Designing A Limited Supply Online Auction

Steven Gao
Department of Computer Science
University of British Columbia
Vancouver, BC V6T 1Z4, Canada
sxgao@cs.ubc.ca

ABSTRACT

We examine the problem of designing a strategy-proof auction in an online setting. This means that agents can arrive and depart dynamically. Thus, the auction mechanism must be able to decide on the optimal allocation of goods based only on bid information from the agents that have already arrived and not on future arrivals. If a good is to be allocated to an agent, the mechanism must do so before that agent departs. We compare two different solutions to the limited supply online auction problem, each with its own set of assumptions and problem specific criteria. In conducting our survey, we note the emerging importance of comparing a mechanism's performance in terms of revenue and efficiency by comparing it to the optimal revenue and efficiency. This emphasis on an algorithmic approach to mechanism design has been denoted as competitive analysis. We conclude our review of online mechanism design by discussing the important avenues for future work.

Categories and Subject Descriptors

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

General Terms

Algorithms, Design, Economics, Performance

Keywords

Online auction, competitive analysis, incentive compatibility, strategy-proof

1. INTRODUCTION

Mechanism design (MD) is concerned with creating rules to enforce specific outcomes within a system of rational and self-interested agents. The classic MD setting is static in that the mechanism makes a single decision in determining the outcome for the system and all agents are present throughout the process. In online mechanism design we consider the same problem but allow agents to arrive and depart dynamically [1]. The mechanism must determine the optimal outcome based on the information limited to only those agents that it has observed up to the current point in time. In other words, the mechanism does not know the preferences of agents that will arrive in the future and it must decide whether to satisfy each bid as the requests are received in real time.

In this paper we examine a subset of mechanisms, known as auctions, in an online situation. In essence, auctions are mechanisms where the final outcome is an allocation of a set of resources to a subset of the agents in the system. The interaction within the system is between an auctioneer (seller) and a set of agents (bidders). As Lavi and Nisan point out in [2], an online auction problem is a contrast to the traditional offline environment where the auctioneer receives *all* bids before generating a set of allocations. The offline setting implies that all agents, as well as the auctioneer, must wait a given amount of time before the transaction can take place. At the end of the waiting period the mechanism performs batch processing on the entire set of bids at once.

The motivation behind studying online auctions is quite evident. Auction design has quickly gained prominence in computer science particularly in the areas such as electronic commerce (e.g. auctions on eBay, electronic catalogues on Amazon), computer and network resource allocation (e.g. network bandwidth allocation), and trading between software agents [2, 3, 4, 5]. These settings are best modeled dynamically since agents enter and depart in a constant stream. In addition, players in the auction are not willing to wait very long for a result [2].

In this paper, we survey work done by Lavi and Nisan [2], Awerbuch et al. [6], Friedman and Parkes [7], and Hajiaghayi et al. [8] on designing limited-supply online auctions. We will attempt to aggregate the analysis done by each of the authors and then summarize the main issues involved in designing a generalized online auction.

In the remainder of this section we discuss research that has contributed to online mechanism design, but is either not directly related to online auction design or could not be included due to space restrictions. In section 2, we formally define an online auction and review some useful terminology. In section 3, we discuss in detail the current work done on designing strategy-proof online auctions for a limited-supply good. In Section 4, we address our contribution to the online mechanism design problem which is to summarize the requirements for a generalized online auction. Finally, in section 5 we conclude with some remarks regarding the direction that online mechanism design is headed.

1.1 Related Work

In addition to the research results we examine in section 3, the following papers may be of interested to those readers who want a broader perspective on online auction design. Common issues

Best Conference Ever '05 Vancouver, British Columbia CA

that are addressed in the papers below are incentive compatibility (or truthfulness), auction competitiveness (amount of revenue produced), the amount of supply of the good being sold (bounded or unbounded), and the distribution over which agent valuations are drawn (known or unknown distribution to the mechanism).

In [3] Goldberg et al. attempt to design truthful competitive mechanisms for the single round, sealed bid auction for an item with *unlimited* supply. Unlimited supplies means the seller can create additional copies of the good at negligible marginal cost. Such auctions can be applicable to real life when selling a digital good, i.e. books, music, videos, and software in downloadable form [3]. They restrict their attention to situations where each agent only wants a single copy of the good. This condition is also known as *unit-demand*.

Although Goldberg et al. consider offline mechanisms their work is important to the online MD problem because it is the first to analyze auctions in terms of their competitiveness. An auction (online or offline) is *competitive* if it produces revenue that is within a constant factor of the offline *optimal fixed price revenue*. We define this formally in section 2. As stated in [3], expressing an online auction's performance in terms of a benchmark offline auction is similar to competitive analysis of online algorithms within computer science, where an online algorithm's performance is compared to the that of the optimal offline algorithm. Such an analysis of an auction is rigorous because it provides theoretical guarantees on the revenue from that auction. The revenue generated by an offline optimal fixed price auction acts as ultimate goal that all online auctions should strive to achieve. They also extend their analysis to a bounded (or limited) supply auction and show that their definition of competitiveness still holds. In fact, the bounded supply case is a generalization of the unlimited supply case, because we can use the former to derive the latter by setting the number of goods equal to the number agents. Goldberg et al. examine single- and multi-priced auctions and argue for the use of a randomized sampling auction instead of a deterministic one in order to achieve competitiveness. Among the papers we examine in section 3, [2, 6, 8] all use a competitive analysis technique similar to that defined in [3], either as a direct result of Goldberg et al. or in parallel.

In [9], Bar-Yossef et al. directly extend the work of Goldberg et al. by considering the same auctions in an online setting. They present a randomized auction which is within a factor $O(e^{\sqrt{loglogh}})$ of the benchmark optimal offline auction, where h is the ratio between highest and lowest bid values. They point out some key differences between their online unbounded auction and an online limited supply auction [2]. Firstly, in a limited supply auction selling a good to the current agent means that it cannot be sold to a future agent. Thus, a competitive mechanism must somehow balance the revenue gained with lost potential revenue. This is usually handled being assuming non-decreasing prices. This is not a consideration in the unbound supply case. Secondly, Lavi and Nisan use the offline Vickrey auction as their benchmark for comparing their online auction in [2], but Goldberg et al. show that it is not competitive in the unbound supply case.

Bagchi et al. [5] also consider an online limited supply auction from an online algorithmic standpoint, but focus on defining optimal strategies for the seller. Their motivation is that many on auction websites (e.g. eBay, uBid) sellers do not use good pricing models and thus lose revenue as a result. The paper shows that a deterministic auction cannot possibly perform within a constant factor of the optimal offline auction, meaning it is not competitive. Thus, a randomized auction is necessary. However, they focus too much on algorithmic design and neglect the importance of making an their mechanism incentive compatible.

Awerbuch et al. [6] use Lavi and Nisan's adversarial model, but their main goal is to bridge the gap between online algorithms and online mechanisms. They describe a method to convert any competitive online algorithm (i.e. without game theoretic considerations) into an online truth-telling mechanism that is strategy-proof in terms of agent valuations. Specifically they derive a total profit that is different from the optimum profit by a bound of $O(p+log\mu)$. The optimum profit is produced by an offline algorithm that knows the true valuations of each agent.

Finally, Gallien [4] examines a limited-supply online model where the goal to is to maximize the expected discounted revenue by selling identical goods to self-interested, time-sensitive agents with unit-demand. He assumes that valuations and arrival-departure times are drawn uniformly from a *known* distribution. He designs a series of strategy-proof mechanisms by. What is interesting is that Gallien's work is done completely separate from the other papers we examine here. In other words there are no common citations between his work and the work of the other research groups we consider. In fact, Gallien uses the term *dynamic* mechanism design to denote the study of online mechanisms. Under the assumptions that the mechanism is stable and is based on discrete time, Gallien derives a strategy-proof mechanism *DP* by turning the mechanism design problem into one centered on dynamic programming.

2. BACKGROUND

In this section we define the formal notation that will be used throughout the rest of the paper as well as desirable properties that we want an online mechanism to have. The notation and properties are presented in a generic fashion so that they may be used for all of the mechanisms we discuss later.

2.1 Formal Definitions

We first define a quasi-linear mechanism and a direct revelation mechanism, since these are important to both offline and online auctions. We then define the class of online mechanisms that we are interested in. Given a set of agents $I = \{1, 2, ..., n\}$, a set of outcomes $O = X \times \Re^n$, where X is some finite set of choices, and a set of strategy spaces for the agents $S = S_1 \times ... \times S_n$:

DEFINITION 1. A quasi-linear mechanism over I and O is defined as $M(q, p, S)$ where $q : S \to X$ is the choice rule, such that $q(s) \in X$ is the choice implemented for strategy $s = (s_1 \times ... \times s_n)$, $p : S \to \Re^n$ is the payment rule, such that $p_i(s) \in \Re$ is the payment made by i, and S is the strategy space for all agents.

To clarify the payment rule can be divided into n payment rules, one for each agent. In the case of an auction, the set of choices X is simply the set of goods that the seller (auctioneer) is selling. Let $|X| = k$ be the number of goods available for sale.

The actions an agent i can take within the mechanism are defined completely by its strategy space S_i. However, based on the *revelation principle* [10, 8] we only need to consider a *direct revelation* mechanism:

DEFINITION 2. A direct revelation mechanism (DRM) is a mechanism $M(q, p, S)$, where the agent's strategy space consists only of disclosing its type $S_i = \Theta_i$ and the choice rule $q : \Theta \to X$ selects an choice $q(\hat{\theta}) \in X$ based on the reported types of the agents $\hat{\theta} = (\theta_1 \times ... \times \theta_n)$.

Thus, within a DRM an agent i's strategy involves either revealing its true type θ_i, which represents i's real preferences, and a false type $\hat{\theta}_i \neq \theta_i \in \Theta_i$. The revelation principle states any social choice function mapping agent types to outcomes that can be implemented by some mechanism in dominant strategies can also be implemented by a direct revelation mechanism that is truthful. Hence, from this point on, we will only look at implementing direct revelation mechanisms and discuss agent strategies only in terms of their types. The specific type space for an agent i is $\theta_i \in \Theta$.

We use a quasi-linear utility function to define the payoffs for each agent. So, an agent i's utility is $u_i(\hat{\theta}) = v_i(q(\hat{\theta}_i, \hat{\theta_{-i}})) - p_i(\hat{\theta}_i, \hat{\theta}_{-i})$, where $\hat{\theta}_{-i}$ is the declared types of all agents other than i, $v_i(q(\hat{\theta}_i, \hat{\theta}_{-i}))$ is i's true valuation for some allocation outcome q of the k goods, and $p_i(\hat{\theta}_i, \hat{\theta}_{-i})$ is the payment that i makes given all reported types $\hat{\theta}$. Note that the type profile can be written as $\hat{\theta} = (\hat{\theta}_i, \hat{\theta}_{-i})$.

For a classic direct revelation mechanism, the type of agent i is simply the set of possible valuations i can have for the outcomes (goods). However, in the online setting we have:

DEFINITION 3. An online mechanism over I and O is defined as a direct revelation mechanism $M(q, p, \Theta)$ where q and p are defined as before, and $\Theta = \Theta_1 \times ... \times \Theta_n$ is the type space for all agents, where the type for agent i is a triple defined as $\theta_i = (v_i, a_i, d_i)$, such that $v_i \in \Re$ is i's true valuation for a set of goods, $a_i \in [0, T]$ is i's arrival time, and $d_i \in [a_i, T]$ is i's departure time.

Only i knows its true type θ_i. Given the reported types $\hat{\theta} = (\hat{\theta_1}, ..., \hat{\theta_n})$ by each agent the mechanism produces an allocation $q_i(t, \hat{\theta})$ at time $0 \leq t \leq T$ and a payment $p_i(\hat{\theta})$ for all i. Here $q_i(t, \hat{\theta})$ is the number of items allocated to i (i.e. 0 or 1) and $p_i(\hat{\theta})$ is the payment for i at its departure time d_i.

We assume throughout this paper that agents are rational and seek to maximize their expected utility.

2.2 Mechanism Properties

We start by defining a dominant strategy (DS) for a player i:

DEFINITION 4. A strategy $\hat{\theta}_i$ for agent i is a dominant strategy if given its true type θ_i and all possible strategies for the other agents,

$$u_i(\hat{\theta}_i, \hat{\theta}_{-i}, \theta_i) \geq u_i(\hat{\theta}'_i, \hat{\theta_{-i}}, \theta_i), \hat{\theta}'_i \neq \hat{\theta}_i \in \Theta_i, \theta_{-i} \in \Theta_{-i} \quad (1)$$

where u_i is i's expected utility.

A dominant strategy maximizes i's expected utility regardless of the strategies (reported types) being played by everyone else. A DS equilibrium occurs when all agents are playing a dominant strategy. Compared to the Nash and Bayes-Nash equilibrium concepts dominant strategy is the most desirable, since it makes no assumptions about the information agents have about each other (e.g. perfect information and common knowledge about each other's preferences). Thus, we are interested in designing mechanisms that implement a dominant strategy equilibrium. This leads to the following definition of *strategy-proof* (SP) as defined in [8]:

DEFINITION 5. A mechanism $M(q, p, \Theta)$ is strategy-proof if given any agent i with true type θ_i and any type vector of the remaining agents' types θ_{-i} it is a dominant strategy for i to play θ_i.

A mechanism that is SP is also known as a *truthful* implementation in dominant strategies, or is simply stated to be *incentive compatible*. Incentive compatibility is important because it guarantees that a player will receive maximum expected utility by revealing its type truthfully. This removes the need for players to consider any strategy other than the truthful one, an important issue in computerized settings, where it may be computationally intractable for a software agent to consider all possible strategies [2].

Revenue is defined as

$$Rev(\hat{\theta}) = \sum_{i \in I} p_i(\hat{\theta}). \quad (2)$$

It is the sum of the payments that each agent gives to the auctioneer based on some allocation q. Thus, revenue is the expected utility of the auctioneer.

For competitive analysis, Goldberg et al. [3] introduce three important values which we adopt. We define the total utility T as the sum of all agent utilities. This is the upper bound on the revenue that a mechanism can generate. Goldberg et al. observe that this is also the revenue that an optimal untruthful multi-price auction would achieve by satisfying all agent requests at their reported bid values. A multi-price auction sells each copy of the good at a different price. Next, we define F as the revenue produced by selling each copy of the good at an optimal fixed price. This is also the revenue generated by an optimal untruthful single-price auction. Finally, we define h as the highest bid value. This will be used to derive lower bound estimates for the mechanisms that we look at. Goldberg et al. show that $F \geq \frac{T}{2 \lg h}$. This can also be stated as F is within a factor of $\Omega(2 \lg h)$ away from T, which restricts the penalty on revenue for using a single-price auction instead of a multi-price one [3].

We define the efficiency of an allocation as

$$Eff(\hat{\theta}) = max \sum_i q_i(\hat{\theta}) v_i. \quad (3)$$

This is the maximal sum of the valuations of all agents.

We use the definition of competitive ratio defined by Lavi and Nisan in [2]. Namely, a mechanism is *c-competitive* with respect to revenue (relative to the Vickrey auction) if for any type profile $\hat{\theta}$ it generates a revenue value that is at least $\frac{1}{c}$ of the revenue produced by the offline Vickrey auction using the same input $\hat{\theta}$.

3. THE LIMITED SUPPLY ONLINE AUCTION

In this section we take a more in-depth look at online mechanisms from [8, 2], where the authors tackle a limited-supply online auction problem, such that the auctioneer is limited to selling k identical items. The reason for comparing these two papers is because they both address the specific auction setting of interest, where [2] is a much cited early attempt at designing a limited supply online auction and [8] is a more recent attempt that tries to make less assumptions. We discuss the problem setting for each paper, which is slightly different than the other, followed by the main results, and finally some remarks about the advantages and disadvantages of their design.

3.1 Incentive Compatible Online Auction

Lavi and Nisan [2] study the online problem in terms of incentive compatibility and competitiveness. The mechanism has an adversary who specifies a set of valuations and a set of arrival-departure time intervals and then randomly assigns valuations to the time intervals for each agent. The mechanism must decide whether to allocate goods or not for each bid as soon as the bid is submitted. Lavi and Nisan use the Vickrey offline auction as a comparison tool to derive a competitive ratio for their mechanism. They also argue for using a worst-case analysis of the generate revenue instead of average-case, as is the standard practice in auction theory.

The motivation is that as auction theory is increasingly applied to computer science domains worst-case bounds provide a stronger guarantee than average-case bounds. This is especially important in situations when the theoretical average-case analysis does not produce the same distribution of results as real world observations.

3.1.1 Problem Formulation

They examined a setting with k identical indivisible goods, with the distinction that as k gets very large, they treat the group of goods as a single divisible good, i.e. k is a continuum. Each agent can request up to m copies of the good. The valuation of a good for each agent i is drawn from a range $[\underline{p}...\overline{p}]$, where $\underline{p}$ also represents the seller's reservation price. Each agent i's valuation v_i is private information. Because an agent can request more than one copy of the good, the paper defines the marginal valuation for agent i as $v_i(q)$, where q is the q-th good. It is assumed in [2] that marginal valuations are non-increasing, i.e. $\forall_i v_i(q+1) \leq v_i(q)$. It is pointed out that this is a necessary assumption that was also made in Vickrey's original paper proposing the Vickrey auction.

When the online auction begins the auctioneer does not know the number of agents, each agent has not been assigned a valuation. At some time t_i agent i *learns* its valuation and must declare its bid right away. As explained earlier we only need to consider DRMs, thus an agent's strategy consists only of declaring it marginal valuation for an additional good. An agent's bid is then defined by a function $v_i : [1...k] \rightarrow R$. Upon receiving a bid, the auctioneer must respond to that request before it can move on to the next bid. The auctioneer responds to an agent i by determining the quantity to allocate to i as well as the amount i must pay.

Lavi and Nisan define revenue as we have in section 2, except that they add in the marginal valuation of any remaining unsold goods using the reservation price $\underline{p}$. They also include the utility of the auctioneer in their formulation of efficiency.

3.1.2 Main Results

Lavi and Nisan define online auctions in terms of supply curves. Namely:

DEFINITION 6. An online auction is "based on supply curves" if for the i-th request $\hat{v}_i(q)$, the mechanism generates a supply curve $p_i(q)$ based on all previous requests, as well as,
(1) the quantity sold to i is $q_i = argmax_q \sum_{j=1}^{q} \hat{v}_i(j) - p_i(j)$ and
(2) the i pays $\sum_{j=1}^{q_i} p_i(j)$.

An important assumption they make throughout is that the supply curves $p_i(q)$ are non-decreasing. On the other hand, the curve representing i's marginal valuation for an additional good is non-increasing. Thus, if we graph the two curves q_i is the maximum value such that $\hat{v}_i(q)$ is at least equal to $p_i(q)$. Note also, that we can express i's expected utility in the more standard form $u_i(q) = \sum_{j=1}^{q_i} \hat{v}_i(j) - P_i$, where $P_i = \sum_{j=1}^{q_i} p_i(j)$ in the discrete case and $P_i = \int_0^{q_i} p_i(q)dq$ in the continuous case.

Using the above definition of an online auction, Lavi and Nisan prove the following theorem:

THEOREM 1. *An online auction is incentive compatible if and only if it is based on supply curves.*

PROOF. If we have an online auction that is based on supply curves, then based on the first condition of the supply curve definition the quantity q_i sold to agent i will be the one that maximizes $\sum_{j=1}^{q} \hat{v}_i(j) - p_i(j)$ which is i's expected utility. Then agent i will have no choice but to disclose its true valuation $v_i(q)$. Thus, the auction is incentive compatible. For the converse, if an online auction is incentive compatible it must be the case that for an agent i its total payment $P_i(q)$ is uniquely determined by q. If not, then supposed i has two different valuations $v_i(q)$, $\hat{v}_i(q)$ that are based on the same allocated quantity, but produce two different total payments P_i, $\hat{P}_i$ respectively. Here $v_i(q)$ is i's true valuation. Assuming $P_i \geq \hat{P}_i$, i will get more utility by declaring the untruthful valuation $\hat{v}_i(q)$. This contradicts the original assumption that the auction is incentive compatible. So given that the total payment is uniquely determined, using the definition of the supply curve we can determine the price for each of the q goods that i receives by using $p_i(j) = p_i(j) - p_i(j-1)$ for $j = 1...q$. This means that the total price and total quantity sold to i is determined by the supply curve. If not, then suppose that q_i is the amount that maximizes i's utility and it is allocate for a bid value of $\hat{v}_i(q)$. Also suppose that for i's true valuation $v_i(q)$ the auction allocates some suboptimal quantity $\hat{q}$. Then i would want to declare the bid value instead of its true valuation. This contradicts the assumption that the auction is incentive compatible. □

Lavi and Nisan derive online mechanism called the Competitive Online Auction (COA) based on the above theorem. They define an online auction as based on a "global supply curve $p_(q)$" it is based on supply curves defined before and if $p_i(q) = p(q + \sum_{j=1}^{i-1} q_j)$, such that q_j is the quantity sold the j-th bidder. This means that the auction uses $p_i(q)$ as the supply curve for the i-th bidder. They present the competitiveness of the COA as the following theorem:

THEOREM 2. *Let* $\phi = \frac{\overline{p}}{\underline{p}}$. *Then the COA is c-competitive in terms of revenue and social efficiency relative to the Vickrey offline auction. Furthermore no other online auction can have a better competitive ratio in terms of revenue and social efficiency. For a randomized auction* $c = \Theta(\ln \phi)$. *For a deterministic auction* $c = \sqrt{\phi}$ *when* $k = 1$, *and* $\phi^{\frac{1}{k+1}} \leq c \leq k\phi^{\frac{1}{k+1}}$ *for* $k \geq 1$.

3.2 Adaptive Online Auction

Hajiaghayi et al. [8] look at a very similar formulation but allow agents to make much broader declarations regarding their valuation $\hat{v}_i$, arrival time $\hat{a}_i$, and departure time $\hat{d}_i$. The only restriction on declaration is that $\hat{a}_i \geq a_i$and $\hat{d}_i \geq \hat{a}_i$. This means that an agent i cannot announce an arrival that is *earlier* than i's actual arrival time. It is obvious that in a real situation this assumption holds because an agent cannot disclose anything to the mechanism if it has not even arrived yet.

The seller derives no utility from keeping a good. Also, an agent i receives no utility from being allocated a good outside of its time interval $[a_i, d_i]$. The valuations are sampled from an independent and identical distribution (i.i.d.), but can be drawn from a distribution that is either known or unknown to the seller. Hajiaghayi et al. basically divides the problem into two cases: (1) for $k = 1$, and (2) $k > 1$. The key contribution of [8] is that it is the first paper to formally address the limited supply online auction problem without assuming anything about the distribution that agent valuations are sampled from, and given this, provide a mechanism that has constant-competitive efficiency and time-SP.

3.2.1 Problem Formulation

The problem formulation is exactly as described in section 2 except with the following additional assumptions. The mechanism knows the number of agents n, and the time-horizon, which is $[0, T]$. Each agent i can demand any number of items, where i is drawn from some unknown, fixed distribution Φ. So i gets its value at time a_i and needs a decision regarding whether the seller will allocate a good to i by time d_i.

They assume *individual rationality*, also known as *voluntary participation*, and *no deficit* for their mechanisms. Voluntary participation states $p_i(\hat{\theta}) = 0$, if $q_i(\hat{\theta}) = 0$ and $p_i(\hat{\theta}) \leq \hat{v}_i$. In other words, each agent i will have a payment of zero if it has not been allocated any goods at the end of the auction. Thus, i's expected utility for participating in the auction is at least equal to i's expected utility outside the auction. No deficit simply means that, $p_i(\hat{\theta}) \geq 0$.

Following the competitive analysis done by Lavi and Nisan [2], Hajiaghayi et al. use the offline Vickrey auction as a benchmark.

They define Vickrey efficiency as

$$Eff_v = \sum_{i \leq k} v^{(i)}. \tag{4}$$

Here $v^{(m)}$ is the m-th highest bid for a copy of the good. Essentially, for efficiency the Vickrey auction attempts to allocate each copy of the good with the agent that has the highest remaining valuation. By remaining we mean the highest valuation given that some agents j may have had a higher evaluation but have either dropped out of the auction or are no longer interested (because they already have the desired number of copies of the good).

3.2.2 Main Results

They solved the $k = 1$ case be dividing into into a number of different situations. In the first situation none of the agents have arrive-departure time intervals that overlap. In other words, there are not two agents that are present at the auction at the same time. In this situation, Hajiaghayi et al. note that the problem is very similar to the well-known secretary problem, where a secretary must interview n different candidates in order to hire the best possible employee. The difficulty arises because the candidates have disjoint arrival-departure times and on top of that the secretary cannot call back candidates after they have left. In other words, the secretary can only choose the current candidate. All previously interviewed candidates are no longer accessible, and all candidates in the future can only be chosen when they arrive. The common difficulty within these two problems is that the mechanism must attempt to choose the optimal agent without having evaluated all the agents.

The general solution to the secretary problem is to interview the initial $t-1$ applicants, record the best "reference candidate" out of this pool, and then hire the first applicant out of the remaining pool that is better than the reference candidate. As noted in [8] as $n \to \infty$ the ratio $\frac{t}{n}$ approaches $\frac{1}{e}$. We offer a simple proof for this result:

PROOF. Consider that the n candidates arrive in a sequence $1, ..., n$. Let us divide the sequence into two pools where the first pool contains m agents that were interviewed and the second pool contains the remaining $n-m$ candidates that have not been interviewed. The secretary can potentially select a suboptimal candidate in two ways: (1) the best candidate was in the first pool, and (2) the best candidate was after an individual that also beat the reference candidate from the first pool. Then the probability for the secretary to succeed is $P = P(m+1$ is the best$) + P(m+2$ is the best$) + P(m+3$ is the best$) + ... + P($n is the best$)$, where $P(m+1) = \frac{1}{n}$, $P(m+2) = \frac{m}{m+1} \times \frac{1}{n}$, $P(m+3) = \frac{m}{m+2} \times \frac{1}{n}$,...,$P(n) = \frac{m}{n-1} \times \frac{1}{n}$. Thus, $P = \frac{m}{n} \times (\frac{1}{m} \times \frac{1}{m+1} \times ... \times \frac{1}{n-1})$. For large n, P converges to $\frac{1}{e} = 0.368$. $\square$

Thus, approximately $\frac{n}{e}$ candidates should be processed before the best one should be considered. For the situation where the arrival-departure times are completely disjoint a mechanism using the above procedure should allow roughly $\frac{n}{e}$ agents to pass before setting a reservation price equal to the highest valuation among those agents. This will result in a SP mechanism because each agent is presented with a price that was not dependent on its valuation and the price does not change while the agent is in the auction. Although the auction setting provides more information than the secretary problem, namely we know the individual bid values of agents in addition to their ordering, Hajiaghayi et al. derived lower bounds on efficiency (2) and revenue (1.5) using the secretary problem solution.

For the multi-item case ($k > 1$), designing a mechanism that is constant competitive as $k \to \infty$ is much harder. Hajiaghayi et al. first point out that it is trivial to achieve an auction in this setting with ek-competitiveness in terms of efficiency and $4k$-competitiveness in terms of revenue by simply discarding $k-1$ goods and performing a single item auction, which degenerates to considering the $k = 1$ case. In the general case they claim to use a modified version of the DSOT auction described by Goldberg et al. in [3]. Note that we did not find an auction in [3] that was denoted as DSOT, so we assume they are using dual-price sampling optimal threshold (DOS) auction. They create a mechanism RM'_k that is c-competitive with the offline Vickrey auction, where $C < 48$. Their proof of its correctness can be found in [8] and is very similar to that of DOS in [3].

4. COMPARING AUCTIONS

In [2] the number of agents is unknown at the beginning and is something the auctioneer knows only at the end. This is in contrast to Hajiaghayi et al. [8] who assume that their mechanism knows the total number of agents that will eventually pass through the system. Lavi and Nisan denote an auction with this assumption as a *partially online* model because it reduces the online nature of the situation. In most real life applications especially electronic commerce the number of participating buyers is unknown. However, the model that Lavi and Nisan present has its own disadvantage because it restricts the strategy space of the agents. Specifically, each agent i is forced to disclose its type to the mechanism at its actual arrival time a_i. Although their mechanism still allows agents to declare false valuations v_i, this is only one of the possible ways that agents in an online setting can be dishonest. Remember that an agent i's reported type also includes $\hat{a}_i$ and $\hat{d}_i$. Hajiaghayi et al. examine a broader strategy space by allowing agents to delay reporting their arrival or declare that they are leaving earlier than they really are. In their closing remarks Lavi and Nisan do attempt to improve their model with a set of modifications that include allowing an agent i to delay its bid for some time $t \geq a_i$, place any number of bids at times following its arrival a_i, or to give i a non-increasing, time-dependent valuation $v_i(q,t)$, which relies on time t in addition to quantity q. They point out that as long as their assumption of a non-decreasing supply curve holds there will be no advantage for an agent to delay its bid or try to make multiple bids. In other words, i will still want to declare its bid as soon as possible, which will be at time a_i upon its arrival. So essentially all of the competitive analysis done by Lavi and Nisan holds.

5. CONCLUSIONS

For Hajiaghayi et al., their primary interest for future work is to derive a mechanism such that it has tighter upper and lower bounds on efficiency and revenue. They also want to examine lower bounds when valuations are drawn from an unknown distribution. An open problem in [2] is to derive constant competitive ratios because their scenario is more difficult to solve since an agent can receive more than 1 item. Finally, Hajiaghayi et al. would also like to consider a generalized version of the problem in [7] where items can be reused (e.g. a wireless connection at an Internet cafe).

Other than the extensions we described in section 4, Lavi and Nison do not present any possible future work to their model. However, a natural extension of the work down by these two groups may be to combine their mechanisms into a more generalized version. In other words, the next step would be to attempt to construct a SP online auction that makes no assumption about the distribution from which the valuations and time intervals are drawn from as in [8], but also allow agents to demand more than one copy of the good [2].

Acknowledgments

I used [11] during the course of writing this paper as a reference for definitions and to get a survey of some mechanisms. As a review paper I found [11] very useful because it summarized a lot of the offline and online mechanisms according to SP conditions. Thus, although the article is a work in progress, I have included it in order to give credit where credit is due.

6. REFERENCES

[1] D. C. Parkes and S. Singh. An mdp-based approach to online mechanism design. In *In Proc. 17th Annual Conference on Neural Information Processing Systems*, 2003.

[2] R. Lavi and N. Nisan. Competitive analysis of incentive compatible on-line auctions. In *EC '00: Proceedings of the 2nd ACM conference on Electronic commerce*, pages 233–241. ACM Press, 2000.

[3] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 735–744. Society for Industrial and Applied Mathematics, 2001.

[4] J. Gallien. Dynamic mechanism design for online commerce. Technical report, Sloan School of Management, MIT, 2002.

[5] Amitabha Bagchi, Amitabh Chaudhary, Rahul Garg, Michael T. Goodrich, and Vijay Kumar. Seller-focused algorithms for online auctioning. In *WADS '01: Proceedings of the 7th International Workshop on Algorithms and Data Structures*, pages 135–147, London, UK, 2001. Springer-Verlag.

[6] B. Awerbuch, Y. Azar, and A. Meyerson. Reducing truth-telling online mechanisms to online optimization. In *Proceedings from thirty-fifth ACM symposium on Theory of computing*, pages 503–510. ACM Press, 2003.

[7] E. J. Friedman and D. C. Parkes. Pricing WiFi at Starbucks: Issues in online mechanism design. In *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pages 240–241. ACM Press, 2003.

[8] M. T. Hajiaghayi, R. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 71–80. ACM Press, 2004.

[9] Ziv Bar-Yossef, Kirsten Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *SODA '02: Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 964–970, Philadelphia, PA, USA, 2002. Society for Industrial and Applied Mathematics.

[10] Y. Shoham and K. Leyton-Brown. Multi agent systems. 2005.

[11] C. Ng. Computational mechanism design: Towards online mechanisms, 2004.