Prompt Mechanisms for Online Auctions

Richard Cole[1]*, Shahar Dobzinski[2] **, and Lisa Fleischer[3]***

[1] Computer Science Department, Courant Institute, New York University
[2] The School of Computer Science and Engineering, the Hebrew University of Jerusalem
[3] Dartmouth College

Abstract. We study the following online problem: at each time unit, one of m identical items is offered for sale. Bidders arrive and depart dynamically, and each bidder is interested in winning one item between his arrival and departure. Our goal is to design truthful mechanisms that maximize the welfare, the sum of the utilities of winning bidders.

We first consider this problem under the assumption that the private information for each bidder is his value for getting an item. In this model constant-competitive mechanisms are known, but we observe that these mechanisms suffer from the following disadvantage: a bidder might learn his payment only when he departs. We argue that these mechanism are essentially unusable, because they impose several seemingly undesirable requirements on any implementation of the mechanisms.

To crystalize these issues, we define the notions of *prompt* and *tardy* mechanisms. We present two prompt mechanisms, one deterministic and the other randomized, that guarantee a constant competitive ratio. We show that our deterministic mechanism is optimal for this setting.

We then study a model in which both the value and the departure time are private information. While in the deterministic setting only a trivial competitive ratio can be guaranteed, we use randomization to obtain a prompt truthful $\Theta(\frac{1}{\log m})$-competitive mechanism. We then show that no truthful randomized mechanism can achieve a ratio better than $\frac{1}{2}$ in this model.

1 Introduction

1.1 Background

The field of algorithmic mechanism design attempts to handle the strategic behavior of selfish agents in a computationally efficient way. To date, most work

* 251 Mercer Street, New York, NY 10012. This work was supported in part by NSF grants IIS0414763 and CCF0515127. cole@cs.nyu.edu.

** Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, by NSF grant CCF-0515127, and by grants from the Israel Science Foundation, the USA-Israel Bi-national Science Foundation. shahard@cs.huji.ac.il.

*** 6211 Sudikoff, Dartmouth College, Hanover, NH 03755. Partially supported by NSF grants CCF-0515127 and CCF-0728869. lkf@cs.dartmouth.edu.

in this field has sought to design truthful mechanisms for *static* settings, e.g., auctions. In reality, however, the setting of many problems is *online*, meaning that the mechanism has no prior information regarding the identity of the participating players, or that the goods that are for sale are unknown in advance. Examples include sponsored search auctions [12], single-good auctions [10], and even pricing WiFi at Starbucks [5].

This paper considers the following online auction problem: at each time unit exactly one of m identical items is offered for sale. The item at time t is called item t. There are n bidders, where bidder i arrives at time a_i and departs at time d_i, both unknown before bidder i's arrival. The interval $[a_i, d_i]$ will be called bidder i's *time window*, and the set of items offered in i's time window will be denoted by W_i. Each bidder is interested in winning at most one of the items within W_i. Let v_i denote the value to the ith bidder of getting an item in W_i. Our goal is to maximize the social welfare: the sum of the values of the bidders that get some item within their time window. As usual in online algorithms, our goal is to optimize the competitive ratio: the worst-case ratio between the welfare achieved by the algorithm and the optimal welfare.

In the full information setting, this problem is equivalent to the online scheduling of unit-length jobs on a single machine to maximize weighted throughput. This online problem and its variants have been widely studied (e.g., $[1, 8, 3]$). The best deterministic algorithm to date guarantees a competitive ratio of ≈ 0.547 [4, 11], while it is known that no deterministic algorithm can obtain a ratio better than $\frac{2}{\sqrt{5}+1} \approx 0.618$ [2]. In the randomized setting, a competitive ratio of $1 - \frac{1}{e}$ is achieved by [1], and no algorithm can achieve a ratio better than 0.8 [2].

This problem provides an excellent example of the extra barriers we face when designing online mechanisms. The only general technique known for designing truthful mechanisms is the VCG payment scheme. In the offline setting we can obtain an optimal solution in polynomial time (with bipartite matching), and then we can apply VCG. In the online setting, however, it is impossible to find an optimal solution, and thus we cannot use VCG. Yet, truthful competitive mechanisms do exist. The competitive ratio of these mechanisms depends on the specific private-information model each mechanism was designed for. This paper considers two different natural models:

- **The Value-Only model:** Here, the private information of bidder i consists of just his value v_i, and the arrival time and the departure time are known to all (but both are unknown prior to the arrival of bidder i).
- **The Generalized Model:** The private information of bidder i consists of two numbers: his value v_i and his departure time d_i. The arrival time is public information (but unknown prior to the arrival of bidder i).

1.2 The Value-Only Model: Is Monotonicity Enough?

The only private information of a bidder in the value-only model is his value, and thus this model falls under the category of *single-parameter* environments – environments in which the private information of each bidder consists of only one number. Fortunately, designing truthful mechanisms for single-parameter environments is quite well understood: an algorithm is truthful if and only if it is monotone. That is, a winning bidder that raises his bid remains a winner.

Using the above characterization, it is possible to prove that the greedy algorithm is monotone [7] (see Section 2.4 for a description). Since [8] shows that greedy is $1/2$ competitive, this gives a truthful mechanism that is $\frac{1}{2}$ competitive.

However, a closer look at this mechanism may make one wonder if it is indeed applicable. The notions of prompt and tardy mechanisms we define next highlight the issue.

Definition 1. *A mechanism for the online auction problem is* prompt *if a bidder that wins an item always learns his payment immediately after winning the item. A mechanism is* tardy *otherwise.*

As we show later in the paper, the tardiness in the greedy mechanism [7, 8] is substantial: there are inputs for which a bidder learns his payment only when he departs. Tardy mechanisms seem very unintuitive for the bidders, and in addition they suffer from the following disadvantages:

- **Uncertainty:** A winning bidder does not know the cost of the item that he won, and thus does not know how much money he still has available. E.g., suppose the mechanism is used in a Las Vegas ticket office for selling tickets to a daily show. A tourist that wins a ticket is uncertain of the price of this privilege, and thus might not be able to determine how much money he has left to spend during his Las Vegas vacation.
- **Debt Collection:** A winning bidder might pay the mechanism long after he won the item. A bidder that is not honest may try to avoid this payment. Thus, the auctioneer must have some way of collecting the payment of a winning bidder.
- **Trusted Auctioneer:** A winning bidder essentially provides the auctioneer with a "blank check" in exchange for the item. Consequently, all bidders must trust the honesty of the auctioneer. Even if the bidders trust the auctioneer, they may still want to verify the exact calculation of the payment, to avoid over-payments that make winning the item less profitable, or even unprofitable. In order to verify this calculation, the bids of all bidders have to be revealed, leading to an undesirable loss of privacy.

Notice that all of these problems are due to the online nature of the setting, and do not arise in the offline setting. To overcome these problems, we present *prompt* mechanisms for the online auction problem. Prompt mechanisms are very intuitive to the bidders as they (implicitly) correspond to take-it-or-leave-it offers: a winning bidder is offered a price for one item exactly once *before* getting the item, and may reject the offer if it is not beneficial for him. We improve upon the greedy algorithm of [7, 8] by showing a different mechanism that achieves the same competitive ratio, but is also prompt.

Theorem: There exists a $\frac{1}{2}$-competitive prompt and truthful mechanism for the online auction problem in the value-only model.

We show that this is the best possible by proving that no prompt deterministic mechanism can guarantee a competitive ratio better than $\frac{1}{2}$.

We also present a *randomized* mechanism that guarantees a constant competitive ratio. The achieved competitive ratio of the latter algorithm is worse than the competitive ratio of the deterministic algorithm. Yet, the core of the proof studies a balls-and-bins problem that might be of independent interest.

1.3 The Generalized Model

While truthful mechanisms for single-parameter settings are well characterized and thus relatively easy to construct, truthful mechanisms for multi-parameter settings, like the generalized model, are much harder to design. The online setting considered in this paper only makes the design of truthful mechanisms a more challenging task.

The online auction problem in the generalized model illustrates this challenge. Lavi and Nisan [9] introduced the online auction problem to the mechanism design community. They showed that no truthful deterministic mechanism for this multi-parameter problem can provide more than a trivial competitive ratio. As a result, Lavi and Nisan proposed a weaker solution concept, set-nash, and provided mechanisms with a constant competitive ratio under this notion. We stress that the set-nash solution concept is much weaker than the dominant-strategy truthfulness we consider.

By contrast with [9], instead of relaxing the solution concept, we use the well-known idea that randomization can help in mechanism-design settings [14]. We provide *randomized* upper and lower bounds in the generalized model for the online auction problem.

Theorem: There exists a prompt truthful randomized $\Theta(\frac{1}{\log m})$-competitive mechanism for the online auction problem in the generalized model.

The main idea of the mechanism is to extend the randomized mechanism for the value-only model to the generalized model. Specifically, we use the random-sampling method introduced in [6] to "guess" the departure time of each bidder, and then we use the above randomized mechanism with these guessed departures. This mechanism is also a prompt mechanism. We notice that it is quite easy to obtain mechanisms with a competitive guarantee of the logarithm of the ratio between the highest and lowest valuations. However, since this ratio might be exponential in the number of items or bidders, this guarantee is quite weak. By contrast, the competitive ratio our mechanism achieves is independent of the ratio between the highest and lowest valuations, and the mechanism is not required to know these valuations in advance.

Theorem: No truthful randomized mechanism for the online auction problem in the generalized model can obtain a competitive ratio better than $\frac{1}{2}$.

The proof of this bound is quite complicated. We start by defining a family of recursively-defined distributions on the input, and then show that no deterministic mechanism can obtain a competitive ratio better than $\frac{1}{2}$ on this family of distributions. We then use Yao's principle to derive the theorem.

The main open question left in the generalized model is to determine whether there is a truthful mechanism with a constant competitive ratio.

Paper Organization In Section 2 we describe prompt mechanisms for the value-only case, and prove that no deterministic tardy algorithms can achieve a ratio better than $\frac{1}{2}$. Due to lack of space, lower and upper bounds for the generalized case are proved only in the full version (see `http://www.cs.huji.ac.il/~shahard`).

2 Prompt Mechanisms and the Value-Only Model

2.1 A Deterministic Prompt $\frac{1}{2}$-Competitive Mechanism

The mechanism maintains a *candidate* bidder c_j for each item j. To keep the presentation simple and without loss of generality, we assume an initialization of the mechanism in which each item j receives a candidate bidder c_j with a value of 0 for winning an item (i.e., $v_{c_j} = 0$).

The mechanism runs as follows: at each time t we look at all the bidders that arrived at time t. We consider these bidders one by one in some arbitrary order (independent of the bids): for each bidder i we look at all the candidates in i's time window, and let c_j be the candidate bidder with the smallest bid (if there are several such candidates, we select one arbitrarily). Formally, $c_j \in \arg\min_{k \in W_i} c_k$. We say that i *competes* on item j. Now, if $v_{c_j} < v_i$, we make i the candidate bidder for item j. After all the bidders that arrived at time t have been processed, we allocate item t to the candidate bidder c_t.

The next theorem proves that this algorithm is monotone, i.e., a bidder that raises his bid is still guaranteed to win. This is also a necessary and sufficient condition for truthfulness. We are still left with the issue of finding the payments themselves. First, observe that the payment of each winning bidder must equal his *critical value*: the minimum value he can declare and still win. Notice that this value is indeed well defined if the algorithm is monotone. For each bidder i this value can be found by using a binary search on the possible values of v_i. Clearly, this procedure takes a polynomial time. See, e.g., [13] for a more thorough discussion. By the discussion above, it is clear that a mechanism is prompt if and only if i's critical value can be found by the time i wins an item. In this case, the payment can also be calculated in polynomial time.

Theorem 1. *The mechanism is prompt and truthful. Its competitive ratio is* $\frac{1}{2}$.

Proof. To show that the mechanism is truthful we have to show that it is monotone: that is, a winning bidder i still wins an item by raising his value v_i to v'_i. First, observe that fixing the declarations of the other bidders, i competes on item j regardless of his value. We now compare two runs of the mechanism, with i declaring v_i and with i declaring v'_i, and show that at each time the candidate for any item j' is the same in both runs. In particular, it follows that the set of winners stays the same, and thus the mechanism is monotone.

First, observe that the two runs are identical until the arrival of i. Look at the next bidder e that arrives after i. For a contradiction, suppose that the candidate for some item changes after bidder e arrives. It follows that i declaring v_i' causes e to compete on an item different than the one that e competes on when i declares v_i. This is possible only if e is competing on j if i declares v_i, but if i declares v_i', e competes on $h \neq j$. It follows that if i declares, v_i' both i and e compete on j, and that i wins j. Thus, $v_i \geq v_e$. When i raises his bid e competes on h. Let c_h be the candidate for h at the time that e arrives. We have that $v_i' > v_{c_h} \geq v_i$, and thus $v_e < v_{c_h}$ so e does not become a candidate on h, and the set of candidates stays the same. To finish the monotonicity proof, look at the rest of the bidders one by one, and repeat the same arguments.

As for the promptness of the mechanism, observe that the identity of the item that i competes on is determined only by the information provided by bidders that had already arrived by the time of i's arrival. The winner of any item j is of course completely determined by the information provided by bidders that arrived by time j. Thus, we can calculate the payment of a winning bidder immediately after he wins an item.

We now analyze the competitive ratio of the mechanism. Let $OPT = (o_1, ..., o_m)$ be the optimal solution, and $ALG = (p_1, ..., p_m)$ be the solution constructed by the mechanism. That is, o_j is the bidder that wins item j in OPT and p_j is the bidder that wins item j in ALG. We will match each bidder i that wins an item in OPT to exactly one bidder l that wins an item in ALG. Furthermore, we will make sure that $v_i \leq v_l$, and that each bidder in ALG is associated with at most two bidders in OPT. This is enough to prove a competitive ratio of $\frac{1}{2}$.

The bidders are matched as follows: for each item j, let $o_{j_1}, \cdots, o_{j_{k_j}}$ be the bidders (ordered by their arrival time) that won an item in the optimal solution and are competing on j. Now match each o_{j_r} to $p_{j_{r+1}}$ for $r < k_j$. Match $o_{j_{k_j}}$ to p_j, the bidder that wins j in ALG (it is possible that $p_j = o_{j_{k_j}}$).

Observe that bidder p_j is associated with at most two bidders that win some item in OPT: bidder $o_{j_{k_j}}$, and at most one bidder, o_{j_i}, that is competing on an item j, where j is the item that o_j $(= o_{j_{i+1}})$ is competing on in ALG. To finish the proof, we only have to show that $v_{o_{j_{k_j}}} \leq v_{p_j}$ and $v_{o_{j_i}} \leq v_{p_j}$. Since $o_{j_{k_j}}$ and p_j both compete for slot j (possibly they are the same bidder) and p_j wins, $v_{o_{j_{k_j}}} \leq v_{p_j}$. Now we show the second claim. When $o_{j_{i+1}}$ arrives, o_{j_i} is already competing on slot j; as $o_{j_{i+1}}$ chooses to compete on slot j rather than slot j' which is also in its interval, thus the current candidate for slot j has value at least $v_{o_{j_i}}$. But the eventual winner of slot j, p_j, can only have a larger value; i.e. $v_{o_{j_i}} \leq v_{p_j}$. □

2.2 A Prompt Randomized Mechanism

We present a randomized prompt $O(1)$-competitive mechanism for the online auction problem in the value-only model. The analysis of the competitive ratio of the mechanism is related to a variant of the following balls-and-bins question:

Balls and Bins (intervals version): n balls are thrown to n bins, where the ith ball is thrown uniformly at random to bins in the interval $W_i = [a_i, d_i]$. We are given that the balls can be placed in a way such that all bins are filled, and each ball i is placed in exactly one bin in $[a_i, d_i]$. What is the expected number of full bins (bins with at least one ball)?

The theorem below proves that, for every valid selection of the a_i's and d_i's, in expectation at least $\frac{1}{10}$ of the bins will be full (notice that in the online auction problem the "balls" have weights). There is a gap between this ratio and the worst example we know: in Subsection 2.3 we present an example in which at most $\frac{11}{24}$ of the bins are full in expectation. Improving the analysis of the balls and bins question will almost immediately imply an improvement in the guaranteed competitive ratio of the mechanism.

The Mechanism

1. When bidder i arrives, assign it to exactly one item in W_i to compete on uniformly at random.
2. At time j conduct a second-price auction on item j among all the bidders that were selected to compete on item j in the first stage.

Theorem 2. *The mechanism is prompt and truthful, and guarantees a competitive ratio of* $\frac{1}{10}$.

Proof. To see that the mechanism is truthful, recall that in the value-only model the arrival time and the departure time of each bidder are public information. It follows that the identity of bidders competing on a certain item is determined only by the outcome of the random coin flips. It is well known that a second-price auction is truthful, and thus we conclude that the mechanism is truthful. Clearly, the mechanism is prompt since the price is determined by the second-price auction which is conducted before allocating the item to the winning bidder.

We now turn to analyzing the competitive ratio of the mechanism. Instead of analyzing this ratio directly, we analyze the competitive ratio of the following process. In addition to the input of the mechanism, the input of the process consists also of "forbidden" sets $S_1 \subseteq W_1, ..., S_n \subseteq W_n$. Later we will see how to construct these sets in a way that guarantees a constant competitive ratio.

1. For each bidder i that *won an item in the optimal solution*, select exactly one item j in W_i to compete on uniformly at random. If $j \in S_i$ then bidder i is not competing on any item at all.
2. At time j allocate item j to one bidder i, where bidder i is selected uniformly at random from the set of all bidders that are competing on item j.

We will compare runs of the mechanism and the process in which the same random coins are used in Step 1. We argue that the competitive ratio of the mechanism is at least as good as the competitive ratio of the process. To see this, observe that in the first step we are restricting ourselves only to bidders that won an item in the optimal solution. Furthermore, some of these bidders are eventually not competing on any item at all. Also, the bidder that is assigned item j is selected uniformly at random from the set of the bidders that are competing on item j, while in Step 2 of the mechanism the bidder with the highest valuation is assigned item j. Obviously, the mechanism does at least as well as the process. We will need the following technical lemma:

Lemma 1. *Let C_j be the random variable that denotes the number of bidders competing on item j (the* congestion *of item j). Let $U_{i,j}$ be the random variable that gets the value of the utility of bidder i from winning item j (that is, v_i if bidder i wins item j, and 0 otherwise). Then,*

$$E[U_{i,j}|i \textit{ is competing on item } j] \geq \frac{v_i}{E[C_j]+1}$$

Proof. We start by bounding from above $E[C_j|i$ is competing on item $j]$. That is, the expected congestion of item j given that bidder i is competing on j. Notice that the expected congestion produced by all other bidders apart from bidder i cannot exceed $E[C_j]$, since the item chosen for each bidder to compete on is selected independently. We are given that bidder i is already competing on item j, and thus we conclude that $E[C_j|i$ is competing on item $j] \leq E[C_j]+1$.

We now prove the main part of the lemma. Notice that $E[U_{i,j}|i$ is competing on item $j] =$ $\Pr[i$ won item $j|i$ is competing on item $j] \cdot v_i$. Let E denote the set of all coin flips in which bidder i is competing on item j (observe that each event $e \in E$ occurs with equal probability). Let $n_j(e)$ be the congestion of item j in $e \in E$.

$$E[U_{i,j}|i \text{ is competing on item } j] = \Sigma_{e \in E} \frac{v_i}{|E| \cdot n_j(e)} \geq \frac{v_i}{E[C_j]+1}$$

where the first equality is by the definition of expectation, and the second inequality is by the convexity of the function $\frac{1}{x}$, and Jensen's inequality. □

As is evident from the lemma, if the expected congestion of all items that are in bidder i's time window is $O(1)$, then bidder i's expected utility is $\Theta(v_i)$. Unfortunately, it is quite easy to construct instances in which for every i, $S_i = \emptyset$ and some items face super-constant congestion. Instead, we will specify for each bidder i a set of items S_i, of size at most half of the size of his time window. We will see that by a proper choice of the S_i's the expected congestion of *every* item is bounded by 4.

Then, as each bidder i (that participates in the optimal solution) has a probability of at least one half of competing on some item, by Lemma 1 bidder i recovers in expectation at least $\frac{1}{2} \cdot \frac{1}{E[C_j]+1}$ of his value; by Lemma 2 this bidder receives in expectation at least $\frac{1}{10}$ of his value. Using the linearity of expectation, we conclude that the mechanism is $\frac{1}{10}$-competitive.

Lemma 2. *There exist sets $S_1, ..., S_n$ such that for each bidder i (that wins an item in the optimal solution), $S_i \subseteq W_i$, and $|S_i| \leq \frac{|W_i|}{2}$, and for each item j, $E[C_j] \leq 4$.*

Proof. The proof of the lemma consists of m stages. In each step we will consider bidders with time windows of length exactly t, where t will take values in descending order from m to 1. We will show for each bidder i with $|W_i| = t$ how to construct his set S_i. By the end of each step, we will be guaranteed that if $|W_i| \geq t$, then for each item in $W_i \setminus S_i$, the expected congestion is at most 4.

We start by handling the case where $t \geq \frac{m}{2}$. Fix some bidder i with $|W_i| \geq \frac{m}{2}$. We are considering only bidders that get an item in the optimal solution, and since there are m items, we need to take into account at most m bidders. Observe that since $W_i \geq \frac{m}{2}$, the average expected congestion of an item in W_i cannot exceed 2. We let S_i be the set of all items in W_i for which the expected congestion is at most 4. By simple Markov arguments, $|S_i| \leq \frac{|W_i|}{2}$. We now have that for every bidder i with $|W_i| \geq \frac{m}{2}$, and for each $j \in W_i \setminus S_i$, $E[C_j] \leq 4$.

Consider now Step t, where $t < \frac{m}{2}$. We first consider the congestion due to bidders with time windows of length at most t. Then we will see that our analysis remains almost the same when including bidders with larger time windows.

Fix some bidder i with $|W_i| = t$. We now bound from above the total congestion of the items in W_i. In the optimal solution, there are at most t bidders that won an item in W_i. Their contribution to the congestion of W_i is bounded from above by assuming that each one is competing on items in W_i time window with probability 1. Hence, the total contribution of these bidders is at most t.

Consider the bidders that won one item j, $a_i - t \leq j \leq a_i - 1$, in the optimal solution. (Our analysis will only improve if $a_i - t \leq 0$.) Clearly, if bidder b won item j in the optimal solution, then that item j is within b's time window. Since a bidder is selected to compete on an item uniformly, it is easy to verify that his contribution to the expected congestion of W_i is maximized when his arrival time is j and his departure time is $j + t - 1$. (Recall that we are only considering bidders with time window of size at most t.) In this case, his contribution to the expected congestion of W_i is $\frac{j+t-a_i}{t}$. Summing over all bidders (with time windows of size at most t) that won one item j, $a_i - t \leq j \leq a_i - 1$, we get that the total contribution of these bidders is at most $\frac{t}{2}$.

Similarly, the total contribution of bidders with time windows of size at most t that won items $d_i + 1$ to $d_i + t$ in the optimal solution is at most $\frac{t}{2}$. It is easy to see that all other bidders with time windows of at most t contribute nothing to the expected congestion of items in W_i. In total, we get that the total expected congestion of items in W_i (due to bidders with time window of length at most t) is at most $\frac{t}{2} + \frac{t}{2} + t = 2t$, and thus the average expected congestion due to these items is at most 2.

As before, we let S_i be the set of all items in W_i for which the expected congestion is at most 4. Again, standard Markov arguments assure that $|S_i| \leq \frac{|W_i|}{2}$. We now have that for every bidder i with $|W_i| = t$, and for each $j \in W_i \setminus S_i$, the average expected congestion incurred by bidders with time windows of size at most t is at most 4. We still need to take into account the congestion incurred by bidders with time windows larger than t. Here we observe that by our construction of the S_i's, these bidders can only contribute to the congestion of items with an expected congestion of at most 4. Therefore, we claim that for each bidder i with $|W_i| \geq t$, and $j \in W_i \setminus S_i$, we have that $E[C_j] \leq 4$. We finish the proof of the lemma by considering smaller values of t, down to $t = 1$. $\square$

$\square$

2.3 A Bad Example

The following example shows that the mechanism presented has a competitive ratio strictly worse than $\frac{1}{2}$. The example is an instance of the balls and bins question presented earlier. For $1 \leq i \leq \frac{n}{3}$, we let $W_i = [i, \frac{2n}{3}]$. For $\frac{n}{3} < i \leq n$, we let $W_i = [\frac{n}{3}+1, i]$. The probability that bin i in $[1, \frac{n}{3}]$ will be empty is:

$$\Pr[\text{no ball falls in bin } i \in [1, \tfrac{n}{3}]\,] = \Pi_{t=1}^{i} \Pr[\text{ball } t \text{ does not fall to bin } i \in [1, \tfrac{n}{3}]\,]$$

$$= \Pi_{t=1}^{i}(1 - \frac{1}{\frac{2n}{3} - t + 1}) = \Pi_{t=1}^{i}(\frac{\frac{2n}{3} - t}{\frac{2n}{3} - t + 1}) = \frac{\frac{2n}{3} - i}{\frac{2n}{3}}$$

We now calculate the expected number of empty bins in the range $[1, \frac{n}{3}]$. Observe that the probability of bin $i \in [1, \frac{n}{3}]$ to be equal to the probability of bin $n-i+1$. Thus, the expected number of empty bins in $[1, \frac{n}{3}]$ is equal to the expected number of empty bins in $[\frac{2n}{3}, n]$:

$$\sum_{t=1}^{\frac{n}{3}} \frac{\frac{2n}{3} - t}{\frac{2n}{3}} = \frac{\frac{n}{3}(\frac{2n}{3} - 1 + \frac{n}{3})}{2 \cdot \frac{2n}{3}} = \frac{n-3}{4}$$

Next we handle bins in the range $[\frac{n}{3}, \frac{2n}{3}]$. By reasoning similar to the previous calculations, the probability that no ball i, $\frac{n}{3} \leq i \leq \frac{2n}{3}$, falls into bin t in this range is $\frac{t - \frac{n}{3} - 1}{\frac{n}{3}}$. The probability that no ball i, $1 \leq i \leq \frac{n}{3}$ falls in bin t is $\Pi_{j=1}^{\frac{n}{3}}(1 - \frac{1}{\frac{2n}{3} - i + 1}) = \frac{\frac{n}{3}}{\frac{2n}{3}} = \frac{1}{2}$. Similarly, the probability that no ball i, $\frac{2n}{3} \leq i \leq n$ falls in bin t is $\frac{1}{2}$. Thus, with probability $\frac{t - \frac{n}{3}}{\frac{n}{3}} \cdot \frac{1}{4}$ no ball falls into bin t, $\frac{n}{3} \leq t \leq \frac{2n}{3}$.

To conclude, the expected number of empty bins in the ranges $[1, \frac{n}{3}]$ and $[\frac{2n}{3}, n]$ together is $\approx \frac{n}{2}$. The expected number of empty bins in $[\frac{n}{3}, \frac{2n}{3}]$ is $\Sigma_{t=\frac{n}{3}}^{\frac{2n}{3}} \frac{t - \frac{n}{3}}{\frac{n}{3}} \cdot \frac{1}{4} \approx \frac{1}{8} \cdot \frac{n}{3}$. In total, about $\frac{13}{24}$ of the bins are empty in expectation. We note that this constant can be somewhat increased to $\frac{4}{7}$ by recursively applying this construction on balls in the middle third (and keeping the other balls' time windows the same). Details are omitted from this extended abstract.

2.4 Limitations of Deterministic Tardy Mechanisms

Here we show that the prompt mechanism of Section 2.1 is optimal. In order to develop some intuition about tardy mechanisms, we start by showing that the greedy mechanism of [7] is tardy.

Recall that the greedy mechanism allocates item t to the bidder with the highest valuation that is present at time t (and has not been assigned any item yet). Consider the following example: two bidders, red and green, arrive at time 1. The red bidder has a value of 10 for winning an item, and his departure time is 5. The green bidder has a value of 6 and a departure time of 1. We consider two scenarios: in the first one, four bidders arrive at time 2, each of them with value 100 and a departure time of 5. In the second scenario, there are no more arrivals. Observe that the greedy mechanism assigns the red bidder the first item. To see that the red bidder cannot learn his payment immediately, recall the following characterization of the payment in single-parameter mechanisms: the payment of a winning bidder is equal to the minimum value he can bid and still win.

In order to win an item in the first scenario, the red bidder must declare a value of at least 6, and therefore this is his payment in this scenario. However, in the second scenario a declaration of 0 will make him win the second item. The mechanism cannot distinguish between the two scenarios when the red bidder wins at time 1, and thus cannot determine the payment at time 1. We conclude that the greedy mechanism is tardy.

The following proposition shows that every prompt deterministic mechanism for the online auction problem achieves a competitive ratio of no better than $\frac{1}{2}$.

Proposition 1. *Every prompt deterministic mechanism for online auctions (even in the value-only model) has a competitive ratio of no better than* $\frac{1}{2}$*.*

Proof. Consider the following setting: two bidders arrive at time 1, each having a value of 1, and a departure time of time 2. Suppose there are no more arrivals of other bidders. Any mechanism that achieves a competitive ratio better than 2 must assign one bidder the first item, and the other item to the second bidder. Let a be bidder that was assigned the first item, and b be the bidder that was assigned the second item.

Claim. Let M be a prompt mechanism with a finite competitive ratio. In the scenario described above, there is no declaration of a value v_b that makes bidder b win the first item.

Proof. Let P_b denote the payment of bidder b for winning the second item with a declaration of 1. Observe that $p_b < 1$[4]. We consider two cases, one in which b declares a value of $w > 1$, and one in which b declares a value of $w < 1$.

Suppose that bidder b raises his bid from 1 to w, and was assigned the first item. The mechanism is prompt, so the payment of bidder b is determined immediately. Suppose, for a contradiction, that this payment is higher than p_b. In this case, if bidder b's true value was w, he could improve his profit by declaring a value of 1, and be assigned the second item. Hence the payment must be at most p_b. Clearly, the payment can not be strictly less than p_b, since otherwise if b's true value is 1, he has an incentive to declare a value of w and increase his profit. Thus the payment must be equal to p_b, but now we will see that this cannot be

[4] If p_i is equal to 1, we add some "noise" to the value to get a strict inequality.

the case. Consider the following setting: b's true value is 1, and therefore he does not win the first item. At time 2 a bidder c with value $w' >> w$ arrives. Bidder c is going to depart immediately. In order to maintain a finite competitive ratio the mechanism must assign bidder c the second item. Thus, if bidder b's true value is 1, he has an incentive to declare a value of w (and therefore win the first item for a payment of p_b), and the mechanism is not truthful.

The other case is where b bids a value w, $w < 1$, and thereby wins the first item (with payment less than 1). As before, if a bidder c with a departure time of 2 and a very high value arrives at time 2, then the mechanism must assign c the second item in order to guarantee a finite competitive ratio. If bidder b's true value is 1, he has an incentive to declare w instead, and win the first item. $\square$

Now alter the scenario described above, and let b's value be $w >> 1$. By the claim, bidder b will not be assigned the first item. However, if at time 2 bidder c with a departure time of 2 and a value of w arrives, the total welfare the mechanism achieves is at most $1 + w$, while the optimal welfare is $2 \cdot w$. $\square$

References

1. Yair Bartal, Francis Y. L. Chin, Marek Chrobak, Stanley P. Y. Fung, Wojciech Jawor, Ron Lavi, Jiří Sgall, and Tomáš Tichý. Online competitive algorithms for maximizing weighted throughput of unit jobs. In *STACS*, pages 187–198, 2004.
2. Francis Y. L. Chin and Stanley P. Y. Fung. Online scheduling with partial job values: Does timesharing or randomization help? *Algorithmica*, 37(3):149–164, 2003.
3. Marek Chrobak, Wojciech Jawor, Jiří Sgall, and Tomáš Tichý. Improved online algorithms for buffer management in QoS switches. In *ESA*, pages 204–215, 2004.
4. Matthias Englert and Matthias Westermann. Considering suppressed packets improves buffer management. In *SODA'07*.
5. Eric J. Friedman and David C. Parkes. Pricing wifi at starbucks: issues in online mechanism design. In *EC '03*.
6. Andrew V. Goldberg, Jason D. Hartline, Anna R. Karlin, Mike Saks, and Andrew Wright. Competitive auctions. *Games and Economic Behavior*, 2006.
7. Mohammad T. Hajiaghayi, Robert Kleinberg, Mohammad Mahdian, and David C. Parkes. Adaptive limited-supply online auctions. In *EC'05*.
8. A. Kesselman, Zvi Lotker, Yishay Mansour, Boaz Patt-Shamir, Baruch Schieber, and Maxim Sviridenko. Buffer overflow management in QoS switches. In *STOC*, pages 520–529, 2001.
9. Ron Lavi and Noam Nisan. Online ascending auctions for gradually expiring items. In *SODA'05*.
10. Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, pages 233–241, 2000.
11. Fei Li, Jay Sethuraman, and Clifford Stein. Better online buffer management. In *SODA'07*.
12. Mohammad Mahdian and Amin Saberi. Multi-unit auctions with unknown supply. In *EC '06*.
13. Ahuva Mu'alem and Noam Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI-02*, 2002.
14. Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, 1999.