A model of a reputation system for incentive engineering

B. Mortazavi and G. Kesidis
CS&E and EE Depts
The Pennsylvania State University
University Park, PA, 16802
mortazav@cse.psu.edu and kesidis@engr.psu.edu

***Abstract*—Reputation systems are used to provide incentives for cooperation among participants of, and generally help to secure, peer-to-peer networks. In this paper, a survey of such systems is provided followed by the description of a model of a reputation framework that can capture the phenomenon of peer nodes misrepresenting reputations for malicious or selfish reasons. For special case, the model is shown to converge in mean to reputations that "reveal" the true propensity of peer nodes to cooperate. The paper concludes with a simulation study that considers weighted voting, hierarchical trust groups and misrepresentations.**

I. Introduction

Peer-to-peer (P2P) overlay systems have been proposed to address a variety of problems and enable new applications. For example, overlays have been proposed to feasibly implement new network protocols, e.g., routing [19], [29] (BGP), DNS, and/or quality-of-service management [47] (including wireless ad hoc networking contexts). In addition, communication applications (instant messaging, and voice-over-IP (VoIP) like Skype), distributed computation (e.g., seti@home), and even collaborative anomaly or intrusion detection systems (IDSs) are being implemented P2P.

A focus of this paper will be the context of P2P content distribution networks (CDNs), i.e., file sharing systems. The peer nodes of a CDN may be connected in either a structured (based on distributed hash tables) or unstructured (random) manner. P2P systems are further categorized into decentralized, centralized, partially centralized, and hybrid centralized [12]. In a decentralized setting, all nodes act equally while in a partial centralized some nodes could serve with more responsibility than the others. In such settings, the interactions are facilitated by a group of servers or supernode peers. Chord [46], CAN [39], Pastry [45] and Tapestry [50] are examples of fully decentralized, structured P2P CDNs. Among unstructured CDNs, Gnutella [3] is decentralized, Kazaa [5] is partially centralized and Napster [7] is hybrid centralized, in which a content directory is maintained in a central node while the peers communicate directly to exchange content. The reader is referred to [44] for a broad overview of P2P systems and [12] for a more detailed survey on this subject. Also, [8] is also a good collection of P2P resources for further reading.

Modern P2P systems need to deal with selfish users (a.k.a. "leechers" or "free-riders") or malicious users [21], [23], [15], P2P worms [10], [4], flash crowds, etc. For example, in the contexts of routing/forwarding in multihop wireless ad-hoc communication networks or P2P CDNs, a goal of reputation systems is to provide *incentives* for future "contributive" cooperation (resource sharing) by all peer nodes that presently benefit; such cooperation is vital to the efficient operation of the system. In certain electronic commerce networks, such as eBay [2], reputations are used to help "secure" individual transactions, i.e., give incentives to users to act responsibly when bidding on or selling merchandise. Both improving performance and reducing implementation costs of reputation systems are challenging typically because of the scale and distributed nature of the networks in which they are deployed. Performance issues include robustness in the presence of malicious and selfish users, including those that, acting alone or with collusion with others, target the reputation system itself by, e.g., lying about their own reputation or that of others when polled[1] or spoofing a fellow peer's reputation information (a tactic that is typically addressed by some kind of authentication system suitable for a large-scale distributed system, see [15], [40], [18] for example).

Again, if no incentives for cooperation are present, the existence of selfish users will diminish the performance of the P2P system [22]. Incentives are typically cumulative in nature in that sustained cooperation on a transaction-by-transaction basis yields significant rewards. These rewards may be explicitly financial (as in the case of micropayments [34], [43]) or reputational in nature where reputations, in some cases, may have implicit financial associations. For example, an eBay seller with a high reputation may garner more and higher bids for their merchandise; also, sellers may reject the bids of buyers with low reputations (or a low percentage of positive feedback). The point of such reputation systems in this context is to promote responsible bidding by potential buyers, accurate representation of merchandise by sellers, and prompt follow-through by both after the auction concludes. Alternatively, incentives could be "rule based," as in the case of Bit-Torrent [1], [37] where a specific amount of upload (cooperation) is *required* before each download. In summary for incentives that are cumulative in nature, the users need to perform a series of contributive transactions in order to receive better service at a later time. The

[1]This activity is sometimes called "shilling", c.f., mention of Kazaa Lite in section II.

higher the number of successful uploads, the better the service received from others, as in Kazaa [5]. The focus of this paper is on a reputation system that essentially is a cumulative incentive mechanism. Related work is surveyed in section II.

At any point in time, a node's reputation should indicate its relative ranking in the network based on its previous transactions with others, i.e., each time a node receives a request for a resource from another node and grants it successfully, its reputation should increase. Moreover, the likelihood that a request is granted is larger for requesting peer nodes with higher reputation from the point of view of the requested peer. This encourages nodes to be presently cooperative in order to receive better service at a later time.

In the following, we will not consider systems wherein the reputation of a requested peer decreases as a result of an unsuccessful transaction. One reason for this is to allow the nodes, without penalty, to ignore the queries that are perceived to be denial of service attacks. Also, a query may have been mistakenly directed to a node that did not have the requested content, in which case the node could not possibly grant the request and should not be penalized. Finally, a node might have been overloaded with queries and is not physically able to grant any more requests. We will see how this assumption facilitates modeling of the reputation system in section IV-A.2.

The reputation information of the nodes could either be advertised through a central entity like in eBay [2], or maintained in a distributed fashion as in Kazaa [5], or some partially decentralized solution. In a centralized reputation system, a server needs to keep the state of each node and the outcome of its transactions, while in a decentralized reputation system lack of a central entity make it challenging to aggregate and maintain nodes' reputations and is prone to manipulations, as mentioned previously. In this paper, we will present and validate a model of reputations and utilize them in an at least partially decentralized manner. Implementation of reputation systems are discussed in section III.

The rest of the paper is organized as follows. A computational model for maintaining reputations together, with analytical convergence results for a special case, are given in section IV. In section V, the results of a simulation study are presented. We conclude in section VI.

II. Related Work on Incentive Mechanisms

As briefly mentioned in the previous section, free riding (the behavior of the selfish users who benefit from communal resources but do not cooperate by sharing theirs') has been shown to cause performance degradation in the P2P network [22], [38], [48]. Specifically, that nearly 70% of Gnutella clients do not share any files and that 1% of the peers return 50% of the responses [11]. Such P2P dynamics are similar to those of "public good" in economics [14], [28]. For example, in the absence of external incentives, the phenomenon of "tragedy of the commons" [24] occurs where consumers only consider maximizing their own utility when making consumption decisions resulting in overall decrease in public utility. This section focuses on related work on incentive mechanisms to avoid such phenomena, with emphasis on those using reputation systems.

In [13], authors propose and compare several economic incentive mechanisms for P2P networks. These transaction-by-transaction incentives are implicitly formed either as a result of monetary compensation or contribution rules. The rules force the peers to share some of their resources while compensations are obtained by the peers upon their contributions. One drawback to rule-based approaches is that, if enforced in a distributed manner, the rules are prone to illegitimate manipulations by the client (requesting) peer. Also, a central entity is required to govern transactions and the monetary benefit associated with them (just as with the micropayments approach [34], [43]). In [20] [31] [32], 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. Clearly every node obtains its maximum possible level of utility in the presence of an equilibrium.

In eBay, peers rate each other after each transaction [41], [42]. The ratings for last 6 months are used to compute the overall reputation of a peer. Prior to each transaction the nodes could retrieve the reputation ranking of their peers in order to make prudent decisions. The incentive to cooperate is indirectly provided by the reputation ranking mechanism. As mentioned above, this is a centralized reputation approach. In Kazaa, the reputation of each peer is stored locally (in the peer client itself). Upon logging in, the reputation of the node is introduced to the system. A centralized approach to Kazaa reputations would create bottleneck at the reputation server in this large-scale network. On the other hand, a fully decentralized approach can be more easily subverted. In the case of Kazaa, its client was cracked and a Kazaa Lite [6] client was made available that permits the client peer to falsely report its own reputation.

A system is introduced in [17] in which each peer maintains a reputation and trust rating for a selected number of peers. The reputation of a peer is, again, a measure of how he/she has conducted transactions in the past and the trustworthiness of a peer is an indicator of how much the reputation information of others received from that peer can be relied on. From time to time, peers advertise their local reputation ratings of others to help modify the reputation information stored other peers. The authors use a Bayesian approach in which users decide whether or not the second hand reputation information should be accepted to modify the reputation ratings. Trust ratings are also updated as a result of receiving this second hand information and comparing it to prior reputation ratings.

In [33], the authors explore a similar approach with a focus on using reputation rankings to isolate malicious users. The resource providers are chosen based on their reputation levels in the system and the reputation of others is maintained at a peer both based on previous interactions (first hand information) and the advertised information from others (second hand information). The authors use a weighted selection procedure to modify local ratings. Their mechanism, however, requires a parallel download from several providers to examine the validity of the resource before locally deciding on the ratings. This could potentially introduce too much overhead which would, in

turn, result in inefficient use of network resources.

An incentive mechanism is introduced In [30] where the resources are distributed among the nodes based on their utility functions, connection types and reputations. The more a node shares resources, the higher its reputation and the better the service its receives from other nodes. Similarly in [49], authors introduce a reputation mechanism (for large peer-to-peer systems) in which the rankings are directly related to the quality of service of the peers. They further discuss aggregation of the rankings through referrals and defense against misrepresented ratings by weighted majority techniques.

The authors of [35] explore enforcement of a reputation-based policy that enables providers to choose among several simultaneous requesters based on their reputation ratings. This encourages peers to enhance their reputations in order to receive desired services. Similar to other approaches, the provider is also selected by the requester based on the reputation ratings. The mechanism is introduced in a partially decentralized setting; some peers are responsible for holding and advertising the reputations of others based on a hash function, and the peers are assumed to advertise the reputations truthfully using techniques described in [36]. The authors also analyze network efficiency for combinations of provider selection methods (such as highest reputation, comparable reputation or black list) and requester selection policies (such as highest reputation and probabilistic reputation).

EigenTrust [25], is a reputation rating mechanism that was originally designed to decrease the number of inauthentic files in P2P sharing network and isolate malicious users. A unique global value is assigned to a peer based on its previous uploads by normalizing and aggregating local trust values from other peers. The local trust values of the "acquaintances" of a peer requesting reputation values are aggregated and weighted based on the trust the peer has in them (the more trustworthy the node, the more reliable the reputation of others advertised by the node). The authors propose a distributed iterative EigenTrust algorithm that calculates a global trust vector at each node. The trust values are used for two purposes, one to isolate malicious users by downloading from reputable nodes and the other is to create incentives for the peers by *rewarding* them. The reward could be in forms of increased connectivity to other reputable peers or increased bandwidth. In [26], the same authors show that these incentives reward cooperative peers and give new peers a fair opportunity to cooperate.

III. Implementation Issues

Before introducing our reputation model, we briefly discuss implementation issues that need to be considered when designing such a model.

Handling the reputations could directly benefit from the structure of the overlay network. For example in an unstructured hybrid or fully centralized P2P overlay, it is clearly beneficial to use a centralized reputation system in which the central entity keeps track of the transactions between the nodes and updates and maintains the reputation of the peers accordingly. Every peer could then refer to the server for further information on its peers (either to retrieve or update). While this approach is less complicated to implement, it would still make the server a single point of failure, prone to denial of service attacks, and generally not scalable.

In a fully decentralized overlay, the reputation rankings of the peers could be distributed across the network in the peers themselves (as in Kazaa [5]). Communicating the reputation values could then be structured or unstructured depending on the network. Considering an unstructured overlay network, for the purpose of scalability, we assume our reputation mechanism will be implemented in a hierarchical manner. That is, peer nodes are assumed to be clustered into trust groups. Each peer node keeps a local vector of the reputation of all the nodes in its own group. Similarly, the supernode of each group keeps a vector of the *group* reputations of other groups with which it will interact (i.e., forward inter-group queries). Inter-group (respectively, intra-group) reputations are in play for inter-group (respectively, intra-group) queries. Inter-group reputations may roughly correspond to the mean intra-group reputations if the propensity of queried nodes to cooperate does not depend on the group of the querying node.

Retrieval of inter-group reputations would be managed by a protocol run only over the supernodes. This way the relevant reputations are gathered and eventually communicated to the queried node by its local supernode. The process of communicating the reputations can, however, be further optimized by periodic advertisements (or aperiodic prefetching) of the reputations instead of retrievals on a transaction-by-transaction basis. This can be done for both intra-group and inter-group reputations. The following reputation system may be separately applied to just one element in one level of a group hierarchy.

IV. Incentive Architecture and Mechanism

In this section, we present a specific model of a reputation mechanism that can account for deliberate misrepresentations of reputation by referred peer nodes. The fact that the reputations of a node eventually reveal its true level of cooperativeness in the network is demonstrated for a special case.

A. Model of a Reputation System for Peer-to-Peer Networks

Consider a group of N peer nodes that subject one another to queries for, say, files. A query (say from 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} will increase. In the following, a general reputation system model will be described and its ability to "reveal" the propensity of peers to cooperate (respond positively to queries) will be established for a special case.

1) Basic definitions of the reputation system: We begin with a set of definitions. Without consideration of reputation, let $\pi_j > 0$ be the fixed probability that peer j will respond positively to a query, i.e., j's propensity to cooperate. In the following, a sequence of transactions is considered with $R_{ij}(n)$ representing the reputation of j from i's point-of-view after the n^{th} transaction. We assume that transactions are independent so that the $(N^2 - N)$-vector $\mathbf{R}$ of reputation states R_{ij} will be a Markov chain on $(\Sigma_{N-1})^N$ where Σ_{N-1} is the N-dimensional simplex, see (1). $\mathbf{R}$ makes a transition upon the conclusion of each transaction. Let

$$\bar{R}_i(n) \quad \equiv \quad \frac{1}{N-1} \sum_{k,\ k \neq i} R_{ki}(n) \tag{2}$$

be the mean reputation of i after the n^{th} transaction and let the *response function*

$$G_j(\pi_j, \bar{R}_i)$$

be the probability that j responds positively to i's query. Generally, response functions G will be assumed to have the following properties:

- G is nondecreasing in both arguments,
- $G(\pi, \bar{R}) = 0$ and $\pi > 0$ imply $\bar{R} = 0$, and
- $G(\pi, \bar{R}) \leq \pi$ for all $\bar{R} \in [0, 1]$.

So, peer j obtains and averages the reputations R_{ki} from all other peer peers k and modifies its probability of responding positively accordingly, i.e., a polling/voting system.

Now a specific mechanism for updating reputations as a result of transactions will be defined. If the n^{th} transaction involves i querying j, then with probability $G_j(\pi_j, \bar{R}_i(n-1))$:

$$R_{ik}(n) = \begin{cases} (R_{ij}(n-1) + C)/(1+C), & k = j \neq i \\ R_{ik}(n-1)/(1+C), & k \neq j, i \end{cases} \tag{3}$$

for some fixed $C > 0$, i.e., R_{ij} becomes relatively larger only when the transaction ij succeeds. With probability $1 - G_j(\pi_j, \bar{R}_i(n-1))$:

$$R_{ij}(n) \quad = \quad R_{ij}(n-1) \text{ for all } i \neq j,$$

i.e., if the transaction fails, there is no change in the reputation. Note that reputations of peers may (relatively) decrease upon positive transactions that do not involve them.

Reputation penalties for unsuccessful transactions can explicitly be incorporated into the above model in a straightforward way. For example, if the n^{th} transaction is ij, the transaction is unsuccessful and $R_{ij}(n-1) \geq \epsilon \geq 0$, peer node i could set

$$R_{ik}(n) \quad = \quad \begin{cases} \frac{R_{ij}(n-1) - \min\{C,\ R_{ij}(n-1) - \epsilon\}}{1 - \min\{C,\ R_{ij}(n-1) - \epsilon\}} & k = j \neq i \\ \frac{R_{ik}(n-1)}{1 - \min\{C,\ R_{ij}(n-1) - \epsilon\}} & k \neq j, i \end{cases}$$

for some fixed ϵ. Recall that we made an argument dissuading the use of such negative feedback in the introductory section of this paper. Such negative reputation dynamics will not be considered in the following.

Similar to [17], [25], [33], [49], we can account for misrepresentation and subsampling of reputations by using the following instead of $\bar{R}_i$ in the reputation model:

$$\bar{R}_{ji}(n) \quad = \quad \frac{\sum_{k,k \neq i} \lambda_{jki} h(R_{jk}(n)) R_{ki}(n)}{\sum_{k,k \neq i} h(R_{jk}(n))} \tag{4}$$

where the terms λ_{jki} can be used to represent how node k may misrepresent when polled by j for i's reputation, i.e., $\lambda_{jki} \in [0, 1/R_{ki}(n)]$ and misrepresentation occurs when $\lambda_{jki} \neq 1$. The h-function parameters can be used to weight reputation information by the reputation of the pollee[2] [33]. Examples are $h(R_{jk}) \equiv \mathbf{1}\{R_{jk} > \theta_j\}$ and $h(R_{jk}) \equiv R_{jk}\mathbf{1}\{R_{jk} > \theta_j\}$ where $\theta_j \geq 0$ is a reputation *threshold* that may be used by node j to define "trust." As a special case, we can model federations that are used by peers for reputation polling, i.e., consider M groups $\{\mathcal{N}_m\}_{m=1}^M$ of peers, where

$$\bigcup_{m=1}^{M} \mathcal{N}_m$$

is the set of all N peers and $|\mathcal{N}_m| \geq 2$ for all m. This is modeled by taking

$$h(R_{jk}) \quad \equiv \quad \begin{cases} 1 & \text{if } j, k \in \mathcal{N}_m \text{ for some } m \\ 0 & \text{else} \end{cases}$$

Note that $\{\mathcal{N}_m\}_{m=1}^M$ need not be a partition, i.e., a single peer j could belong to more than one group. Finally, note that we clearly need to assume that the denominator of (4) is always strictly positive for all i, j, n. Recall from section III that "complete" polling could be conducted on a group basis, i.e., a hierarchical (partially decentralized) reputation system.

2) Modeling the Transaction Process: In the introductory section, we argued why reputations should not be reduced as a result of a failed transaction. Under this assumption, a model of reputation dynamics does not therefore need to consider the *circumstances* of a failed query. Moreover, we can combine the probability that the a queried peer refuses to comply together with the probability that that peer is not logged onto the peer-to-peer system. That is, let ρ_{ij} be the probability that a transaction involves i querying j so that

$$\sum_i \sum_{j,\ j \neq i} \rho_{ij} \quad = \quad 1.$$

The quantity ρ_{ij} and/or the quantity π_j could also account for the probability that the user j is "on" the system. Thus, we can simplify the analysis of reputation dynamics by not explicitly modeling peer-node arrivals and departures to the reputation system and associated effects on the query resolution system.

For the purposes of subsequent analysis, we further assume that the successive transaction attempts are independent. With specific regard to content that is extremely popular for a period of time: peer nodes j with such files will experience high query rates (i.e., high ρ_{ij}). Rather than attempting to model time-varying parameters ρ_{ij}, we will assume that these parameters

[2]Such reputation "weightings" can also be used in more general voting systems for distributed decision making as, e.g., applied to anomaly detection.

are constant and will simply tend to be higher for those nodes j that have highly desirable content. Again, our point in the following analysis is to determine whether the reputation process does indeed reveal the *long term* (over many transactions) "propensity to cooperate" of the peer nodes.

In [31], [32], the authors set-up a game model for a P2P content distribution network (CDN) that accounts for uplink/downlink bandwidth resources made available to the CDN. Also, the reputation ("cumulative contribution") of the peers is incrementally modified not by a fixed amount (C) but by an amount equal to the size of the file involved in the transaction. The following game, a significant modification of that in [31], [32][3], is given here to concretely show the role of the previously described reputation system in a game model of a P2P CDN.

Suppose that each peer i has a fixed bidirectional access capacity κ_i to the Internet. This capacity is divided between an uplink u and downlink d capacity for each peer, i.e., $u + d = \kappa$. Consider a P2P CDN operating in discrete time with synchronized queries by the peers. 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 our simplified transactions model); each requestee j is also informed of the downlink capacity d_i of the requester i. The size of the file requested is r_i. In the following, when the 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 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. Node j transmits to peer $i \in M_j$ at rate $x_{ij} = \min\{r_i,\ R_{ij}u_j\}$ if $\sum_{k\in M_j} \min\{r_k,\ R_{kj}u_j\} < u_j$, otherwise

$$x_{ij} \quad = \quad \frac{\min\{r_i,\ R_{ij}u_j\}}{\sum_{k\in M_j} \min\{r_k,\ R_{kj}u_j\}}. \tag{5}$$

3) Each requester i adjusts their reputation of requestee j, e.g., by adding $R_{ji} + cx_{ij}$, for some constant $c > 0$, and then normalizes all reputations stored at i.

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. For the specific iteration described above, note that the allocated rates x are continuous functions of the requests r and reputations R; such continuity is necessary in order to invoke Brouwer's result on existence of a fixed-point of the iteration [16]. Moreover, $x_{ij} \leq r_i$, i.e., requesting peer i will not receive at a rate larger than requested. Also, note that free riders ($u \equiv 0$) will ultimately receive zero reputations in this system and thus get $x = 0$ under (5). Note the intuitive incentive effects that modified reputations in step 3 will have on subsequent steps (5). Also note that 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., $U_i(x_i)$. 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 the net utility $U_i(x_i) - \kappa_i$, see, e.g., [27].

We reiterate that modeling research of 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^{th} peer's demand for a file, it will prefer to query peers j (who possess that file) which have larger reputations R_{ji} and uplink capacities u_j. So, instead of fixed values, in a slightly more realistic setting we would expect the ρ_{ij} to be increasing functions of R_{ji} and u_j.

In the remainder of this paper, we will study the reputation framework described above without consideration of network resources, P2P query resolution, and dynamic user demand.

3) Mean Reputation Process: For $n \geq 1$, we can now directly derive for "complete and honest" polling:

$$\begin{aligned}
\mathsf{E}(R_{ij}(n) \mid \mathbf{R}(n-1)) \ &= \ \left(1 - \sum_{k,\ k\neq i} \rho_{ik}\right) R_{ij}(n-1) \\
&\quad + \rho_{ij}[R_{ij}(n-1)(1 - G_j(\pi_j, \bar{R}_i(n-1)) \\
&\qquad + \frac{R_{ij}(n-1)+C}{1+C} G_j(\pi_j, \bar{R}_i(n-1))] \\
&\quad + \sum_{k,\ k\neq i,j} \rho_{ik}[R_{ij}(n-1)(1 - G_k(\pi_k, \bar{R}_i(n-1)) \\
&\qquad + \frac{R_{ij}(n-1)}{1+C} G_k(\pi_k, \bar{R}_i(n-1))] \\
&= \ \left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G_k(\pi_k, \bar{R}_i(n-1))\right) R_{ij}(n-1) \\
&\quad + \frac{C}{1+C} \rho_{ij} G_j(\pi_j, \bar{R}_i(n-1)).
\end{aligned} \tag{6}$$

In the first equation of this display, the n^{th} transaction: does not involve i querying in the first term, involves i querying $j \neq i$ in the second term, and involves i querying $k \neq j$ in the third term.

Since the terms $\lambda_{jki} h(R_{jk}(n)) \geq 0$ depend on reputations only through $\mathbf{R}(n)$, we can generalize the model (6) to account for misrepresentation and subsampling of reputations by replacing in (6) the $\bar{R}_i$ as defined in (2) with $\bar{R}_{ji}$ as defined in (4), i.e.,

$$\begin{aligned}
&\mathsf{E}(R_{ij}(n) \mid \mathbf{R}(n-1)) \\
&= \ \left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G_k(\pi_k, \bar{R}_{ki}(n-1))\right) R_{ij}(n-1) \\
&\quad + \frac{C}{1+C} \rho_{ij} G_j(\pi_j, \bar{R}_{ji}(n-1)).
\end{aligned} \tag{7}$$

4) Accumulation points of the reputation system: If we take expectation of both sides of (7) and set

$$\begin{aligned}
\mathsf{E}R_{ij}(n-1) \ &= \ \mathsf{E}R_{ij}(n) \\
&= \ \mathsf{E}(\mathsf{E}(R_{ij}(n) \mid \mathbf{R}(n-1))),
\end{aligned}$$

[3] In particular, their assumptions about user utilities and their reputations, the latter being cumulative *net* contributions to the P2P system that decline with received content.

we see that the sample paths of the ergodic $\mathbf{R_{ij}}(n)$ have marginal distributions whose limiting accumulation points are distributions satisfying:

$$\begin{aligned} &\mathsf{E}^*(R_{ij} \sum_{k,\ k\neq i} \rho_{ik} G_k(\pi_k, \bar{R}_{ki})) \\ &= \mathsf{E}^*(\rho_{ij} G_j(\pi_j, \bar{R}_{ji})) \quad \text{for all } i \neq j \end{aligned} \tag{8}$$

where E^* is expectation under any one of the limiting distributions under consideration.

5) Convergence of mean reputations for a case of complete and honest polling : In this subsection, assume the *common* response function G satisfies the following property: there is a strictly positive $\varepsilon \ll 1$ such that, for all $0 \leq \pi, \bar{R} \leq 1$,

$$\varepsilon\pi \ \leq\ G(\pi, \bar{R}) \ \leq\ \pi. \tag{9}$$

Also assume the response function G is *separable*, i.e.,

$$G(\pi, \bar{R}) \ = \ \pi g(\bar{R})$$

for nondecreasing g. So, by (9), $g(0) \geq \varepsilon$ and $g(1) \leq 1$. For example, $g(\bar{R}) \equiv \varepsilon + \bar{R}(1-\varepsilon)$ or $g(\bar{R}) \equiv \max\{\varepsilon,\ \min\{c\bar{R},\ 1\}\}$ for some constant $c > 1$. Note that by arranging that the response functions are all strictly positive (when $\pi > 0$), a new cooperative node with low initial reputation can obtain content in order draw queries and thereby its own reputation. The following result could apply to a system intra-group reputations or a system inter-group reputations as discussed in section III.

Theorem 1: If G is separable and the statement of (9) holds, then for complete and honest polling (6):

$$\lim_{n\to\infty} \mathsf{E}R_{ij}(n) \ = \ \frac{\rho_{ij}\pi_j}{\sum_{k,k\neq i}\ \rho_{ik}\pi_k} \quad \text{for all } i \neq j.$$

Proof:
For separable G, define

$$X_{ij}(n) \ \equiv \ \frac{\rho_{ij}\pi_j}{\sum_{k,k\neq i}\ \rho_{ik}\pi_k} - R_{ij}(n)$$

for all $i \neq j$ and $n \geq 0$. By (6),

$$\begin{aligned} &\mathsf{E}(X_{ij}(n) \mid \mathbf{R}(n-1)) \\ &= \left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G(\pi_k, \bar{R}_i(n-1))\right) \\ &\quad \times \left(\frac{\rho_{ij} G(\pi_j, \bar{R}_i(n-1))}{\sum_{k,k\neq i}\ \rho_{ik} G(\pi_k, \bar{R}_i(n-1))} - R_{ij}(n)\right) \\ &= \left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G(\pi_k, \bar{R}_i(n-1))\right) \\ &\quad \times X_{ij}(n-1) \end{aligned}$$

where (9) allows division by $g(\bar{R}_i(n-1)) > 0$ for the second equality. Thus,

$$\begin{aligned} &\mathsf{E}X_{ij}(n) = \mathsf{E}(\mathsf{E}(X_{ij}(n) \mid \mathbf{R}(n-1))) \\ &= \mathsf{E}\left(\left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G(\pi_k, \bar{R}_i(n-1))\right) X_{ij}(n-1)\right). \end{aligned}$$

Since

$$\begin{aligned} 1 \ &> \ \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G(\pi_k, \bar{R}_i(n-1)) \\ &\geq \ \frac{C\varepsilon}{1+C} \sum_{k,\ k\neq i} \rho_{ik}\pi_k \ \equiv \ \alpha \ > \ 0, \end{aligned}$$

we can easily show that

$$\mathsf{E}|X_{ij}(n)| \ \leq \ (1-\alpha)\mathsf{E}|X_{ij}(n-1)|.$$

This argument can be used successively to show

$$\begin{aligned} \mathsf{E}|X_{ij}(n)| \ &\leq \ (1-\alpha)^2 \mathsf{E}|X_{ij}(n-2)| \\ &\leq \ (1-\alpha)^n \mathsf{E}|X_{ij}(0)| \end{aligned}$$

from which the theorem statement immediately follows. □

Interpreting this theorem, the mean reputation $\mathsf{E}R_{ij}(n)$ becomes proportional to the mean rate $\rho_{ij}\pi_j$ of successful transactions ij. Further suppose all possible $N(N-1)$ queries ij are equally likely (query load perfectly balanced among the peers), i.e., $\rho_{ij} = 1/(N^2 - N)$ for all $i \neq j$. In this case, the theorem implies

$$\lim_{n\to\infty} \mathsf{E}R_{ij}(n) \ = \ \frac{\pi_j}{\Pi_{-i}} \text{ for all } i \neq j$$

where

$$\Pi_{-i} \ \equiv \ \sum_{j,\ j\neq i} \pi_j.$$

So, for each peer i, $\mathsf{E}R_{ij}$ is proportional to π_j in steady state and this reputation system will reveal the propensities to cooperate of the peers.

V. Simulation Study

In this section, we describe the results of preliminary simulation studies of the sample paths of $\mathbf{R}(n)$. We used a common, separable response function $G(\pi, \bar{R}) = \pi * \min(1, (N/2)\bar{R})$. Since a "typical" value of the normalize reputations R are about $1/N$, this choice ensured that the values of G were on the order of the values of π.

Earlier experiments led us to choose the value of C (the reward given to a node for a successful transaction) to be $3/N$. If C is chosen too high, several successful transactions would raise a node's reputation so high that it would make it sufficient to receive service from others for quite sometime without the need to cooperate. On the other hand, small values for C may not create enough incentives for the nodes to consistently cooperate.

Our simulation consisted of a series of rounds. In each round, two nodes were randomly selected to perform one transcation, say ij, and then the reputation of i was computed as in (4). The response function, G, of j was computed, i.e., node j decided to grant the request with probability G. As a result of a successful transaction, node i then updated its reputation values

based on (3). A first order autoregressive estimator of the mean reputations was then updated using:

$$\tilde{\mathbf{R}}(n) \quad = \quad \beta\tilde{\mathbf{R}}(n-1) + (1-\beta)\mathbf{R}(n) \qquad (10)$$

for forgetting factor $1 > \beta > 0$. In section V-A, we show how the choice of β affected the dynamics of the reputation values of a node, i.e., it reduced oscillations of (low-pass filters) the reputation processes $\mathbf{R}$. The modified reputation values of node j were then used in subsequent rounds as described above.

We performed our simulations in three phases. In phase 1, the transactions between any two nodes depended on the individual reputations of the nodes via the mean reputation of the requestee assuming all other nodes are polled, i.e., a nonhierarchical structure. In phase 2, we studied a hierarchical structure in which the nodes are arranged in groups; the intra-group transactions occurred as in the first phase (with only intra-group polling), while the inter-group transactions involved the group reputations instead of the individual ones. Finally, reputation misrepresentations were considered in phase 3, i.e., $\lambda_{jki} \neq 1$ for some j, k, i. In phases 1 and 3, $N = 100$ nodes were considered, while there were $N = 20$ nodes consisting of 5 groups of 4 in phase 2. Each transaction was assumed equally likely (i.e., $\rho_{ij} = 1/(N^2 - N)$ for all i, j), and each trial consisted of about 100 transactions per pair (i.e., approximately $100(N^2 - N)$ transactions). We assumed different probabilities of cooperation among the nodes; more specifically, the π_i were selected independently at random from the interval $[0.5, 1]$ according to a uniform distribution. Finally, the initial values $\mathbf{R}(0)$ were independently selected at random according to a uniform distribution on $[0, 1]$ and then normalized.

A. Non-hierarchical structure

We first assumed a non-hierarchical (flat) reputation system. Figures 1(a) through 1(b) are a typical observed sample path of $R_{ji}(n)$, the reputation of a specific node j from the perspective of a node i. The abrupt increases (the size of which depend directly on C) in the reputation value indicate successful transactions ij, while the more gradual reductions occur when node i has successful transactions with other nodes (recall that unsuccessful transactions do not impact the reputations). For visualization purposes, π_j/Π_{-i} (≈ 0.01 in this case) is depicted as a straight horizontal line on the figures. This quantity was observed to the the mean time-averaged value of the reputation processes for all trials, as expected from the theorem in section IV-A.5.

In figure 1(a), the value of the forgetting factor β was set to 0.85 while it was set to 0.15 in figure 1(b). As a result, the former shows significantly less variation about the exptected mean π_j/Π_{-i} though it has a slightly longer transient phase (before reaching steady state about the mean).

We also examined the use of local reputations R_{ji} (figure 1(c)) instead of mean reputations $\bar{R}_{ji}$ (figure 1(a)) in the response function G. Again, a typical sample path is depicted. We see that both cases result in a similar marginal variance about the mean, but the use of local reputations shows has lower "frequency," i.e., individual excursions away from the mean have longer time-durations; so, over a given fixed interval of time (subset of consecutive transactions), it is more probable that a crossing of the true mean value of the process will not be observed.

Finally, in this phase we further investigated weighting the reputations received from other nodes with their current reputation values. Referring to (4) for all nodes, we took $h(R) \equiv R$ in one set of trials and $h(R) = \mathbf{1}\{R > \theta\}$ with $\theta = 0.01$. For both cases, the reputation dynamics were observed to be similar to those when no weighting was used. This was expected since all transactions were equally likely and the π_j values were selected independently using the same distribution.

B. Hierarchical structure

Again, in this phase, 20 nodes were partitioned into 5 groups of 4. Depending on which two nodes are selected at each round, the transactions either fell within their group (intra-group) or between different groups (inter-group). For intra-group transactions, individual reputations were used and normalized within the group (summing over all other nodes j in i's group, $\sum_{j,\ j\neq i} R_{ji} = 1$). For inter-group transactions, group reputations were maintained through the assumed supernodes (one per group).

In figure 2(a), a sample path of one group's reputation from the point of view of another is depicted. Since the transactions between any two groups occur much more frequently than between a pair of nodes, the group reputation sample path appears smoother and has a shorter transient phase than the individual reputation sample path depicted in figure 2(b). The average cooperation levels π of the individual members of the group are also depicted with a horizontal line in figure 2(a). It can be seen that the sample path fluctuates around this expected mean value and the time-average was observed to converge to it, as in the non-hierarchical experiments.

In Figure 2(b), the individual reputations of one specific node from the point if view of another in the same group is shown. The sample path decreases less rapidly than its non-hierarchical counterpart with similar parameters because the transactions across groups do not impact the individual reputations. However, the time-average reputations still reveal a node's propensity to cooperate. Again, both sample paths depicted were typical of those observed during our trials.

Note one reason why inter-group transactions do not utilize or update the individual reputations is that, if two *different* propensities to cooperate (for inter-group vs. intra-group transactions) are present, nodes that cooperate well within the group should not get rewarded universally. Studying such behavior (in particular, for special-interest groups) is among our plans for future work.

C. Misrepresentations

Finally, we used the λ terms in (4) to model misrepresentation of reputations by certain pollees. We assumed 25% of the nodes lied unfavorably (about all other nodes to all other nodes), 10% lied in favor, and the rest were honest. The corresponding λ terms were chosen to be 0.75, 1.25 and 1 respectively. Since all requests were equally likely ($\rho_{ij} = 1/(N^2 -$

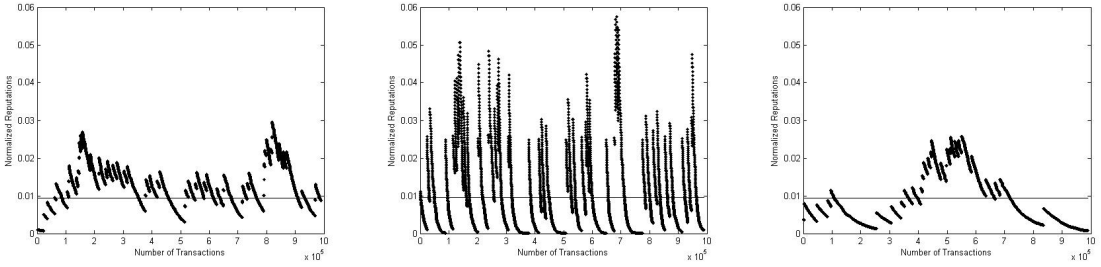

(a) β =0.85 & $\bar{R}$ is used in the G function

(b) β = 0.15 & $\bar{R}$ is used in the G function

(c) β = 0.85 & local reputations are used in the G function

Fig. 1. 100 honest nodes with 100 transactions per pair

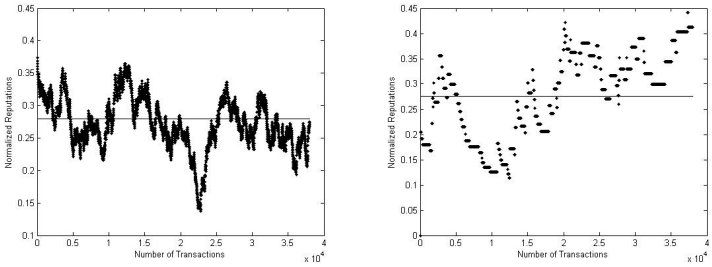

(a) Group Reputations

(b) Individual Reputations

Fig. 2. 20 nodes in 5 groups of 4, 100 transactions per pair of nodes, β = 0.95 and $\bar{R}$ is used in G

N)), the value of the mean reputation of the j^{th} node, from the perspective of the i^{th} node, was simply predicted to be

$$(0.25 \cdot 0.75 + 0.10 \cdot 1.25 + 0.65 \cdot 1)\frac{\pi_j}{\Pi_{-i}} = 0.9625\frac{\pi_j}{\Pi_{-i}}.$$

This was consistent with the results of our simulation trials where we observed nodes receiving about 4% fewer successful transactions compared to honest reporting (with the same settings for G, C, etc). Clearly, in this model, when the aggregated reputations are weighted using $h(R) = R$, the highly reputable nodes can lie more effectively.

VI. Summary and Future Work

In summary, we surveyed reputation frameworks as used by peer-to-peer overlay systems, especially to provide incentives for persistent contributive cooperation by the peer nodes. We formulated a normalized reputation model and proved, for a special case, that it ultimately revealed the nodes' propensities to cooperate under honest reporting. This model was studied by further by simulation.

In the future, we plan to expand our model to account for selfish/malicious peers that collude. Simulations will possibly be conducted on the platform [9].

We will also study the possibility of the nodes having different propensities to cooperate for inter-group versus intra-group transactions where the latter is typically signficantly larger.

Finally, we will also explore ways to integrate routing and reputations for specific architectures. Recall mention of potential problems with this idea in the introductory section of this paper. Note in particular that we will have to explicitly model peer arrivals and departures (join and leaves) for such an integrated system, i.e., the simplifying assumption of section IV-A.2 would not be made.

References

[1] Bit-torrent. http://www.bittorrent.com/.
[2] eBay. http://ebay.com.
[3] Gnutella. http://www.gnutella.com.
[4] Igloo. http://vil.nai.com/vil/content/v_100046.htm.
[5] Kazaa. http://www.kazaa.com.
[6] Kazaa lite. http://www.kazaalite.nl.
[7] Napster. http://www.napster.com.
[8] P2P resources. http://www.cs.dartmouth.edu/~zhaom/research/marianas/resource.html.
[9] p2psim. http://pdos.csail.mit.edu/p2psim/index.html.
[10] VBS.gnutella worm. http://securityresponse.synmantec.com/avcenter/venc/data/vbs.gnutella.html.
[11] E. Adar and B. A. Huberman. Free riding on gnutella. *First Monday magazine*, Sept. 2000.
[12] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys (CSUR)*, 36(4), December 2004.
[13] P. Antoniadis, C. Courcoubetis, and R. Mason. Comparing economic incentives in peer-to-peer networks. *Computer Networks*, 46(1):133–146, 2004.
[14] A. Asvanund, K. Clay, R. Krishnan, and M. D. Smith. An empirical analysis of network externalities in peer-to-peer music-sharing networks. *Information Systems Research*, 15(2):155–174, June 2004.
[15] M. Blaze, J. Feigenbaum, and A. D. Keromytis. The role of trust management in distributed systems security. In *Secure Internet Programming*, pages 185–210, 1999.

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

[17] 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.

[18] D. Clarke, J.-E. Elien, C. Ellison, M. Fredette, A. Morcos, and R. L. Rivest. Certificate chain discovery in SPKI/SDSI, 1999. To be published, November 1999.

[19] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and K. van der Merwe. The Case for Separating Routing from Routers. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, Portland, OR, September 2004.

[20] 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.

[21] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems, 2004.

[22] 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.

[23] S. Ganeriwal and M. B. Srivastava. Reputation-based framework for high integrity sensor networks. In *Proc. of the 2nd ACM workshop on Security of ad hoc and sensor networks (SASN '04)*, pages 66–77, 2004.

[24] G. Hardin. The tragedy of the commons. *Science*, 162:1243–48, 1968.

[25] 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.

[26] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. Incentives for combating freeriding on P2P netwroks. Technical report, Stanford University, 2003.

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

[28] R. Krishnan, M. D. Smith, and R. Telang. The economics of peer-to-peer networks. *Journal of Information Technology Theory and Application (JITTA)*, 5(3):31–44, 2003.

[29] Z. Li, P. Mohapatra, and C.-N. Chuah. Virtual multi-homing: On the feasibility of combining overlay routing with BGP routing. In *Proc. of Networking 2005*, pages 1348–1352, Waterloo, Canada, 2005.

[30] 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.

[31] 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.

[32] 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.

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

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

[35] 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.

[36] T. G. Papaioannou and G. D. Stamoulis. Enforcing credible reporting in peer-to-peer enviroments, working paper. In *Athens University of Economics and Business*, January 2004.

[37] D. Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proc. of SIGCOMM*, Portland, Oregon, 2004.

[38] 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.

[39] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A scalable content-addressable network. In *Proceedings of the ACM Conference of the Special Interest Group on Data Communication (SIGCOMM)*, pages 161–172, New York, NY, August 2001.

[40] M. Reiter and S. Stubblebine. Toward acceptable metrics of authentication. In *Proc. of IEEE Symposium on Security and Privacy*, pages 10–20, 1997.

[41] P. Resnick, K. Kuwabara, R. Zeckhauser, and E. Friedman. Reputation systems. *Communications of the ACM*, 43(12):45–48, 2000.

[42] P. Resnick and R.Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay s reputation system. In *Proc. of NBER workshop on empirical studies of electronic commerce*, 2000.

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

[44] K. Ross and D. Rubenstein. Tutorial on P2P systems. *http://cis.poly.edu/˜ross/papers/P2PtutorialInfocom.pdf*, 2004.

[45] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, Nov. 2001.

[46] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Trans. on Networking*, 11(1):17–32, 2003.

[47] L. Subramanian, I. Stoica, H. Balakrishnan, and R. Katz. OverQoS: An Overlay Based Architecture for Enhancing Internet QoS. In *1st Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, March 2004.

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

[49] 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.

[50] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. D. Kubiatowicz. Tapestry: A resilient global-scale overlay for service deployment. *IEEE Journal on Selected Areas in Communications*, 22(1):41–53, Jan. 2004.