Model and simulation study of a peer-to-peer game with a reputation-based incentive mechanism

Bita Mortazavi* and George Kesidis
CS&E and EE Depts
The Pennsylvania State University
University Park, PA, 16802
mortazav@cse.psu.edu and kesidis@engr.psu.edu
*Also a Member of Technical Staff at Verizon Wireless

***Abstract*— A distributed, cumulative reputation system is used to encourage resource sharing and cooperation among participants in a peer-to-peer system, specifically a content distribution (file-swapping) network. The reputation system that is employed is a variation of one that was shown to be able to reveal true propensities to cooperate of the peer nodes. The actions of peer nodes are herein modeled with a game in which users modify the amount of uplink bandwidth they allocate to the network. In principle, increasing uplink bandwidth will result in improved reputation that will, in turn, result in improved downloading performance. This was demonstrated by our simulation results.**

I. Introduction

Peer-to-peer (P2P) systems are currently receiving a significant amount of attention. The traffic generated through such applications has been shown to be dominating in the Internet [5]. The attraction of these systems, when compared to client/server frameworks, is in their robustness, reliability and cost efficiency. Among P2P networks, content distribution (file swapping) networks (CDNs), such as Kazaa (or "KaZaA") [4], Gnutella [3], eDonkey [2] and Bit-Torrent [1] are very popular essentially because of their scalability. However, the rationally selfish behavior of users[1] who benefit from communal resources but do not cooperate by sharing theirs' (i.e., free riding) has been shown to cause performance degradation in peer-to-peer CDNs [12], [24], [26], [27].

Reputation systems [9], [20], [17], [28], [23], [14], [22], monetary incentives [21], [25] and rule-based incentives [1], [10], [7], are different approaches proposed to address the free riding phenomena in P2P networks. Participating peers rank each other in response to directly perceived or referred behavior. Ideally for a CDN, the most cooperative nodes are those with highest reputation rankings and receive better service in return. Reputations are cumulative in nature and provide indirect incentives. On the other hand, monetary incentives more directly reward those who cooperate by sharing resources (e.g., micropayments approach [21], [25]). In rule based incentives, the system itself governs the behavior of the peers by imposing rules that force peers to share some of their resources while compensations are obtained by the peers upon their contributions.

In [11], [18], [19], peers play a game in hopes of maximizing their own utility (their "cumulative contribution" acts as a reputation). The game is designed so that the peers need to maintain a level of cooperation in sharing their bandwidth resources for an equilibrium to exist. We recently [22] introduced a cumulative reputation framework that, for special cases, can be proved to converge in mean to reveal the true propensity of peer nodes to cooperate. Based on this framework, a game was formulated in which users play to maximize the success rate of their queries for files by adjusting their cooperation level and thereby improving their reputation. A simulation study (considering weighted voting, hierarchical trust groups and misrepresentations) showed that all users with a desired goal of a query success rate $\geq 50\%$ were encouraged to cooperate.

In this paper we study the behavior of the users in a "continuous" game where the quality of the service received from peers is measured by the download rate of the desired files. The peer-reputation ranking system is a simplified version of that in [22]. We show that the users who cooperate more by sharing more uplink bandwidth [18], [19] achieve a better reputation ranking and eventually a higher net utility. In section II, we introduce our game model and highlight the differences between it and previously proposed models. The model is further studied through simulation in section III. Conclusions are drawn in section IV.

II. An uplink resource management game

Consider a group of N peer nodes that subject one another to queries for resources, specifically files. A query (say from peer i to j) together with a response (j's response to i's query) form a *transaction*. For $i \neq j$, let R_{ij} be the *normalized reputation* of peer j from the point of view of peer i, i.e., for all peers i, it will always be the case that

$$\sum_{j,\ j\neq i} R_{ij} \quad = \quad 1. \tag{1}$$

As transactions occur, these reputation states will change. In particular, if i queries j and the subsequent response is that j gives i the requested file (i.e., responds *positively*), then R_{ij}

This work was supported in part by NSF ANIR grant 0219747 and CyberTrust grants 0549389 and 0524202.

[1] In this paper, the terms "peer", "user" and "node" are used interchangeably.

will increase. In [22], we survey approaches to securing such a reputation system.

In [18], [19], the reputation ("cumulative contribution") of the peers is incrementally modified not by a fixed amount (C) but by an amount equal to the rate at which the file has received. The following game is a significant modification of that in [18], [19]; in particular, with regard to their assumptions about user utilities and their reputations, the latter being cumulative *net* contributions to the P2P system that decline with received content.

Suppose that each peer i has a fixed aggregate download capacity D_i and aggregate upload capacity U_i. Each peer i can designate an amount of their uplink capacity $u_i \leq U_i$ that is used for contributive/cooperative interaction with the P2P CDN system under consideration. We do not model the user's "control variable" u as a fixed total access capacity minus the downlink [18], [19], but as simply bounded by a fixed upper bound on uplink capacity. The system is taken to operate in discrete time according to the query epochs. That is, consider the following dynamics for each point in discrete time:

1) Each peer i queries at most one peer, say $j \neq i$ (this would occur with probability ρ_{ij} in a simplified transactions model); each requester i is also informed of the current uplink rate u_j of the requestee j. The size of the file requested is r_i. In the following, when the downlink capacity terms d are directly compared with the file size r terms (or the r terms are called "rates"), the r terms will be implicitly assumed to be divided by the unit of discrete time. We assume all requested file sizes $r_i \leq d_i$.
2) As a result of the previous step, each peer j is in receipt of a set M_j of queries where we note M_j may be empty or may have more than one query. Each requester $i \in M_j$ "deserves" an allocation

$$\delta_i \equiv \frac{R_{ji}}{\sum_{k \in M_j} R_{jk}} u_j.$$

The excess demand is

$$\varepsilon \equiv \left[\sum_{k \in M_j} r_k - u_j \right]^+$$

where $[z]^+ \equiv z$ if $z > 0$, otherwise $[z]^+ \equiv 0$. The idea is to distribute a penalty for excess demand ε only among those users which are requesting more than what they deserve, i.e., among

$$M_j^* \equiv \{i \in M_j \mid r_i > \delta_i\}.$$

That is, for all requesters $i \in M_j^*$ allocate

$$x_{ji} = r_i - \frac{g(R_{ji})}{\sum_{k \in M_j^*} g(R_{jk})} \varepsilon, \quad (2)$$

otherwise (i.e., for $i \notin M_j^*$)

$$x_{ji} = r_i, \quad (3)$$

where g is a positive nonincreasing function.
3) Each requester i adjusts their reputation of requestee j, e.g., by adding $R_{ij} + cx_{ji}$, for some constant $c > 0$, and then normalizing all reputations stored at i. Note this adjustment of reputation is a simplified version of the one introduced in [22].

Obviously, there are many other possible variations of the above method of employing reputations to decide on resource disbursement and then adjusting reputations in response to those decisions. With regard to the former, we assumed that $x_{ji} \leq r_i$ is desired, i.e., requesting peer i ought not receive at a rate larger than requested. Also, under overload conditions, i.e., $\sum_{k \in M_j} r_k > u_j$: $x_{ji} = r_i$ for all peers i that have high reputations R_{ji} and low request r_i relative to other peers in M_j; and $x_{ji} \ll r_i$ for all peers i that have low reputations R_{ji} and high request r_i relative to other peers in M_j. Note that under equations (2) and (3), if there is excessive demand, i.e., $\varepsilon > 0$, then

$$\sum_{i \in M_j} x_{ji} = u_j,$$

otherwise $x_{ji} = r_i$ for all $i \in M_j$ (the latter even when $M_j^* \neq \emptyset$). Another benefit of (2) and (3) is "initialization," i.e., a peer with no reputation will be granted a positive uplink capacity, when $g(0) < \infty$; so, we suggest that selection of the form $g(R) \equiv \alpha / R$, for fixed $\alpha > 0$, *not* be used. In our simulation study, we used

$$g(R) \equiv \exp(-\alpha R). \quad (4)$$

Note the grants x are continuous in the requests r. Such continuity is necessary for the *existence of a fixed-point of the iteration* by Brouwer's theorem [8].

Intuitively, if a user reduces their uplink-rate u, their reputation R will eventually be reduced, and their allocations x will also be reduced during periods of excessive demand, i.e., $\varepsilon > 0$; this is the "incentive" effect that step 3 will subsequently have under equations (2) and (3).

At the end of step 3, a requestee can assess the value obtained from the CDN from the result x of the current transaction (or by an accumulation of past and present results) via a utility function, i.e., $V_i(x_i/r_i)$ [2]. The utility is typically assumed to be nondecreasing but concave (law of diminishing returns) with $V(0) = 0$ [15]. This utility can be compared against the cost of participating in the CDN that, in this case, could be a function of the access capacity κ_i. A game can then be formulated where, e.g., peers iteratively adjust typically constrained *control variables*, say $u_i \leq \kappa_i$ with κ_i fixed (and therefore $d_i = \kappa_i - u_i$). This adjustment would occur immediately after step 3 with a "greedy" local objective to maximize their *unimodal net utility*. More precisely, over the interval $[nT, (n+1)T - 1]$ (i.e., T transactions), each user i assesses the effect of their choice of uplink-rate $u_i(n)$ at time nT on their net utility

$$\nu_i(n) \equiv V(X_i(n)) - au_i(n) \quad (5)$$

[2] Alternatively the utility function can be an increasing function of just the received rate, i.e., $V_i(x_i)$.

where the uplink-rate is constant over the interval ($= u_i(n)$),

$$X_i(n) \quad = \quad \sum_{k=nT}^{(n+1)T-1} x_i(k)$$

is the total download over the interval (alternatively, a mean download or mean download divided by requested rate, could be used for example), and $x_i(k)$ is the download of user i at time/transaction k (note that clearly $x_i(k) = 0$ if the $k^{\rm th}$ transaction does not involve peer i querying another peer), and a is a constant capturing the user's relative valuation of utility V and uplink u allocated to the peer-to-peer network.

Writing X_i as a function u_i (clearly, X_i is also a function u_j for all $j \neq i$), we also assume that

$$\frac{\partial V_i}{\partial u_i}(0) = V_i'(X_i(0))\frac{\partial X}{\partial u_i}(0) - a \quad > \quad 0.$$

for all i so that no user will select zero uplink-rate.

The interval T needs to be assumed sufficiently large to observe enough transactions to accurately make these assessments; clearly, the choice of T could be user-dependent but in the following simplified simulation we assume that it is common to all users and that the uplink-rate updates are synchronized among the users. We assumed that each user i sequentially modify their constrained uplink-rate $u_i \in [1, U_i^{\max}]$ to try to maximize their net utility ν_i using a distributed annealing strategy [16], [13], [6].

Modeling research on the *demand* processes of P2P CDNs is only in its preliminary stages and that games allow us to model dynamic user iteration with the P2P system (including the query-resolution/routing protocol). Clearly, given the $i^{\rm th}$ peer's demand for a file, it will prefer to query peers j (who possess that file) which have larger reputations R_{ij} and uplink capacities u_j. So, instead of fixed values, in a more realistic setting we would expect the ρ_{ij} to be increasing functions of R_{ij} and u_j. For simplicity in the following simulation study, we assumed that successive query-response transactions were mutually independent and the specific transaction ij occurred with probability ρ_{ij}.

III. A SIMULATION STUDY

In our simulation study of $N = 100$ peer nodes with equally likely transactions, nodes assessed their selection of contributed uplink-rate at the end of each interval by calculating their net utility (5). We considered two cases for the utilities V with $V_i^{\max} = U_i^{\max} = 1000$kbps for all i:

- unbounded $V_i(\xi) = V_i^{\max} \log(1 + \xi)$
- bounded $V_i(\xi) = V_i^{\max}(2/\pi)$arctan$(\xi)$

where ξ is the total amount received, X, by i divided by the total amount requested, $\sum_k r_i(k)$, by i in an interval. Peers dynamically chose uplinks between 1kbps and $U^{\max}$ when the parameter a specifying the relative cost of uplink bandwidth in (5) was approximately 0.01 ($a = 0.007$ in our reported simulations). To ensure congestion in the simulated peers, we either randomly chose the requested transfer rate $r_i \in [u_j/(2P),\ u_j/P]$ (where P was the number of concurrent requesing peers at all nodes on average) or $r_i = u_j/2$.

The reputation values for each peer from the point of view of another were all randomly initialized and normalized (1). Each peer i also randomly initialized its uplink bandwidth $u_i \in [1,\ U_i^{\max}]$. Intervals T were chosen to be either 1000 and 2500 transactions long.

The simulations were performed for a total of 200 *rounds*, each divided into a "query process" and "response process". At the beginning of each round, F unique pairs of nodes were selected at random and one peer in each pair queried the other; these are ordered pairs chosen from a set of size $N(N - 1)$. Note that query resolution was not covered in this work and peers were assumed to be made aware of the location of the files they desired. This process resulted in an average of F/N queries per node per round. Then, in the response process, the queried peers evaluated their excess demand and prepared to download files to the requesters at rates calculated based on step 2 in section II. The querying peers i, upon receipt of the requested resource at rate x_{ji} from the queried/provider peer j, modified j's reputation according to step 3 above. We chose $C = 0.0001$ because we found that this value was not too large so that a peer derived too much reputation benefit from a single transaction (so that it would not need to be as cooperative for a long time thereafter) and not too small so giving negligle incentive to cooperate. Note that unlike the reputation model in [22], we used only local reputation values (instead of mean pooled reputations) to govern future transactions.

Successive rounds were grouped into *intervals* of a total of T transactions. At the end of each interval, all nodes evaluated their net utilities (5) to decide how to modify their uplink bandwidths u for the next interval. These decisions were assumed to be performed synchronously. Utilities were calculated based on the cumulative received rate x and cumulative request rate r during the past interval. The following distributed annealing strategy [16], [13], [6] was used where the "temperature" $\zeta > 0$, the uplink step-size $\gamma > 0$, and the interval index $n \geq 1$.

0.1 Uplink rates are initialized: $u_i(0) \in [1, U_i^{\max}]$ for all peers i.

0.2 After an initial interval of T transactions, the initial net utility $\nu_i(0)$ is calculated for all peers i.

$n.1$ A candidate uplink rate $u'(n)$ is assigned to each node to be used in the $n^{\rm th}$ interval: if $1+\gamma \leq u_i(n-1) \leq U^{\max} - \gamma$,

$$u_i'(n) = \begin{cases} u_i(n-1) & \text{with probability 1/3} \\ u_i(n-1) + \gamma & \text{with probability 1/3} \\ u_i(n-1) - \gamma & \text{with probability 1/3} \end{cases}$$

$n.2$ After an interval of T transactions, the net utilities $\nu_i'(n)$ and the differences $\delta_i \equiv \nu_i'(n) - \nu_i(n-1)$ are calculated.

$n.3$ If $\delta \geq 0$, $u_i'(n)$ is accepted, i.e., set $u_i(n) = u_i'(n)$; otherwise, it is only accepted with probability $\exp(\delta_i/\zeta)$.

$n.4$ Increment n and go to step $n.1$.

In figures 1 through 4, $\gamma = 5$, $\zeta = 1$ and the unbounded utilities were used by all peers. Also, the requested rates obeyed $r_i = u_j/2$ to create congestion at the nodes (random selection of $r_i \in [u_j/(2P),\ u_j/P]$ with $P = 5$ was shown to create less congestion and hence fewer peers were encouraged to increase their uplink-rates (cooperate)).

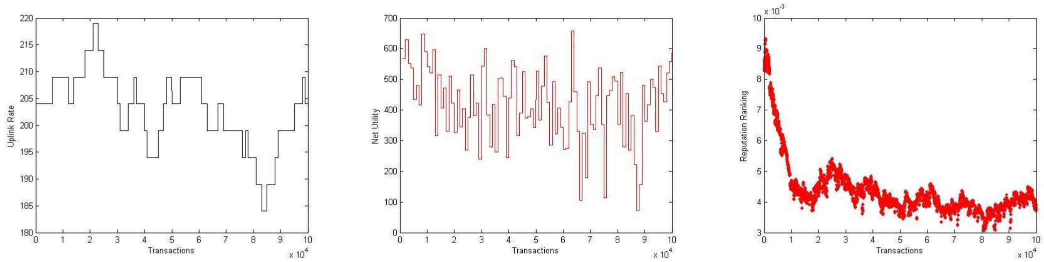

(a) Uplink Dynamics (b) Net Utility (c) Reputation Ranking

Fig. 1. A user's sample path for the simulation using $T = 1000$

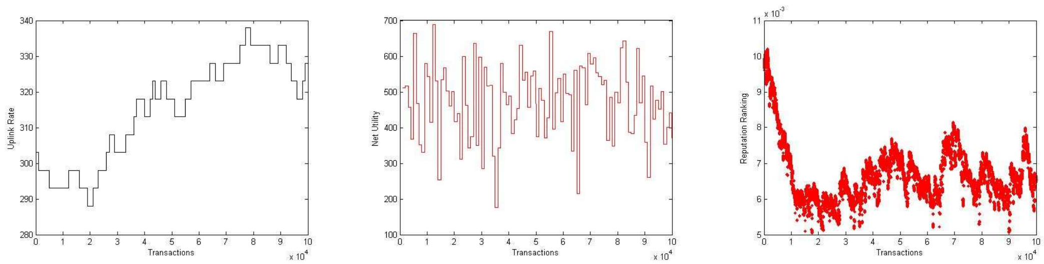

(a) Uplink Dynamics (b) Net Utility (c) Reputation Ranking

Fig. 2. Another user's sample path for the simulation using $T = 1000$

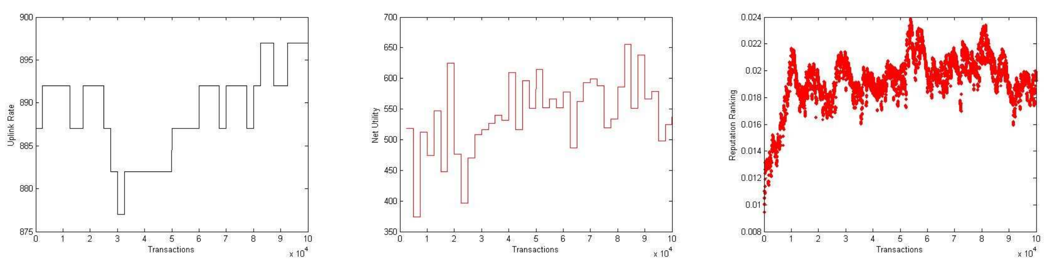

(a) Uplink Dynamics (b) Net Utility (c) Reputation Ranking

Fig. 3. A user's sample path for the simulation using $T = 2500$

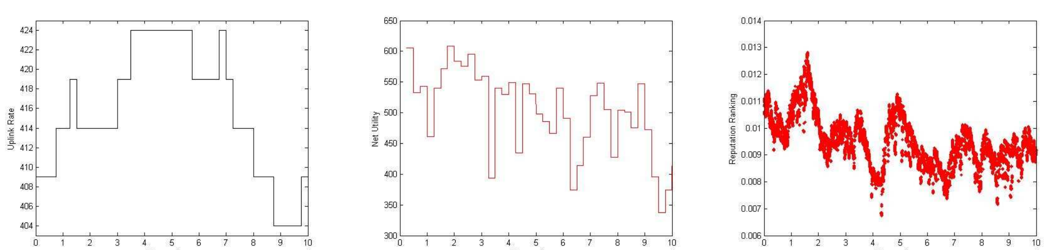

(a) Uplink Dynamics (b) Net Utility (c) Reputation Ranking

Fig. 4. Another user's sample path for the simulation using $T = 2500$

Figures 1 and 2 depict the sample path behavior of peers in this dynamic network environment with $T = 1000$ total transactions per interval consisting of 2 rounds of $F = 500$ parallel transactions per interval. The peer in Figure 1, generally reduced his uplink bandwidth resulting in an only slightly reduced net utility. However, the mean reputation value of the peer decreased very sharply at the beginning mainly because: the initial uplink choice of 205kbps was fairly low ($U^{\max} = 1000$kbps) and the adjustment size γ was small. After about 10^4 transactions, the reputation values changed less drastically. In figure 2, this node also began with a low uplink value (similarly, a sudden decrease in the reputation ranking resulted) but more constantly adjusted to higher uplink values. Net utilities were slightly higher on average than those of the aforementioned user while reputations in the steady state tended to be larger.

Examples more clearly demonstrating these "incentive" trends of the reputation system are given in Figures 3 through

4 wherein intervals of $T = 2500$ (5 rounds of F = 500 parallel transactions per interval) were used instead. Less drastic initial changes in the reputation values are due to the moderate initial uplink value for the user of figure 4. In figure 3, the reputation value initially increases rapidly because the initial uplink of this user was relatively high.

Simulation results using bounded utilities were similar to those reported above. Also, other sets of experiments were performed with different parameters. For larger temperatures, e.g., $\zeta = 100$, changes in the uplink resulting in lower net utility are more likely to be accepted. As might be expected, the nodes responded much more slowly to the network dynamics and faced either gradual increase or decrease in the uplink-rate instead of oscillations. this resulted in corresponding changes in net utilities and reputation values but more gradually and with less oscillations. The sudden initial changes in the reputation values were observed as before if the initial uplink was chosen too high or two low. Finally, we also used a gradient ascent approach, instead of distributed annealing, with significantly poorer results. We postulate that the reason for this is that a given peer's net utility depends on the uplinks of all other peers, all of which potentially change each interval, and that performance of gradient ascent degrades more significantly when the optimized function itself (the net utility) changes (possibly related problems involving poor *local* maxima are also a factor).

IV. Summary and Conclusions

A distributed, cumulative reputation system was used to encourage cooperation among participants in a peer-to-peer file sharing system. The actions of peer nodes were modeled with a continuous game in which users modified the amount of uplink bandwidth they allocate to the network. We demonstrated by simulation that increasing uplink bandwidth (cost) results in improved reputation that will, in turn, result in improved downloading performance (benefit).

In the future, we plan to explore the potential benefits of using reputation values averaged over peers [22] in order to reduce volatility. In addition, we will study distributed and automated tuning of the many parameters that are in play so as to make the reputation system more responsive thereby reducing the delay between a peer's action (uplink bandwidth) and its appropriate effect on the peer's utility. Finally, we will consider other applications of peer-to-peer systems such as distributed computation.

References

[1] Bit-torrent. http://www.bittorrent.com/.

[2] edonkey. http://www.edonkey.com/.

[3] Gnutella. http://www.gnutella.com.

[4] Kazaa. http://www.kazaa.com.

[5] Top applications (bytes) for subinterface 0[0]: Sd-nap traffic. http://www.caida.org/analysis/workload/byapplication/sdnap/.

[6] E. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines*. Wiley, 1989.

[7] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, 46(1):133–146, 2004.

[8] K. Border. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, London, 1985.

[9] S. Buchegger and J.-Y. L. Boudec. Robust reputation system for P2P and mobile ad-hoc networks. In *Second Workshop on Economics of Peer-to-Peer Systems*, June 2004.

[10] B. Cohen. Incentives Build Robustness in BitTorrent. In *Workshop on Economics of Peer-to-Peer Systems*, Berkeley, CA, USA, May 2003.

[11] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *5th ACM conference on Electronic commerce*, pages 102 – 111, New York, NY, 2004.

[12] D. R. Figueiredo, J. K. Shapiro, and D. Towsley. A public good model of availability in peer-to-peer systems. Technical Report 04-27, CSE Dept, Michigan State University, 2004.

[13] B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, Vol. 13, No. 2:pp. 311–329, 1988.

[14] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in P2P networks. In *Proc. of the 12th international conference on World Wide Web (WWW)*, pages 640–651, New York, NY, 2003.

[15] F. Kelly. Charging and accounting for bursty connections. In *Internet Economics (Editors L.W. McKnight and .P. Bailey)*, pages 253–278, 1997.

[16] G. Kesidis and R. Rao. Purposeful mobility for relaying and surveillance in ad-hoc sensor networks. *IEEE Transactions on Mobile Computing*, 2004.

[17] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau. Incentive p2p networks: a protocol to encourage information sharing and contribution. *SIGMETRICS Performance Evaluation Review*, 31(2):23–25, 2003.

[18] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau. A game theoretic approach to provide incentive and service differentiation in p2p networks. In *Proc. of the joint international conference on Measurement and modeling of computer systems*, pages 189–198, New York, NY, 2004.

[19] R. T. B. Ma, S. C. M. Lee, J. C. S. Lui, and D. K. Y. Yau. An incentive mechanism for p2p networks. In *Proc. of the 24th International Conference on Distributed Computing Systems (ICDCS)*, pages 516–523, Washington, DC, USA, 2004.

[20] S. Marti and H. Garcia-Molina. Limited reputation sharing in P2P systems. In *Proc. of the 5th ACM conference on Electronic commerce*, May 2004.

[21] S. Micali and R. L. Rivest. Micropayments revisited. In *Lecture Notes in Computer Science*, pages 149–163. Springer-Verlag, 2002.

[22] B. Mortazavi and G. Kesidis. Incentive-compatible cumulative reputation systems for peer-to-peer file-swapping. In *Proc. IEEE CISS*, Princeton, March 2006.

[23] T. G. Papaioannou and G. D. Stamoulis. Effective use of reputation in peer-to-peer environments. In *Fourth International Scientific Workshop on Global and Peer-to-Peer Computing*, April, 2004.

[24] L. Ramaswamy and L. Liu. Free riding: A new challenge to peer-to-peer file sharing systems. In *36th Hawaii International Conference On System Sciences (HICSS)*, 2003.

[25] R. L. Rivest. Peppercoin micropayments. In *Lecture Notes in Computer Science*, pages 2–8. Springer-Verlag, 2004.

[26] G. d. Veciana and X. Yang. Fairness, incentives and performance in peer-to-peer networks. In *Allerton Conference on Communication, Control and Computing*, 2003.

[27] X. Yang and G. d. Veciana. Service capacity of peer to peer networks. In *IEEE INFOCOM*, 2004.

[28] B. Yu, M. P. Singh, and K. Sycara. Developing trust in large-scale peer-to-peer systems. In *Proc. of First IEEE Symposium on Multi-Agent Security and Survivability*, 2004.