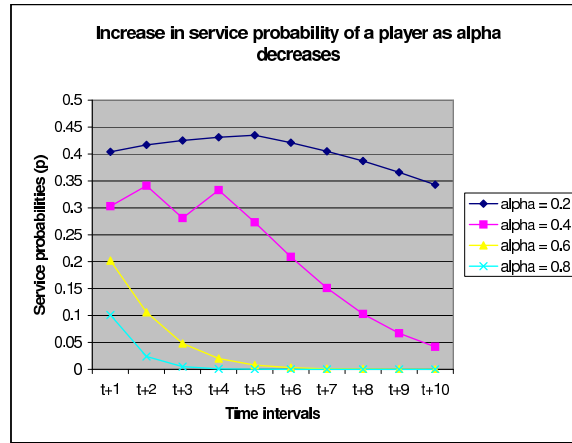


Fig. 3. A decrease in value of α require players to serve each other with greater probability. An initial reputation value of 0.5 is used in all the cases for different values of α . Also, we set $U/C = 100$.

α determines how much importance is given to a peer's current performance as compared to its past service record. A low value of α (i.e. giving more importance to nodes past actions up to the current time period) means that peers need to continually provide service to be able to maintain high reputation and access service from the system. On the other hand, if α is high peers can easily increase their reputation in any period in which they provide service, irrespective of how cooperative they have been in past with regards to providing service to others.

Thus, a simple way to improve the system efficiency is to set α as low as possible. As shown in Figure 3, as we decrease α the value of p tends to 0.5 (which is the maximum possible service probability in a system with selfish peers when $C \ll U$; the same result can be obtained directly, if we set $C = 0$ and $\alpha = 0$ in Equation 3).

V. RELATED WORK

Almost all the current research in P2P systems to overcome the free-riding problem relies on one of the following two incentive approaches for nodes to cooperate - reputation ([11], [12], [13], [14], [15], [16], [17]) and monetary schemes ([18], [19]). Most of the reputation based schemes assume that there are only a small set of malicious peers that do not provide service. It is assumed that most of the nodes are good and they would monitor the contribution of others to ensure that free-riding does not take place.

In a P2P system, a model where all the peers behave selfishly, however, appears to be more appropriate. This is true as P2P systems get used for more sophisticated application, such as distributed/grid computing, rather than simply sharing of music files, for example. In such scenarios, users have incentive to behave selfishly as stakes are typically higher. Recently game theory has been used to model the behavior of nodes in P2P systems. Below we describe two such recent papers. The first paper uses reputation and the second paper uses money as a form of incentive to motivate peers to cooperate.

A. [Chiranjeeb et. al.]

In [9] the authors use game theory to study the interaction of strategic and rational peers, and propose a differential service based incentive scheme to improve the system's performance. The authors first consider a simplified setting of homogeneous peers, where all peers derive equal benefit from everybody else. In this case it is shown that there are exactly two Nash equilibria, and there are closed form analytic formulae for these equilibria. The stability properties of these equilibria were investigated and it was

shown that in a repeated game setting, the equilibrium with the better system welfare is realized. The authors use the symmetry of a homogeneous system to reduce the interaction of a peer as a two-player game. This game is modeled as a Cournot-duopoly and the Nash equilibrium contribution of the players is calculated. The result is then extended to a N -player game.

For a system with heterogeneous peers, it was concluded that no closed form solution is possible and so simulation was used to study such a setting. The main findings were that the qualitative properties of the Nash equilibrium are impervious to (1) exact form of the probability function used to implement differential service, (2) perturbations like users leaving and joining the system, (3) non-strategic or non-rational players, who do not play according to the rules.

We believe that the a major drawback of the proposed differential service mechanism is the difficulty of its implementation. This is because of the following two reasons: (1) It is assumed that a peer wishing to join the system first determines the benefit that it can derive from the system. If the benefit is larger than a critical benefit, then the peer's best option is to join the system and operate at the Nash equilibrium value of contribution. If on the other hand the benefit is less than a critical value, the peer is better off not joining the system. The benefit for a peer is the amount of resource contributed by other peers as well as the utility of those resources to the peer. Obtaining such information beforehand can be difficult and the issue is not adequately addressed in the paper. (2) Every request from a peer contains a metadata describing the contribution of the peer to the system. The differential service received by the peer is dependent on this information. Since peers have incentive to manipulate this information, the authors propose to use a *neighbor audit scheme*, in which peers continually monitor the contributions of their neighbors. We feel that such altruism on part of the peers is not reasonable to assume, especially when the underlying game model is uncooperative and peers behave selfishly.

B. [Philippe et. al.]

The authors of [10] examine the design implications of the assumption that users selfishly act to maximize their own rewards. The authors construct a game theoretic model of the system and analyze equilibria of user strategies under several payment mechanisms. The idea is to encourage users to balance what they take from the system with what they contribute to the system. This is done by charging users for every download and rewarding them for every upload. The payment mechanisms differ primarily in the way that this amount of money is determined.

A micro-payment mechanisms was described in which a centralized server at the end of each time period charges an agent the amount of money proportional to the difference between the number of downloads and uploads. Since users do not like micro-payments (having to decide before each download if a file is worth a few cents imposes mental decision costs), two variations of the basic scheme were also proposed. First is a quantized micro-payment mechanism in which users pay for downloads in blocks of b files, where b is a fixed parameter. At the end of a time period, the number of files downloaded by a user is rounded up to the next multiple of b , and the user is charged for this many blocks. Second is a point-based mechanism, which uses an internal currency called "points" instead of micro-payments. Like quantized micro-payments, this mechanism lets users trade a fixed amount of dollars for a block of b points. Even though files are paid for with points on a per-file basis, where dollars are concerned, the mechanism is essentially quantized.

The proposed micro-payment mechanism relies on a centralized entity (like an index server in Napster [1]) to keep track of all the file transactions in the network and accordingly charge the agents or nodes. In a completely distributed P2P network how the above payment mechanisms can be implemented is not clear. Also, the authors at present consider only three levels of user strategies for download (uploads) - no downloads (no uploads), moderate downloads (moderate uploads), and heavy downloads (heavy

uploads). It was concluded that both heavy downloads and uploads constitute a unique Nash equilibrium strategy for the agents. But how the users should behave exactly within each level of strategy for both downloads and uploads is not clear.

VI. CONCLUSION AND FUTURE WORK

In this paper we have proposed a simple mechanism based on nodes' reputation to overcome the free-riding problem prevalent in P2P systems. The model presented here addresses the free-riding problem as it predicts that even for selfish users serving others is the best strategy. Game theory is used to predict the optimum (Nash equilibrium) strategies of selfish nodes such that their profits are maximized. Game theory is also used to provide valuable insight into the behavior of individual nodes, as well as the performance of the overall system. Interestingly, game theory provided us a proof for some of the intuitive results, such as the strategy of serving less than 50 percent of the times when the cost of providing service is negligible.

The proposed game theoretic solution of the free-riding problem has several significant advantages - fairness, simple implementation, and ease of calculating optimum strategies.

For simplicity we assumed that the cost and utility attached with serving and obtaining services, is the same for all service types and for all users. In future, we want to develop more elaborate models that takes into account the heterogeneity of service types and users. Also, we would investigate the applicability of game theory in P2P systems when heterogeneity is the over-riding factor in designing of protocols for system operations.

REFERENCES

- [1] Napster. <http://www.napster.com/>.
- [2] Gnutella. <http://gnutella.wego.com/>.
- [3] KaZaA. <http://www.kazaa.com/>.
- [4] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A Scalable Peer-to-peer Lookup Protocol for Internet Applications. In *Proceedings of the 2001 ACM SIGCOMM Conference*, 2001.
- [5] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proceedings of ACM SIGCOMM (San Diego, 2001)*, 2001.
- [6] E. Adar, and B. A. Huberman. Free Riding on Gnutella. *First Monday*, vol. 5, No. 10, Oct. 2000.
- [7] M. J. Osborne. A course in game theory. Cambridge, Mass. : MIT Press, c1994.
- [8] I. Foster, and C. Kesselman. The Grid: Blueprint for a New Computing Infrastructure, 2nd Edition, Morgan Kaufmann, 2004.
- [9] C. Buragohain, D. Agrawal, S. Suri. A Game Theoretic Framework for Incentives in P2P Systems. In *Proc. of the Third International Conference on Peer-to-Peer Computing (P2P'03)*, 2003.
- [10] P. Golle, K. Leyton-Brown, I. Mironov, and M. Lillibridge. Incentives for sharing in peer-to-peer networks. In *Proc. of 2001 ACM Conference on Electronic Commerce*, 2001.
- [11] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. EigenRep: Reputation Management in P2P Networks. In *Proc. of ACM World Wide Web Conference, Budapest, Hungary*, May 2003.
- [12] K. Aberer and Z. Despotovic. Managing trust in a peer-2-peer information system. In *Proc. of the Tenth International Conference on Information and Knowledge Management (CIKM '01)*, pages 310-317, 2001.
- [13] H. T. Kung, and C. H. Wu. Differentiated Admission For Peer-To-Peer Systems: Incentivizing Peers To Contribute Their Resources. In *Proceedings of the 2003 Workshop on Economics of Peer-to-Peer Systems, Berkeley CA*, 2003.
- [14] T. W. "Johnny" Ngan, D. S. Wallach, and P. Druschel. Enforcing Fair Sharing of Peer-to-Peer Resources. In *Proceedings of IPTPS 2003, Berkeley*, February 2003.
- [15] M. Gupta, P. Judge, and M. Ammar. A reputation system for peer-to-peer networks. In *ACM 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, 2003.
- [16] E. Damiani, S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Managing and sharing servents' reputations in P2P systems. *IEEE Transactions on Data and Knowledge Engineering*, 15(4):840-854, July/August 2003.
- [17] S. Marti, and H. Garcia-Molina. Limited reputation sharing in P2P systems. In *Proc. of the 5th ACM conference on Electronic commerce, New York, NY, USA*, 2004.
- [18] V. Vishumurthy, S. Chandrakumar, and E. G. Sirer. KARMA: A Secure Economic Framework for Peer-to-Peer Resource Sharing. In *Proceedings of the 2003 Workshop on Economics of Peer-to-Peer Systems, Berkeley CA*, 2003.
- [19] Mojo Nation. <http://www.mojojunction.net/>