Cumulative Reputation Systems for Peer-to-Peer Content Distribution

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
*also a member of technical staff at Verizon Wireless

***Abstract*— Reputation systems can be used to provide incentives for cooperation among participants of peer-to-peer networks. In this paper, a survey of such systems is first provided. We then focus on content distribution networks that have honest but rationally selfish users. A reputation framework is proposed 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 is designed in which users play to maximize the files received from the system by adjusting their cooperation level and gaining a better reputation as a result.**

I. Introduction

Peer-to-peer (P2P) overlay systems have been proposed to address a variety of problems and enable new applications. The attraction of these systems, when compared to client/server frameworks, is in their robustness, reliability and cost efficiency. They have been proposed to feasibly implement new network protocols, e.g., routing [17], [27] (BGP), DNS, and/or quality-of-service management [46] (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. The focus of this paper is in the context of P2P content distribution networks (CDNs), i.e., file sharing systems [9], [6].

The modern P2P systems need to deal with selfish (a.k.a. "leechers" or "free-riders") or malicious users[1] [19], [21], [12], P2P worms [7], [3], Byzantine faults and Sybil attacks [14], Eclipse attacks [44], [45], flashcrowds, etc. Some of these problems are particularly challenging for large-scale, peer-to-peer systems. Reputation systems have been proposed to address some of these issues. 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., create incentives for 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 (acting alone or with collusion with others) target the reputation system itself by, e.g., lying about their own reputation[2] or that of others when polled (such as badmouthing). A detailed description of different attacks on reputation systems can be found in [23]. In general, if no incentives for cooperation are present, the existence of rationally selfish users will diminish the performance of the P2P system [20]. 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 [32], [43]) or reputational in nature where reputations, in some cases, may have implicit financial associations such as in eBay. For example, an eBay seller with a high reputation[3] 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], [38], [16] where a specific amount of upload (cooperation) is typically[4] *required* before each download, i.e., incentives are on a transaction-by-transaction basis.

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 higher the number of successful uploads, the better the service received from others, as in Kazaa [4]. The focus of this paper is to develop a novel reputation framework, in which in the absence of misrepresentation of the reputation values, can reveal the propensity to cooperate of the peers. This reputation system

[1] In this paper, we will use the terms "users", "peers", "nodes", and "peer nodes" interchangeably.

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

[3] It also helps to accept payments in credit card form as this offers the buyer some measure of purchase protection.

[4] Under Bit-Torrent, not all transactions are necessarily bidirectional; this allows new users with little desirable content to participate.

is essentially a cumulative incentive mechanism that tries to encourage cooperation among rationally selfish users. Related work on incentives is surveyed in section II.

The rest of the paper is organized as follows. In section III, an incentive-compatible convergence result for a proposed cumulative reputation system is given; we note that the convergence theorem holds only for a special case. Section concludes with survey of some techniques used to scale and secure the reputation system itself. In section III-G, a simple game model is introduced wherein peers are encouraged by the reputation system to adjust their "propensities to cooperate" in order to receive a certain level of satisfaction (downloading transactions success) from the P2P CDN. We conclude in section IV with a brief summary.

II. Related Work on Incentive Mechanisms

As mentioned above, 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 P2P networks [20], [39], [47], [8]. Specifically, that nearly 70% of Gnutella clients do not share any files and that 1% of the peers return 50% of the responses [8]. Such P2P dynamics are similar to those of "public good" in economics [11], [26], where in the absence of external incentives, the phenomenon of "tragedy of the commons" [22] occurs; 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 [10], 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 [32], [43]). In [18] [29] [30], 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.

eBay provides a centralized reputation server through which 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. This centralized approach indirectly provides incentive to cooperate. 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 distributed in the clients. 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. This fully decentralized approach can be subverted as seen in Kazaa Lite [5]. In the case of Kazaa, its client was cracked and a Kazaa Lite client was made available that permits the client peer to falsely report its own reputation.

A system is introduced in [13] in which each peer maintains reputation and trust ratings 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 at other peers. Through a Bayesian approach users decide whether or not the second hand reputation information should be accepted to modify local information. Trust ratings are updated upon receiving the second hand information and comparing them to prior reputation ratings.

In [31], 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). A weighted selection procedure is used to modify reputation 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 [28] 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 it receives from other nodes. Similarly in [48], in the reputation mechanism (for large peer-to-peer systems) the rankings are directly related to the quality of service of the peers. The paper also discusses aggregation of the rankings through referrals and defense against misrepresented ratings by weighted majority techniques.

The authors of [36] explore a partially decentralized reputation system in which both the requesters and the providers are selected based on their reputation rankings. 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 [37]. 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 [24], is one of the initial proposals on aggregated reputation rankings that was originally designed to decrease the number of inauthentic files in P2P sharing network and to isolate malicious users. A unique global value is assigned to a peer and is updated 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 algorithm that calculates a global trust vector at each node which are then used for both isolating malicious users and 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 [25], authors further show that these incentives reward cooperative peers and give new peers a fair opportunity to cooperate.

Finally the authors of [35], introduce two incentive based mechanisms one considering storage as a resource the other bandwidth. Authors also show that playing a game gives incentive to the peers to be honest about their resources.

III. Incentive Architecture and Mechanism

In this section we present a specific reputation system that in the absence of misrepresentation of the reputation values eventually reveals the true level of cooperativeness in the network. This has also been proved for a special case. Since reputation systems typically rely on referrals, they are vulnerable to attacks such as badmouthing. While the focus of the paper is not to address different attacks on P2P reputation systems but rather to propose a novel reputation model, our reputation system also accounts for deliberate misrepresentations of the reputations by the referred nodes.

A. Model of a "Flat" Reputation System for P2P 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.

B. 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[5] 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 is 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 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 $G_j(\pi_j, \bar{R}_i(n-1))$:

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

for some fixed $C > 0$, i.e., R_{ij} becomes relatively larger only when the transaction ij succeeds. Note that if (1) holds for $\mathbf{R}(n-1)$, then it will also hold for $\mathbf{R}(n)$ under (3).

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.

C. Modeling the Transaction Process

As mentioned above and reasoned in the introductions, a model of reputation dynamics does not need to consider the *circumstances* of a failed query. Moreover, we can combine the

[5] We consider the cooperativeness of a user a behavior and not related to the amount of resources it has cached. A node that does not possess a resource would simply not receive a query for it.

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

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 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 the remainder of this paper, we will study the reputation framework described above without consideration of network resources [29], [30], [34]), P2P query resolution [9], [6], and dynamic user demand.

D. Mean Reputation Process

For $n \geq 1$, we can now directly derive for "complete and honest" polling:

$$\begin{aligned}
\mathsf{E}(R_{ij}(n) \mid \mathbf{R}(n-1)) \ &= \ \left(1 - \sum_{k,\ k\neq i} \rho_{ik}\right) R_{ij}(n-1) \\
&\quad + \rho_{ij}[R_{ij}(n-1)(1 - G_j(\pi_j, \bar{R}_i(n-1)) \\
&\qquad + \frac{R_{ij}(n-1) + C}{1+C} G_j(\pi_j, \bar{R}_i(n-1))] \\
&\quad + \sum_{k,\ k\neq i,j} \rho_{ik}[R_{ij}(n-1)(1 - G_k(\pi_k, \bar{R}_i(n-1)) \\
&\qquad + \frac{R_{ij}(n-1)}{1+C} G_k(\pi_k, \bar{R}_i(n-1))] \\
&= \ \left(1 - \frac{C}{1+C} \sum_{k,\ k\neq i} \rho_{ik} G_k(\pi_k, \bar{R}_i(n-1))\right) R_{ij}(n-1) \\
&\quad + \frac{C}{1+C} \rho_{ij} G_j(\pi_j, \bar{R}_i(n-1)). \qquad (4)
\end{aligned}$$

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.

E. 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 $0 \leq \pi, \bar{R} \leq 1$,

$$\varepsilon\pi \ \leq \ G(\pi, \bar{R}) \ \leq \ \pi. \qquad (5)$$

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

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

for nondecreasing g. So, by (5), $g(0) \geq \varepsilon$ 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 the response functions in such a way that are all strictly positive (when $\pi > 0$), a new cooperative node with low initial reputation can still obtain content in order to later satisfy queries and increase its reputation. The following result could apply to a system intra-group reputations or a system inter-group reputations as discussed in subsection III-F below.

Theorem 1: If G is separable and the statement of (5) holds, then for complete and honest polling (4):

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

The proof can be found in the appendix.

So, the mean reputation $\mathsf{E}R_{ij}(n)$ is a consistent estimator of the mean rate $\rho_{ij}\pi_j$ of successful transactions ij. In particular, 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 $\Pi_{-i} = \sum_{j,\ j\neq i} \pi_j$. 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.

F. Trust groups for scalability and reliability

In this subsection, we survey well-known techniques for deploying and securing large-scale reputation systems. [23]. Potential attacks on cumulative distributed reputation systems include but, are not limited to, badmouthing and ballot box stuffing which are variations of Byzantine attacks (false reputation referrals and associated collusions).

Similar to [13], [24], [31], [48], we can account for misrepresentation and subsampling of reputations by using the following instead of $\bar{R}_i$ in the reputation model:

$$\bar{R}_{ji}(n) \ = \ \frac{\sum_{k,k\neq i} \ \lambda_{jki} h(R_{jk}(n)) R_{ki}(n)}{\sum_{k,k\neq i} \ h(R_{jk}(n))} \qquad (6)$$

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[6] [31]. 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." Note that the consistency theorem I can be extended to the more complex situation of equation 6.

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

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

in (6). i.e.,

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

As a special case, we can model federations that are used by peers for *feasible and reliable* reputation polling (in the presence of both lying and spoofing of reputation referrals), 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 $\{\mathcal{N}_m\}_{m=1}$ need not be a partition, i.e., a single peer j could belong to more than one group. Also note that we clearly need to assume that the denominator of (6) is always strictly positive for all i, j, n. Such local trust repositories can be securely implemented using one of a variety of existing light-weight authentication techniques within each group in a distributed or centralized fashion [12], [40], [15]. For example, this could involve a trusted server or "supernode" that maintains reputations and handles polling for the entire group. Obviously, the implicit assumption any "trust" group m (assumed small enough to allow complete sharing of reputation information among its peers) is that, for all appraised nodes i, $\lambda_{jki} \approx 1$ for all nodes $j, k \in \mathcal{N}_m$.

finally we consider an example of attacks targeting the P2P CDN itself, not just reputation systems, namely Sybil attacks [14] wherein one end-user employs many different identities. A typical solution to Sybil attacks involves a centralized *registration* server that authenticates a unique identifier upon registration for each peer. If authenticated reputations are based only on positive feedback and reputation referrals occur only among trusted peers, then multiple identities will dilute the reputation of the end-user thereby providing a natural disincentive for Sybil attacks. The effect of Sybil attacks would then be contained within their trust groups.

G. A simple game model

We now briefly describe a simple game model in which peers file-swap and adjust their reputation estimates. In addition, they may modify their own cooperation level. The game is organized into rounds where, in each round, peer nodes request files from each other. Whether the requestee (server peer of a transaction) grants or denies a request is based on the requestee's cooperation level and the requester's (client peer's) reputation ranking in the system. Upon a successful transaction, the requester increases the requestee's reputation level as described above.

All peers are engaged in number of transactions in any given round. At the end of a round, peers evaluate their success rate (number of successfully received files versus the number of requested files) for the current round. If a peer's ratio is above a threshold (the peer's "satisfaction" level), they reduce their cooperation level π to conserve resources (uplink bandwidth in particular [29], [30], [34]); otherwise, they become more cooperative in the hopes of receiving better service in the future rounds through improved reputation.

More specifically, the reputation system can be set-up in different ways. For example, a non-hierarchical (flat) framework can be used where transactions between any two nodes depend on the individual reputations of the nodes via the mean reputation of the requestee assuming all other nodes are polled. Alternatively, a hierarchical framework in which the nodes are arranged in (trust) groups can be used: the intra-group transactions occur as in the first phase (with only intra-group polling), and the inter-group transactions involve group reputations instead of the individual ones.

In the simulation study described in [33], we showed how the reputation system encouraged all peers to cooperate to receive better service. We also showed how different satisfaction levels impacted peers' decisions to cooperate. In both hierarchical and non-hierarchical reputation framework, as expected from the theorem in section III-E, the simulation results showed that the reputation values of a node converged to the nodes propensity to cooperate. When reputation values were misrepresented by some peers, a corresponding reduction in successful transactions by the "victimized" client peers was observed, compared to previous experiments assuming completely honest reporting. This study also showed that peers with medium to high expectations were encouraged by the reputation system to increase their propensities to cooperate. Finally, in [34], we modeled the actions of peer nodes were modeled with a game in which users modify the amount of uplink bandwidth they allocate to the network [29], [30]. Increasing uplink bandwidth resulted in improved reputation that, in turn, resulted in improved downloading performance.

IV. Summary

In summary, we surveyed cumulative 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, i.e., assuming rationally selfish but honest end-users. The paper concluded with a description of game theoretic models that employed these reputation systems.

References

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

[2] eBay. http://ebay.com.

[3] Igloo. http://vil.nai.com/vil/content/v_100046.htm.

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

[5] Kazaa lite. http://www.kazaalite.nl.

[6] P2P resources. http://www.cs.dartmouth.edu/~zhaom/research/marianas /resource.html.

[7] VBS.gnutella worm. http://securityresponse.synmantec.com/avcenter/venc /data/vbs.gnutella.html.

[8] E. Adar and B. A. Huberman. Free riding on gnutella. *First Monday magazine*, Sept. 2000.

[9] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys (CSUR)*, 36(4), December 2004.

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

[11] 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.

[12] 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.

[13] 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.

[14] M. Castro, P. Druschel, A. Ganesh, A. Rowstron, and D. S. Wallach. Security for structured peer-to-peer overlay networks. In *Proceedings of the 5th USENIX Symposium on Operating Systems Design and Implementation (OSDI '02)*, Boston, Massachusetts, December 2002.

[15] 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.

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

[17] 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.

[18] 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.

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

[20] 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.

[21] 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.

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

[23] A. Jsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. 2005. http://security.dstc.edu.au/papers/JIB2005-DSS.pdf.

[24] 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.

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

[26] 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.

[27] 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.

[28] 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.

[29] 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.

[30] 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.

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

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

[33] B. Mortazavi and G. Kesidis. Incentive-compatible reputation systems for P2P CDNs with rationally selfish but honest users. submitted.

[34] B. Mortazavi and G. Kesidis. Model and simulation study of a peer-to-peer game with a reputation-based incentive mechanism. In *Proc. Information Theory and Applications Workshop (IEEE), UC San Diego*, Feb. 2006.

[35] T.-W. J. Ngan, A. Nandi, A. Singh, D. S. Wallach, and P. Druschel. On designing incentives-compatible peer-to-peer systems. In *2nd Bertinoro Workshop on Future Directions in Distributed Computing (FuDiCo II: S.O.S.)*, Bertinoro, Italy, June 2004.

[36] 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.

[37] 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.

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

[39] 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.

[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] A. Singh, M. Castro, A. Rowstron, and P. Druschel. Defending against eclipse attacks on overlay networks. In *Proceedings of the 11th ACM SIGOPS European Workshop*, Leuven, Belgium, September 2004.

[45] A. Singh, T.-W. J. Ngan, P. Druschel, and D. S. Wallach. Eclipse attacks on overlay networks: Threats and defenses. In *IEEE Infocom 2006*, Barcelona, Spain, Apr. 2006. To appear.

[46] 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.

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

[48] 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.

APPENDIX: PROOF OF THEOREM 1

For separable G, define

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

for all $i \neq j$ and $n \geq 0$. By (4),

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

where (5) 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. □