

Strategyproof Mechanisms for Ad Hoc Network Formation

C. Jason Woodard* and David C. Parkes†
Division of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138

April 1, 2003

Abstract

Agents in a peer-to-peer system typically have incentives to influence its network structure, either to reduce their costs or to increase their ability to capture value. The problem is compounded when agents can join and leave the system dynamically. This paper explores the use of mechanism design techniques to offset the incentives for strategic behavior and facilitate the formation of networks with desirable global properties.

1 Introduction

Consider the numerous efforts underway to create community wireless networks in urban areas. Today, many are simply lists of “hot spots” operated by individuals and businesses. Others have more ambitious goals in the spirit of the Rooftops project at MIT:

The overall system should be financially and technologically self-sufficient. It should allow people to join the Internet generally without recourse to wireline carriers, and without substantial work beyond the purchase and installation of the node.¹

These goals imply networks in which most packets receive transit across one or more wireless nodes to reach the Internet, with the cost of the relatively few wireline connections either donated or shared among the users.

Despite the enormous value such networks have the potential to create, they present classic incentive problems that may inhibit their growth. At one extreme, if costs cannot be shared at all, nodes at the edges of the network will be free riders, benefiting disproportionately from others’ connectivity and willingness to provide transit. But if nodes can charge arbitrary prices, those with wireline connections or at network “bottlenecks” will be able to extract rents from their less favorably connected neighbors, which may in turn lead to investment in technically unnecessary links to mitigate the hold-up problem. An ideal mechanism would balance these opposing forces, not only for a fixed network structure but also given agents’ choices about how and when to connect to the system.

Network formation has been studied by a growing body of research in economics and game theory, e.g., papers by Jackson and Wolinsky [5] and Dutta and Mutuswami [1]. This literature focuses on situations in which agents choose which links to form, thereby creating a network whose value—and each agent’s ability to capture it—depends on its structure. A common theme of this work is the tension between stability and

*Also Harvard Business School, Boston, MA 02163. Email: jwoodard@hbs.edu.

†Email: parkes@eecs.harvard.edu.

¹“Rooftops,” <http://rooftops.media.mit.edu>, accessed March 24, 2003.

efficiency: high-value networks may be vulnerable to agents’ incentives to modify their structure, while networks in which all agents are doing as well as they can for themselves may not maximize social welfare. The network formation literature generally assumes an environment of perfect information, and agents are often assumed to interact through decentralized bargaining.

A related body of research has considered networks from a mechanism design perspective. The field of mechanism design explores conditions under which system designers can achieve various economic and computational properties in the presence of strategic agents (see Parkes [8] for a survey). Many of these results can be applied directly to systems that involve networks, for example, mechanisms that allocate costs in a multicast tree [2], and mechanisms that address an optimal routing problem [3]. This work typically assumes that the set of nodes in the system is *fixed*, a simplification that contrasts with the more dynamic approach of the network formation literature. On the other hand, agents delegate the choice of network structure to a central mechanism and it becomes critical to ensure that their incentives support truthful revelation of preference information. This focus on private information poses significant additional challenges.

One notable exception to the focus on static systems within mechanism design is the recent work of Friedman and Parkes [4], which introduces the problem of *online* mechanism design for dynamic systems with agents that arrive and leave over time. The model that we introduce in this paper can be viewed as a relaxation of this earlier work, as we describe below.

We believe that bringing these streams of research together may shed light on issues of particular relevance to the peer-to-peer systems community, as suggested by the wireless networking scenario above and elaborated in the remainder of the paper.

2 Problem Definition

Let $N = \{1, 2, \dots, n\}$ be a set of agents. For any $S \subset N$, let G_S be the set of all possible networks involving only links between members of S . Let each agent have private information summarized by a type, $\theta_i \in \Theta_i$, that determines its preferences over networks. These preferences are represented by value functions $v_i : G_N \times \Theta_i \rightarrow \mathbb{R}^+$, with an empty network normalized to a value of 0, where G_N is the set of all possible networks.² The mechanism design approach uses payments to implement networks with good properties, even though the types are private information. As is standard, we assume agents have quasilinear utility functions $u_i(g, p_i, \theta_i) = v_i(g, \theta_i) - p_i$, where $g \in G_N$ is a network and $p_i \in \mathbb{R}$ is a monetary amount. The outcome from a mechanism is a pair (g, p) , where $p = (p_1, \dots, p_n)$ defines the payment by each agent.

The cornerstone of traditional mechanism design is the class of *direct* mechanisms, in which each agent announces a type, $\hat{\theta}_i$, then the mechanism responds with an outcome. The mechanism is assumed to be capable of enforcing the outcome, perhaps through some prior contract with the agents. For this reason, mechanisms that violate individual rationality, i.e., that may make an agent worse off by participating, are viewed with suspicion.³ In a *strategyproof* direct mechanism, it is a dominant strategy for each agent to report its true type. In other words, whatever any other agent reports and whatever the types of other agents, an agent can be truthful. This provides useful robustness to a system, and simplicity for participants.

If we hold the set of agents fixed, we can apply the standard repertoire of direct mechanisms, including the celebrated Vickrey-Clarke-Groves (VCG) mechanism, to the network formation problem. The VCG mechanism is both strategyproof and efficient, meaning in this context that it always chooses a network that maximizes the total value of the system to the agents. Unfortunately, although it is individually rational and

²It is convenient to assume a finite set of agents. This can be easily relaxed, by restating an agent’s valuation in terms of the structure of the network, with either indifference to agent identity or the ability to capture classes of agent identities.

³In the context of a peer-to-peer system, this may nevertheless require some method to verify that peers are choosing to perform as instructed.

revenue maximizing among efficient mechanisms [6], the VCG mechanism is not budget balanced; money may need to be pumped into it to retain strategyproofness. This tradeoff is pervasive in mechanism design. Alternative approaches, e.g., Moulin and Shenker [7], take strategyproofness and budget balance as central criteria, and seek mechanisms that minimize inefficiency.

If we allow agents to join and leave the system dynamically, we must look beyond classical direct mechanisms, and towards *online* mechanisms [4]. We need a mechanism that takes a sequence of decisions, possibly one each time the state of the system changes. This is akin to the online algorithm problem, where choices must be made on-the-fly without the benefit of information about the agents yet to arrive.

Moreover, in our context these online decisions must be *mutually consistent* to achieve useful incentive properties. This is clear enough in a resource allocation context: the mechanism cannot give to one agent a resource it promised another in a previous period. The notion of consistency is trickier in a network formation context. First, the mechanism can only form networks that are feasible with the agents that are present in the system at a given time; if an agent leaves, the existing network becomes infeasible by definition. This may be a mere technicality, if the departure does not affect the network’s value to the remaining agents—or it could be a catastrophe. (Consider the wireless networking scenario again, and imagine there is only one agent with a wireline Internet connection!) Second, even if all agents remain in the system after they arrive, adding a new agent may impose externalities on the others, such as additional network congestion. These externalities may render the current network no longer value maximizing, necessitating a change in network structure to maintain efficiency. The payments required to maintain incentive compatibility might also change.

The online mechanism design situation is complicated by the possibility that agents may want to strategically manipulate their arrival or departure times. In the context of VCG-based mechanisms, an online mechanism can achieve strategyproofness in this broader sense if the online choice rule is *perfectly competitive* with an optimal offline choice with complete information about all future arrivals [4]. When this is possible, one can simply update the state of the network every time an agent arrives or leaves, with payments computed to reward each agent with the marginal utility it contributes to the system. But this solution might be computationally and informationally infeasible—not to mention financially expensive.⁴ In this paper, we will sidestep the issue by assuming agents’ arrival and departure times are exogenous to the mechanism and truthfully reported (or observed) and focus on the mutual consistency issue identified above. Even under these more restrictive assumptions, the problem remains challenging.

3 The Online Mechanism

Consider agents that arrive sequentially, letting t_i denote the arrival time of agent i . Assume an ordering over agents, such that $i < j$ if and only if $t_i < t_j$. Suppose that as soon as agent i arrives into the system and announces its type, $\hat{\theta}_i$, the mechanism must choose a network g_i^* and collect a payment p_i . We propose a choice rule similar to that of the Groves family of mechanisms:

$$g_i^*(\hat{\theta}_i) = \arg \max_{g \in \Phi_i} \left[v_i(g, \hat{\theta}_i) + w_i(g, \hat{\theta}_1, \dots, \hat{\theta}_{i-1}) \right], \quad (1)$$

where $w_i \geq 0$ is an arbitrary function that can depend on the network selected and the previously reported types, and Φ_i defines a subset of the feasible networks $G_{N(t_i)}$, given agents $N(t_i)$ in the system at time t_i . In a setting in which the designer’s goal is to maximize the total value of the system, the function w_i can be selected to be an estimate of the value of the system for future agents, taking into account the commitments implied by the mechanism’s choice at time t_i .

⁴In the context of strategic arrivals, we can instead consider mechanisms with *expected* optimal online choice rules, which provides Bayesian Nash incentive compatibility instead of strategyproofness [4].

In order to guarantee the strategyproofness of the mechanism, such that agents choose to announce types truthfully, we require that the value of the network to an agent stay constant as long as it remains in the network.⁵ The function Φ_i expresses these consistency constraints:

$$\Phi_i = \left\{ g \in G_{N(t_i)} : v_j(g, \hat{\theta}_j) = v_j(g_j^*, \hat{\theta}_j) \text{ for all } j < i, j \in N(t_i) \right\}.$$

With this, the payment by agent i is computed at time t_i with a method similar to the Groves mechanisms, with w_i taking the place of the usual linear summation term:

$$p_i(\hat{\theta}_i) = h_i(\hat{\theta}_1, \dots, \hat{\theta}_{i-1}) - w_i(g_i^*(\hat{\theta}_i), \hat{\theta}_1, \dots, \hat{\theta}_{i-1}). \quad (2)$$

For the sequential mechanism defined by choice rule (1) and payment rule (2) to be strategyproof, it is sufficient that the consistency constraints can be satisfied. With this, and by a standard Groves argument, the utility to agent i for an announced type $\hat{\theta}_i$ is

$$\begin{aligned} u_i(\hat{\theta}_i) &= v_i(g_i^*(\hat{\theta}_i), \theta_i) - p_i(\hat{\theta}_i) \\ &= v_i(g_i^*(\hat{\theta}_i), \theta_i) + w_i(g_i^*(\hat{\theta}_i), \hat{\theta}_1, \dots, \hat{\theta}_{i-1}) - h_i(\hat{\theta}_1, \dots, \hat{\theta}_{i-1}). \end{aligned} \quad (3)$$

By inspection, since i can only influence its utility through the effect of its announced type on g_i^* , the agent should announce its true type to make the choice rule (1) explicitly maximize the first two terms of (3).

In comparison with standard Groves mechanisms, it is interesting that the choice rule in (1) is more relaxed because w_i need not be the total value to all other agents. But, we retain a greedy form of this utilitarian requirement in Groves via the consistency constraints, which capture the intuition that future decisions must be optimal for the new agent *subject to* the requirement that they in no way hurt any previous agents still in the system.

Compared with Friedman and Parkes [4], we have relaxed the assumption of strategic arrivals, and moved away from implementing VCG payments. We find this interesting, because we can achieve strategyproofness with respect to announced types without placing stringent requirements, such as perfect-competitiveness with an offline decision, on the online choice rule of the mechanism. Instead, of central importance in our mechanism is that the consistency requirements are satisfied by sequential choices.

4 Ongoing Work

In closing, we discuss what can go wrong when consistency requirements cannot be satisfied, and then propose some approaches to address this issue.

First, what if the consistency constraints cannot be satisfied? For one thing, individual rationality may be violated. As suggested by the example of a wireless network in which only one agent has a wireline connection, if there are agents whose presence is essential for the system to have value, and those agents leave—or never arrive—then agents that have made positive payments may be left with negative utility overall. One way to avoid this, of course, would be to ensure that such agents do arrive and prevent them from leaving. If this is impractical, it may still be possible to guarantee *interim* individual rationality (i.e., positive expected utility for every agent type) if we have information about the distribution of arrivals and departures. But strategyproofness may be violated even if individual rationality is not, simply by breaking the requirement that $v_i(g, \hat{\theta}_i)$ remain constant throughout agent i 's time in the system.

Several directions seem particularly promising for further exploration, to address the consistency issue but retain incentive properties:

⁵This can be relaxed if the departure time of an agent is known, with an agent instead receiving a guaranteed *total* value over its time in the system.

restrictions Place restrictions on the set of feasible networks to ensure sequential consistency.

oblivious scaling Dynamically scale back the values of agents dynamically, when faced with consistency problems, with a method that is oblivious to the announced type of an agent.

friendly agents Inject a number of friendly agents into a system, that can step in and ensure the feasibility of consistency requirements when necessary.

In addition, there may be restrictions on agent types that lead out positive results. Finally, if the class of strategyproof mechanisms proves too restrictive, we will naturally explore weaker implementation concepts.

To further understand the implications of the consistency requirements, we need to investigate the kinds of constraints that arise for various classes of value functions. Simple value functions (e.g., c if connected, 0 otherwise) should yield simpler constraints than complex ones, but they are also less expressive. We would like to discover more complex value functions that nonetheless yield simple and easily satisfied constraints. We also recognize that the computational issues are no less important than the economic ones—an ideal mechanism would be distributed, requiring only local information (e.g. [9])—and that theory will only get us so far: the inevitable design tradeoffs will need to be assessed in experiments and working systems.

5 Conclusions

There is a fascinating and important research program in developing economically motivated computational methods to facilitate the formation of networks with desirable properties in peer-to-peer systems. We have provided a formal model in which to study mechanism design in a dynamic setting with agents that can arrive and leave dynamically. We identified the requirement that online decisions are *mutually consistent*. This is of paramount importance in extending traditional techniques from mechanism design to the context of ad hoc network formation.

References

- [1] Dutta, Bhaskar, and Suresh Mutuswami (1997). “Stable Networks.” *Journal of Economic Theory* 76, 322–344.
- [2] Feigenbaum, Joan, Christos Papadimitriou, and Scott Shenker (2001). “Sharing the Cost of Multicast Transmissions.” In *Journal of Computer and System Sciences* 63, 21–41.
- [3] Feigenbaum, Joan, Christos Papadimitriou, Rahul Sami, and Scott Shenker (2002). “A BGP-based Mechanism for Lowest-Cost Routing.” In *Proc. 21st Symposium on Principles of Distributed Computing*, ACM Press, 173–182.
- [4] Friedman, Eric J., and David C. Parkes (2003). “Pricing WiFi at Starbucks – Issues in Online Mechanism Design.” Poster paper, to appear in *Proc. 4th ACM Conf. on Electronic Commerce (EC’03)*.
- [5] Jackson, Matthew O., and Asher Wolinsky (1996). “A Strategic Model of Social and Economic Networks.” *Journal of Economic Theory* 71, 44–74.
- [6] Krishna, Vijay, and Motty Perry (2000). “Efficient Mechanism Design.” Working paper.
- [7] Moulin, Hervé, and Scott Shenker (2001). “Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency.” *Economic Theory* 18, 511–533.
- [8] Parkes, David C. (2001). *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*, Ph.D. dissertation, University of Pennsylvania, Chapter 2.
- [9] Shneidman, Jeffrey, and David C. Parkes (2003). “Rationality and self-interest in peer to peer networks.” To appear in *Proc. 2nd Int. Workshop on Peer-to-Peer Systems (IPTPS’03)*.