

Online Auctions with Re-usable Goods

Mohammad T. Hajiaghayi*†

Robert D. Kleinberg*‡

Mohammad Mahdian §

David C. Parkes ¶

ABSTRACT

This paper concerns the design of mechanisms for online scheduling in which agents bid for access to a re-usable resource such as processor time or wireless network access. Each agent is assumed to arrive and depart dynamically, and in the basic model require the resource for one unit of time. We seek mechanisms that are truthful in the sense that truthful revelation of arrival, departure and value information is a dominant strategy, and that are online in the sense that they make allocation decisions without knowledge of the future. First, we provide two characterizations for the class of truthful online allocation rules. The characterizations extend beyond the typical single-parameter settings, and formalize the role of restricted misreporting in reversing existing price-based characterizations. Second, we present an online auction for unit-length jobs that achieves total value that is 2-competitive with the maximum offline value. We prove that no truthful deterministic online mechanism can achieve a better competitive ratio. Third, we consider revenue competitiveness and prove that no deterministic truthful online auction has revenue that is constant-competitive with that of the offline Vickrey-Clarke-Groves (VCG) mechanism. We provide a randomized online auction that achieves a competitive ratio of $O(\log h)$, where h is the ratio of maximum value to minimum value among the agents; this mechanism does not require prior knowledge of h . Finally, we generalize our model to settings with multiple re-usable goods and to agents with different job lengths.

Categories and Subject Descriptors

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

1. INTRODUCTION

*CSAIL, MIT, 32 Vassar Street, Cambridge, MA 02139, USA, Emails: {hajiagh, rdk}@csail.mit.edu.

†Part of this work was done while the author was an intern in the Microsoft Research Theory Group.

‡Supported by a Fannie and John Hertz Foundation Fellowship.

§Microsoft Research, One Microsoft Way, Redmond WA 98052, U.S.A., Email: mahdian@microsoft.com.

¶Department of Electrical Engineering and Comp Science, Harvard University, Cambridge, MA 02139, USA. Supported in part by NSF grant IIS-0238147. Email: parkes@eecs.harvard.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.

Copyright 2005 ACM 1-59593-049-3/05/0006 ...\$5.00.

1.1 Motivation

Online mechanism design concerns the design of mechanisms for markets in which agents arrive and depart over time, and the mechanism must compute allocation and payment decisions online without knowledge of the agents who will subsequently arrive. Such problems arise in many practical applications of mechanism design (e.g., pricing access to a WiFi port at Starbucks [7], or scheduling computer jobs on a shared server.) These problems are generally quite difficult because they combine the challenges of mechanism design (i.e., ensuring truthfulness¹) with the challenges of designing online algorithms (i.e., dealing with uncertainty about future inputs). As an example, one of the most important techniques for designing truthful mechanisms (the Vickrey-Clarke-Groves (VCG) scheme) is inapplicable in most online problems because it requires computing an optimal allocation, which is generally impossible in the online setting [7]. In this paper we will analyze truthful online mechanisms in terms of their competitive ratio with the efficiency and revenue of an (off-line) VCG scheme.

The setting we will consider is a simple scheduling problem, in which agents bid for access to a re-usable resource over a sequence of time slots. We will assume that the resource has a finite capacity, most often 1. An agent has an arrival and departure time, and in the basic model a value for receiving one unit of the resource during this interval. Our objective is to design a mechanism for this problem which is strategyproof with respect to not only the values, but also the arrival and departure times. The requirement of strategyproofness with respect to arrival and departure times makes the online auction problem difficult since it places constraints on the timing of allocations. This is demonstrated by Lavi and Nisan [15], who prove that without any restriction on the types of possible misreports, it is impossible to achieve a bounded competitive ratio on efficiency. In this paper, we study the problem with the assumptions of *no early arrivals* and *no late departures*, i.e., we assume that agents cannot report an arrival time earlier than their true arrival time or a departure time later than their true departure time. This model was also adopted by Porter [20] for his work on online auctions with re-usable goods.

The assumptions of no early arrival and no late departure are reasonable in situations where it is possible to verify that agents are physically present throughout their reported arrival-departure interval (i.e., agents can hide their presence but not create a phantom presence). Also, no early arrivals makes sense if one considers that the arrival time is the first time an agent learns about the auction. No late departures makes sense, for example, in markets for a resource such as processor time where it is possible to delay granting an agent the benefit from using the resource (even though it is ready) until departure, for instance by waiting until that time to report the outcome of the processor's computation.²

Finally, it is worth mentioning that we consider the *no early ar-*

¹We use truthfulness and strategyproofness interchangeably in this paper.

²Porter [20] adopted this motivation in his work on online auctions for processor scheduling.

rivals assumption in this paper because it is quite natural and because the characterization of truthful mechanisms for this model is simpler to state. However, almost all of our mechanisms remain strategyproof without assuming no early arrivals.

1.2 Our contributions

In this paper, we give a characterization for the online allocation rules that are truthfully implementable. The characterization is interesting because the online auction problem studied here is neither *single-parameter* [1, 16], nor *order-based* [13], and a complete characterization was previously not available for this model. Furthermore, our result explicates the importance of limited misreports in the design of truthful online auctions. We see this, for instance, in the *no early arrivals* and *no late departures* assumptions. The characterization is stated in terms of monotonicity and timing properties on allocation rules, but does not imply the W-MON condition [13], which is known to be necessary for truthful implementation in settings with unrestricted misreports. We are also able to extend our characterization to address randomized mechanisms. A parallel characterization is provided in terms of price-based auctions, that demonstrates that while fully agent-independent prices are not necessary, prices must be value-independent and satisfy monotonicity requirements with respect to report arrival and departure times. Fully agent-independent prices are sufficient for truthfulness, but only necessary when there are unrestricted misreports.

For unit-length jobs, we use the monotonicity characterization to develop a truthful online auction that is 2-competitive with the optimal offline efficiency. The allocation problem is a special case of the *maximum-weight matching problem*. Without the constraint of truthfulness, the best known upper bounds on competitive ratio are 2 in the deterministic case and $e/(e-1)$ in the randomized case [3, 10], while the best known lower bounds are 1.618 in the deterministic case and 1.25 in the randomized case [6, 8]. With the truthfulness constraint, we demonstrate that no deterministic mechanism can obtain a better competitive ratio than 2, closing the gap. We also extend the result to an asynchronous model in which time is a continuous parameter, agents need not arrive and depart at integer times, and jobs are interruptible, obtaining a 5-competitive mechanism in this case. The auctions also extend to the case of k re-usable goods available in each period, or equivalently a re-usable resource with capacity k .

This simple 2-competitive online auction can have arbitrarily bad revenue with respect to the offline VCG revenue. We prove that no deterministic mechanism is constant-competitive with VCG for revenue by giving a lower-bound of $\Omega\left((\log h / \log \log h)^{1/4}\right)$, where h is the ratio of the maximum value to minimum value of agents. Allowing for randomization in the mechanism, we can nearly match this bound, obtaining a competitive ratio of $O(\log h)$ even without knowledge of h .

Finally, we generalize our model to a setting with jobs with varying job length, introducing a fourth parameter to the private type of an agent. Porter [20] had provided a deterministic auction that satisfies an efficiency competitive ratio bound of $(1 + \sqrt{k})^2 + 1$, where k is the ratio of maximum to minimum *value density* (that is, value divided by job length). We provide a randomized truthful auction with a competitive ratio $O(\log(l))$ where l is the ratio of the maximum job length to the minimum job length of agents, assuming that upper and lower bounds on the job lengths are known in advance to the mechanism designer. The auction is based on the asynchronous unit-length auction. On the other hand, when the ratio l is not known in advance we can achieve the same competitive ratio with efficiency, using an online auction that is truthful with respect to all parameters except job length.

1.3 Prior work

Online mechanism design has been the focus of several recent papers. Many of these papers (e.g. [5, 11]) assume that the agents

arrive in a predetermined order which is not under their control, and that an agent's only private information is her value. Designing truthful mechanisms is much easier in these single-parameter models. Some of the online mechanisms which have appeared (e.g. [2, 14]) are strategyproof against agents misstating their arrival or departure time because they are based on prices which are monotonically non-decreasing over time. However, such approaches do not lead to competitive online mechanisms in our setting because the non-decreasing price requirement is too restrictive.

Hajiaghayi, Kleinberg and Parkes [9] present constant-competitive online mechanisms for auctioning identical goods. Unlike the present paper, they assume that the agents arrive in random order and study a setting without re-usable goods in which goods can be allocated at any particular time.

Friedman and Parkes [7] initiated the study of VCG-based online mechanisms. Such mechanisms are (dominant-strategy) truthful in the rare instances where the underlying allocation problem admits an online algorithm with competitive ratio 1. Recently, Parkes and Singh [18, 19] have studied VCG-based online mechanisms with Bayesian-Nash equilibrium, adopting the framework of Markov Decision Processes. The setting for this work is quite general, but the solution concept is weaker than dominant strategy truthfulness.

Porter [20] presents a truthful mechanism for the variation on our model with different length jobs. In order for an agent to derive positive utility, it must be granted the resource for a total duration equal to its job length. Porter achieves a competitive ratio of $(1 + \sqrt{k})^2 + 1$ with respect to the optimal efficiency, where k is the ratio of maximum to minimum *value density* (value divided by processing time) of a job, and proves that this ratio is optimal for deterministic mechanisms. Here, we provide a randomized mechanism whose efficiency competitive ratio is $O(\log(l))$ where l is the ratio of maximum to minimum processing time. This significantly improves Porter's result, except for those instances in which the amount of variation in job lengths is exponentially greater than the variation in value densities.

Lavi and Nisan [15] study a closely-related online auction problem, in which items have expiration times and may be allocated any time at or before their expiration. (Our model can be regarded as a special case of theirs when the number of items is equal to the number of time slots and items may only be allocated at their exact expiration time.) Assuming unrestricted misreports, they prove strong negative results for deterministic truthful auctions (no such mechanism can achieve a competitive ratio better than the number of items) and this leads them to consider a weaker solution concept called *Set-Nash equilibrium* which admits constant-competitive mechanisms. Here, rather than modifying the solution concept, we achieve constant-competitiveness by restricting the set of allowable misreports. As we have argued earlier, this restriction (the *no late departures* assumption) is natural in the context of online mechanism design.

Earlier, Ng et al. [17] presented a generalization of the model in which jobs have both *length* and *size*, and in which there are multiple units of a reusable good available in each period. Their VIRTUALWORLDS mechanism generalizes the online auction presented in this paper for the unit-length model, but was presented without competitive analysis and for a model in which agents cannot misreport their arrival time or patience.

1.4 Outline

The outline of the paper is as follows. Section 2 introduces the model we are discussing. Section 3 provides the main characterization results for truthful online allocation rules. Further generalizations of this characterization, together with a price-based characterization that applies for a model of restricted misreporting and points beyond the results in this paper, are postponed until Section 8. Section 4 presents several competitive truthful online mechanisms in synchronous and asynchronous models for unit length jobs. In Section 5, we prove that no deterministic strategyproof mechanism

can achieve a competitive ratio better than 2 for efficiency in the synchronous setting. In Section 6 we consider truthful online mechanisms that are competitive with respect to revenue and prove that there is no deterministic strategyproof online mechanism whose revenue is constant-competitive with that of the Vickrey-Clarke-Groves (VCG) mechanism, but there is a randomized online mechanism which achieves a competitive ratio of $O(\log h)$, where h is the ratio of maximum value to minimum value of agents. In Section 7, we generalize our results to the case in which the agents have different job lengths and can lie about their processing times. Finally, we conclude with some open problems in Section 9.

2. THE MODEL

We will consider mechanism design problems for scheduling a single re-usable resource over a finite time interval $[0, T]$. In this section we define a simple model for studying such problems. Generalizations of this basic model will also be considered.

In our basic model, there are n agents, and the type of an agent i is characterized by an ordered triple $v_i = (a_i, d_i, w_i)$ with $0 \leq a_i \leq d_i \leq T$ and $0 \leq w_i < \infty$. We refer to a_i, d_i as the agent's arrival and departure time, respectively, and we refer to w_i as the agent's value. The set of all possible agent types is denoted by V . An allocation is a function $x = (x_1, \dots, x_n)$ which assigns to each agent a subset of $[0, T]$ which is a finite union of subintervals, such that distinct agents are assigned disjoint subsets.³ The set of all such allocations is denoted by A . For an agent i and an allocation x , let $q_i(x) = 1$ if the set $x_i \cap [a_i, d_i]$ contains an interval of length at least 1, otherwise $q_i(x) = 0$. The value of agent i for allocation x is equal to $q_i(x)w_i$, i.e. the agent gets value w_i if it is allocated at least one time unit between its arrival and departure times, and otherwise its value is zero.

We will be studying direct revelation mechanisms, in which each agent participates by simply announcing its type. A deterministic mechanism consists of a social choice rule (also called an allocation rule) $f : V^n \mapsto A$ and a payment rule $p : V^n \mapsto \mathbb{R}_{\geq 0}^n$. Here, $f(v_1, \dots, v_n)$ represents the allocation which is chosen when the vector of reported types is (v_1, \dots, v_n) , and $p_i(v_1, \dots, v_n)$ represents the amount agent i must pay. We will sometimes summarize the allocation rule f by specifying the function $q : V^n \mapsto \{0, 1\}^n$ whose i -th component is $q_i(f(v_1, \dots, v_n))$. Note that the value of each agent for the allocation $f(v_1, \dots, v_n)$ is completely determined by the value of $q(v_1, \dots, v_n)$; for this reason, by abuse of notation we will sometimes also refer to q as the allocation rule. We will assume that agents have quasi-linear utilities, so the utility of agent i for outcome x and payment p_i is $q_i(x)w_i - p_i$. The efficiency of a mechanism is the combined value of all agents for the allocation, i.e. the quantity $\sum_i q_i(v)w_i$. The revenue of a mechanism is the sum of the payments collected from agents, i.e. the quantity $\sum_i p_i(v)$. We will be comparing the efficiency and revenue of our mechanisms against the standard VCG mechanism, which computes the allocation and payments off-line. See Krishna [12] for a definition of VCG mechanisms.

For now, we adopt a model in which we allow only late reports of arrivals and early reports of departures. We consider relaxations of this definition in Section 8. As discussed in the introduction, this assumption is justifiable in many practical applications, and also necessary to obtain constant-competitive mechanisms [15].

We are interested in mechanisms that satisfy the following two properties: *voluntary participation*, i.e., for every $v \in V^n$, and every agent i with $v_i = (a_i, d_i, w_i)$, $p_i(v) \leq q_i(v)w_i$, and *strategyproofness* (also known as *truthfulness*), that is, for every $v \in V^n$ with $v_i = (a_i, d_i, w_i)$ (think of v as the true types of agents), and every bid $\hat{v}_i = (\hat{a}_i, \hat{d}_i, \hat{w}_i)$ for agent i satisfying $\hat{a}_i \geq a_i$ and $\hat{d}_i \leq$

³We use this general model to capture both the synchronous and asynchronous settings, as well as job lengths greater than one. An agent may receive multiple subintervals, but only has value for the allocation if one of the subintervals is long enough to satisfy its demand.

d_i , we have $q_i(v)w_i - p_i(v) \geq q_i((\hat{v}_i, v_{-i}))w_i - p_i((\hat{v}_i, v_{-i}))$, where $v_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ (i.e., the utility of i is maximized if she bids truthfully).

3. CHARACTERIZING TRUTHFULNESS

In this section, we provide two simple characterization theorems for truthful mechanisms. The first establishes a monotonicity criterion that is necessary and sufficient for the existence of a payment rule that truthfully implements a given allocation rule; this can be considered a generalization of well-known results concerning truthful mechanisms for one-parameter agents (see, for example, [1]). The second is a necessary and sufficient price-based characterization for truthful mechanisms. Both theorems can be substantially generalized; such generalizations will be treated in Section 8.

DEFINITION 1 (MONOTONICITY). We say that a type $v_i = (a_i, d_i, w_i)$ dominates the type $v'_i = (a'_i, d'_i, w'_i)$, denoted $v_i \succ v'_i$, if $a_i \leq a'_i$, $d_i \geq d'_i$, and $w_i > w'_i$. An allocation rule $q : V^n \mapsto \{0, 1\}^n$ is called monotone if for every agent i and every $v, v' \in V^n$ with $v_i \succ v'_i$ and $v_j = v'_j$ for $j \neq i$, we have $q_i(v) \geq q_i(v')$.

THEOREM 2. Let $q : V^n \mapsto \{0, 1\}^n$ be an allocation rule. There is a payment rule p such that the mechanism (q, p) is strategyproof if and only if q is monotone.

The proof of this theorem is omitted, since it will be subsumed by the more general Theorem 6, which applies to randomized mechanisms as well. Notice that Theorem 2 is existential. Whether the payment rule can be computed efficiently depends on the specifics of the model and the allocation rule.

We can now provide a price-based characterization for truthful online auctions in terms of a value-independent price schedule. Consider price schedule, $ps_i(a, d, v_{-i})$, that will define the price to agent i for an allocation, given that the agent announced interval $[a, d]$. Notice this price can depend on its reported arrival and departure, but not on its reported value. We can define a price-based online auction:

DEFINITION 3 (PRICE-BASED ONLINE AUCTION). An online auction is price-based if and only if there exists a value-independent price schedule $ps_i(a, d, v_{-i})$, such that for any $v \in V$ and any agent i , $q_i(v) = 1$ if and only if $ps_i(a_i, d_i, v_{-i}) \leq v_i$, and payments $p_i(v) = ps_i(a_i, d_i, v_{-i})$ in this case, or zero otherwise.

DEFINITION 4 (PRICE MONOTONICITY). Prices are monotonic if $ps_i(a, d, v_{-i}) \leq ps_i(a', d', v_{-i})$, for all $a' \geq a$ and $d' \leq d$.

THEOREM 5. An online allocation rule f is truthfully implementable if and only if there is a truthful price-based auction with value-independent and monotonic prices $ps_i(a, d, v_{-i})$ that implements f .

PROOF. (\Rightarrow) We refer to the proof of Theorem 6, which appears in the next section. If f is truthfully implementable, then the corresponding function q is monotone, and the proof of Theorem 6 establishes that the prices $p_i(v)$ defined in (7) truthfully implement f . Note that the right side of (7) depends only on a_i, d_i , and v_{-i} , and that it is monotonically non-decreasing in a_i and non-increasing in d_i , so that we may use it to define the value-independent and monotonic price schedule $ps_i(a, d, v_{-i})$.

(\Leftarrow) Given an allocation rule f which is implemented by prices $ps_i(a, d, v_{-i})$, we must prove that f is truthfully implementable, i.e. that the corresponding function q is monotone. This follows from the monotonicity of prices: if $v_i = (a_i, d_i, w_i)$ dominates $v'_i = (a'_i, d'_i, w'_i)$, then $ps_i(a_i, d_i, v_{-i}) \leq ps_i(a'_i, d'_i, v_{-i})$. Now, suppose $w_i > w'_i$, then if $q_i(v'_i, v_{-i}) = 1$ we have $ps_i(a'_i, d'_i, v_{-i}) \leq w'_i$, which implies $ps_i(a_i, d_i, v_{-i}) < w_i$. This in turn implies $q_i(v) = 1$, as desired. \square

3.1 Truthful randomized mechanisms

A randomized mechanism is a probability distribution over deterministic mechanisms. We assume agents are risk-neutral, so that their utility for a probability distribution over outcomes is equal to the expected utility of a random sample from this distribution. Note that this means it is immaterial, from the standpoint of determining the utility of agent i , whether the price charged to i is equal to the random variable $p_i(v)$ or to its expectation. Accordingly, for the rest of this section we will assume that the payment $p_i(v)$ is deterministic, and that only the allocation $f(v)$ is random. As before, we will summarize the allocation rule f using a function $q = (q_1, \dots, q_n)$, where $q_i(v)$ is now defined to be the probability that the time interval allocated to i has length at least 1, given $f(v)$. Note that q_i now takes values in the interval $[0, 1]$ rather than the two-element set $\{0, 1\}$. As before, the utility of agent i depends only on the value of $q_i(v)$ and not on any other feature of the distribution of $f(v)$. For this reason, by abuse of notation we will sometimes refer to q (rather than f) as the allocation rule. The word ‘‘monotone’’, when applied to such a function q , is still interpreted according to Definition 1.

THEOREM 6. *Let $q : V^n \mapsto [0, 1]^n$ be an allocation rule. There is a payment rule p such that the mechanism (q, p) is strategyproof and satisfies voluntary participation if and only if q is monotone.*

PROOF. First, we prove the ‘‘if’’ part. Let q be a monotone allocation rule and $v \in V^n$ with $v_i = (a_i, d_i, w_i)$. We define the payment rule as

$$p_i(v) = q_i(v)w_i - \int_0^{w_i} q_i((a_i, d_i, x), v_{-i}) dx. \quad (1)$$

We show that the allocation rule q in combination with the payment rule p constitute a strategyproof mechanism. It is also clear from the above definition that this mechanism satisfies the voluntary participation property. If the mechanism is not strategyproof, there is an agent i , a vector $v \in V^n$ of true types with $v_i = (a_i, d_i, w_i)$, and a non-truthful bid $\hat{v}_i = (\hat{a}_i, \hat{d}_i, \hat{w}_i)$ with $\hat{a}_i \geq a_i$ and $\hat{d}_i \leq d_i$ such that the utility $q_i(\hat{v}_i, v_{-i})\hat{w}_i - p_i(\hat{v}_i, v_{-i})$ of agent i if she bids \hat{v}_i is strictly greater than the utility $v_i q_i(v) - p_i(v)$ that she derives from being truthful. Using Equation 1, this can be written as follows:

$$\begin{aligned} (w_i - \hat{w}_i)q_i(\hat{v}_i, v_{-i}) + \int_0^{\hat{w}_i} q_i((\hat{a}_i, \hat{d}_i, x), v_{-i}) dx \\ > \int_0^{w_i} q_i((a_i, d_i, x), v_{-i}) dx. \end{aligned} \quad (2)$$

By monotonicity of q and the inequalities $\hat{a}_i \geq a_i$ and $\hat{d}_i \leq d_i$, we have

$$\int_0^{w_i} q_i((a_i, d_i, x), v_{-i}) dx \geq \int_0^{w_i} q_i((\hat{a}_i, \hat{d}_i, x), v_{-i}) dx. \quad (3)$$

Equations 2 and 3 imply

$$(w_i - \hat{w}_i)q_i(\hat{v}_i, v_{-i}) > \int_{\hat{w}_i}^{w_i} q_i((\hat{a}_i, \hat{d}_i, x), v_{-i}) dx. \quad (4)$$

We now consider two cases: if $w_i > \hat{w}_i$, then by dividing both sides of the above inequality by $w_i - \hat{w}_i$ we obtain that $q_i(\hat{v}_i, v_{-i})$ is strictly greater than the average of $q_i((\hat{a}_i, \hat{d}_i, x), v_{-i})$ over $x \in [\hat{w}_i, w_i]$, which contradicts the monotonicity of q . Similarly, if $w_i < \hat{w}_i$, then by dividing both sides of inequality (4) by $w_i - \hat{w}_i$ we obtain that $q_i(\hat{v}_i, v_{-i})$ is strictly less than the average of $q_i((\hat{a}_i, \hat{d}_i, x), v_{-i})$ over $x \in [w_i, \hat{w}_i]$, which again contradicts the monotonicity of q . This contradiction establishes the strategyproofness of the mechanism with allocation rule q and payment rule p .

Conversely, assume q is an allocation rule for which there is a payment rule p such that (q, p) is strategyproof. Consider an agent i and types $v, v' \in V^n$ with $v_i \succ v'_i$ and $v_{-i} = v'_{-i}$ such that $q_i(v) < q_i(v')$. Let $v_i = (a_i, d_i, w_i)$ and $v'_i = (a'_i, d'_i, w'_i)$. Since $v_i \succ v'_i$, we have $w_i > w'_i$. Now consider a scenario where the true types of

the agents are given by v . In this scenario, if i bids truthfully, she will have a utility of $q_i(v)w_i - p_i(v)$, but if she bids v'_i , then her utility will be $q_i(v')w_i - p_i(v')$. Therefore, strategyproofness of (q, p) implies

$$q_i(v)w_i - p_i(v) \geq q_i(v')w_i - p_i(v'). \quad (5)$$

Now, consider a scenario where the true types of the agents are v' and agent i might lie by announcing v_i . A similar argument for this scenario implies

$$q_i(v')w'_i - p_i(v') \geq q_i(v)w'_i - p_i(v). \quad (6)$$

By adding inequalities 5 and 6 and using the inequality $w_i > w'_i$, we obtain $q_i(v) \geq q_i(v')$. Therefore, q is monotone. \square

It is worth remarking that the payment rule $p_i(v)$ defined in (1) is equivalent to the following simpler definition in the case of deterministic mechanisms:

$$p_i(v) = \min\{w'_i : q_i((a_i, d_i, w'_i), v_{-i}) = 1\}, \quad \text{if } q_i(v) = 1, \quad (7)$$

and $p_i(v) = 0$ otherwise. In words, an agent that is allocated pays the smallest value it could have reported and still received an allocation.

4. MECHANISMS FOR JOBS OF UNIT LENGTH

4.1 The synchronous model

In this section we consider the basic setting defined in Section 1: Agents arrive and depart at the beginning and end of time slots numbered $0, \dots, T$ (i.e., $V = \{(a, d, w) : a, d \in \{0, \dots, T\}, a \leq d, w \in \mathbb{R}_{\geq 0}\}$), and there is only one re-usable good that can be allocated to at most one agent in each time slot. An agent wins if she receives this good in one of the time slots between her arrival and departure. The mechanism must decide which agent (if any) receives the good at time slot t , based on the information available at time t , i.e., the arrival times and the values of all agents that have arrived at a time $\leq t$, and the departure time of all agents that have left at a time $\leq t$. The payment of an agent must be computed based on the information available at her departure (i.e. it can be delayed past the period in which an allocation decision is made).

The following theorem shows that there is a 2-competitive strategyproof mechanism for this problem. The proof of this theorem is based on a simple greedy allocation rule that in each interval selects the bidder with the highest value that has not received the good yet. It is worth noting that this mechanism is strategyproof even if bidders are allowed to announce an arrival time that is before their actual arrival time.

THEOREM 7. *There exists a strategyproof online 2-competitive mechanism in the synchronous model with a single re-usable good.*

PROOF. We use the following greedy allocation rule: At any time step, allocate the good to the bidder with the highest value that is present at that time and has not received the good yet. For the payment rule, we define $p_i(v)$ by equation (7). It is clear that both our allocation rule and our payment rule are computable in polynomial time and in an online fashion. It is also easy to see that this allocation rule is monotone. The reason for this is that an agent i who loses does not affect the state of the algorithm; therefore, if i loses when bidding (a_i, d_i, w_i) , she would still lose if she arrives at a time later than a_i , leaves before d_i , or announces a value less than w_i . Thus, by Theorem 6 (and its proof), this mechanism is strategyproof and satisfies voluntary participation.

Now, we show that this allocation rule is 2-competitive. We do this by a charging argument. Consider an off-line optimal solution OPT . For any agent i who wins in OPT , we charge her value to an agent who wins in the greedy solution. If i herself is a winner in the greedy solution, we charge her value to herself. Otherwise, let t be

the time at which i wins the auction. Since i never wins in the greedy solution, she is present at time t , and therefore the algorithm must pick a winner j at time t whose value is not less than the value of i . We charge the value of i to j . It is not hard to see that this charging scheme charges each agent j in the greedy solution at most twice, each time for a value less than the value of j . Therefore, the value of OPT is at most twice the value of the greedy solution. \square

4.2 The asynchronous model

In the synchronous model the arrival and departure times of agents were restricted to be integers. In the asynchronous model, we let agents have types in $V = \{(a, d, w) : a, d, w \in \mathbb{R}_{\geq 0}, 0 \leq a \leq d - 1 \leq T\}$, and allow the allocation of the good to an agent to begin at any time. Furthermore, we allow the mechanism to revoke an allocation before it is completed. This is necessary for constant-competitiveness, since if an agent i with $d_i = a_i + 1$ and a high value w_i arrives at a time that the mechanism has started but not finished an allocation to another agent then the mechanism must interrupt this first allocation. An agent derives a value w_i if she is allocated the good for one unit of time continuously, i.e., it is not possible to preempt and then resume jobs.

For this model, the following theorem gives a 5-competitive mechanism. This mechanism is similar in idea to the greedy mechanism in Theorem 7, except here once an allocation starts, we increase the value of the corresponding agent by a multiplicative factor that increases continuously over time at an exponential rate, thereby giving her an advantage over newly arriving agents. This allows us to use a charging argument to prove that the mechanism is 5-competitive. We also use this algorithm in Section 7 to handle the case in which we have different job lengths.

THEOREM 8. *There exists a strategyproof online 5-competitive mechanism in the asynchronous model with a single re-usable good.*

PROOF. We modify the greedy algorithm in the proof of Theorem 7 as follows. We call a point t in time *critical* if a new agent arrives at time t , or an allocation completes at time t . At any critical point t , we compare the values of all agents that are present at time t , and have not been allocated the good for one unit of time continuously. If there is an agent in this set who has already had the good for $\delta < 1$ units of time before time t (i.e., she has received an allocation at time $t - \delta$ and has not been interrupted after that), then we multiply the value of this agent by 2^δ before comparing it to the values of other agents. After the comparison, the agent with the highest value receives the allocation. If this agent is different from the agent who has had the allocation since $t - \delta$, the latter agent is interrupted.

For the payment rule, we again use the general rule given by equation (7). This rule can be implemented efficiently as follows: For each agent i who wins, we run the algorithm without this agent, and let T denote the set of all critical points $t \in [a_i, d_i]$ in this run. For every $t \in T$, we define the value at time t as the maximum over the values of all agents that are present in the system at t and have not received the allocation for one continuous unit of time before t . We now calculate the price of allocation for i at time t by taking the maximum, over all $t' \in (t, t+1)$, of the value at t' divided by $2^{t'-t}$. Also, if there is an allocation to an agent j that is started at time $t - \delta$ for some $0 \leq \delta < 1$ and continues beyond t , we take the maximum of the above value and 2^δ times the value of j , and let the price for i at time t be the maximum of these values. The payment of i is the minimum, over all $t \in [a_i, d_i - 1]$ of the price of allocation to i at time t . It is clear that the allocation and payment rules are both efficiently computable in an online fashion. Furthermore, it is easy to see that the payment computed by the above algorithm is the same as the one given in the proof of Theorem 6, and therefore in order to establish the strategyproofness of our mechanism, we only need to prove its monotonicity.

Monotonicity of our allocation rule is easy to see for the departure time and the value: If an agent i does not win when the types of

agents are v , then she will not win if she lowers her value or leaves earlier than her true departure time. Consider a situation where i announces an arrival time a'_i after her true arrival time a_i . If agent i does not receive an allocation in the interval $[a_i, a'_i]$ in the truthful scenario, then the behavior of the algorithm is the same in both scenarios and therefore i does not win in the non-truthful scenario. If she starts an allocation in this interval, then the behavior of the algorithm in these scenarios might be different, since during the time intervals that were allocated to i in the truthful scenario, other agent might receive the allocation in the non-truthful scenario. The value of these agents cannot be higher than the value of i . But we know that any allocation to i in the truthful scenario is interrupted before it completes. Hence, allocations made during the same time intervals in the non-truthful scenario must also be interrupted. Therefore, the state of the algorithm at time a'_i is the same in both scenarios. Thus, i does not win in the non-truthful scenario. This establishes the monotonicity of our allocation rule.

We now prove that this algorithm is 5-competitive. This is done by charging the value of any winning agent in an optimal solution OPT to a winning agent in our algorithm. We assume, without loss of generality, that OPT does not interrupt any allocation. For any winning agent i in OPT , if she is also a winner in our algorithm, then her value is charged to herself. Otherwise, consider the time t at which i is allocated the good in OPT . At this time, our algorithm has allocated the good to an agent j_0 . This agent might be interrupted in our algorithm. If she is interrupted, let j_1 be the agent that interrupts her. We continue this chain until we reach an agent j_k who is not interrupted, and charge the value of i to this agent. We now calculate the maximum total value charged to an agent j with value w_j who wins at time t in our algorithm. If j wins in OPT , there is a charge of w_j . Consider an agent i in OPT whose value is charged to j . Let $t' = t - \delta$ be the time at which i receives an allocation in OPT . It is clear from the algorithm that $\delta > -1$ and the value of i is at most $2^{-\delta}w_j$. Also, the value of δ for any two such i 's must be apart by at least one. Therefore, the total charge to j is at most $w_j + \sum_{l=-1}^{\infty} 2^{-l}w_j = 5w_j$. This shows that our algorithm is 5-competitive. \square

There is an example which shows that the above analysis is tight, i.e., the competitive ratio of the above algorithm is not better than 5.

4.3 Multiple re-usable goods

Finally, we show that both mechanisms proposed above can be generalized to the case where there are k identical re-usable goods instead of one, achieving the same competitive ratios. (Formally, the setting of k identical re-usable goods is defined using the same model as in Section 2, except that the set of allocations A is enlarged to encompass all functions mapping agents to subintervals of $[0, T]$, such that each $t \in [0, T]$ belongs to at most k such subintervals.)

THEOREM 9. *There is a strategyproof online mechanism which is 2-competitive (5-competitive) for efficiency, in the synchronous (asynchronous) model with k re-usable goods.*

PROOF. The proof is essentially the same as the proof of Theorem 7, except that the greedy allocation rule gives the goods at time t to the unassigned bidders with the k largest values at time t . \square

5. COMPETITIVE RATIO LOWER BOUND

In this section we prove that no deterministic strategyproof mechanism can achieve a competitive ratio strictly better than 2 in the synchronous setting. This shows that the result of Theorem 7 is tight. Note that if we do not care about truthfulness, the best known lower bound for the problem is the Golden ratio (≈ 1.618) [8].

THEOREM 10. *No truthful online deterministic mechanism can obtain a $(2 - \epsilon)$ -approximation for efficiency in the synchronous model with a single re-usable good for any constant $\epsilon > 0$.*

PROOF. Suppose, there is a truthful online deterministic mechanism A which can obtain a $2 - \epsilon$ approximation for efficiency. We design a set of scenarios for which we obtain a contradiction on the behavior of the algorithm.

First scenario is as follows. Assume that we have two agents. The first one called x has type $v_x = (1, 2, 1 + \delta)$, where $0 < \delta < \frac{\epsilon}{1-\epsilon}$. Assume there is another agent y whose type is $v_y = (1, 3, 1)$. In this case, if agents are truthful, mechanism A should assign x to the first time slot and assign y to the second time slot. In addition, mechanism A should always charge x less than $1 + \delta$, otherwise x has motivation to lie about its value. It means for positive δ , x has a positive utility in this scenario.

In the next scenario, we have an agent y' which behaves the same as agent y does in the previous scenario, i.e., $v_{y'} = v_y$. We have an agent x' of type $v_{x'} = (1, 3, 1 + \delta)$. Finally, we have an agent z' of type $v_{z'} = (2, 3, \infty)$. In this case, we claim that mechanism A should always assign x' to the first time slot, since otherwise x' can lie about its type to announce it the same as that of x in the previous scenario and still get a positive value. Note that in this case, x' cannot be assigned to the second time slot because of agent z , whose value is very large, and at time 1, mechanism A does not know whether such an agent exists or not. In addition, in this scenario, agent y' cannot be assigned to any time slot.

Finally, we consider the third scenario in which we have two agents x'' and z'' for which $v_{x''} = v_x$ and $v_{z''} = v_{z'}$, and another agent y'' for which $v_{y''} = (1, 2, 1)$. In this case, mechanism A should not assign y'' to the first time slot. The reason is that in this case, y'' has a positive value (since mechanism A should charge y'' less than one) and thus in the second scenario, agent y' could lie about its type to announce it the same as that of y'' in this scenario and thus get a positive value (agent y' had zero value in the second scenario.) It means in this scenario, if agent z'' does not appear at all, since mechanism A does not know about it at time 1, still agent y'' should not be assigned to any time slot. Thus the mechanism obtain efficiency at most $1 + \delta$, though the optimum efficiency is $2 + \delta$ in this case. Since $\frac{2+\delta}{1+\delta} > 2 - \epsilon$ by the choice of δ , we have the desired contradiction to the existence of mechanism A . \square

6. REVENUE OF THE AUCTION

Theorems 7 and 8 established the existence of mechanisms whose efficiency is constant-competitive with that of the VCG mechanism. In contrast to these positive results, there is no strategyproof mechanism whose *revenue* is constant-competitive with that of the VCG mechanism, if we insist on mechanisms which satisfy the following axiom.

DEFINITION 11. *An impatient bidder is an agent with arrival and departure times (a_i, d_i) satisfying $d_i = a_i + 1$, i.e. an agent who can only accept an allocation at time a_i . We say that a mechanism M considers impatient bidders anonymously if it has the following property: whenever M assigns a time slot t to an impatient bidder x with value w , there is no impatient bidder y with arrival time t whose value is strictly greater than w .*

THEOREM 12. *If M is a truthful, deterministic online mechanism satisfying voluntary participation, and M considers impatient bidders anonymously, then the competitive ratio of M with respect to the VCG revenue is not bounded above by any constant C .*

PROOF. Suppose M is a truthful deterministic online mechanism whose revenue is C -competitive with the off-line VCG mechanism, for some constant C , and suppose M considers impatient bidders anonymously. We will derive a contradiction using a series of scenarios, as in the proof of Theorem 10. All of our scenarios will involve a timeline with time slots numbered $0, 1, \dots, T$, a set of patient bidders $\{A_1, A_2, \dots, A_q\}$ each with value 3, arrival time 0, and departure time T , and pairs of impatient bidders $B_0, C_0, B_1, C_1, \dots, B_T, C_T$, such that B_i and C_i both arrive at time i and depart at time $i + 1$, and whose values satisfy $w(B_i) <$

$w(C_i)$. In any such scenario, the agent winning time slot i in the off-line VCG mechanism has a VCG payment which is bounded below by $w(B_i)$, hence the VCG revenue is bounded below by $\sum_{i=0}^T w(B_i)$.

This lower bound proof is a bit more complicated than the proof of Theorem 10, so we will first give an outline of the main ideas. By considering a series of scenarios $\{\Upsilon_k\}$, we will prove the following fact: if agents $\{A_1, A_2, \dots, A_q\}$ (defined as above) arrive at time 0 along with impatient bidders B_0, C_0 with values 1 and 2, then M can not allocate the time slot to B_0 or C_0 . Since the VCG revenue in this case is bounded below by 1, we know that M must collect at least $1/C$ in revenue. Next we will consider a carefully constructed series of scenarios $\{\Psi_q\}$, in which the set of agents at time 0 is as above, and all agents arriving after time 0 are impatient and have values much smaller than $1/C$. Using Ψ_q , we will show that it is impossible for M to guarantee at least $1/C$ in revenue. Roughly speaking, this is because truthfulness prevents M from extracting much revenue from the patient bidder who wins time slot 0, and (for q large enough) there will not be enough time after this to make up the difference.

Scenario Υ_k ($1 \leq k \leq T$) is specified as follows. Choose a sufficiently large number x (any number greater than $2C + 3Cq$ will suffice), and put $w(B_i) = x^i$, $w(C_i) = 2x^i$ for $0 \leq i \leq k$, $w(B_i) = w(C_i) = 0$ for $i > k$. The VCG revenue in this case is at least $1 + x + \dots + x^k$. If M does not allocate time slot k to B_k or C_k , then its revenue is bounded above by $3q + (1 + x + \dots + x^{k-1})$ (due to the voluntary participation constraint) and this is less than $1/C$ times the VCG revenue, by our choice of x . Thus we may conclude that in scenario Υ_k , M allocates time slot of k to an impatient bidder. Since M is an online mechanism, and Υ_k is indistinguishable from Υ_T until time $k + 1$, we may conclude that in scenario Υ_T , each time slot $k > 1$ is allocated to an impatient bidder. If M also allocates time slot 0 to an impatient bidder, then each of the patient bidders A_1, \dots, A_q derives zero utility in scenario Υ_T . This contradicts truthfulness, since A_i may derive positive utility by announcing value $2 + \epsilon$, arrival time 0, and departure time 1. (In this case, A_i will win time slot 0 and pay at most $2 + \epsilon$, since M considers impatient bidders anonymously and satisfies voluntary participation.)

Now let y be a very small number (any number smaller than $1/48C^3$ will suffice) and consider a scenario with agents A_1, \dots, A_q and B_0, C_0 defined as before, and with impatient bidders B_i, C_i at time i whose values are $y, 2y$ respectively. The VCG revenue is unbounded as a function of T , hence if M is C -competitive it must eventually allocate a time slot $i > 0$ to one of the impatient bidders B_i, C_i . Let τ_q be the first such time slot. Define a series of scenarios $\{\Psi_q : 0 \leq q \leq 6C^2\}$ as follows: in scenario Ψ_q , there are agents A_1, \dots, A_q with arrival time 0, departure time T , and value 3, there are impatient bidders B_0, C_0 at time 0 with values 1, 2, respectively, there are impatient bidders B_i, C_i at time i ($1 \leq i < \tau_{q-1}$) with values $y, 2y$, respectively, and there is an impatient bidder B_i with value y at time $i = \tau_{q-1}$. (When $q = 0$, we interpret τ_{q-1} to mean $\tau_0 + 1$.) By considering Ψ_0 we can get an explicit upper bound on τ_0 . The VCG revenue in scenario Ψ_0 is $1 + \tau_0 y$, whereas the revenue of M is bounded above by $2 + 2y$, hence $C(2 + 2y) \geq 1 + \tau_0 y$, i.e. $\tau_0 \leq 2C(1 + 1/y)$.

Now we will consider Ψ_q for $q > 0$, obtaining a lower bound on $\tau_{q-1} - \tau_q$. The argument proceeds as follows. By the definition of τ_q , we know that M doesn't allocate any time slots to impatient bidders until time τ_q , so the revenue from impatient bidders is at most $2y \max(0, \tau_{q-1} - \tau_q + 1)$, by voluntary participation. For each patient bidder A_i , it is possible to receive time slot τ_{q-1} at a price of at most $2y$ by declaring an arrival time of τ_{q-1} , departure time $\tau_{q-1} + 1$, and value $2y$. In such a case, time slot τ_{q-1} will be allocated to A_i because M considers impatient bidders anonymously, and the price will be at most $2y$ by voluntary participation. As M is truthful, it must be the case that A_i derives at least as much utility from truthfully announcing her type.

Thus M can not charge A_i a price greater than $2y$, so the combined revenue from patient bidders is at most $2qy$. Using the trivial lower bound of 1 on the VCG revenue, and the C -competitiveness of M , we now have $1 \leq 2Cqy + 2Cy \max(0, \tau_{q-1} - \tau_q + 1)$. We have $q \leq 6C^2$ and $y < 1/48C^3$, hence $2Cqy < 1/4$. Thus $3/4 < 2Cy \max(0, \tau_{q-1} - \tau_q + 1)$, i.e., $\tau_{q-1} - \tau_q > \frac{3}{8Cy} - 1$. Summing over $q = 1, 2, \dots, 6C^2$ we obtain $\tau_0 > \frac{9C}{4y} - 6C^2$. Combining this with the upper bound $\tau_0 \leq 2C(1 + 1/y)$ from earlier, we see that $\frac{2C}{y} + 2C > \frac{9C}{4y} - 6C^2$, i.e. $6C + 2 > \frac{1}{4y} > 12C^3$, a contradiction since $C \geq 1$. \square

By closely examining the proof of Theorem 12, we can strengthen it to the following result.

COROLLARY 13. *Suppose that the bids w_i are constrained to belong to an interval $[a, b]$ whose endpoints are known to the mechanism designer, and let $h = b/a$. If M is a truthful, deterministic online mechanism which considers impatient bidders anonymously and satisfies voluntary participation, then the competitive ratio of M with respect to the VCG revenue is at least $\Omega\left(\left(\frac{\log h}{\log \log h}\right)^{1/4}\right)$.*

PROOF. Let us explicitly determine an interval $[a, b]$ containing all the bid values considered in the scenarios $\{\Upsilon_k\}$ and $\{\Psi_q\}$ which arose in the proof of Theorem 12. The smallest bid considered was y , in scenario Ψ_q for all values of q . Note that y may be taken to be equal to $1/49C^3$. The largest bid considered was $2x^T$ in scenario Υ_T . Here x may be taken to be $3C(1+q) < 20C^3$, since $q \leq 6C^2$. The largest value of T considered was $T = \tau_0 \leq 2C(1 + 1/y) \leq 100C^4$, in scenario Ψ_0 . Thus $x^T = O(C^{300C^4})$. Therefore the bid interval $[a, b] = [y, 2x^T]$ satisfies $b/a = O(C^{O(C^4)})$, and the right side is at most h , for some $C = O\left(\left(\frac{\log h}{\log \log h}\right)^{1/4}\right)$. \square

Next we show that a randomized mechanism can nearly match the lower bound established in Corollary 13, even if the bid interval $[a, b]$ is not known in advance. Specifically, the revenue of our mechanism is $O(\log h)$ -competitive with the VCG revenue, where h is the ratio of maximum value to minimum value of agents. This result can be considered parallel to a result of Lavi and Nisan [14], who derived the same revenue competitive ratio in a different online auction setting. However, it is worth noting that their result requires foreknowledge of the bid interval $[a, b]$.

THEOREM 14. *There is a randomized online mechanism which achieves a competitive ratio of $O(\log h)$ when all bids belong to an interval $[a, b]$ satisfying $b/a = h$. The mechanism need not know the values of a and b .*

PROOF. For simplicity, we will work in the synchronous model. The same competitive ratio can be achieved in the asynchronous model by incorporating the mechanism from Theorem 8 into this proof.

Suppose first that the bid interval is known to be $[1, h]$. Then the mechanism is extremely simple: at time 0, guess a random power of 2 between 1 and h and define this to be the reserve price p . In each period, if there is at least one bidder present whose bid is above p and has not yet been satisfied, choose one such bidder at random and allocate the time slot to that bidder, charging a price of p . For each period t , if the VCG mechanism allocates t to a bidder x and charges q to that bidder, then with probability $1/\lceil \log_2(h) \rceil$ the random price p satisfies $q/2 \leq p < q$. If so, then either the mechanism charges p to x , or it charges p to an agent who is allocated time slot t . Using a charging scheme as in the proof of Theorem 7, this implies that the mechanism's competitive ratio is bounded above by $4\lceil \log_2(h) \rceil$.

If the bid interval is not known initially, then it is a bit harder to design a $O(\log h)$ -competitive mechanism. We will define such a mechanism, which combines the random-reserve-price notion introduced above, the greedy mechanism analyzed earlier, and a random

partitioning technique which is often useful in designing competitive auctions. At time 0, the mechanism samples a random number ξ_k independently from the uniform distribution on $[0, 1]$, for each integer k . (Of course, in an actual implementation the numbers ξ_k will be determined by lazy evaluation, i.e. we sample ξ_k the first time we need to examine its value and not earlier.) A random partition of the set of agents into two sets A, B is computed online, by assigning each agent (at its arrival time) randomly, independently, and equiprobably to A or B . For each agent i , we now determine a threshold price schedule $\alpha_t(i)$ as follows. If $i \in B$, then $\alpha_t(i) = \infty$ for all t . Let A_t (resp. B_t) denote the set of agents arriving at or before time t , and assigned to set A (resp. B). If $i \in A$ and B_t is the empty set, $\alpha_t(i) = 0$. Otherwise, let $w_{\min}(B_t), w_{\max}(B_t)$ be the minimum and maximum bids reported by agents in B_t . Among all integers k such that $w_{\min}(B_t)/2 \leq 2^k < w_{\max}(B_t)$, choose the one for which ξ_k is minimum, and set $\alpha_t(i) = 2^k$ in $i \in A$.

At time t , the mechanism computes its allocation as follows. It first defines a set of eligible agents, by taking the set A_t and removing all agents who have been allocated a time slot in a previous period. If the set of eligible agents is non-empty, then an eligible agent i with maximum bid value is chosen (randomly and uniformly, if there is more than one eligible agent with maximum bid value) and is declared the *winner* at time t . Letting w_i denote the bid value of this winning agent i , the mechanism allocates time slot t to agent i if $\alpha_t(i) < w_i$; otherwise t is unallocated.

The pricing rule is defined as follows. For each agent i , a price schedule $p_t(i)$ is computed by simulating the same allocation rule with agent i absent. (All other random choices, including the partition of the remaining agents into sets A, B , are unchanged in the simulation.) Letting $\beta_t(i)$ denote the bid of the agent who is the winner at time t in this simulation (or $\beta_t(i) = 0$ if there is no winner), we set

$$p_t(i) = \max\{\alpha_t(i), \beta_t(i)\}.$$

The payment for agent i is determined by the value-independent, monotone price schedule $ps_i(a, d, v_{-i}) = \min_{a \leq t \leq d} p_t(i)$. It is easy to check that this price schedule implements the allocation rule described above, hence the mechanism is strategyproof by Theorem 5.

To prove that the mechanism is $O(\log h)$ -competitive with the VCG revenue, we begin by identifying a set of agents whose contribution to the VCG revenue may be easily bounded. An agent i with type (a_i, d_i, w_i) is *pitiful* if the VCG mechanism charges a positive price to agent i , yet there exists an integer k_i such that $2^{k_i-1} < w_i \leq 2^{k_i}$ and every other bid received at or before time a_i is greater than 2^{k_i} . (We call k_i the *index* of pitiful agent i .) Note that for distinct pitiful agents i, j with arrival times $a_i \leq a_j$, the indices k_i, k_j are also distinct and satisfy $k_i > k_j$. Thus the sum of the bids of all pitiful agents is bounded above by 2^{k^*+1} , where k^* is the index of the earliest-arriving pitiful agent, if there is any such agent.

Let t_0 be the earliest time at which more than one agent arrives. (If there is no such time, then the VCG revenue is zero and there is nothing to prove.) Let $w_1 \geq w_2$ be the two largest bids arriving at time t_0 (corresponding to agents i_1, i_2 , respectively) and note that $w_2 > 2^{k^*-1}$, since t_0 is no later than the arrival time of the earliest pitiful agent. With probability $1/4$, $i_1 \in A, i_2 \in B$. If so, the probability that our mechanism allocates time slot t_0 to an agent at a price of at least 2^{k^*-1} is at least $1/\lceil \log_2(h) \rceil$. It follows that the total amount charged to pitiful agents by the VCG mechanism is at most $16\lceil \log_2(h) \rceil$ times the expected revenue of our mechanism.

It remains to bound the total amount charged to non-pitiful agents by the VCG mechanism. Assume that VCG allocates time slot t to non-pitiful agent i at a price $p > 0$. Let 2^m be the largest power of 2 less than w_i . We claim there exist agents $j \neq i, k \neq i$, arriving at or before time a_i , such that $w_j \geq p$ and $w_k/2 \leq 2^m$. The existence of agent k follows from the fact that i is not pitiful. The existence of agent j follows from the truthfulness of the VCG mechanism;

otherwise, i could improve its utility by claiming that its departure time is $a_i + 1$ and its value is $p - \varepsilon$, for some sufficiently small positive ε .

With probability $1/2$, $i \in A$. If so, the allocation rule satisfies one of the following properties:

1. There exists a time at which agent i is a winner.
2. The winner at time t is not agent i , but this agent bids at least as much as agent i .

In the first case, let $i' = i$ and let t' be the time at which i wins. In the second case, let i' be the winner at time t and let $t' = t$. (Note that $t' \geq a_i$ in both cases, since i is ineligible in time slots earlier than a_i . Note also that $i' \in A$, since agents in B are never eligible.) Defining agents j, k as above, and conditioning on the event $i \in A$, the probability that $j, k \in B$ is at least $1/4$. (It is equal to $1/4$ unless $j = k$.) By the properties of agents j and k , the interval $[\min\{w_j, w_k\}/2, \max\{w_j, w_k\})$ contains a power of 2 between $p/2$ and $w_{i'}$, say 2^ℓ . Conditional on the event that $i \in A$, $j, k \in B$, the probability is at least $1/\lceil \log_2(h) \rceil$ that $\alpha_{t'}(i') = \alpha_T(i') = 2^\ell$. Thus, while the VCG mechanism charges p to agent i , our mechanism charges at least $(16 \lceil \log_2(h) \rceil)^{-1} p$ in expectation to agent i' . This argument credits a given agent i' at most twice: once when $i = i'$, and once when t is the time slot which our mechanism allocates to i' . Thus the expected amount charged to non-pitiful agents by the VCG mechanism is at most $32 \lceil \log_2(h) \rceil$ times the revenue of our mechanism. \square

It may seem that the revenue of the VCG mechanism is a rather weak benchmark against which to compare our mechanism's revenue. However, as illustrated by Theorem 12, deterministic truthful mechanisms can not be constant-competitive even against this benchmark. Moreover, the VCG revenue is at least as large as the maximum total value of a feasible allocation to the set of agents that are disjoint from those satisfied by the optimal allocation. This follows from the fact that the sets of agents who can win the auction simultaneously form a matroid [4, 21].

It is also worth mentioning that the greedy mechanism used in the proof of Theorem 7, which is 2-competitive for efficiency, can have an arbitrarily bad competitive ratio with respect to the VCG revenue. Consider $n+2$ agents, 1 to $n+2$, as follows. Agent i , $2 \leq i \leq n+1$, has type $v_i = (i-1, i+1, 2)$. Agent 1 has type $v_1 = (1, 2, 1)$ and agent $n+2$ has type $v_{n+2} = (n+1, n+2, 1)$. It is easy to observe that off-line VCG charges each of agents $2, \dots, n+1$ a price of 1 and thus collects n for revenue. However, the greedy algorithm only charges agent 2 a price of 1 and all others 0. It means the competitive ratio of the greedy mechanism for revenue can be arbitrarily large.

7. DIFFERENT JOB LENGTHS

In this section we consider the case where jobs are allowed to have different lengths. In other words, each agent's type is now characterized by *four* values: a_i, d_i, w_i , and L_i , where a_i, d_i are interpreted as before, and L_i is a positive real number specifying the *length* of job i . The value of a given outcome x for agent i is equal to w_i if the subset of $[0, T]$ allocated to i contains an interval whose length is at least L_i , otherwise x has zero value for i .

First, we start by assuming that all job lengths are in an interval $[L_{min}, L_{max}]$ which is known to the mechanism beforehand. Using a technique similar to that adopted in Section 6, we show that there exists a strategyproof mechanism that achieves a competitive ratio of $O(\log(L_{max}/L_{min}))$ for efficiency.

THEOREM 15. *There is a randomized strategyproof mechanism that achieves a competitive ratio of $O(\log(L_{max}/L_{min}))$ for efficiency when all job lengths are in an interval $[L_{min}, L_{max}]$ known to the mechanism, even if agents are allowed to lie about their length.*

PROOF. The mechanism is as follows: let L be a random power of two in the interval $[L_{min}, 2L_{max}]$. Every job of length more than

L is rejected. Jobs of length less than or equal to L are treated like jobs of length L . Finally, the mechanism in Theorem 8 is used to schedule these jobs. We observe that this allocation rule has a truthful implementation. First, notice that since the mechanism rounds all job lengths that are less than L to L and removes jobs of length greater than L , no agent can benefit by lying about their length. Therefore, we only need to prove truthfulness with respect to other parameters. By Theorem 6, it is enough to show that the allocation rule is monotone. Monotonicity with respect to values and departure times is obvious. Monotonicity with respect to arrival times follows from the monotonicity of the mechanism in Theorem 8.⁴

We now show that the above allocation rule achieves a competitive ratio of $O(\log(L_{max}/L_{min}))$ for efficiency. Consider an optimal solution OPT . Let L_1, L_2, \dots, L_k denote all powers of two in the interval $[L_{min}, 2L_{max}]$. We partition the set of jobs served in the solution OPT into k subsets OPT_1, \dots, OPT_k , where OPT_i is the set of jobs in OPT that have length more than L_{i-1} and less than L_i (we let $L_0 := L_{min}$). For each i , we number the jobs in OPT_i with consecutive natural numbers in the order that they are served in the solution, and let OPT'_i be the set of odd numbered jobs or the set of even numbered jobs in OPT_i , whichever has higher total value. Clearly, the sum of the values of jobs in OPT'_i is at least half the sum of values of jobs in OPT_i . Furthermore, if we round the length of all jobs in OPT'_i up to L_i , we obtain a feasible solution of the instance constructed in the mechanism if the value L picked by the mechanism is L_i . Therefore, by Theorem 8, the value of the solution found by the allocation rule conditioned on $L = L_i$ is at least $1/5$ times the value of jobs in OPT'_i . Since for each i , the probability that $L = L_i$ is $1/\lceil \log(L_{max}/L_{min}) \rceil$, the expected value of the solution found by our mechanism is at least $O(1/\log(L_{max}/L_{min}))$ times the sum of values of OPT_i for $i = 1, \dots, k$, or $O(1/\log(L_{max}/L_{min}))$ times the value of the optimal solution. \square

When the interval $[L_{min}, L_{max}]$ is not known in advance, the following theorem shows that we can still achieve a competitive ratio of $O(\log(L_{max}/L_{min}))$ using a mechanism that is strategyproof with respect to all parameters except job lengths. The proof is very similar to the proofs of Theorems 14 and 15, and is deferred to the full version of this paper.

THEOREM 16. *There is a randomized strategyproof mechanism that achieves a competitive ratio of $O(\log(L_{max}/L_{min}))$ for efficiency when all job lengths are in an interval $[L_{min}, L_{max}]$. The mechanism need not know the values of L_{min} and L_{max} , but assumes that the agents do not lie about their length.*

8. A GENERAL FRAMEWORK FOR TRUTHFUL ONLINE AUCTIONS

In this section, we generalize the characterization in Section 3 to other models of misreporting in online auctions and extend the existing theory on price-based characterizations of truthful auctions to models with restricted misreporting. The standard theory of truthful mechanisms states that a truthful mechanism must be implemented in terms of an agent-independent price function [13], where the price to an agent is independent of its reported type. This need not be the case in truthful online auctions, where a patient agent can make a smaller payment than if it was less patient, even when it would receive the good in the same period. This can occur, for instance, in the simple unit-length synchronous auction. In online auctions the price can depend on an agent's reported arrival and departure, although not on its reported value. Prices must satisfy a monotonicity property with respect to reported arrivals and departures and allocations must be carefully timed for some misreporting

⁴Notice that here we are using the fact that job lengths that are less than L are all rounded up to L . Without this, the allocation rule would not be monotone with respect to arrival times and can not be truthfully implemented.

models. Auctions that are based on agent-independent price schedules, $ps_i(t, v_{-i})$ in period t , are truthful with appropriate timing requirements. However, the existence of a simple price schedule of this kind is not necessary for truthfulness.

8.1 Restricted Misreports

In the main model in this paper, we considered agents that can only report late arrivals and early departures. This is an example of a domain in which there is restricted misreporting of agent type.

Let $L(v_i) \subseteq V_i$ denote the available misreports (or *lies*) available to agent i with value v_i . We assume *transitivity*, so that $v'_i \in L(v_i)$ and $v''_i \in L(v'_i)$ implies $v''_i \in L(v_i)$. The standard mechanism design model has $L(v_i) = V_i$ and the standard multiagent model with cooperative agents has $L(v_i) = v_i$. A general definition of truthfulness in this model is as follows:

DEFINITION 17 (TRUTHFULNESS). *A mechanism (f, p) for social choice rule, $f : V^n \rightarrow A$, and payment rule, $p : V^n \rightarrow \mathbb{R}_{\geq 0}^n$, is truthful if for any agent i and any $v_{-i} \in V_{-i}$, and any $v_i \in V_i$ and $\hat{v}_i \in L(v_i)$, we have $v_i(f(v)) - p_i(v) \geq v_i(f(\hat{v}_i, v_{-i})) - p_i(\hat{v}_i, v_{-i})$.*

The following lemma provides a price-based characterization of truthful mechanisms in domains with restricted misreporting. The price can depend on an agent's reported type, but in a limited way.

LEMMA 18. *For available lies L , a function f is truthful if and only if there exists a price function, $\tilde{p}_i : A \times V_i \times V_{-i} \rightarrow \mathbb{R} \cup \infty$, such that,*

- (B1) *For any agent i , any $v \in V$, any $x \in A$, if there exists $v'_i \in L(v_i)$ such that $f(v'_i, v_{-i}) = x$ then price $\tilde{p}_i(x, v_i, v_{-i}) = \min\{\tilde{p}_i(x, v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\}$, otherwise $\tilde{p}_i(x, v_i, v_{-i}) = \infty$.*
- (B2) *For any agent i , any $v \in V$, $f(v) \in \arg \max_{x \in A} \{v_i(x) - \tilde{p}_i(x, v_i, v_{-i})\}$.*

PROOF. (\Leftarrow) By contradiction, suppose (B1) and (B2) hold but that f is not truthful. We can define payments $p_i(v) = \tilde{p}_i(x, v_i, v_{-i})$ for $x = f(v)$. Consider some v and $v'_i \in L(v_i)$, for which $v_i(x) - p_i(v) < v_i(y) - p_i(v'_i, v_{-i})$ and $x = f(v)$, $y = f(v'_i, v_{-i})$. By (B1) we have $p_i(v'_i, v_{-i}) = \tilde{p}_i(y, v'_i, v_{-i}) \geq \tilde{p}_i(y, v_i, v_{-i})$, and $v_i(x) - \tilde{p}_i(x, v_i, v_{-i}) < v_i(y) - \tilde{p}_i(y, v_i, v_{-i})$, a contradiction with (B2).

(\Rightarrow) Given a truthful f , we know there is a payment function $p_i(v)$. To show (B1) and (B2), we construct the price function $\tilde{p}_i : A \times V_i \times V_{-i} \rightarrow \mathbb{R} \cup \infty$, as follows. For any i , $v \in V$, and $x \in A$, if there exists $v'_i \in L(v_i)$ such that $f(v'_i, v_{-i}) = x$ we set $\tilde{p}_i(x, v_i, v_{-i}) = \min\{p_i(v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\}$, otherwise $\tilde{p}_i(x, v_i, v_{-i}) = \infty$. To show (B1), fix some x , some v_i and some v_{-i} for which there is a $v'_i \in L(v_i)$ for which $f(v'_i, v_{-i}) = x$. First, there can be no $\hat{v}_i \in L(v_i)$ with $f(\hat{v}_i, v_{-i}) = x$ for which $\tilde{p}_i(x, v_i, v_{-i}) > \tilde{p}_i(x, \hat{v}_i, v_{-i})$, because this would give,

$$\begin{aligned} & \min\{p_i(v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\} \\ & > \min\{p_i(v''_i, v_{-i}) : v''_i \in L(\hat{v}_i), f(v''_i, v_{-i}) = x\}, \end{aligned}$$

which contradicts the transitivity of misreporting. So, we have

$$\begin{aligned} & \tilde{p}_i(x, v_i, v_{-i}) \leq \\ & \min\{\tilde{p}_i(x, v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\}, \end{aligned}$$

which together with

$$\begin{aligned} & \tilde{p}_i(x, v_i, v_{-i}) \geq \\ & \min\{\tilde{p}_i(x, v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\}, \end{aligned}$$

gives (B1) (since $v'_i = v_i \in L(v_i)$).

Now, we show the same price function also satisfies (B2). By contradiction, consider some v such that $f(v) = x$ and $v_i(x) - \tilde{p}_i(x, v_i, v_{-i}) < v_i(y) - \tilde{p}_i(y, v_i, v_{-i})$. Let $v'_i \in L(v_i)$ be the type

that determined $\tilde{p}_i(y, v_i, v_{-i})$, i.e. for which $f(v'_i, v_{-i}) = y$ and for which $p_i(v'_i, v_{-i}) = \tilde{p}_i(y, v_i, v_{-i})$. When other agents declare v_{-i} and agent i has true type v_i it is better to declare $v'_i \in L(v_i)$, a contradiction with truthfulness. \square

Given price functions $\tilde{p}_i(x, v_i, v_{-i})$ satisfying conditions (B1) and (B2) we can define a payment rule $p_i(v) = \tilde{p}_i(f(v), v_i, v_{-i})$ that implements f truthfully. Thus, the agent's payment is the price for the optimal alternative, as chosen by the allocation rule.

Note that payments can depend on the reported type of an agent but only in a restricted way (B1). In words, the price on alternative x given report v_i must be the lowest price that is available to the agent given misreports that are possible for a type of v_i . When misreports are unrestricted we recover the standard agent-independent price-based characterization:

COROLLARY 19. *For unrestricted reports, then for any truthful function f , there exists an agent-independent price function, $\tilde{p}_i : A \times V_{-i} \rightarrow \mathbb{R} \cup \infty$, such that for any $v \in V$ and any agent i , $f(v) \in \arg \max_{x \in A} \{v_i(x) - \tilde{p}_i(x, v_{-i})\}$.*

PROOF. Observe that if $L(v_i) = V_i$, then for any $v_{-i} \in V_{-i}$, and any $x \in A$, we have

$$\begin{aligned} \tilde{p}_i(x, v_i, v_{-i}) &= \min\{\tilde{p}_i(x, v'_i, v_{-i}) : v'_i \in L(v_i), f(v'_i, v_{-i}) = x\} \\ &= \min\{\tilde{p}_i(x, v'_i, v_{-i}) : f(v'_i, v_{-i}) = x\} \\ &= \tilde{p}_i(x, v_{-i}) \end{aligned}$$

for all $v_i \in V_i$, where $\tilde{p}_i(x, v_{-i})$ is an agent-independent price function that does not depend on the reported type of agent i . \square

8.2 Application to online auctions

In this section we state generalizations of the monotonicity and price-based characterizations in Section 3 to additional models of misreporting. These characterizations generalize well-known results in the case of one-parameter agents (see, for example, [1]) and strengthen Theorem 3 in [7]. For space reasons, the proofs in this section are omitted; they will appear in the full version of this paper. We consider the following three models of misreporting:

- (A1) $L(v_i) = \{v'_i = (a'_i, d'_i, w'_i) : a'_i \geq a_i, d'_i \leq d_i, \text{ any } w'_i\}$
- (A2) $L(v_i) = \{v'_i = (a'_i, d'_i, w'_i) : a'_i \geq a_i, \text{ any } d'_i, \text{ any } w'_i\}$
- (A3) $L(v_i) = \{v'_i = (a'_i, d'_i, w'_i) : d'_i \leq d_i, \text{ any } a'_i, \text{ any } w'_i\}$

Although we focus on Model A1 in this paper (since it seems more realistic), all of our truthful mechanisms in Sections 4 and 6 can be also stated for the more general Model A3.⁵

The notion of *critical value* will be important in defining the characterization.

DEFINITION 20 (CRITICAL VALUE). *The critical value $v^c(a, d, v_{-i})$ is the minimal value w'_i , with $v'_i = (a, d, w'_i)$ for which $q_i(v'_i, v_{-i}) = 1$, or ∞ if $q_i(v'_i, v_{-i}) = 0$ for all w'_i .*

For models A2 and A3 we define the notion of a critical period, $I_c(a, d, v_{-i}) \subseteq T$, which will limit when allocations can be made.

DEFINITION 21 (CRITICAL PERIOD FOR MODEL A2). *In model A2, $I_c(a, d, v_{-i}) = [t, d]$ where period t is defined as the earliest $t \in [a, d]$ for which $v^c(a, t, v_{-i}) = \min_{\tau: a \leq \tau \leq d} v^c(a, \tau, v_{-i})$.*

DEFINITION 22 (CRITICAL PERIOD FOR MODEL A3). *In model A3, $I_c(a, d, v_{-i}) = [a, t]$ where period t is defined as the latest $t \in [a, d]$ for which $v^c(t, d, v_{-i}) = \min_{\tau: a \leq \tau \leq d} v^c(\tau, d, v_{-i})$.*

An online allocation rule is said to be *supported* by a price schedule $ps_i(a, d, v_{-i})$ in the sense of the *price-based online auctions* (Definition 3), so that the rule allocates a good to the agent if and only if the agent's reported value $w_i \leq ps_i(a, d, v_{-i})$ and at price $ps_i(a, d, v_{-i})$.

⁵Note that we cannot go beyond Model A3, because of a recent result of Lavi and Nisan [15] mentioned in the introduction.

THEOREM 23. (Model A1) *The following are equivalent statements in a domain with no early arrival and no late departure:*

1. An allocation rule has a truthful implementation.
2. An allocation rule is monotonic.
3. An allocation rule is supported by a monotonic and value-independent price schedule $ps_i(a, d, v_{-i})$.

THEOREM 24. (Model A2) *The following are equivalent statements in a domain with no early arrivals:*

1. An allocation rule has a truthful implementation.
2. An allocation rule is monotonic, and when making an allocation assigns the item in the critical period.
3. An allocation rule is supported by a monotonic and value-independent price schedule $ps_i(a, d, v_{-i})$, and when making an allocation assigns the item no earlier than the first period $t : a_i \leq t \leq d_i$ for which $ps_i(a_i, t, v_{-i})$ is minimal.

THEOREM 25. (Model A3) *The following are equivalent statements in a domain with no late departures:*

1. An allocation rule has a truthful implementation.
2. An allocation rule is monotonic, and when making an allocation assigns the item in the critical period.
3. An allocation rule is supported by a monotonic and value-independent price schedule $ps_i(a, d, v_{-i})$, and when making an allocation assigns the item no later than the last period $t : a_i \leq t \leq d_i$ for which $ps_i(t, d_i, v_{-i})$ is minimal.

In each case, the payment rule that makes the allocation rule truthful is as defined in Eq. 7 in Section 3.1.

The price-based characterizations provide a useful complement to the monotonicity-based characterization on allocation rules. Rather than define monotonic allocation rules, one can define monotonic and value-independent price schedules for which the supported allocation rule is feasible.

9. OPEN PROBLEMS

In this paper, we considered online auctions with re-usable goods. We presented several upper and lower bounds on competitiveness for both revenue and efficiency of truthful online mechanisms. In a general framework, we also provided necessary and sufficient characterizations for allocation rules that can be implemented in a truthful online auction. Here, we present several open problems.

The main open question is to determine whether the known lower bounds for deterministic mechanisms can be extended to apply to randomized mechanisms. We are referring here to our Theorem 10 (for efficiency) and Theorem 12 (for revenue), as well as the lower bound by Lavi and Nisan [15] (for efficiency, in misreporting model A2). We also conjecture that the lower bound of Theorem 12 can be improved to $\Omega(\log h)$ and extended to asymmetric mechanisms.

Another open problem is to determine whether there is a deterministic mechanism whose revenue is $O(\log h)$ -competitive with the VCG mechanism, at least in the case where the interval $[1, h]$ is known to the mechanism. (The paper by Lavi and Nisan [14], which addresses a different type of online auction problem, achieves this bound using a deterministic mechanism with known bid interval.)

Acknowledgements. We would like to thank Jason Hartline, Nicole Immorlica, and Ron Lavi for fruitful discussions.

10. REFERENCES

- [1] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
- [2] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proceedings of the thirty-fifth ACM symposium on Theory of computing*, pages 503–510. ACM Press, 2003.
- [3] Y. Bartal, F. Chin, M. Chrobak, S. Fung, R. Lavi, J. Sgall, and T. Tichy. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *Proceedings of the 21st Symposium on Theoretical Aspects of Computer Science*, pages 187–198, 2004.
- [4] S. Bikhchandani, S. de Vries, J. Schummer, and R. Vohra. Linear programming and vickrey auctions. *IMA Volumes in Mathematics and its Applications, Mathematics of the Internet: E-Auction and Markets*, 127:75–116, 2001.
- [5] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online learning in online auctions. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003)*, pages 202–204. New York, 2003. ACM.
- [6] F. Chin and S. Fung. Online scheduling for partial job values: Does timesharing or randomization help? *Algorithmica*, 37:149–164, 2003.
- [7] E. J. Friedman and D. C. Parkes. Pricing WiFi at starbucks: issues in online mechanism design. In *The 4th ACM Conference on Electronic Commerce*, pages 240–241, 2003.
- [8] B. Hajek. On the competitiveness of on-line scheduling of unit-length packets with hard deadlines in slotted time. In *Proceedings of the 2001 Conference on Information Sciences and Systems*, 2001.
- [9] M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 71–80. ACM Press, 2004.
- [10] A. Kesselman, Z. Lotker, Y. Mansour, B. Patt-Shamir, B. Schieber, and M. Sviridenko. Buffer overflow management in qos switches. In *Proceedings of the 33rd ACM Symposium on Theory of Computing*, pages 520–529, 2001.
- [11] R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *44th Annual Symposium on Foundations of Computer Science (Cambridge, 2003)*. IEEE Computer Soc., Los Alamitos, CA, 2003.
- [12] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [13] R. Lavi, A. Mu'alem, and N. Nisan. Towards a characterization of truthful combinatorial auctions. In *Proc. 44th Annual Symposium on Foundations of Computer Science*, 2003.
- [14] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *The 2nd ACM Conference on Electronic Commerce*, pages 233–241, 2000.
- [15] R. Lavi and N. Nisan. Online ascending auctions for gradually expiring items. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2005. to appear.
- [16] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49(5):577–602, September 2002.
- [17] C. Ng, D. C. Parkes, and M. Seltzer. Virtual Worlds: Fast and Strategyproof Auctions for Dynamic Resource Allocation. In *Fourth ACM Conf. on Electronic Commerce (EC'03)*, 2003. Shorter version. Extended version available at <http://www.eecs.harvard.edu/~parkes/pubs/virtual.pdf>.
- [18] D. C. Parkes and S. Singh. An MDP-based approach to Online Mechanism Design. In *Proc. 17th Annual Conf. on Neural Information Processing Systems (NIPS'03)*, 2003.
- [19] D. C. Parkes and S. Singh. Approximately efficient online mechanism design. In *Proc. 18th Annual Conf. on Neural Information Processing Systems (NIPS'04)*, 2004.
- [20] R. Porter. Mechanism design for online real-time scheduling. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 61–70. ACM Press, 2004.
- [21] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proceedings of the 20th Symposium on Theoretical Aspects of Computer Science*, 2003.