

Practice and Theory of Incentives in Networked Systems (PINS): Workshop Report

Rahul Sami Dina Katabi Peyman Faratin John Wroclawski
sami@csail.mit.edu dk@mit.edu peyman@mit.edu jtw@csail.mit.edu

1 Introduction

The development of large, open networked systems without a single controlling organization has led to much interest in using tools from Game Theory and Mechanism Design to tackle incentive issues in such systems. Recently, this subject has inspired research in many different fields, including AI, Economics, E-commerce, CS Theory, and Networking. However, there appears to be a wide gulf between the theoretical research and practical realities of networked systems, as the research has had very little impact on mainstream networked-system design. The SIGCOMM PINS workshop, held in collaboration with SIGecom, was envisaged as a forum to initiate direct dialogue between game theorists and networking researchers.

With this in mind, the workshop had a strong focus on *practical* incentives, and the scope of the workshop included the question of whether incentive-based techniques are necessary (or useful) in the first place. The goal of the workshop was to promote an exchange of ideas on the true applicability, range, and validity of game-theoretic and economic models for analysis and design of the Internet and Internet-based systems.

2 Workshop Format

The workshop started with a tutorial by Ramesh Johari (Stanford University) outlining basic game theory and Mechanism Design, as well as potential applications to networking problems. This served as a common foundation for all workshop participants, many of whom were not familiar with this field. The tutorial was followed by nine submitted paper presentations. The workshop concluded with a panel discussion on the role of incentives in networks.

3 Tutorial

The tutorial used examples based on interdomain routing and peer-to-peer systems to introduce many game theoretic concepts and results: static and dynamic games; Nash equilibrium and Dominant Strategy Equilibrium; Mechanism Design; the revelation principle; impossibility results; and the Vickrey-Clarke-Groves family of mechanisms. The tutorial was lively, with the audience posing many questions for the speaker. It helped the audience understand the remainder of the workshop better – many of the topics the tutorial touched upon were raised again in later talks.

4 Paper Presentations

Session 1: Experimental Research

Daniel Friedman presented research from an experimental economics perspective [1]. A laboratory game of downloading in a congested network was set up, and human subjects were told that they would be paid an amount dependent on their downloading success. In some cases, the subjects were pitted against bots following simple strategies. The subjects' behavior was then compared to theoretical predictions based on the "social optimum" and the Nash equilibrium of the downloading game. Friedman reported that the actual behavior was significantly different from the theoretical predictions – the agents downloaded at a rate that was much higher than that explained of "rational utility maximizers."

The paper modeled the network as an M/M/1 queue. Members of the audience with a networking background pointed out that an M/M/1 queue is a poor model of actual networks but it is not clear how to accurately model the network in such studies.

Session 2: Incentives in Practice

This session had three presentations addressing the question of whether incentives are useful in practical networked systems.

Ratul Mahajan talked about his group's experience with attempting to use game theory in networks [2]. He talked about two problems that appear to be natural candidates for a solution based on game theory: improving the performance of inter-domain routing via route negotiation and discouraging cheating in wireless multi-hop networks. In both cases, they examined game-theoretic solutions, but found them to be impractical. Specifically, in the ISP route negotiation case, they found that existing proposals based on Mechanism Design require ISPs to reveal too much information. Further, they are based on optimizing a single objective function, whereas ISPs want more flexibility in their optimization criteria. In the wireless multi-hop network problem, they found that incentive schemes to guarantee fairness, such as barter or currency methods, can actually reduce cooperation in a heterogeneous network, because users who serve more than they need would quickly lose incentive to continue serving traffic. Further, the majority of users do not seek to manipulate the scheme, and so the researchers decided instead to use a low-overhead scheme that tries to catch cheaters, but is not provably strategyproof. The audience asked if they had considered using other game-theoretic models for the ISP negotiation case, in particular, cooperative game theory; Ratul answered that the information revelation was still a problem, as cooperative game models required ISPs to know each others' values.

Next, Jon Crowcroft also argued that incentives were unnecessary in mobile ad hoc networks [3]. His main point was that the "costly" resource in this case, *viz.* battery power, is of very low value; thus, it does not provide an incentive for large-scale deviation from the cooperative protocol. He also suggested that the best approach is to catch cheating nodes by having nearby nodes monitor their transmissions. He further said that new technologies are often first adopted by pioneers who want the technology to succeed and visionaries who can see the long-term benefit of cooperation. Selfish users are likely to join only after the technology is well established, and hence putting in incentive mechanisms from the beginning will unnecessarily slow the adoption of the technology. Thus, he argued for putting in the incentive mechanism after the technology is established. This raised an interesting issue— an audience member asked if it would be possible to make changes to a system once it had become widely deployed and backward compatibility became an issue. Jon replied that we could put in hooks for an incentive system from

the start, and introduce the actual incentive system later.

In the last talk of the session, Mike Afegan argued *for* introducing explicit pricing information – into the Border Gateway Protocol (BGP) [4]. He argued that incentives are already important in determining the routing structure of the Internet – they determine the commercial relationships that ISPs form with each other. An ISP's choice of route is largely driven by the price it must pay to other ISPs for carrying traffic, not the quality that these routes deliver to its users. This had led users to get around the ISP choices with overlay networks, for example, making the entire system unnecessarily complex. Mike argued that the pricing information should be coupled with the routing information and be able to change on the same timescale, so that users can make their own quality-cost tradeoff; hence, it would be best to include it in the routing protocol. The audience raised several questions in the following discussion. When asked if pricing would reveal ISPs' private information, Mike responded that the information would not be globally revealed, but only shared in a pairwise manner. This much information revelation already occurs (at a different timescale) during ISP peering negotiations. Then, there was a question raised as to whether it is better to separate the routing and negotiation into two protocols, so as to keep the routing protocol simple. Mike replied that as he was advocating exchanging pricing information at the same granularity and timescale as routing information, it seemed like a good idea to put them in the same protocol.

Session 3: Work-in-Progress Talks

Yang-hua Chu considered the problem of incentivizing peers to contribute in a peer-to-peer streaming system [5]. The valuable resource here is bandwidth; some peers are on Ethernet links and are relatively rich, whereas peers behind DSL or cable modem links are relatively poor. Yang-hua makes the case for using a taxation scheme to decide how much bandwidth each peer should contribute, instead of the usual bit-for-bit scheme in which peers must contribute as much as they receive. With a bit-for-bit scheme, DSL or cable modem peers are severely limited by their upstream bandwidth, which is often much lower than their receiving bandwidth. With a taxation scheme, resource-rich peers subsidize resource-poor peers, and in doing so improve the overall social welfare.

Nicholas Christin talked about the choice of solution concept to model users in networked systems [6]. He gave three examples: network topology formation, protection against security threats, and TCP congestion control. In each example, he showed that the Nash equilibrium so-

lution concept predicts an unrealistic outcome, and cannot explain the observed behavior. The perfectly optimizing behavior assumed in the Nash equilibrium seems too demanding; instead, we need a model of “near-rational” behavior. He demonstrated that the observed behavior matched the set of ϵ -Nash equilibria for each example. The audience asked about the loss of predictive power in moving from Nash equilibria to ϵ -Nash equilibria: Every Nash equilibrium is also an ϵ -Nash equilibrium, and hence the original unrealistic predictions are not ruled out. The speaker agreed, but suggested that by making other assumptions about reasonable behavior, such as assuming that the transmission rate must increase at a bounded rate in the TCP example, the unrealistic equilibria might be ruled out.

Jeff Sheidman presented a model for faithfulness in Internet algorithms [7]. His departure point was the observation that Internet protocols allow players a wider range of strategic possibilities than traditional economic models assume: Not only can they lie about their private information or utility, players can also manipulate the computation and communication that the protocol expects them to follow. He argued that protocols should be checked for incentive-compatibility, to determine whether users will benefit from hacking and changing them. As an example, he outlined a manual backtracing procedure for the BitTorrent file-sharing protocol, and reported that several points for successful manipulations were discovered. The discussion centered on his program analysis procedure – Jeff clarified that his manual procedure was sound though not complete, *i.e.*, it did not always report all successful manipulations, but any reported manipulations are correct.

Session 4: Theory and Models

Ion Stoica talked about modelling free-riding and white-washing in peer-to-peer systems [8]. Observing that user behavior often exhibits a mixture of altruism and strategizing, he presented a simple model of users parametrized by their “generosity” level. Simulations based on this model were then used to derive insight about the system performance, measure in terms of overall social welfare; these simulations were also used to evaluate potential incentive mechanisms. The results show that when the average generosity levels are high, incentive mechanisms give little gain. When the generosity level is low, a mechanism for punishing free-riders can produce significant improvements; in this case, however, a lack of permanent identities can result in a significant social cost. The audience asked whether it is possible to incorporate into the model

the cost incurred by generous users to punish a free-rider. Ion replied that it was possible, but would result in a more complicated model, and it is much harder to get insight from a complex model.

Hervé Moulin gave the last presentation of the day. He spoke about the possibility of extending strategyproof Mechanism Design to allow for the merging and splitting strategies that are often possible in the Internet [9]. Hervé’s canonical description was in terms of scheduling jobs at a single server, in which users may merge multiple jobs into a single job or split a single job into smaller jobs; the results can also be applied to other shared resources such as bandwidth. The common strategies of longest-job-first or shortest-job-first are easily attacked in such a setting. Hervé mentioned that a completely random job ordering would be merge-proof and transfer-proof, but is unfair: Small job owners may have to wait arbitrarily long to have their requests satisfied. He then presented a randomized mechanism that is merge-proof (but not split-proof) and meets several natural fairness properties, and showed that it is nearly optimal with respect to these properties. The audience asked if it could be generalized to multiple resources – many identical servers, or even heterogeneous resources. Hervé replied that even moving to two servers made the problem much harder, so it was unclear what could be proven for these models.

5 Panel Discussion

The panelists were David Clark (MIT), Joan Feigenbaum (Yale University), John Ledyard (Caltech), and David Wetherall (U. Washington). The panelists have different backgrounds. Clark and Wetherall are network researchers, Ledyard is an economist, and Feigenbaum is a CS theorist. The panel was loosely structured as a series of questions, to which each panelist responded; however, the discussion was allowed to flow on its own as the audience also got actively involved. Here, we give a brief description of some of the topics discussed.

The first topic was on control as an alternative to incentives – *i.e.*, the authority might use control mechanisms to limit user choices and force the user to adopt a strategy that improves the social welfare. For example, as an alternative to incentive-compatible TCP, it is conceivable that ISPs could put a box in front of every user that limits the user sending rate. Dave C. and John pointed out that there is not a black-and-white distinction between control and incentives; rather, the ability of the controller to punish users can create strong disincentives. Joan suggested that control meant something that participants cannot feasibly violate, *i.e.*, a real restriction of the strategy space. Her

view was that if techniques such as cryptography could reduce the strategy space, they could be tried before, or alongside, incentive mechanisms. David W. raised the example of Distributed Denial-of-Service (DDoS) prevention as a problem in which control was *necessary*, and incentives could never work. Dave C. raised the point that it is hard to distinguish DDoS traffic from regular flash crowd traffic – the difference is only in the intent behind the traffic, and intent is hard to determine. The audience asked if payments or liabilities could be used as a solution. John responded that payments are tricky, because the liability legal structure is mainly in the US, and the Internet has a very international character.

An audience member asked why incentives must be so tightly coupled with money, when people’s motivations are often not translatable into monetary terms. John responded that it was a common misconception to think that economics was about maximizing wealth. This is the goal only of some of economics, which assumes quasilinear utilities, etc. However, there is a huge literature on experimental economics which suggest that people are often altruistic. He suggested that adding a little “grease” to an economic model can dramatically improve its explanatory power. The art is to identify what additional features are important to consider in any particular problem setting.

The next topic centered on the tradeoff between simplicity and models: Are simple (perhaps inaccurate) game models sufficient to draw useful conclusions, or are more complex models necessary? David W. felt that in practice, there is very little cheating against a huge backdrop of altruism, and hence, it is best to include this in a model rather than start with the assumption that everyone is a selfish maximizer. Joan suggested that strategic behavior is costly, and hence will be used only if it is clear that there is a significant benefit; when that is clear, then people *will* be strategic. Dave C. pointed out that spam is a domain in which strategic behavior was unanticipated, but profitable.

An audience member asked about relaxations of the standard equilibrium notions, such as ϵ -equilibrium: Are they useful models of actual behavior, or are they too simplistic? John responded that ϵ -equilibria are useful for explaining behavior observed in the lab; however, they blow up the number of equilibria, and thus have lower predictive power. Price-taking behavior is not game-theoretic at all, but is a great model of real behavior. He listed other solution concepts, like quantal response equilibrium [10] and ϵ -Nash [6], as concepts that *sort of* work in a lab setting, but said we should stay tuned to experimental economics for further concepts. Joan mentioned new con-

cepts from the computer science side, such as tolerable manipulability, which says that something can be non-strategyproof, but you do not care as long as it improves whatever you are trying to optimize. She also asked John what his favorite equilibrium concept was. He suggested that ex-post Nash [11] could be useful, or Bayes Nash [12] for one-shot games. For network games, he said that a Nash equilibrium is good enough. An audience member asked why, given that Nash is a bad predictor in repeated games. John responded that it depends on the scale of the game - for a two-player game, it is bad, but for a game with ten million players, when you do not know who you are playing against, the Nash equilibrium is pretty good.

6 Conclusion and Experiences

The PINS workshop got a very good turnout – there were over 60 registered attendees. Overall, the audience felt that the workshop was informative and fun, with a lot of active discussion and debate. We were very fortunate to have several economists attend; they were able to give very valuable pointers to the economic literature, and correct some common misconceptions. This is a very exciting area of research – it may take a long process to bring the theory and practice together, but it is useful for each to keep an eye on the other.

References

- [1] Daniel Friedman and Bernardo Huberman. Internet Congestion: A Laboratory Experiment. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 177 – 182, 2004.
- [2] Ratul Mahajan, Maya Rodrig, David Wetherall, and John Zahorjan. Experiences Applying Game Theory to System Design. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 183 – 190, 2004.
- [3] Elgan Huang, Jon Crowcroft, and Ian Wassell. Rethinking Incentives for Mobile Ad Hoc Networks. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 191 – 196, 2004.
- [4] Mike Afergan and John Wroclawski. On the Benefits and Feasibility of Incentive Based Routing Infrastructure. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 197 – 204, 2004.

- [5] Yang-hua Chu, John Chuang, and Hui Zhang. A Case for Taxation in Peer-to-Peer Streaming Broadcast. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 205 – 212, 2004.
- [6] Nicolas Christin, Jens Grossklags, and John Chuang. Near rationality and competitive equilibria in networked systems. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 213 – 219, 2004.
- [7] Jeff Shneidman, David Parkes, and Laurent Massoulié. Faithfulness in Internet Algorithms. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 220 – 227, 2004.
- [8] Michal Feldman, Christos Papadimitriou, Ion Stoica, and John Chuang. Free-Riding and Whitewashing in Peer-to-Peer Systems. In *Proceedings of the SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pp. 228 – 235, 2004.
- [9] Hervé Moulin. On scheduling fees to prevent merging, splitting and transferring of jobs. Paper available at <http://www.ruf.rice.edu/~econ/papers/2004papers/schedfees4.pdf>.
- [10] Richard D. McKelvey, and Thomas R. Palfrey. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior*, 10, pp. 6 – 38, 1995.
- [11] Bengt Holmstrom, and Roger B. Myerson. Efficient and Durable Decision Rules with Incomplete Information. *Econometrica*, pp. 1799 – 1819, 1983.
- [12] J. C. Harsanyi. Games with incomplete information played by Bayesian players. *Management Science*, 14, 1967.