Strategyproof Mechanisms for Ad Hoc Network Formation

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

May 21, 2003

Abstract

Agents in a peer-to-peer system typically have incentives to influence its network structure, either to reduce their costs or increase their ability to capture value. The problem is compounded when agents can join and leave the system dynamically. This paper proposes three economic mechanisms that 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]

A system that achieves these goals will need to induce the formation of 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.

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

[1] "Rooftops," http://rooftops.media.mit.edu, accessed May 19, 2003.

Network formation has been studied by a growing body of research in economics and game theory, e.g., papers by Jackson and Wolinsky [6] 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 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, e.g., mechanisms that allocate costs in a multicast tree [2], and mechanisms that address an optimal routing problem [3]. In this work, agents typically delegate the choice of network structure to a central mechanism, making it crucial to ensure that they have incentives to reveal their preferences truthfully. This focus on private information poses significant additional challenges. On the other hand, most of this work 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. One notable exception is the recent work of Friedman and Parkes [5], which introduces the problem of *online* mechanism design for 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. The goal of this paper is to state the network formation problem in the language of mechanism design, and begin to identify its special features. Section 2 does this for a single time period, and Section 3 extends the single-period model to a repeated setting, under strong assumptions. Section 4 relaxes these assumptions and considers the problem in a fully online setting, where an additional requirement arises—that the mechanism's decisions be *consistent* across time, in a sense to be explained below. Section 5 discusses the possibility that this requirement may not be satisfiable, and suggests ways to overcome it, which we are exploring in ongoing experiments.

Before proceeding, we note that the mechanisms proposed here are centralized. That is, we assume that agents express their preferences to a trusted central authority, which then takes a decision that applies globally to the system. An ideal mechanism would be distributed (see [4]), requiring only local information and admitting a more flexible trust model. These issues are explored further by Shneidman and Parkes [10], and we hope to address them in future work.

2 The Static Problem

Let $N = \{1, 2, \ldots, n\}$ 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 $v_i : G_N \times \Theta_i \to \mathbb{R}^+$, where G_N is the set of all possible networks involving any of the agents, and an empty network is normalized to a value of 0.[2] In the mechanism design approach, agents are assumed to maximize a utility function that depends on the decision taken by the mechanism, also called an *outcome*. We make the standard assumption that agents have quasilinear utility functions, i.e., $u_i(g, p_i, \theta_i) = v_i(g, \theta_i) - p_i$, where $g \in G_N$ is a network and $p_i \in \mathbb{R}$ is a monetary amount, with the sign convention that $p_i > 0$ is a payment from the agent to the mechanism and vice versa. This allows the mechanism to compare utility across agents, and balance it by transferring money among them. Thus, we define the outcome of a mechanism as a pair (g, p), where g is a network, as above, and $p = (p_1, \ldots, p_n)$ is a vector of payments.

The cornerstone of traditional mechanism design is the class of *direct* mechanisms, in which each agent announces a type, $\widehat{\theta}_i$ (perhaps untruthfully), 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 than not, 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 obtain no higher utility than by being truthful itself. This property provides a useful robustness to the system, and simplifies the decision-making process 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.[4] The choice and payment rules for the VCG mechanism are, respectively:

$$g^*(\widehat{\theta}_i) = \arg\max_{g \in G_N} \sum_{j \in N} v_j(g, \widehat{\theta}_j), \tag{1}$$

$$p_i(\widehat{\theta}_i) = \sum_{j \in N \setminus i} v_j(g^*_{-i}, \widehat{\theta}_j) - \sum_{j \in N \setminus i} v_j(g^*, \widehat{\theta}_j) \tag{2}$$

where $g^*_{-i} = \arg\max_{g \in G_{N \setminus i}} \sum_{j \in N \setminus i} v_j(g, \widehat{\theta}_j)$, i.e., the highest-valued network that could be formed among agents excluding i. The mechanism defined by (1) and (2) is efficient by inspection, and strategyproof by the standard Groves argument: since agent i can only influence its utility through the effect of its announcement of its type on the chosen network, it should announce truthfully so as to cause the mechanism to choose the network that is truly most valuable for that agent.

We note that although the static problem is economically straightforward—at least, it maps cleanly into the standard mechanism design framework—it may be computationally complex. As an indication, there are $2^{n(n-1)/2}$ possible networks with undirected links among n nodes, each of which may in principle be valued differently by each agent. Moreover, to compute the VCG payments, the choice rule

[2] It 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.

[3] Nevertheless, in the context of a peer-to-peer system, the distributed interaction among agents may necessitate some means to verify that peers are choosing to perform as instructed.

[4] Unfortunately, although it is individually rational and revenue maximizing among efficient mechanisms [7], the VCG mechanism is not budget balanced in general; 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 [8], take strategyproofness and budget balance as central criteria, and seek mechanisms that minimize inefficiency. In particular, a lack of budget balance can exist in the present setting because it is akin to a double auction, with agents playing the roles of both buyers and sellers.

must be applied again n times, once with each agent removed. In practice, however, agents will have (or be restricted to express) much simpler preferences, and the optimization problem will be solvable using standard techniques.

3 The Dynamic Problem: Repeated VCG

If we allow agents to join and leave the system dynamically, we need a mechanism that takes a sequence of decisions, possibly one each time the state of the system changes. A natural way to construct such a thing would be to simply rerun the static VCG mechanism whenever an agent joins or leaves the system. It is easy to write down the choice and payment rules for this *repeated* VCG mechanism:

$$g^*(\widehat{\theta}_i) = \arg\max_{g \in G_{N(t_i)}} \sum_{j \in N(t_i)} v_j(g, \widehat{\theta}_j), \tag{3}$$

$$p_i(\widehat{\theta}_i) = \sum_{j \in N(t_i) \setminus i} v_j(g^*_{-i}, \widehat{\theta}_j) - \sum_{j \in N(t_i) \setminus i} v_j(g^*, \widehat{\theta}_j) \tag{4}$$

where $g^*_{-i} = \arg\max_{g \in G_{N(t_i)\setminus i}} \sum_{j \in N(t_i)\setminus i} v_j(g, \widehat{\theta}_j)$. This is identical to the static VCG mechanism except for the replacement of N by $N(t_i)$.

While this mechanism is attractive in its simplicity, it has at least two drawbacks. First, strong assumptions are required to retain the incentive properties of the static version. Specifically, the repeated VCG mechanism remains strategyproof if we assume that agents' preferences are time independent (i.e., the valuation functions v_i are not affected by future events, or an agent's expectations or beliefs about them) and there are no costs for changing the configuration of the network.[5] Second, the state of the system must change slowly enough that the mechanism can keep up with the computation and implement the chosen outcome before needing to do it again. To the extent that these assumptions are problematic, we must look beyond classical direct mechanisms, and toward *online* mechanisms.

4 The Dynamic Problem: Sequential VCG

In an online mechanism, a decision is taken for each agent only once, when it announces its arrival. That is, at time t_i the mechanism chooses a network structure and a payment in response to agent i's announced type. This is akin to the online algorithm problem in that choices are made "on the fly," without the benefit of information about agents yet to arrive.

Often in these kinds of problems, the system faces constraints that cause its decisions to be interdependent across time. For example, in a resource allocation context, a resource that is given to an agent in one period may not subsequently be given to another (perhaps until it is returned). The situation is more subtle in a network formation context. We continue to assume, for simplicity, that in each period the mechanism can form any feasible network $g \in G_{N(t_i)}$, regardless of the system's current state. But each agent that remains in the system receives utility *in each period*, which depends on the network formed in that period. Therefore, even though the mechanism must determine a payment for each agent as it arrives, it must contend with the fact that the agents' total utility will generally depend

[5] Slightly weaker assumptions may suffice.

on the networks formed in subsequent periods as well.[6] In this section, we identify a set of *consistency constraints* that address this issue for VCG-based mechanisms, ensuring that they remain strategyproof in an online setting.

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 [5]. 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.[7] 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.

4.1 The Sequential Mechanism

Consider agents that arrive sequentially, letting t_i denote the arrival time of agent i. Suppose that as soon as agent i arrives into the system and announces its type, $\widehat{\theta}_i$, the mechanism must choose a network configuration g_i^* and collect a payment p_i. The network configuration is subject to change in subsequent periods, but mechanism commits to collect the same payment from the agent in each period it remains in the system.

We propose a choice rule similar to that of the Groves family of mechanisms:

$$g_i^*(\widehat{\theta}_i) = \arg\max_{g \in \Phi_i} \left[v_i(g, \widehat{\theta}_i) + w_i(g, \widehat{\theta}_1, \ldots, \widehat{\theta}_{i-1}) \right], \tag{5}$$

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.

In order to guarantee the strategyproofness of the mechanism, so that rational agents choose to announce their types truthfully, we require that the value of the network to an agent stay constant as long as it remains in the network.[8] If this were not the case, an agent might benefit by misstating its type, hoping to be assigned a low payment in the period it arrives, while expecting to benefit from higher-valued network configurations in future periods. This is true even if we assume truthful arrival and departure. The function Φ_i expresses these consistency constraints:

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

[6] We assume in our current work that agents' utility is linearly additive across time. Some relaxation is possible here, also.

[7] In 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 [5]. These rules may require correct assumptions about the distribution of agent types to be effective, but they always exist in principle.

[8] This 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.

With this, the payment by agent i is computed at time t_i using the general formula of the Groves mechanisms, with w_i taking the place of the usual linear summation term:

$$p_i(\widehat{\theta}_i) = h_i(\widehat{\theta}_1, \ldots, \widehat{\theta}_{i-1}) - w_i(g_i^*, \widehat{\theta}_1, \ldots, \widehat{\theta}_{i-1}). \tag{6}$$

For the sequential mechanism defined by choice rule (5) and payment rule (6) 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 $\widehat{\theta}_i$ is

$$\begin{aligned} u_i(\widehat{\theta}_i) &= v_i(g_i^*, \theta_i) - p_i(\widehat{\theta}_i) \\ &= v_i(g_i^*, \theta_i) + w_i(g_i^*, \widehat{\theta}_1, \ldots, \widehat{\theta}_{i-1}) - h_i(\widehat{\theta}_1, \ldots, \widehat{\theta}_{i-1}). \end{aligned} \tag{7}$$

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 (5) explicitly maximize the first two terms of (7).

In comparison with standard Groves mechanisms, it is interesting that the choice rule in (5) is more relaxed, because w_i need not be the total value to all other agents. Intuitively, this is possible because agents arrive sequentially and cannot change the order of their arrival. But we retain a greedy form of this utilitarian principle 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 (or even help!) any previous agents still in the system.

4.2 Vickrey Payments

To facilitate comparison with the repeated VCG mechanism, and because we use it in our computational experiments, let us state a version of the sequential mechanism that sets aside the extra flexibility it affords, and uses the standard VCG payment rule as a particular choice of h_i and w_i:

$$\begin{aligned} g_i^*(\widehat{\theta}_i) &= \arg\max_{g \in \Phi_i} \left[v_i(g, \widehat{\theta}_i) + \sum_{j \in N(t_i) \setminus i} v_j(g, \widehat{\theta}_j) \right] \\ &= \arg\max_{g \in \Phi_i} \sum_{j \in N(t_i)} v_j(g, \widehat{\theta}_j) \end{aligned} \tag{8}$$

where Φ_i is the same as above, and

$$p_i(\widehat{\theta}_i) = \sum_{j \in N(t_i) \setminus i} v_j(g_{-i}^*, \widehat{\theta}_j) - \sum_{j \in N(t_i) \setminus i} v_j(g_i^*, \widehat{\theta}_j), \tag{9}$$

where $g_{-i}^* = \arg\max_{g \in \Phi_{-i}} \sum_{j \in N(t_i) \setminus i} v_j(g, \widehat{\theta}_j)$, and

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

Note that in writing out the payment rule explicitly, we uncover a new subtlety: when computing agent i's marginal impact on the network, we cannot simply "wind the clock back" from time t_i to t_{i-1}, because one or more of the agents present at that time may have left the system by the time i arrives. Instead, we need to compute the sum of the values of the agents present at time t_i, excluding i, for the highest-valued network they could form *at time* t_i, subject to the constraints imposed by those agents. One implication of this is that the consistency constraints may be unsatisfiable in this marginal impact computation, even if they are satisfied in the computation of the network that actually forms, and even if they were satisfied in all previous periods. This issue can be addressed by modifying g^*_{-i}, for example to relax the constraints if they cannot be satisfied. This would not affect the incentive properties of the mechanism, which only require that g^*_{-i} not depend on $\widehat{\theta}_i$.

5 Discussion

We have now stated the network formation problem as a problem of mechanism design, building up from a static direct mechanism to a repeated version of that mechanism, to a fully online mechanism. We identified an important requirement of sequential VCG-based mechanisms, namely that each agent's value and payment remain constant for every period that agent remains in the system. However, our conjecture is that for nontrivial preferences and agent dynamics, there is some probability that the consistency constraints for any given period will be unsatisfiable, at least in a network formation context. We now discuss what can go wrong when consistency requirements cannot be satisfied, and then propose some approaches to address this issue.

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 the value of the network to agent i remain constant throughout its time in the system.

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

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

In addition, there may be restrictions on agent types that lead to 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., 1 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. For example, even if the central mechanism's optimization can be made tractable—perhaps in an approximate version—an ideal mechanism would be distributed, requiring only local information (e.g., [10]). Finally, we recognize that theory will only get us so far; the inevitable design tradeoffs will need to be assessed in experiments and working systems, which we are beginning to develop.

In conclusion, we believe 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 be *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] Feigenbaum, Joan, and Scott Shenker (2002). "Distributed Algorithmic Mechanism Design: Recent Results and Future Directions." In *Proc. 6th Int. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, ACM Press, 1–13.

[5] Friedman, Eric J., and David C. Parkes (2003). "Pricing WiFi at Starbucks – Issues in Online Mechanism Design." Poster paper, forthcoming in *Proc. 4th ACM Conf. on Electronic Commerce (EC'03)*.

[6] Jackson, Matthew O., and Asher Wolinsky (1996). "A Strategic Model of Social and Economic Networks." *Journal of Economic Theory* 71, 44–74.

[7] Krishna, Vijay, and Motty Perry (2000). "Efficient Mechanism Design." Working paper.

[8] Moulin, Hervé, and Scott Shenker (2001). "Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency." *Economic Theory* 18, 511–533.

[9] Parkes, David C. (2001). *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*, Ph.D. dissertation, University of Pennsylvania, Chapter 2.

[10] 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)*.