Pricing WiFi at Starbucks: Issues in Online Mechanism Design

Citation

Friedman, Eric, and David C. Parkes. 2003. Pricing WiFi at Starbucks: Issues in online mechanism design. In EC '03: Proceedings of the 4th ACM Conference on Electronic Commerce: June 9-12, 2003, San Diego, C.A., 240-241. New York: ACM Press.

Published Version

doi:10.1145/779928.779978

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:4101235>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Pricing WiFi at Starbucks – Issues in Online Mechanism Design

Eric J. Friedman*

David C. Parkes†

ABSTRACT

We consider the problem of designing mechanisms for online problems in which agents arrive over time and truthfully announce their arrival. These problems are becoming extremely common in a wide variety of problems involving wireless networking and web serving. We show how the standard results of mechanism design can be modified to apply to this setting, provide conditions under which efficient and incentive compatible mechanisms exist and analyze several important online models including wireless networks and web serving.

Categories and Subject Descriptors

F.2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*.

General Terms

Algorithms, Economics.

1. INTRODUCTION

Mechanism design has a long and successful history in economics and has recently become an active area of research in computer science [7, 1, 2]. However, most research to date has focused on static problems, with the notable exceptions related to repeated auctions [6, 4]. In this paper we focus on online mechanisms in which agents arrive over time.

This is a classic mechanism design problem with a twist since all computations and analysis must be done online. Thus, in addition to trying to infer agents' valuations, we must also try to infer agents' true arrival time, which may not correspond to their announced arrival time. Intuitively, if the current prices in a system are high it might be advantageous for a user to delay submitting their request in the hope that prices might fall.

*School of Operations Research and Industrial Engineering, Cornell University, Ithaca, NY 14853. Email: friedman@orie.cornell.edu. This material is based upon work supported by the National Science Foundation under Grant No. ANI-9730162.

†Division of Engineering and Applied Sciences, 33 Oxford Street, Harvard University, Cambridge, MA 02138. Email: parkes@eecs.harvard.edu. Supported in part by National Science Foundation grant IIS-0238147.

Consider the following motivating scenarios:

WiFi at Starbucks: Consider the pricing of WiFi at a cafe¹ with a limited number of connections. Customers arrive would like to use the WiFi network while enjoying a cup of coffee. What pricing plan will maximize social welfare (or profit)?

Web serving: Consider the pricing of access to a web server where page requests have varied service times. In what order should we serve requests to minimize delay costs?

Ad-hoc Networks: Consider a group of agents who dynamically create their own network. How should we dynamically configure the network to preserve battery life?

In this paper, we will address the incentive issues in online mechanisms, and seek to bring the announcement of truthful arrival times and truthful value information into equilibrium. In particular, we consider online variations of VCG mechanisms, which reduce the mechanism design problem into an algorithmic problem.

2. MODEL AND BASIC RESULTS

We introduce a model to capture the problem of online mechanism design. A mechanism defines a strategy space and a mapping from agent strategies into outcomes. An outcome $\omega = (k, p)$, defines a choice, $k \in \mathcal{K}$, and payments $p = (p_1, p_2, \dots)$ where agent i makes payment p_i to the mechanism. Choice k provides a complete description of the state of the system at all points in time. Agent i has type, $\theta_i = (a_i, v_i)$, which defines both the arrival time of the agent, $a_i \in \mathbb{R}^+$, and a valuation function $w_i(k; \theta_i)$ on choices. We assume *quasilinear* utility functions, with payoff $u_i(\omega; \theta_i) = w_i(k; \theta_i) - p_i$. Let θ to denote the joint type space and θ_{-i} the same space without agent i .

A *social choice function* (SCF), f , defines a mapping from types to outcomes, that we wish to implement within the mechanism. An online mechanism, $\mathcal{M}$, is said to *implement* SCF, f , if outcome $f(\theta)$ is selected by the mechanism in equilibrium, for any joint types θ .

An online *direct-revelation* mechanism (online DRM) is a mechanism in which the strategy space for agent i consists of a single announcement of a reported value, $\hat{v}_i \in V$, at a time, $\hat{a}_i$, chosen by the agent. An *incentive-compatible* online DRM is a mechanism in which the truthful revelation of value, v_i , immediately upon arrival, is an equilibrium of the system. In the special case that truth-revelation is a *dominant strategy*, an online mechanism is said to be *strategyproof*.

A variation of the revelation principle allows us to limit attention to *incentive compatible* and *direct-revelation* online mecha-

¹Starbucks cafe has recently begun offering wireless access in their stores with a variety of pricing plans. (See, <http://www.starbucks.com/retail/wireless.asp>.)

nisms. As with classical mechanism design, this serves to simplify the mechanism design problem by constraining the design space.

THEOREM 1 (ONLINE REVELATION PRINCIPLE). *For both dominant-strategy and Bayesian-Nash equilibrium, if a social choice function can be implemented in an online mechanism, then it can be implemented in an incentive-compatible online direct-revelation mechanism.*

For the remainder of the paper, we consider the problem of making efficient online choices, that maximize the long-term total value of the system. In particular, we consider the problem of implementing a choice rule $k_{\text{eff}}(\theta) \in \mathcal{K}$, that maximizes the *total* value, $V(k; \theta) = \sum_i v_i(k; \theta_i)$.

Our goal is to find online choice rules $k(\theta)$ for which there are there payments that make the overall SCF incentive-compatible.

In order to address incentive issues, we consider the family of Groves [5] mechanisms. In particular, we focus on the Vickrey-Clarke-Groves (VCG) mechanism. Let $\hat{\theta} = (\hat{v}, \hat{a})$ denote the announced types of agents, and $k(\hat{\theta})$ denote the choice implemented with all agents and $k(\hat{\theta}_{-i})$ denote the choice that would be implemented without agent i . Let $V_k(\hat{\theta})$ denote shorthand for $V(k(\hat{\theta}), \hat{\theta})$, and let $V_k(\hat{\theta}_{-i})$ denote shorthand for $V(k(\hat{\theta}_{-i}), \hat{\theta}_{-i})$, the reported value of the choice without i . An online VCG mechanism implements choice $k(\hat{\theta})$ and computes payment, $p_{\text{vcg},i} = v_i(k(\hat{\theta}); \hat{a}_i) - V(k(\hat{\theta}), \hat{\theta}) + V(k(\hat{\theta}_{-i}), \hat{\theta}_{-i})$, to agent i .

It is well known that if the choice rule is offline optimal then the VCG mechanism is strategy proof. But, perfectly-competitive online algorithms only exist in a few very special cases. (We present stylized versions of the motivating scenarios in the full length version of this paper [3].) Indeed, there are information-theoretic lower-bounds for many online choice problems. However, all is not lost and we can find weaker conditions for incentive compatibility due to the physical constraint that agents can't announce their arrival before they actually arrive.

We define an online choice rule, $k(\theta)$, to be *time monotonic*, if the total long-term value of the decision $k(\theta)$ cannot be improved by a unilateral deviation, $\hat{a}_I \geq a_I$ and $\hat{v}_I \neq v_I$, for any agent I , in any state of the system. In particular, an agent must not be able to increase the total value of the system by stalling and reporting some later arrival time.

THEOREM 2. *An online VCG mechanism is strategyproof if and only if the online choice rule is time-monotonic.*

At this point, it is interesting to highlight a fundamental mismatch between the worst-case competitive approach of computer science to online algorithm problems, and what is required here to get an incentive-compatible online mechanism. It is *not* sufficient to use an online choice rule with an *optimal* competitive ratio (i.e. equal to a lower-bound), because this is not consistent with the *expected-utility* maximizing model of a self-interested agent. Rather, what is required is expected-optimal online choice rules. Defining *expected* time-monotonicity in the obvious way, we have the following main result:

THEOREM 3. *The online VCG mechanism is Bayesian-Nash incentive compatible if the online choice rule is expected time-monotonic.*

This is useful, because at least it is clear that expected-optimal online choice rules, that maximize the average-case long-term efficiency of a system, always exist (in contrast to perfectly-competitive online choice rules). Thus, we reduce the online mechanism design problem to the problem of implementing online choice rules that maximize expected long-term value.

The VCG mechanism is Bayesian-Nash, but not dominant strategy, incentive-compatible, because the Bayesian optimality of the online algorithm depends on correct beliefs about the distribution over agent types, which requires that agents report types truthfully.

3. EXAMPLE

In the full length version of this paper [3] we analyze a specific model of the Starbucks problem. Two conclusions that we draw from that analysis are: (1) Even simple dynamic models can be computationally intractable. (2) Simple, but suboptimal, procedures can obtain a high percentage of the maximum efficiency and as such might be of practical importance.

4. CONCLUSIONS

We have introduced the problem of online mechanism design, and demonstrated how to incorporate online problems into the mechanism design framework. We adopted the VCG framework, and showed a mapping between perfectly-competitive online choice rules and expected-optimal online choice rules and Bayesian-Nash incentive compatibility. We also highlighted a basic mismatch between traditional worst-case computer science analysis, and the expected-utility maximizing models of self-interested agents within mechanism design. The most obvious direction for future work is to construct strategyproof but approximate online mechanisms, and online choice rules that are time-monotonic for problems in which there can be no perfectly-competitive rules. We also intend to adopt paradigms from stochastic dynamic optimization to construct Bayesian-Nash incentive compatible online mechanisms.

5. REFERENCES

- [1] A Archer and E Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, 2001.
- [2] Joan Feigenbaum, Christos H Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.
- [3] Eric Friedman and David C. Parkes. Pricing WiFi at Starbucks— Issues in online mechanism design. <http://www.eecs.harvard.edu/parkes/pubs/online.pdf>, 2002.
- [4] Jeremie Gallien. Dynamic mechanism design for online commerce. Technical report, Sloan School of Management, MIT, Cambridge, 2002.
- [5] Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.
- [6] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *Proc. 2nd ACM Conf. on Electronic Commerce (EC-00)*, 2000.
- [7] Noam Nisan and Amir Ronen. Algorithmic mechanism design. *Games and Economic Behavior*, 35:166–196, 2001.