Designing Incentive Mechanisms for Peer-to-Peer Systems

John Chuang
School of Information Management and Systems, UC Berkeley
chuang@sims.berkeley.edu

From file-sharing to mobile ad-hoc networks, community networking to application layer overlays, the peer-to-peer networking paradigm promises to revolutionize the way we design, build and use the communications network of tomorrow, transform the structure of the communications industry, and challenge our understanding of markets and democracies in a digital age.

The fundamental premise of peer-to-peer systems is that individual peers voluntarily contribute resources to the system. However, the inherent tension between individual rationality and collective welfare produces a misalignment of incentives in the grassroots provisioning of P2P services.

What makes free-riding a particularly difficult problem is the unique set of challenges that P2P systems pose: large populations, high turnover, asymmetry of interest, collusion, hidden actions, and zero-cost identities.

This paper discusses some of the research opportunities and challenges in the design of incentive mechanisms for P2P systems.

P2P System Characteristics

Peer-to-peer networks exhibit several economic characteristics that are different from those of traditional provider-centric networks. First, P2P networks derive their utility solely from the contribution of resources from cooperating nodes. End-nodes have to self-organize and assume the roles of routing, message forwarding, and network repair in order to maintain connectivity (at the network layer and/or the application layer). At the same time, nodes have every incentive to free-ride or defect, e.g., contribute less to the common resource pool than they consume from it, or contribute nothing at all. There is therefore a tension between the maximization of individual utility and the achievement of global optimality. Economic theory suggests that, in the absence of proper incentives, private provisioning of public goods is always sub-optimal, and this may lead to the "tragedy of the commons". Indeed, empirical evidence of free-riding behavior on P2P file-sharing networks has led to speculation of the inevitable collapse of such systems due to non-cooperation [1,23].

0-7803-8525-X/04/$20.00 ©2004 IEEE

Second, there is a general lack of accountability in these networks, even when anonymity is disallowed. First, the lack of a service provider or a centralized authority means no entity can be held responsible for poor performance. Second, the actions of individual nodes are often unobservable [9,11]. If a message fails to reach its intended destination, it is very difficult to identify the point of failure in a communication path. It is even more difficult to attribute the cause to either selfish or even malicious behavior on the part of a node, or to more mundane reasons such as network congestion, channel interference, node movement, or node joins/leaves. Therefore, nodes in an application layer overlay network or a mobile ad hoc network may choose to drop packets to conserve energy, and nodes in a file-sharing network may refuse to respond to queries to avoid the cost of uploading a file. This problem of *hidden action* or *moral hazard* is further exacerbated by the high turnover rates and the availability of cheap pseudonyms [13]. An attacker can launch a Sybil attack [6] by repeatedly inserting large numbers of resource-hungry nodes (with unique identities) into the system, rapidly draining the system of its resources.

Third, cooperation is difficult to establish in a network where there are few repeat transactions, i.e., nodes are continually interacting with nodes that they have not met before, and will most likely not meet again in the future. Consider the prisoners' dilemma problem in game theory: two persons accused of conspiring in a crime can either cooperate with each other (by staying silent) or defect (by testifying against the other). In the iterated prisoner's dilemma game, mutual cooperation is the Nash equilibrium if the game is played an infinite number of rounds, but defection becomes the dominant strategy when the game is played for only a finite number of rounds. Put another way, cooperation is rational only when there is a chance of payback (reciprocation) in the future. In the ad-hoc and overlay network context, where a node's connected neighbors may be frequently changing, it may be difficult for reciprocity, and therefore cooperation, to take hold.

Fourth, even with repeated transactions, there may be asymmetry of capabilities and interests between the nodes, leading to persistent trade imbalances and a potential breakdown in reciprocity. In a wireless ad-hoc network, the location of nodes in relation to the access point may result in an asymmetric resource contribution-consumption pattern. For example, in Figure 1, node X, located closest to the access point (AP), will likely carry more traffic than node Y, located further away. Node Y might have few opportunities to reciprocate to X. However, node Y might itself be carrying traffic for node Z. Should Z be treated as a free-rider in this case? In a P2P file-sharing network, Alice may be interested in downloading files from Bob's collection, but Bob may be uninterested in any file in Alice's collection. There might even exist a unidirectional cycle between Alice, Bob and Carol (Figure 2) that provides utility to all three persons without any direct reciprocity.

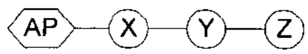

Figure 1: Example wireless ad-hoc network

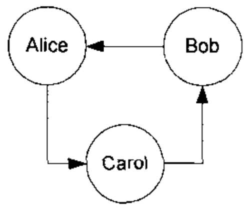

Figure 2: P2P file-sharing example

Research Opportunities and Challenges

Are P2P systems doomed to failure due to non-cooperation? How can we design incentive mechanisms to encourage cooperation in P2P networks? These questions lead to a set of interdisciplinary research opportunities and challenges for P2P system design, including:

- quantifying the incentives and disincentives for cooperation
- quantifying the impact of free-riding on system performance
- encouraging direct and indirect reciprocity
- leveraging peer selection (and other topology considerations)
- dealing with strangers and whitewashers
- overcoming information asymmetries (e.g., hidden action and hidden information)

Tackling these challenges will require cross fertilization of knowledge between computer science and various economics disciplines, including cooperative and non-cooperative game theory, theories of evolution and learning in games, mechanism design, theory of incentives and contracts, agency theory, public and club goods theory, and the economics of altruism and giving, reciprocity and reputation, and network formation.

First of all, what are the costs and benefits associated with participation in a P2P network [4]? To what extent do the costs and benefits depend on the behavior of the individual peer? Clearly, if cooperative behavior results in significant costs but negligible benefits for the peer, there is little incentive for the utility-maximizing peer to cooperate. The system designer can tackle this problem from both ends. First, various token-based and reputation-based schemes have been proposed to provide incentives for cooperative behavior in different network settings

[5,11,14,17,25,26]. In some settings, it may even make sense to create a market such that nodes compete with one another to contribute resources [18]. Second, any disincentives for cooperative behavior should be minimized if not eliminated [10]. Often, perceived disincentives are sufficient to discourage cooperation even when there is little or no actual disincentive. Techniques such as prioritization may be employed to reduce these disincentives.

On the other hand, it is often unnecessary, and sometimes undesirable, to aim for 100% cooperation in a P2P network. First, many networks can tolerate some degree of free-riding. As long as there are a sufficient fraction of cooperative nodes in the network that provides an acceptable (though not necessarily optimal) level of service, there may be no need for an elaborate and circumvention-prone mechanism to exclude free-riders or enforce 100% cooperation. In some cases, simply setting "cooperate" as the default behavior may be sufficient to produce an acceptable level of cooperation. Second, in networks with heterogeneous node capabilities, it may even be desirable for the resource-poor nodes to free-ride. There are a number of open questions. What is the optimal level of free-riding for a given network or application? What characteristics or parameters influence this level? What mechanisms are effective in selectively encouraging cooperation in some nodes and discouraging cooperation in others?

The key to overcoming the individual rationality of defection or free-riding is through encouraging *reciprocity* among the peers. Reciprocity can be facilitated through *repetition* and/or *reputation*.

Repetition allows cooperation to flourish via *direct* reciprocity. However, the lack of repeat transactions and asymmetry in node capabilities and interests in P2P networks means that we may need to turn to reputation to facilitate cooperation via *indirect* reciprocity. In a direct reciprocity scheme, node A will cooperate with node B only if node B has cooperated with (or not defected against) node A in the past. On the other hand, in an indirect reciprocity scheme, node A will cooperate with node B if node B has cooperated with other nodes in the system (subject to some threshold) [20,22]. Therefore, even though node B is a stranger to node A, node A can account for node B's transactional history (or reputation) in its decision-making. While this scheme can dramatically reduce the *strangeliness* of the network, it comes at a significant cost. Whereas direct reciprocity schemes only require the storage of local state, indirect reciprocity schemes require some shared state, either in the form of transaction logs, redeemable tokens, or reputation scores. Furthermore, significant challenges arise in the design of shared states storage and retrieval mechanisms that are scalable and robust against various forms of attacks including collusion [11]. The extent of state-sharing (from regional to global) will dictate the completeness and accuracy of state information on one hand, and the cost of maintaining such state on the other hand.

The *strangeliness* of the network can also be reduced via peer selection and topology formation. Instead of random neighbor selection, nodes with similar interests can self-organize to form affinity groups and preferentially transact with intra-cluster nodes [21]. We can think of clustering as a technique for encouraging direct reciprocity via repetition. Alternatively, we can consider clustering as transforming the P2P network from a public good into a club good [2], thereby avoiding the free-riding and tragedy-of-the-commons problems. On a cautionary note, uncoordinated matching by strategic nodes may lead to sub-optimal outcomes, with the efficiency loss quantified as the *price of anarchy* in studies of congestion games [19,24].

Peer selection itself can be used as a reward for cooperation. For example, in a P2P media streaming system, a receiver's performance is dependent on the number, location, and capabilities of its supplying peers. Therefore, a rational node may respond to an incentive scheme where contributors are rewarded with preferential selection of high quality supplying peers [15]. More generally, service differentiation offers an intuitive and compelling framework for designing P2P incentive mechanisms [3].

Strategic topology formation [7,27] is an important design consideration for P2P networks. P2P topologies can be constructed independently from the underlying network topology, transcending geographical and administrative boundaries. In many cases, the costs and benefits of participating in a P2P network is influenced not just by the behavior of the peers, but also by the topology of the network, the distribution of node capabilities, and the traffic matrix [4]. To the extent that any asymmetries in network topology can be matched with the distribution of node capabilities, the resulting variance of costs and benefits across different peers could be minimized, and strategic churning behavior could be avoided.

We have discussed various approaches to reduce the *strangeliness* of the network. Yet, even with perfect information (i.e., globally accessible, complete details of all transactions), nodes will still have to deal with strangers in the system. Specifically, there will always be newcomers that are joining the network for the first time, and therefore have no history whatsoever. In some cases, a newcomer may be a one-timer (e.g., an out-of-town visitor in a mobile ad-hoc network). In other cases, the newcomer may become an old-timer and take up an extended stay in the network. While the newcomers may turn out to be cooperators or defectors, they are usually given the benefit of the doubt in their very first transaction. For example, the tit-for-tat (TFT) strategy always cooperates with a stranger and reciprocates the last action of a non-stranger. Unfortunately, this cooperate-with-stranger policy fails miserably in an environment with cheap pseudonyms. In particular, a *whitewasher* can repeatedly enter the network with new identities, posing as newcomers, and exploit the nodes that always cooperate with strangers. The other alternative is to always defect on a stranger. While this strategy is robust against whitewashers, it produces a social cost in the form of a tax imposed on newcomers, and may exclude some nodes who would otherwise join the network [12,13]. One way to balance this tension between welcoming newcomers and protecting against whitewashers is to adopt a stranger-adaptive strategy [11]. A node will dynamically adjust its probability of cooperating with a stranger based on observations on the level of whitewashing in the system. This way, the tax is imposed only when necessary.

Last but not least, we need to tackle the various forms of information asymmetry in P2P networks, including *hidden action* and *hidden information*. To address the problem of hidden action, the objective is to design mechanisms such as contracts to provide incentives for peers to exert high effort even when their actions may be unobservable [9]. To address the problem of hidden information, the objective is to design distributed and scalable protocols and mechanisms to induce truthful revelation of private information by strategic peers [8]. Significant research challenges and opportunities remain in these very important areas.

Acknowledgements

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

References

[1] E. Adar and B. A. Huberman. Free riding on gnutella. *First Monday*, 5(10), October 2000.

[2] A. Asvanund, S. Bagla, M. Kapadia, R. Krishnan, M. D. Smith, R. Telang. Intelligent Club Management in Peer-to-Peer Networks. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

[3] C. Buragohain, D. Agrawal, S. Suri. A Game-Theoretic Framework for Incentives in P2P Systems. 3rd IEEE International Conference on Peer-to-Peer Computing, September 2003.

[4] N. Christin and J. Chuang. On the Cost of Participating in a Peer-to-Peer Network. 3rd International Workshop on Peer-to-Peer Systems (IPTPS'04), February 2004.

[5] R. Dingledine, M.J. Freedman, D. Molnar. The Free Haven Project: Distributed Anonymous Storage Service. Workshop on Design Issues in Anonymity and Unobservability, July 2000 (LNCS 2009).

[6] J. Douceur. The Sybil Attack. 1st International Workshop on Peer-to-Peer Systems, March 2002.

[7] B. Dutta and M.O. Jackson. (eds.) *Networks and Groups*. Springer Verlag, 2003.

[8] J. Feigenbaum and S. Shenker. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications, 2002.

[9] M. Feldman and J. Chuang. Hidden Action in Multi-hop Routing. 2004.

[10] M. Feldman, K. Lai, J. Chuang, I. Stoica. Quantifying Disincentives in Peer-to-Peer Networks. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

[11] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust Incentive Techniques for Peer-to-Peer Networks. ACM E-Commerce Conference (EC'04), May 2004.

[12] M. Feldman, C. Papadimitriou, J. Chuang and I. Stoica. Free-riding and Whitewashing in Peer-to-Peer Systems. 3rd Workshop on Economics and Information Security, May 2004.

[13] E. Friedman and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2), 173--199, 2001.

[14] P. Golle, K. Leyton-Brown, I. Mironov. Incentives in Peer-to-Peer File Sharing. ACM Conference on Electronic Commerce, 2001.

[15] A. Habib and J. Chuang. Service Differentiated Peer Selection: An Incentive Mechanism for Peer-to-Peer Media Streaming. 12th IEEE International Workshop on Quality of Service, June 2004.

[16] G. Hardin. The Tragedy of the Commons. *Science*, 162:1243--1248, 1968.

[17] S. Kamvar, M. Schlosser, H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. 12th International World Wide Web Conference, May 2003.

[18] S. Kamvar, B. Yang, H. Garcia-Molina. Addressing the Non-Cooperation Problem in Competitive P2P Systems. 1st Workshop on Economics of Peer-to-Peer Systems, Jun 2003.

[19] E. Koutsoupias, C. Papadimitriou. Worst-Case Equilibria. 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404-413, 1999.

[20] M. Nowak and K. Sigmund. Evolution of Indirect Reciprocity by Image Scoring. *Nature*, 393:573-- 577, 11 June 1998.

[21] V. N. Padmanabhan and K. Sripanidkulchai. The Case for Cooperative Networking. 1st International Workshop on Peer-to-Peer Systems, March 2002.

[22] G. Pollock and L. Dugatkin. Reciprocity and the Evolution of Reputation, *Journal of Theoretical Biology* 159, 25-37, 1992.

[23] S. Saroiu, K. Gummadi and S. Gribble. Measuring and analyzing the characteristics of Napster and Gnutella hosts. *Multimedia Systems* 9(2): 170-184, 2003.

[24] S. Suri, C. Toth, Y. Zhou. Uncoordinated Load Balancing and Congestion Games in P2P Systems. 3rd International Workshop on Peer-to-Peer Systems, February 2004.

[25] V. Vishnumurthy, S. Chandrakumar, E.G. Sirer. KARMA: A Secure Economic Framework for P2P Resource Sharing. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

[26] B. Wilcox-O'Hearn. Experiences Deploying a Large-Scale Emergent Network. 1st International Workshop on Peer-to-Peer Systems, March 2002.

[27] C. J. Woodard and D. Parkes. Strategyproof Mechanisms for Ad Hoc Network Formation. 1st Workshop on Economics of Peer-to-Peer Systems, June 2003.

Designing Incentive Mechanisms for Peer-to-Peer Systems

John Chuang
chuang@sims.berkeley.edu
http://p2pecon.berkeley.edu/

in collaboration with: Michal Feldman, Kevin Lai, Nicolas Christin, Ahsan Habib, Ion Stoica

Economics of P2P?

P2P is more than just illegal file sharing!

2

Agenda

- P2P systems and characteristics
- Designing incentive mechanisms: challenges and opportunities

3

P2P Systems

- Distributed storage and retrieval
  - File-sharing: Napster, gnutella, kaZaA, Overnet, bitTorrent, ...
  - Persistence: Eternity, Intermemory, Freehaven
  - Anonymity: FreeNet, Publius
  - DHTs: Chord, CAN, Pastry, Tapestry, OpenHash, ...
- Distributed computation: Grid, SETI@Home, etc.
- Communications
  - Connectivity: mobile ad-hoc networks, "rooftop" networks
  - Redundancy: resilient overlay networks
  - Scalability/deployability: end system multicast
  - Anonymity: onion-routing, MIX-net, Crowds

4

P2P System Characteristics (1)

- Utility derived solely from contributions by participating nodes
  - Individual rationality vs. collective welfare
  - Free-riding prevalent (70% in gnutella)
  - Tragedy of the commons

5

P2P System Characteristics (2)

- General lack of accountability, even when anonymity is disallowed
  - Absence of service provider
  - Highly dynamic node memberships
  - Difficult/costly to detect defections (hidden action)
  - Exacerbated by cheap pseudonyms

6

P2P System Characteristics (3)

- Few repeat transactions
  - Hard to establish cooperation through direct reciprocity
  - Iterated vs. One-shot Prisoners' Dilemma
- Asymmetries in capabilities and interests

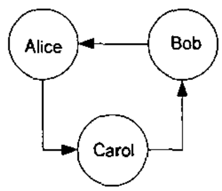

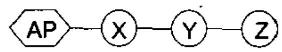

7

Questions

- Are P2P systems doomed to failure due to non-cooperation?
  - What are the cost and benefit of participating in a P2P system? [CC04]
  - What are the disincentives for cooperative behavior? [FLCS03]
- How do we incentivize cooperation given the set of system characteristics?
  - History (private vs. shared history) [FLSC04]
  - Dealing with strangers & whitewashers [FLSC04, FPCS04]
  - Service differentiation (peer selection) [HC04]
  - Overcoming information asymmetries (hidden information, hidden action) [FC04]

8

To serve or not to serve?

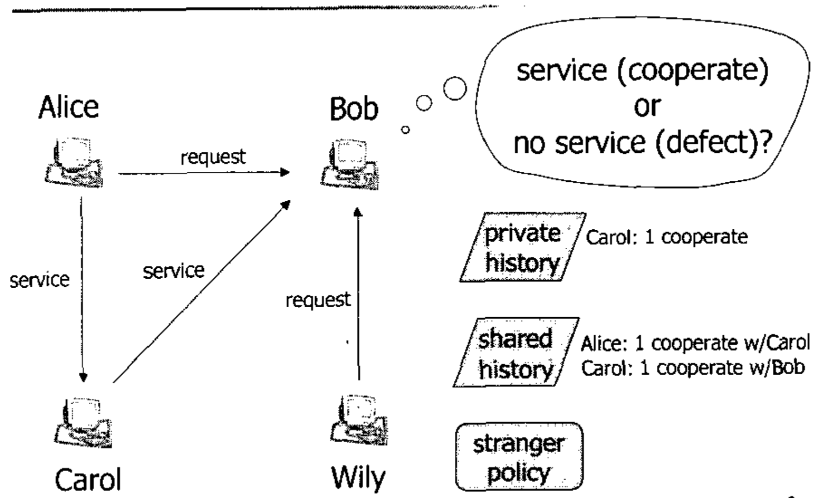

9

Private History

- Advantages
  - Implementation is simple and decentralized
  - Immune to collusion
- Disadvantages
  - Requires repeat transactions
    - e.g., low rate of turnover, small populations
  - Deals poorly with asymmetry of interest

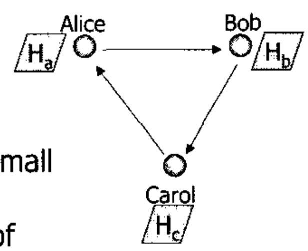

10

Shared History

- Advantages
  - Tolerates few repeat transactions (large populations, high turnover)
  - Tolerates asymmetry of interest
- Disadvantages
  - Susceptible to collusion
    - Subjective shared history via max-flow algorithm [FLSC04]
  - Implementing write-once abstraction requires overhead or centralization:
    - e.g., DHT-based storage w/replication

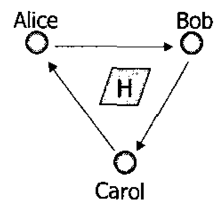

11

Dealing with Strangers

- Availability of cheap pseudonyms gives rise to whitewashing attack
  - Use new identity for each transaction
  - Always defect
- Whitewashers indistinguishable from legitimate newcomers
- Tit-for-tat strategy fully exploited by whitewashers

12

Stranger Policies

- Always cooperate
  - allows exploitation by whitewasher
  - Randomly (P = 50%) cooperate doesn't help
- Always defect
  - forces newcomers to allow exploitation by existing players
  - "social cost of cheap pseudonyms"
- Adaptively cooperate
  - separately estimate stranger friendliness

$$\min\left(\frac{cooperation\_strangers\_have\_given}{cooperation\_strangers\_have\_received}, 1\right)$$

  - only taxes newcomers when necessary

13

P2P Economics

- Some pointers:
  - Workshops
    - p2pecon'03: http://www.sims.berkeley.edu/p2pecon/
    - p2pecon'04: http://www.eecs.harvard.edu/p2pecon/
    - SIGCOMM PINS'04 http://pins.csail.mit.edu/
  - Project: p2pecon@berkeley
    - http://p2pecon.berkeley.edu/

14