Overcoming Free-Riding Behavior in Peer-to-Peer Systems

MICHAL FELDMAN and JOHN CHUANG
School of Information Management and Systems
University of California, Berkeley

While the fundamental premise of peer-to-peer (P2P) systems is that of voluntary resource sharing among individual peers, there is an inherent tension between individual rationality and collective welfare that threatens the viability of these systems. This paper surveys recent research at the intersection of economics and computer science that targets the design of distributed systems consisting of rational participants with diverse and selfish interests. In particular, we discuss major findings and open questions related to free-riding in P2P systems: factors affecting the degree of free-riding, incentive mechanisms to encourage user cooperation, and challenges in the design of incentive mechanisms for P2P systems.

Categories and Subject Descriptors: C.2.4 [**Computer-Communication Networks**]: Distributed Systems; J.4 [**Social And Behavioral Sciences**]: Economics

General Terms: Algorithms, Design, Economics, Performance, Reliability

Additional Key Words and Phrases: incentives, cooperation, peer-to-peer, hidden-action, hidden-information, game-theory

1. INTRODUCTION

The peer-to-peer (P2P) communications model has emerged a widely deployed alternative to the traditional client-server model for many distributed systems. In a typical P2P system, each node is owned and operated by an independent entity, and the nodes collectively form a self-organizing, self-maintaining network with no central authority. As a result, P2P system performance is highly dependent on the amount of voluntary resource contribution from the individual nodes.

Traditional system design assumes *obedient* users – users who adhere to a specified protocol without consideration of their own utility. However, this obedience assumption appears unrealistic in P2P settings where individual participants may interact with one another with varying degrees of collaboration and competition. Therefore, researchers have turned in recent years to a model of *rational* users – users who act to maximize their own utility, including deviating from the proto-

Authors' address: M. Feldman and J. Chuang, School of Information Management and Systems, 102 South Hall, Berkeley, CA 94720-4600.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.
© 2005 ACM 1529-3785/2005/0700-0041 $5.00

Samsung v. Truesight
IPR2024-01477
Exhibit 1015

Exhibit 1015
Page 01 of 10

col specification if they could increase their utility by doing so. In P2P systems, where cooperation may incur significant communication and computation costs, rational users may refuse to contribute their fair share of resources. Thus, individual rationality is in conflict with social welfare.

Users who attempt to benefit from the resources of others without offering their own resources in exchange are termed "free-riders." In 2000, a measurement study of the Gnutella file-sharing network [Adar and Huberman, 2000] found that approximately 70% of peers provide no files and that the top 1% of the peers provide approximately 37% of the total files shared. Similar patterns have been observed in subsequent studies of Napster and Gnutella networks [Saroiu et al., 2002]. In 2005, [Hughes et al., 2005] found free-riders have increased to 85% of all Gnutella users.

The free-riding phenomenon is by no means unique to P2P systems. However, the characteristics of P2P systems present interesting challenges and opportunities for the design of incentive-compatible systems. Some of these characteristics include: lack of central authority, highly dynamic memberships, availability of cheap identities ("pseudonyms"), hidden or untraceable actions, and collusive behavior. This paper surveys different approaches, e.g., based on payments or reciprocity, to alleviate or overcome this free-riding problem.

While free-riding is the focus of this paper, rational behavior manifests itself in many other distributed system settings, including strategic network formation [Fabrikant et al., 2003; Christin and Chuang, 2005], selfish routing [Cole et al., 2003] and interconnection [Ferreira and Sirbu, 2005], congestion control [Akella et al., 2002], multicast cost sharing [Feigenbaum et al., 2001], and selfish caching [Chun et al., 2004], just to name a few. At the conclusion of this paper, we will outline some of the important open questions and fruitful areas of research in this area.

2. RESEARCH QUESTIONS

There is growing recognition among distributed system designers that the ultimate success of their system depends not just on traditional technical considerations such as performance, robustness and scalability, but also on economic considerations such as incentive-compatibility. Failure to address vulnerabilities to rational manipulation may lead to the failure of an otherwise well-designed system.

In designing incentive-compatible P2P systems, it is important to understand the characteristics and factors that influence the degree of free-riding, and use this understanding to propose and analyse incentive schemes that mitigate free-riding behavior by rational users. In particular, researchers in the field are concerned with the following research questions:

—What factors determine the pervasiveness of free-riding?

—What is the effect of free-riding on system performance?

—How do system characteristics impose challenges on the design of incentive schemes?

—What mechanisms can be used to facilitate cooperative behavior in P2P systems?

—How does the availability of cheap identities affect user behavior in P2P systems, and what are its implications on the desired generosity toward strangers?

Exhibit 1015
Page 02 of 10

—What is the effect of invisible (hidden) actions on the design of incentive mechanisms?

—What are appropriate assumptions about the rationality of users and their behavior in P2P systems?

3. PROPOSED INCENTIVES FOR COOPERATION

Users contribute resources to a P2P system for a number of reasons. Some may offer resources in exchange for receiving other resources in the present time or in the future. Others may do so under threat of retaliation. Still others do so out of altruism. In this section, we outline some of the incentive schemes that have been proposed in recent years to encourage user cooperation within P2P systems.

3.1 Inherent Generosity

Empirical research in behavioral economics demonstrates that purely self-interested models usually fail to explain the observed behavior of people [Camerer, 2003]. The *warm-glow* model [Andreoni, 1989] is based on the insight that some users gain utility from the mere act of giving. Consistent with this approach, [Feldman et al., 2004b] devise a modeling framework that studies the phenomenon of free-riding in P2P systems while taking user *generosity* into account. Every user decides whether to contribute or free-ride based on how her generosity compares to the current contribution cost in the system, which is the inverse of the total percentage of contributors. The model analytically determines the resulting percentage of free-riders in the system based on the distribution of generosity in the population. We find that if the societal generosity is below a certain threshold, then there are too many selfish rascals around and the system collapses. But if it exceeds the threshold, the contribution level increases in the societal generosity with diminishing returns.

3.2 Monetary Payment Schemes

Monetary payment schemes dictate that the service recipients simply pay the service providers for resources they consume. One of the first studies that considered payment schemes in P2P transactions is [Golle et al., 2001], which uses a game theoretical model to study the potential benefits of introducing micropayment methods into P2P file-sharing systems. Since then, various pricing schemes have been proposed in the context of file-sharing networks, and more generally, packet forwarding in multi-hop networks.

Monetary schemes allow for rich and flexible economic mechanisms but suffer the notable drawback of seeming highly impractical since they require an infrastructure for accounting and micropayments. Much of the research in this field is less concerned with the feasibility of micropayments but instead considers problems that remain under the assumption that monetary exchanges are possible. In what follows, we outline some of the major problems that are studied in this area and approaches that have been proposed to solve them.

(1) **The individual costs of service providers may be hidden from the service recipient. How can we encourage truthful revelation of costs?** The problem of hidden-information has long been of interest in the economics literature. The theory of *mechanism design* (MD) provides an elegant mathematical

Exhibit 1015
Page 03 of 10

framework for game design, whereby the behavior of strategic players results in the socially desired outcome. It accomplishes the alignment of incentives by means of a monetary exchange between the participating entities. A rich body of literature has been dedicated to the application of MD to software agents and Internet applications. [Rosenschein and Zlotkin, 1994] and [Varian, 1995] argue that the "hyper-rational" view of MD may in fact be more appropriate for software agents, which have much better computational power than do human beings.

Classical MD, however, (i) ignores the computational complexity of the proposed algorithms and (ii) assumes a central model in which all nodes can communicate directly with a "center." In an attempt to transform existing mechanisms into protocols for use in computational environments, these two issues should be addressed. To that end, Nisan and Ronen [Nisan and Ronen, 1999] have introduced the notion of *algorithmic mechanism design* (AMD), which creates a formal model of centralized computation that combines incentive compatibility and computational tractability. AMD has been further extended to *distributed algorithmic mechanism design* (DAMD) by [Feigenbaum and Shenker, 2002]. DAMD seeks solutions within this framework that are both fully distributed and computationally tractable.

VCG (Vickrey-Clarke-Groves) mechanisms [Clarke, 1971; Groves, 1973; Vickrey, 1961] are a class of efficient and *strategyproof* mechanisms (i.e., encourage truthful revelation in *dominant strategy*). [Feigenbaum et al., 2002] propose a strategyproof mechanism (of the VCG family) for shortest paths in BGP routing that induces truthful revelation of transit costs. [Zhong et al., 2003] design a payment mechanism for packet forwarding in mobile ad-hoc networks in which the payment to intermediate nodes who carry the packets is based on reports sent by the nodes to a central center. The payments are designed in a way that encourages truthful revelation about the service that has been provided to and consumed by users.

(2) **The actions of the service providers may be *hidden* from the service recipient. How can we encourage cooperative behavior despite hidden action?**

The *principal-agent* (PA) model has been widely used in the economics literature to study situations in which a task is delegated to an agent, whose action is hidden from the principal. The model seeks to determine the optimal way for the principal to cope with this type of *information asymmetry*.

[Feldman et al., 005b] apply the PA model to routing in multi-hop networks, where the actions taken by individual intermediate nodes, i.e., routers, are typically hidden from the communicating endpoints. This work shows how the hidden-action problem can be overcome through appropriate design of contracts, in both the direct (the endpoints contract with each individual router) and recursive (each router contracts with the next downstream router) cases. It further demonstrates that per-hop monitoring does not necessarily improves the utility of the principal or the social welfare in the system. In addition, it generalizes existing mechanisms that deal with hidden-information [Feigenbaum et al., 2002] to handle scenarios involving both hidden-information and hidden-action.

(3) **How should the payment be delivered from the recipient to the providers?**

[Jakobsson et al., 2003] propose an architecture for micropayments in multi-hop

Exhibit 1015
Page 04 of 10

celluar networks in which senders upload tokens into their packets that are later redeemed by the forwarding nodes according to the reports of the nodes in the system. This scheme relies on a trusted accounting center that makes the payments to nodes based on the received tokens.
[Buttyan and Hubaux, 2002] propose two decentralized payment schemes based on virtual currency. In the first, the sender uploads the required payment, in full, into the forwarded packet and intermediate nodes each "earn" a portion of the payment as the packet traverses through them. In the second, the receiver pays for the packet delivery. Each intermediate node earns profits by buying the packet from its previous node and then selling it to the subsequent node at a higher price; eventually the receiver buys the packet, thus paying for the total forwarding cost.

3.3 Reciprocity-Based Schemes

In reciprocity-based schemes, users maintain histories of past behavior of other users and use this information in their decision making processes. These schemes can be based on direct reciprocity or indirect reciprocity. In direct-reciprocity schemes, user A decides how to serve user B based solely on the service that B has provided to A in the past. In contrast, in indirect-reciprocity schemes, the decision of A also depends on the service that B has provided to other users in the system.

Direct-reciprocity schemes are appropriate for applications with long session durations, as they provide ample opportunities for reciprocation between pairs of users. For example, the BitTorrent file-distribution system employs a tit-for-tat incentive mechanism to encourage cooperative behavior between a set of nodes performing coordinated exchange of large digital files [Cohen, 2003]. A measurement study of several BitTorrent networks found increased levels of cooperation over other P2P file-sharing networks [Andrade et al., 2005]. Even then, it is still possible to identify vulnerabilities to rational manipulation in its specifications [Shneidman and Parkes, 2004]. A recent study demonstrates, via analysis and PlanetLab experiments, that a free-rider can still realise the same file download completion time as a contributor [Jun and Ahamad, 2005].

Direct-reciprocity is also appropriate for different flavors of P2P multicast streaming applications. In [Ngan et al., 2004], the multicast tree is periodically rebuilt in order to facilitate direct reward and retaliation by neighbor nodes. In [Castro et al., 2003], multiple delivery trees are used for each streaming session, and no participant can occupy a leaf position on all trees and avoid forwarding traffic to others. As a generalization of the tit-for-tat scheme, [Chu et al., 2004] propose a *taxation* scheme for P2P streaming applications. In accounting for the heterogeneity in upstream bandwidths among users, a "tax schedule" specifies the amount of subsidy between resource rich and resource poor nodes.

Many indirect-reciprocity schemes have been proposed in the literature, and they are often called reputation-based schemes. They differ from one another primarily in the computation of reputation scores and the mapping of scores to strategies (e.g., [Dellarocas and Resnick, 2003; Gupta et al., 2003; Jurca and Faltings, 2005; Kamvar et al., 2003; Yu and Singh, 2000]). Indirect-reciprocity schemes are more scalable than direct-reciprocity schemes, especially for P2P systems with large population sizes, highly dynamic memberships, and infrequent repeat transactions [Feldman et al., 2004a]. However, indirect-reciprocity schemes rely on second-

Exhibit 1015
Page 05 of 10

hand observations and thus must confront trust issues that do not arise in direct-reciprocity schemes.

A growing recognition of the potential that lies in reciprocity-based schemes has led to an extensive amount of research in this area. Here are some of the main issues that have been explored.

(1) **How should *newcomers* be treated in reciprocity-based schemes?** The assumption of persistent identities breaks down in systems where identities are cheap and can be obtained or replaced at almost no cost. This is desirable for network growth as it encourages newcomers to join the system. However, this also allows misbehaving users to escape the consequences of their actions by repeatedly switching to new identities (i.e., *whitewashing*), while being indistinguishable from legitimate newcomers.

There are two ways to counter whitewashing attacks. The first is to require the use of free but irreplaceable pseudonyms, e.g., through the assignment of strong identities by a trusted central authority [Castro et al., 2002]. In the absence of such mechanisms, it may be necessary to impose a penalty on all newcomers, including both legitimate newcomers and whitewashers, which exacts a social cost [Friedman and Resnick, 1998; Feldman et al., 2004b; Feldman and Chuang, 2005]. However, it may be possible to adopt a stranger policy that adapts to the observed behavior of previous strangers, thereby avoiding unncessary incurrences of the social cost [Feldman et al., 2004a].

(2) **Strategies that are based on indirect reciprocity are vulnerable to collusive behavior such as false accusation and false praise.**

Several studies propose statistical methods to detect misbehavior and mitigate attacks on reputation systems. [Buchegger and LeBoudec, 2004] propose a Bayesian approach, in which nodes monitor their neighbors' behavior and periodically exchange this information. Nodes update their view based on their neighbors' reports if the information passes a Bayesian deviation test. This scheme reflects a more general approach calling for the storage of two reputation scores for each user: one that indicates the user's contribution level in the system, and a second one that reflects upon the user's trustworthiness in reporting information about other users. [Dellarocas, 2000] takes a different approach and proposes to alleviate these attacks through methods that are based on anonymity and cluster filtering.

Yet other researchers propose to base the computation of the reputation scores on the weighted graph that represents the interactions between the entities. The nodes of the graph signify the interacting entities and the weighted edges signify direct interactions between the entities along with their quality. This approach is often used to analyze false praise attacks. [Kamvar et al., 2003] devise the EigenTrust algorithm that calculates the PageRank on the recommendation graph, where recommendations from trusted nodes are worth more. [Feldman et al., 2004a] propose to apply the maxflow algorithm on the graph. The maxflow is the greatest level of reputation the source can give to the sink without violating "reputation capacity" constraints. This method alleviates collusion behavior in which a group of non-cooperative users dishonestly report high reputation values on each other to subvert the reputation system.

The effect of collusion is magnified in systems with cheap pseudonyms, where users

Exhibit 1015
Page 06 of 10

can engage in the *sybil attack* [Douceur, 2002] – create fake identities and collude with their own multiple identities. [Cheng and Friedman, 2005] formalize the notion of *sybilproofness* in reputation schemes. They demonstrate that if reputation is determined only by the structure of the graph and the edge values, then no sybilproof mechanism exists. Conversely, if reputation values are computed with respect to a fixed node in the graph, sybilproofness is guaranteed when several conditions are satisfied. In particular, some flow-based algorithms (similar to the algorithm presented in [Feldman et al., 2004a]) satisfy the conditions of sybilproofness.

4. FUTURE DIRECTIONS

Overcoming free-riding behavior is central to the performance and robustness of P2P systems. Many incentive mechanisms have been proposed and implemented with the goal of inducing cooperation from strategic peers. The different studies and approaches presented in this survey suggest several important avenues of research.

— Game theory is perhaps the most pervasive tool in the study of incentive schemes for P2P systems, and different variants of the prisoner's-dilemma (PD) model are used in this area. Is game theory the right tool? Is PD the right model? What characteristics do these tools capture and what are missing?

— In applying game theory to system design, a fundamental question is "what constitutes rational behavior?" The predominant concept of rationality is that of Nash equilibrium. Yet, many other *solution concepts* exist. For example, models of *near* rationality may better predict empirically observed behavior than models of *perfect* rationality[Christin et al., 2004]. Other solution concepts may be appropriate for environments with low information setting and/or highly dynamic player populations [Friedman and Shenker, 1997]. It is important to understand the appropriateness of a range of solution concepts to different applications and how various techniques can affect the solution concepts achieved.

— As suggested by [Feigenbaum and Shenker, 2002], one may envision an environment with different types of agents, ranging from rational and bounded-rational agents to obedient, malicious or faulty ones, having goals other than utility maximization. How would these different user behavior models affect the design of incentive schemes?

— Is it possible to establish a common framework, including the formalization of adversarial models, definitions of robustness, etc., so as to facilitate a fair and meaningful evaluation of the the large number of reputation systems [Morselli et al., 2004]? Is it possible to prove robustness using rigorous analytic tools? What are the limitations of an analytic model in the context of reputation systems? Both possibility and impossibility results would be valuable.

— Contracts can be used to encourage desired behavior of agents in equilibrium. What is the computational complexity of optimal contracts in different tasks and network structures?

In summary, user behavior can have potentially devastating effects on P2P system performance, and so must be explicitly accounted for in P2P system design. Great strides have been made and valuable lessons have been drawn by the research community in the last few years. For example, we now appreciate the importance

Exhibit 1015
Page 07 of 10

of identifying and tackling rational manipulation at all stages, including revelation of private information, computation of algorithms, and execution of tasks. At the same time, it is clear that further progress is needed on many fronts, e.g., insights on appropriate models and tools, so that the intellectual foundation is in place for the practice of economics-informed design of P2P systems.

ACKNOWLEDGMENTS

This work is supported in part by the National Science Foundation under ITR awards ANI-0085879 and ANI-0331659.

REFERENCES

Adar, E. and Huberman, B. A. (2000). Free Riding on Gnutella. *First Monday*, 5(10).

Akella, A., Karp, R., Papadimitriou, C., Seshan, S., and Shenker, S. (2002). Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP. In *SIGCOMM 2002*.

Andrade, N., Mowbray, M., Lima, A., Wagner, G., and Ripeanu, M. (2005). Influences on cooperation in bittorrent communities. In *To appear, ACM SIGCOMM workshop on the Economics of Peer-to-Peer Systems (P2PECON'05)*.

Andreoni, J. (1989). Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence. *Journal of Political Economy*, 97(6):1447–58.

Buchegger, S. and LeBoudec, J.-Y. (2004). A Robust Reputation System for Peer-to-Peer and Mobile Ad-Hoc Networks. In *Workshop on Economics of Peer-to-Peer Systems*.

Buttyan, L. and Hubaux, J. (2002). Enforcing Service Availability in Mobile Ad-Hoc WANs. *IEEE/ACM Workshop on Mobile Ad Hoc Networking and Computing (MobiHOC)*.

Camerer, C. F. (2003). *Behavioral Game Theory*. Princeton University Press.

Castro, M., Druschel, P., Ganesh, A., Rowstron, A., and Wallach, D. S. (2002). Security for Structured Peer-to-Peer Overlay Networks. In *Proceedings of Multimedia Computing and Networking 2002 (MMCN '02)*.

Castro, M., Druschel, P., Kermarrec, A.-M., Nandi, A., Rowstron, A., and Singh, A. (2003). Splitstream: High-bandwidth multicast in cooperative environments. In *Proceedings of 19th ACM Symposium on Operating Systems Principles*.

Cheng, A. and Friedman, E. (2005). Sybilproof reputation mechanisms. In *To appear, ACM SIGCOMM workshop on the Economics of Peer-to-Peer Systems (P2PECON'05)*.

Christin, N. and Chuang, J. (2005). A cost-based analysis of overlay routing geometries. In *(IEEE INFOCOM*.

Christin, N., Grossklags, J., and Chuang, J. (2004). Near Rationality and Competitive Equilibria in Networked Systems. In *Proc. SIGCOMM workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*.

Chu, Y.-H., Chuang, J., and Zhang, H. (2004). A Case for Taxation in Peer-to-Peer Streaming Broadcast. In *Proc. SIGCOMM workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*.

Chun, B., Chaudhuri, K., Wee, H., Barreno, M., Papadimitriou, C., and Kubiatowicz, J. D. (2004). Selfish caching in distributed systems. In *ACM Principles of Distributed Computing (PODC)*.

Clarke, E. (1971). Multipart Pricing of Public Goods. *Public Choice*, 1:17–33.

Cohen, B. (2003). Incentives build robustness in bittorrent. In *Workshop on Economics of Peer-to-Peer Systems*.

Cole, R., Dodis, Y., and TimRoughgarten (2003). Pricing Networks with Selfish Routing. In *Workshop on Economics of Peer-to-Peer Systems*.

Dellarocas, C. (2000). Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behavior. In *Second ACM Conference on Electronic Commerce*.

Dellarocas, C. and Resnick, P. (2003). Online Reputation Mechanisms: A Roadmap for Future Research. Summary Report of the First Interdisciplinary Symposium on Online Reputation Mechanisms. In *Workshop on economics of peer-to-peer networks*.

Exhibit 1015
Page 08 of 10

Douceur, J. R. (2002). The Sybil Attack. In *Electronic Proceedings of the International Workshop on Peer-to-Peer Systems*.

Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., and Shenker, S. (2003). On a network creation game. In *ACM Symposium on Principles of Distriubted Computing (PODC)*.

Feigenbaum, J., Papadimitriou, C., Sami, R., and Shenker, S. (2002). A BGP-based Mechanism for Lowest-Cost Routing. In *Proceedings of the ACM Symposium on Principles of Distributed Computing*.

Feigenbaum, J., Papadimitriou, C., and Shenker, S. (2001). Sharing the Cost of Multicast Transmissions. In *Journal of Computer and System Sciences*, volume 63, pages 21–41.

Feigenbaum, J. and Shenker, S. (2002). Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In *Proceedings of the International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*.

Feldman, M. and Chuang, J. (2005). The Evolution of Cooperation Under Cheap Pseudonyms. In *Proc. of the seventh international IEEE conference on E-Commerce Technology*.

Feldman, M., Chuang, J., Stoica, I., and Shenker, S. (2005b). Hidden-Action in Multi-Hop Routing. In *ACM Conference on Electronic Commerce (EC'05)*.

Feldman, M., Lai, K., Stoica, I., and Chuang, J. (2004a). Robust Incentive Techniques for Peer-to-Peer Networks. In *ACM Conference on Electronic Commerce (EC'04)*.

Feldman, M., Papadimitriou, C., Stoica, I., and Chuang, J. (2004b). Free-Riding and Whitewashing in Peer-to-Peer Systems. In *Proc. SIGCOMM workshop on Practice and Theory of Incentives and Game Theory in Networked Systems*.

Ferreira, P. and Sirbu, M. (2005). Interconnected Communication Networks Provisioned Selfishly. In *ACM Conference on Electronic Commerce (EC'05)*.

Friedman, E. and Resnick, P. (1998). The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199.

Friedman, E. and Shenker, S. (1997). Learning and implementation on the internet. In *Manuscript. New Brunswick: Rutgers University, Department of Economics*.

Golle, P., Leyton-Brown, K., Mironov, I., and Lillibridge, M. (2001). Incentives For Sharing in Peer-to-Peer Networks. In *Proceedings of the 3rd ACM conference on Electronic Commerce, October 2001*.

Groves, T. (1973). Incentives in Teams. *Econometrica*, 41:617–663.

Gupta, M., Judge, P., and Ammar, M. (2003). A reputation system for peer-to-peer networks. In *Proceedings of the 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*.

Hughes, D., Coulson, G., and Walkerdine, J. (2005). Freeriding on gnutella revisited: The bell tolls? In *Submitted to IEEE Distributed Systems Online*.

Jakobsson, M., Hubaux, J.-P., and Buttyan, L. (2003). A Micro-Payment Scheme Encouraging Collaboration in Multi-Hop Cellular Networks. In *Financial Cryptography*.

Jun, S. and Ahamad, M. (2005). Incentives in bittorrent induce free riding. In *To appear, ACM SIGCOMM workshop on the Economics of Peer-to-Peer Systems (P2PECON'05)*.

Jurca, R. and Faltings, B. (2005). Reputation-based pricing of p2p services. In *To appear, ACM SIGCOMM workshop on the Economics of Peer-to-Peer Systems (P2PECON'05)*.

Kamvar, S. D., Schlosser, M. T., and Garcia-Molina, H. (2003). The EigenTrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the Twelfth International World Wide Web Conference*.

Morselli, R., Katz, J., and Bhattacharjee, B. (2004). A game-theoretic framework for analyzing trust-inference protocols. In *2nd Workshop on Economics of Peer-to-Peer Systems*.

Ngan, T.-W. J., Wallach, D. S., and Druschel, P. (2004). Incentive-compatible peer-to-peer multicast. In *Workshop on Economics of Peer-to-Peer Systems*.

Nisan, N. and Ronen, A. (1999). Algorithmic Mechanism Design. In *Proceedings of the 31st Symposium on Theory of Computing*.

Rosenschein, J. S. and Zlotkin, G. (1994). *Rules of Encounter*. MIT Press.

Exhibit 1015
Page 09 of 10

Saroiu, S., Gummadi, P. K., and Gribble, S. D. (2002). A Measurement Study of Peer-to-Peer File Sharing Systems. In *Proceedings of Multimedia Computing and Networking 2002 (MMCN '02)*.

Shneidman, J. and Parkes, D. C. (2004). Overcoming rational manipulation in mechanism implementation.

Varian, H. R. (1995). Economic Mechanism Design for Computerized Agents. In *Proc. of Usenix Workshop on Electronic Commerce*.

Vickrey, W. (1961). Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16:8–37.

Yu, B. and Singh, M. P. (2000). A Social Mechanism of Reputation Management in Electronic Communities. In *Proceedings of fourth international workshop in cooperative information*.

Zhong, S., Chen, J., and Yang, Y. R. (2003). Sprite: A Simple, Cheat-Proof, Credit-Based System for Mobile Ad-Hoc Networks. In *22nd Annual Joint Conference of the IEEE Computer and Communications Societies*.

Received June 2005; revised July 2005; accepted July 2005

Exhibit 1015
Page 10 of 10