

Clairvoyant Mechanisms for Online Auctions

Philipp Brandes¹, Zengfeng Huang², Hsin-Hao Su³, and Roger Wattenhofer¹

¹ ETH Zurich, Switzerland
{pbrandes,wattenhofer}@ethz.ch

² UNSW, Australia,
zengfeng.huang@unsw.edu.au

³ MIT, USA
hsinhao@csail.mit.edu

Abstract. In this paper we consider online auctions with buyback; a form of auctions where bidders arrive sequentially and the bidders have to be accepted or rejected immediately. Each bidder has a valuation for being allocated the good and a preemption price. Sold goods can be bought back from the bidders for a preemption price. We allow unbounded valuations and preemption prices independent from each other. We study the clairvoyant model, a model sitting between the traditional offline and online models. In the clairvoyant model, a sequence of all potential customers (their bids and compensations) is known in advance to the seller, but the seller does not know when the sequence stops. In the case of a single good, we present an algorithm for computing the difficulty Δ , the optimal ratio between the clairvoyant mechanism and the pure offline mechanism (which knows when the sequence stops, and can simply sell the good to the customer with the highest bid, without having to pay any compensations). We also present an optimal clairvoyant mechanism if there are multiple goods to be sold. If the number of goods is unbounded, however, we show that the problem in the clairvoyant model becomes \mathcal{NP} -hard. Based on our results in the clairvoyant model, we study the Δ -online problem (where the sequence is unknown to the mechanism, but the difficulty Δ of the input sequence is known). We show that there is a tight gap of $\Theta(\Delta^5)$ between the offline and the online model.

1 Introduction

Traditional auctions have a rich theory but only make sense in the presence of at least two bidders. In reality, however, many auctions have a rather low demand, and bidders do not compete concurrently. Instead, bidders appear *online*, one after the other.

A familiar example is booking a seat in an airplane. Prices for a flight fluctuate over time, a known pattern is that seats become more expensive as a flight fills up, because the airline starts to learn that there is demand for the flight. Selling seats in an airplane is not a traditional auction since customers are not bidding against each other. Rather, potential customers check the price well in

advance of a flight. If the price is right, they book a seat, sealing the deal with the airline. Airlines generally try to marginally overbook flights, i.e., they sell more tickets than available, assuming that not all customers will actually show up at the gate. Sometimes there are more customers than seats, and the airline must get some customers off the plane. This is usually achieved by having them fly later and giving them some cash as compensation. We believe that such compensations are easily covered by the high premium of late customers.⁴

In this paper we analyze these online auctions. Our bidders come in an online fashion and name their price for a good. The seller can choose to sell the good for that price, or not sell the good (and hope for a better bid to come in later). Bidder and seller also establish a compensation, in case the good is sold to the customer but the deal is later canceled (in the case of a better bidder showing up, worth paying the compensation). These online auctions need two ingredients: First, a good with a price that may fluctuate over time. Second, customers which want to receive the good (or a reservation for the good) quickly. In particular, the time between the arrivals of two customers should generally be larger than the time a customer is willing to wait for the outcome of her bid. In this case online auctions seem to be a better suitable model than traditional auctions. We believe that such online auctions happen often in practice. Booking flights is the running example in this paper, but there are plenty of other examples. Selling ad slots on web pages is a popular one. Since the number of page views is not known beforehand, some sold slots might not be served and thus those slots need to be bought back. More examples are real estate sales, selling network services with quality of service guarantees, or concert tickets.

A simple example will show that online auctions become academically interesting for a worst case analysis only if reasonable compensations are present: Let us assume that a first customer offers a low price but a prohibitively high compensation. If the seller accepts the deal, a next customer offering a much higher price will show up. On the other hand, if the seller does not accept the deal, no other customer will show up. No matter how the seller decides regarding the first customer, the mistake could be devastating.

The starting point for our analysis is what we call the *clairvoyant model*, a hybrid online/offline model. In the clairvoyant model, a sequence of all potential customers (their bids and compensations) is known in advance to the seller, but the seller does not know when the sequence stops, i.e., who the last customer of the sequence is. No matter who the last customer is, the seller wants to do a good job, i.e., the seller wants to sell the good to a customer with a high bid *and* keep compensations that accumulated so far low. It turns out that the clairvoyant model is a stepping stone for a deeper understanding of online auctions, sitting

⁴ In reality, airlines do not implement online auctions in the clean form described in this paper. Airlines do not seem to maximize their profits with this mechanism, probably for psychological reasons. As such, on web pages, flights still can be sold out, instead of just asking for a higher and higher premium for an unexpectedly popular flight.

nically between the pure online and offline models. It introduces a novel technique for analyzing online auctions from a theoretical point of view.

Our contributions are as follows: After introducing the clairvoyant model, we present an optimal mechanism for it in the case of a single good. The result of that mechanism is a factor Δ worse than a pure offline mechanism (that knows when the sequence stops, and can simply sell the good to the customer with the highest bid, without having to pay any compensations). In other words, the parameter Δ tells us how nasty the compensations are. It directly tells us the difficulty of an input sequence. If compensations are minimal (just return the money to canceled customers), then we have by definition $\Delta = 1$. We also show an optimal clairvoyant mechanism if there are multiple goods to be sold. If the number of goods is unbounded, however, we prove that the clairvoyant model becomes \mathcal{NP} -hard. Based on the results in the clairvoyant model, we study the pure online problem (where the sequence is unknown to the mechanism) in a deterministic setting. If Δ is known, we show that there is a tight gap of $\Theta(\Delta^5)$ between the online and the offline model.

2 Related Work

There has been a lot of research of traditional (“offline”) auctions, inspired by the seminal papers of Vickrey, Clarke, and Groves (“VCG”) [5, 12, 27]. They introduce the notion of truthfulness, which means that no bidder has an advantage if she is not telling the truth about her valuation. There is a large amount of work on traditional auctions, for an overview see, e.g., Nisan, Roughgarden, Tardos, and Vazirani [23].

Online mechanisms have been introduced in [8, 19]. In those online mechanisms, the bidders have an arrival and departure time and a valuation for the good. It is assumed that the good expires after a certain period of time, and that a replacement becomes available. In this setting, it was shown that something similar to VCG style second price auctions is still a viable allocation strategy. The initial motivation behind these kind of online auctions is the WiFi at Starbucks [8]. Customers arrive and then depart some time later with each customer having a valuation for the WiFi. Many papers on online mechanisms mainly focus on truthfulness or other incentive compatible solution concepts, e.g., [13, 20, 24, 25]. An overview of online auctions can be found in [23].

Somewhat related to our online auctions are not even auctions, but the secretary problem [21]. In the classic setting one employer interviews n secretaries, with the goal to hire the best secretary. The employer has to decide right after an interview whether to hire or discard a secretary. Unlike our model, previous decisions cannot be recalled. If secretaries are interviewed in random order, it has been shown that the optimal strategy is to first interview n/e secretaries, and then simply hire the first secretary that is better than all previously interviewed secretaries [21]. It has also been shown that, if the input is adversarial (as in our work), the situation is hopeless; the best strategy is to just hire a random secretary, without any interview process [11]. This setting has been adapted to

the online auctions in [14]. Instead of secretaries, there are buyers and instead of a job there is a single indivisible good. They present a mechanism that is, if the buyers appear in random order – as in the original problem – $e+o(1)$ -competitive for efficiency and $e^2+o(1)$ competitive for revenue. Since we have the possibility to cancel previous decisions with financial compensations, our model allows more freedom.

The work closest to ours considers online auctions with buyback, introduced independently by Babaioff et al. and Constantin et al. [3, 6]. Both limit the preemption price (paid to reacquire the good) to a constant fraction of the valuation v of a bidder and this fraction is independent of the individual bidder. Lower and upper bounds for deterministic and randomized algorithms depending on the fraction of the preemption price are presented in their work. Our work allows arbitrary values for the preemption price (that can depend on the specific customer) and we analyze how to deal with this very heterogeneous set of customers. This kind of auction is not truthful since a buyer can overstate her preemption price and thus gain if her good is bought back [6]. In [2] the goods cannot be allocated to any subset of bidders, but bidders form a matroid, This is extended to an intersection of matroids in [1], while still limiting the buyback factor. The concept of buyback has also been applied to the knapsack problem [3, 15, 18] where the goods appear in an online fashion and can be removed later on from the knapsack. Buyback is also used in scheduling with eviction [9].

Online algorithms often face two different types of problems: First, they do not know the future, and second, they have to deal with past mistakes. Hartline and Sharp [16, 17] formalized the two types of problems. When problems are analyzed in this framework, they are called incremental problems. This approach has been applied to various problems, e.g., to maximum flow, online median, facility location, and clustering [4, 7, 22, 26]. Our setting is different as we can potentially fix past mistakes with compensations. Nevertheless, our clairvoyant analysis is a relative of incremental problems.

3 Model

We consider an online auction. There are r indivisible and identical *goods*. Each bidder b_i is willing to buy exactly one good, and has a *valuation* v_i for being allocated a good. The bidders arrive one after another; whether to allocate a good to a bidder must be decided immediately. Bidders that are not allocated a good cannot be recalled, but bidders that are allocated a good can be recalled. A recalled bidder b_i is willing to return her good if she receives adequate compensation. We call the value *preemption price*, which is paid if the good is *bought back*. The preemption price of bidder b_i is denoted by π_i . In summary, bidder b_i is fully specified by $b_i = (v_i, \pi_i)$. Neither v_i nor π_i are bounded, any value in \mathbb{R}^+ is allowed. We assume that the input sequence of bidders b_1, \dots, b_n is created in advance by an adversary who knows the mechanism that is used to allocate the goods. As described above, if the good of a bidder b_i is bought back, the mechanism has to pay the preemption price. For now, we assume that

the mechanism retains the initial valuation v_i of the bidder. We denote this the *retaining* model. In this model we assume that $v_i \leq \pi_i$ for every bidder b_i . We will show later that this is not necessary and in fact use the model when the value is not retained, which is called the *non-retaining* model.

Let us concentrate on the case of a single good ($r = 1$). Let $\text{offline}(\ell)$ denote the highest valuation of the first ℓ bidders, i.e., $\text{offline}(\ell) = \max_{1 \leq i \leq \ell} v_i$. Since the pure offline mechanism knows the whole input sequence and when it stops, it can sell the good just to one single bidder, the bidder with the highest valuation.

As discussed in the introduction, the online mechanism cannot be competitive with the offline model. Essentially, an online mechanism has to deal with two different issues: First, it does not know the future, and second, it needs to offer a solution at all times. We will now introduce the *clairvoyant* model, a model between pure online and offline. The clairvoyant model knows the whole sequence b_1, \dots, b_n of future potential bidders, but does not know when the sequence stops, i.e., who the last bidder of the sequence is. Because of this, a clairvoyant mechanism must offer a solution at all times.

Both pure online and clairvoyant mechanisms may need to accept more than one bidder (and hence buy the good back). Let S be the set of all bidders that have been accepted during the course of a mechanism and let $[\ell]$ denote the set of the first ℓ bidders, i.e., $\{b_1, \dots, b_\ell\}$. We define $\text{gain}(S, \ell) = \sum_{b_i \in S \cap [\ell]} (v_i - \pi_i) + \max_{b_i \in S \cap [\ell]} \pi_i$. It is the sum of valuations of bidders in S up to bidder b_ℓ , minus the preemption prices for the bidders whose good were bought back. Since we retain the value of a bidder, we have $v_i \leq \pi_i$ for every bidder b_i and thus the bidder with the highest preemption price is also the last accepted bidder.

Since the mechanism does not know when the input sequence stops and it thus can stop anytime, we evaluate any mechanism in its worst round. Specifically, given S , the *gain competitiveness* is defined to be $\max_{1 \leq \ell \leq n} \frac{\text{offline}(\ell)}{\text{gain}(S, \ell)}$. If we now minimize this over the best mechanism (the set S of accepted bidders), we get the optimal gain competitiveness

$$\Delta = \min_S \max_{1 \leq \ell \leq n} \frac{\text{offline}(\ell)}{\text{gain}(S, \ell)}.$$

This can be interpreted as the difficulty of the input sequence. In other words, our mechanisms are evaluated in their worst round, i.e., the round in which it has the highest competitive ratio compared to the pure offline mechanism. This forces our mechanisms into accepting bidders early, and possibly repeatedly, thus paying preemption prices repeatedly. The task is to design mechanisms that choose a set S and thereby allocate the goods to the bidders minimizing gain competitiveness.

We will clarify the terms defined above by presenting a simple example. Let the input sequence be $(1, 2), (4, 100), (50, 60)$. A pure offline mechanism will accept $b_3 = (50, 60)$ since this is the bidder with the highest valuation. A clairvoyant mechanism must always accept the first bidder since it could also be the last one. Assume that it also accepts the third bidder. We now calculate the gain

competitiveness for this set as

$$\max \left\{ \frac{\text{offline}(1)}{v_1} = \frac{1}{1}, \frac{\text{offline}(2)}{v_1} = \frac{4}{1}, \frac{\text{offline}(3)}{v_1 + v_3 - \pi_1} = \frac{50}{1 + 50 - 2} \right\} = 4.$$

Note that this is also optimal since accepting bidder b_2 prevents the mechanism from choosing b_3 , hence $\Delta = 4$. This gives us a theoretical insight on the input sequence. No online mechanism could have done better. As explained, the clairvoyant model sits between pure offline and online models. It turns out that it is comparable to both pure models, even though the pure models are not comparable to each other.

4 Auctioning Off a Single Good

We start our analysis by considering the special case of just a single good being sold, i.e., $r = 1$.

4.1 Clairvoyant Mechanism

We now present a mechanism that optimally solves the clairvoyant model, giving us insights into what is possible for an online mechanism.

Theorem 1. *There exists a clairvoyant mechanism that calculates the set of bidders that should be accepted to solve the online auction for one good optimally, i.e., it calculates Δ . If the inputs are integers, its runtime is polynomial; otherwise it is a FPTAS.*

We now formalize and extend the impossibility result from the introduction. Due to space limitations, the proofs have been moved to the full version.

- Lemma 1.** (1) *The value of Δ depends on the input sequence and is unbounded.*
(2) *The gain competitiveness of the pure online mechanism is unbounded and independent of Δ .*
(3) *No randomized online mechanism can achieve bounded gain competitiveness if the number r of items is in $o(n)$, i.e., $r \in o(n)$.*

4.2 Bounded Preemption Prices

The impossibility results from the introduction and the previous section exploited that the preemption price could be arbitrarily large. Thus, in the following we restrict the previously arbitrarily large preemption prices to be at most ρ times as large as the valuation, i.e., $\rho \geq \frac{\pi_i}{v_i}$ for all $1 \leq i \leq n$. Intuitively, this can either be seen as a simple, reasonable constraint for the customers. If someone values a seat on an airplane with some value v , then losing this seat should not be arbitrarily larger than v . One could also model this scenario in such a way that every customer also has to buy an insurance whose compensation depends on the premium. If she loses her seat, then the insurance will pay her the preemption

price. Now the price of the insurance is closely related to the preemption price. This interpretation also guarantees us that at most a factor of ρ between v_i and π_i for every bidder b_i . The following results resemble closely those in [3, 6]. The factor ρ allows us to design a mechanism that is 4ρ gain competitive. It accepts a bidder if her valuation is at least by a factor 2 larger than the preemption price of the bidder that is currently allocated the good.

Theorem 2. *There exists a mechanism that has 4ρ factor gain competitiveness.*

Corollary 1. *If $\rho \geq \frac{\pi_i}{v_i}$ for every bidder b_i , then $\Delta \leq 4\rho$.*

4.3 Online Mechanism with Δ

This raises the question whether restricting the preemption price is the only way to go. We already know that Δ contains valuable information about the input sequence. But does it contain all the necessary information for an online mechanism to be competitive? We now provide the mechanisms with this information and denote them Δ -online mechanisms. These more powerful online mechanisms can achieve a $\mathcal{O}(\Delta^5)$ factor approximation of the clairvoyant mechanisms. Note that this information is not as strong as knowing that the preemption price of every bidder is at most a factor of ρ larger. The clairvoyant mechanism might accept someone whose preemption price is much larger than its valuation. We briefly describe the mechanism. Simply put, this mechanism accepts bidders with a sufficiently small preemption price (and a high enough valuation to pay back the last bidder). Furthermore, it also accepts bidders that have such a high valuation that the clairvoyant mechanism also had to accept it.

We denote the current bidder with $b = (v, \pi)$. We call the last accepted bidder $b^* = (v^*, \pi^*)$. The online mechanism accepts the first bidder for sure, so initially $b^* = (v_1, \pi_1)$. After the first bidder, the current bidder b is accepted for two different reasons: We call bidders *good* if $\pi \leq 2\Delta^2 v$; if a bidder is not good, it is *bad*. The mechanism will accept a good bidder if its valuation $v > 2\pi^*$. We call bidders *crucial* if $v > 2\Delta v^{**}$, where $v^{**} \geq v^*$ is the largest valuation seen so far. The mechanism will accept a crucial bidder if its valuation $v > \pi^*/(1 - \frac{1}{\Delta^2})$. The pseudocode is shown in Algorithm 1.

In this section a Δ -online mechanism is presented that is $\mathcal{O}(\Delta^5)$ competitive. But first, we need some additional notation.

Theorem 3. *Given the value of Δ , there exists a mechanism that has gain competitiveness $\mathcal{O}(\Delta^5)$ compared to the offline solution.*

Proof. Notice that the clairvoyant mechanism will accept every crucial bidder. Let $\bar{b}_1 = (\bar{v}_1, \bar{\pi}_1), \bar{b}_2 = (\bar{v}_2, \bar{\pi}_2), \dots$ be the subsequence of bidders who are crucial, and let $\bar{b}_0 = b_1$ be the very first bidder, who will also be accepted by the clairvoyant mechanism. We will prove the theorem by induction over the crucial bidders. Our induction hypothesis is that before \bar{b}_i came, the gain competitiveness of the mechanism is at most $8\Delta^5$, we then prove that before \bar{b}_{i+1} came, the

ALGORITHM 1: A Δ -online mechanism

```
accept the first bidder and set  $(v^*, \pi^*) = (v_1, \pi_1)$  and  $v^{**} = v_1$  ;  
while there is a new bidder  $b_i$  do  
    if  $\pi_i \leq 2\Delta^2 v_i$  and  $v_i > 2\pi^*$  then  
        buy good back and give it to bidder  $b_i$ ;  
         $\pi^* \leftarrow \pi_i$  and  $v^* \leftarrow v_i$  ;  
    end  
    else if  $v_i \geq 2\Delta v^{**}$  and  $v_i > \pi^*/(1 - \frac{1}{\Delta^2})$  then  
        buy good back and give it to bidder  $b_i$ ;  
         $\pi^* \leftarrow \pi_i$  and  $v^* \leftarrow v_i$  ;  
    end  
     $v^{**} = \max\{v^{**}, v_i\}$ ;  
end
```

gain competitiveness remains $8\Delta^5$. Before we can continue our proof, we need two helper lemmas.

As before, let $b^* = (v^*, \pi^*)$ be the last bidder our mechanism has accepted.

Lemma 2. *If the clairvoyant mechanism accepts a bad bidder $\hat{b} = (\hat{v}, \hat{\pi})$, then the next bidder it will accept must be the first crucial bidder that comes afterward.*

Proof. Let $\bar{b} = (\bar{v}, \bar{\pi})$ be the next bidder clairvoyant mechanism accepts after $\hat{b} = (\hat{v}, \hat{\pi})$, and v^{**} be the maximum valuation of all bidders before \bar{b} . Then we must have $\bar{v} > \hat{\pi} > 2\Delta^2 \hat{v}$. Note that $v^{**} \leq \Delta \hat{v}$, since otherwise the gain competitiveness of the clairvoyant mechanism will be larger than Δ , and thus we have $\bar{v} > 2\Delta v^{**}$, and therefore \bar{b} must be crucial. As clairvoyant mechanism needs to accept all crucial bidders, \bar{b} must be the first crucial bidder after $(\hat{v}, \hat{\pi})$.

Lemma 3. *If b^* is bad, then the next bidder our mechanism accepts must be the first crucial bidder $\bar{b} = (\bar{v}, \bar{\pi})$ that comes afterward. Furthermore, the gain after accepting \bar{b} is at least $\frac{1}{\Delta^2} \bar{v}$.*

Proof. Let \bar{b} be the next crucial bidder after b^* . If b^* is bad, then b^* must be crucial since our mechanism only accepts bad bidders that are crucial. So the clairvoyant mechanism will also accept b^* since it accepts every crucial bidder. By Lemma 2, the next bidder after b^* the clairvoyant mechanism will accept is \bar{b} . So $\bar{v} - \pi^* \geq \frac{1}{\Delta} \bar{v}$, since otherwise the gain of clairvoyant mechanism will be less than $\frac{1}{\Delta} \bar{v}$. This implies that $\bar{v} \geq \pi^* + \frac{1}{\Delta} \bar{v} > \pi^* + \frac{1}{\Delta^2} \bar{v}$ and therefore $\bar{v} > \pi^*/(1 - \frac{1}{\Delta^2})$. Thus, our mechanism will also accept \bar{b} . Let v^{**} be the maximum valuation before \bar{b} , then $v^{**} \leq \Delta v^*$. So between b^* and \bar{b} , our mechanism will not accept any bidder.

By our assumption, b^* is last bidder our mechanism accepts before \bar{b}_i , so if b^* is bad, \bar{b}_i must be the first crucial bidder after b^* , and our mechanism will accept \bar{b}_i . The gain after accepting \bar{b}_i is at least $\frac{1}{\Delta^2} \bar{v}_i$, and the gain competitiveness is at most Δ^2 .

If our mechanism does not accept \bar{b}_i , then $\bar{v}_i < \pi^* + \frac{1}{\Delta^2}\bar{v}_i$. Moreover, by Lemma 3, if our mechanism does not accept \bar{b}_i , then b^* is good, and thus $\bar{v}_i < \pi^* + \frac{1}{\Delta^2}\bar{v}_i \leq 2\Delta^2v^* + \frac{1}{\Delta^2}\bar{v}_i$. Thus, we have $\bar{v}_i - \frac{1}{\Delta^2}\bar{v}_i < 2\Delta^2v^*$ or equivalently $\bar{v}_i < 2\Delta^2v^*/(1 - \frac{1}{\Delta^2}) \leq 3\Delta^2v^*$ (wlog assuming $\Delta > 2$, otherwise we can achieve constant factor competitiveness by treating Δ as two in the mechanism). This implies that the current gain competitiveness is at most $6\Delta^2$ using that b^* is a good bidder.

Based on the above analysis and a simple induction we conclude that if our mechanism accepts a bad bidder $b^* = (v^*, \pi^*)$, the gain is at least $\frac{1}{\Delta^2}v^*$ at this moment. It is also easy to see, if b^* is good, then the gain is at least $v^*/2$ (analogue to the proof of Theorem 2).

We now combine the previous observations. Let $c = \{\bar{b}_i, c_1, c_2, \dots, c_t\}$ be the sequence of bidders that arrive between \bar{b}_i and \bar{b}_{i+1} (excluding \bar{b}_{i+1}). Let $b' = (v', \pi')$ be the last bidder the clairvoyant mechanism accepts. If the clairvoyant mechanism only accepts good bidders in c , then the gain competitiveness between our online mechanism and the clairvoyant mechanism is at most $4\Delta^4$, because $v' \leq 2\pi^* \leq 4\Delta^2v^*$ holds at all time (otherwise, our online mechanism will accept (v', π')) and the gain of our online mechanism is at least v^*/Δ^2 , which implies the gain competitiveness is at most $4\Delta^4$.

Thus, we only need to consider the case when clairvoyant mechanism accepts at least one bad bidder in c (possibly \bar{b}_i). By the above analysis, we know that if the clairvoyant mechanism accepts some bad bidder $\hat{c} = (\hat{v}, \hat{\pi})$, then the next bidder it accepts is \bar{b}_{i+1} . Furthermore, $v^{**} \leq \Delta\hat{v}$, where v^{**} is maximum valuation before \bar{b}_{i+1} .

Before accepting $\hat{c} = (\hat{v}, \hat{\pi})$ the clairvoyant mechanism only accepts good bidders. Now suppose we are at the time right before \hat{c} comes. Suppose, at this time, our online mechanism accepts $b^* = (v^*, \pi^*)$ and clairvoyant mechanism accepts (v', π') . We first consider the case when b^* is good. Then we have $v' \leq 2\pi^* \leq 4\Delta^2v^*$. Let m be the maximum valuation before \hat{c} . We have $m \leq \Delta v'$, and $\hat{v} \leq 2\Delta m$ (otherwise $\bar{b}_{i+1} = \hat{c}$). Remember that v^{**} is the maximum valuation before \bar{b}_{i+1} . Hence, $v^{**} \leq \Delta\hat{v} \leq 2\Delta^2m \leq 2\Delta^3v' \leq 8\Delta^5v^*$.

We now conclude this proof with a simple case distinction. If $b^* = (v^*, \pi^*)$ is good, then the gain competitiveness of our mechanism will never be worse than $8\Delta^5$ after it accepts $b^* = (v^*, \pi^*)$, as the gain is at least $v^*/2$. Moreover, before \hat{b} (with \hat{v}) came, both our mechanism and clairvoyant mechanism only accept good bidders, so the gain competitiveness of our mechanism is at most $8\Delta^5$ before this time. So the gain competitiveness of our mechanism is at most $8\Delta^5$ before \bar{b}_{i+1} comes.

On the other hand, if $b^* = (v^*, \pi^*)$ is bad, which implies that $\hat{c} = b^* = \bar{b}_i$, and that the clairvoyant mechanism does not accept any bidder before \bar{b}_{i+1} . This implies that $v^{**} \leq \Delta v^*$, and the gain competitiveness of our mechanism in this period is at most Δ^3 , since the gain is at least $\frac{1}{\Delta^2}v^*$.

The bound from Theorem 3 is tight. We proceed by showing the matching lower bound for any deterministic mechanism.

$$\begin{array}{l} \textcircled{\textcircled{b_{j-1}}} \quad b_{j-1} = (v_j/(2d), v_j d/4) \\ \textcircled{b_j} \quad b_j = (v_j, v_j d^2/2) \\ \textcircled{\textcircled{b_{j+1}}} \quad b_{j+1} = (v_j d^2/2, v_j d^2) \end{array}$$

(a) The bidders b_{j-1} and b_{j+1} are accepted by the clairvoyant mechanism. The bidders b_j and b_{j+1} are accepted by the online mechanism resulting in negative gain.

$$\begin{array}{l} \textcircled{\textcircled{b_{j-1}}} \quad b_{j-1} = (v_j/(2d), v_j d/4) \\ \textcircled{b_j} \quad b_j = (v_j, v_j d^2/2) \\ \textcircled{\textcircled{b_{j+1}}} \quad b_{j+1} = (v_j d^2/2, v_j d^2) \\ \textcircled{b_{j+2}} \quad b_{j+2} = (2v_j d^3, v_j d^{1000}) \end{array}$$

(b) The bidders b_{j-1} and b_{j+1} are accepted by the clairvoyant mechanism. The bidders b_j and b_{j+2} are accepted by the online mechanism. Thus, a bidder $b_{j+3} = (v_j d^{999}, v_j d^{1337})$ would inevitably lead to a gain competitiveness of $\omega(\Delta^5)$.

$$\textcircled{b_{j+3}} \quad b_{j+3} = (v_j d^{999}, v_j d^{1337})$$

$$\begin{array}{l} \textcircled{\textcircled{b_{j-1}}} \quad b_{j-1} = (v_j/(2d), v_j d/4) \\ \textcircled{b_j} \quad b_j = (v_j, v_j d^2/2) \\ \textcircled{\textcircled{b_{j+1}}} \quad b_{j+1} = (v_j d^2/2, v_j d^2) \\ \textcircled{\textcircled{b_{j+2}}} \quad b_{j+2} = (2v_j d^3, v_j d^{1000}) \\ \textcircled{\textcircled{b_{j+3}}} \quad b_{j+3} = (2v_j d^4, v_j d^{1337}) \end{array}$$

(c) The bidders b_{j-1} , b_{j+1} , and b_{j+2} are accepted by the clairvoyant mechanism. The bidders b_j and b_{j+3} are accepted by the online mechanism. Thus, a bidder $b_{j+4} = (v_j d^{1336}, v_j d^{2000})$ would inevitably lead to a gain competitiveness of $\omega(\Delta^5)$.

$$\textcircled{b_{j+4}} \quad b_{j+4} = (v_j d^{1336}, v_j d^{2000})$$

$$\begin{array}{l} \textcircled{\textcircled{b_{j-1}}} \quad b_{j-1} = (v_j/(2d), v_j d/4) \\ \textcircled{b_j} \quad b_j = (v_j, v_j d^2/2) \\ \textcircled{\textcircled{b_{j+1}}} \quad b_{j+1} = (v_j d^2/2, v_j d^2) \\ \textcircled{\textcircled{b_{j+2}}} \quad b_{j+2} = (2v_j d^3, v_j d^{1000}) \\ \textcircled{\textcircled{b_{j+3}}} \quad b_{j+3} = (2v_j d^4, v_j d^{1337}) \end{array}$$

(d) The bidders b_{j-1} , b_{j+1} , and b_{j+3} are accepted by the clairvoyant mechanism. The bidders b_j and b_{j+4} are accepted by the online mechanism. Thus, a bidder $b_{j+5} = (v_j d^{1999}, v_j d^{2000})$ would inevitably lead to a gain competitiveness of $\omega(\Delta^5)$.

$$\textcircled{b_{j+5}} \quad b_{j+5} = (4v_j d^{1999}, v_j d^{2000})$$

Fig. 1: The bidders accepted by the clairvoyant mechanism are marked with (thinly) dashed lines. The online mechanism accepts by definition b_j . If the online mechanism accepts the bottom left bidder, the bidder on the bottom right appears; resulting in a $\omega(\Delta^5)$ gain competitiveness.

Theorem 4. *Any deterministic Δ -online mechanism has gain competitiveness of $\Omega(\Delta^5)$ compared to the offline solution.*

Proof. For any $d > 0$, we will present a sequence of bidders, for which the gain competitiveness between the offline mechanism and the clairvoyant mechanism is at most $2d$, but for any online mechanism, the gain competitiveness is at least $4d^5$. Given Δ , we can set $d = \Delta/2$. Thus, any online mechanism is at least $\Omega(\Delta^5)$ worse than the offline mechanism. The input sequence is depicted in Figure 1.

The input sequence starts with bidder b_1 with $(v_1, \pi_1) = (1, 1)$, then the adversary inserts a sequence of bidders $b_{i+1} = (v_{i+1}, \pi_{i+1})$, for $i = 1, 2, \dots$, where $(v_{i+1}, \pi_{i+1}) = (2d^i, d^{i+2})$. Let b_j be the first bidder in this sequence that the online mechanism accepts. Notice that the online mechanism has to accept one, since otherwise the gain competitiveness is infinity. The clairvoyant mechanism accepts bidder b_{j-1} , but not b_j . The adversary then sets $(v_{j+1}, \pi_{j+1}) = ((d^2/2)v_j, d^2v_j)$, so that the online mechanism cannot accept this bidder because the new gain would be at most $v_j d^2/2 - v_j d^2/2 - \pi_1 < 0$. The clairvoyant mechanism accepts bidder b_{j+1} to maintain gain competitiveness $\mathcal{O}(\Delta)$.

The next bidder b_{j+2} that comes has $(v_{j+2}, \pi_{j+2}) = (d^3v_j, d^{1000}v_j)$, so the online mechanism cannot accept this bidder either, since otherwise the adversary can make the next bidder have a valuation of $d^{999}v_j$, which makes the gain competitiveness much larger than $4d^5$. The clairvoyant mechanism does not accept bidder b_{j+2} and still maintains gain competitiveness $\mathcal{O}(\Delta)$.

Bidder b_{j+3} is then $(v_{j+3}, \pi_{j+3}) = (2d^4v_j, d^{1337}v_j)$. For the same reason, the online mechanism cannot accept this one. The clairvoyant mechanism accepts bidder b_{j+3} to maintain gain competitiveness $\mathcal{O}(\Delta)$. If the online mechanism accepts this bidder, then the clairvoyant mechanism accepts bidder b_{j+2} , but not bidder b_{j+3} (see Figure 1).

Bidder b_{j+4} is $(v_{j+4}, \pi_{j+4}) = (4d^5v_j, d^{2000}v_j)$, and again the online mechanism cannot accept this one. The clairvoyant mechanism does not accept bidder b_{j+4} and still maintains gain competitiveness $\mathcal{O}(\Delta)$.

At this point, the online mechanism accepted (v_j, π_j) , and the gain competitiveness is at most $\frac{4d^5v_j}{v_j} = 4d^5$. Thus, the claim follows.

5 Auctions with Several Goods

In this section we consider auctions with r goods. The pure offline mechanism chooses the best r bidders and never has to pay a preemption price. If r is constant, then we show the following constructive result.

Theorem 5. *Checking whether there is a solution with gain competitiveness of δ in an online auction is r goods can be computed in $\mathcal{O}(n^{r+1})$.*

Similar to the problem of checking whether there is a k -clique in a graph, the general version of this problem is \mathcal{NP} -hard.

Theorem 6. *Checking whether there is a solution with gain competitiveness of δ in an online auction is \mathcal{NP} -hard.*

References

1. B. V. Ashwinkumar. Buyback Problem - Approximate Matroid Intersection with Cancellation Costs. In *ICALP*, pages 379–390, 2011.
2. B. V. Ashwinkumar and R. Kleinberg. Randomized Online Algorithms for the Buyback Problem. In *WINE*, pages 529–536, 2009.
3. M. Babaioff, J. D. Hartline, and R. Kleinberg. Selling ad campaigns: online algorithms with cancellations. In *EC*, pages 61–70, 2009.
4. M. Chrobak, C. Kenyon, J. Noga, and N. E. Young. Incremental medians via online bidding. *Algorithmica*, 2008.
5. E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
6. F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pál. An online mechanism for ad slot reservations with cancellations. In *SODA*, pages 1265–1274, 2009.
7. S. Dasgupta and P. M. Long. Performance Guarantees for Hierarchical Clustering. *J. Comput. Syst. Sci.*, 70(4):555–569, June 2005.
8. E. J. Friedman and D. C. Parkes. Pricing WiFi at Starbucks: Issues in Online Mechanism Design. In *EC*, pages 240–241, 2003.
9. S. P. Y. Fung. Online Scheduling of Unit Length Jobs with Commitment and Penalties. In *TCS*, pages 54–65, 2014.
10. M. R. Garey and D. S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.
11. J. Gilbert and F. Mosteller. Recognizing the Maximum of a Sequence. *Journal of the American Statistical Association*, 61(313):pp. 35–73, 1966.
12. T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
13. M. T. Hajiaghayi. Online Auctions with Re-usable Goods. In *EC*, 2005.
14. M. T. Hajiaghayi, R. D. Kleinberg, and D. C. Parkes. Adaptive limited-supply online auctions. In *EC*, pages 71–80, 2004.
15. X. Han, Y. Kawase, and K. Makino. Online Knapsack Problem with Removal Cost. In *COCOON*, pages 61–73, 2012.
16. J. Hartline and A. Sharp. An Incremental Model for Combinatorial Maximization Problems. In *WEA*, pages 36–48, 2006.
17. J. Hartline and A. Sharp. Incremental flow. *Networks*, 50(1):77–85, 2007.
18. Y. Kawase, X. Han, and K. Makino. Unit Cost Buyback Problem. In *ISAAC*, 2013.
19. R. Lavi and N. Nisan. Competitive Analysis of Incentive Compatible On-line Auctions. In *EC*, pages 233–241, 2000.
20. R. Lavi and N. Nisan. Online Ascending Auctions for Gradually Expiring Items. In *SODA*, pages 1146–1155, 2005.
21. D. V. Lindley. Dynamic Programming and Decision Theory. *j-APPL-STAT*, 10(1):39–51, 1961.
22. R. R. Mettu and C. G. Plaxton. The Online Median Problem. In *FOCS*, 2000.
23. N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
24. D. C. Parkes and S. P. Singh. An MDP-Based Approach to Online Mechanism Design. In *NIPS*, 2003.
25. D. C. Parkes, S. P. Singh, and D. Yanovsky. Approximately Efficient Online Mechanism Design. In *NIPS*, 2004.
26. C. G. Plaxton. Approximation Algorithms for Hierarchical Location Problems. In *STOC*, pages 40–49, 2003.
27. W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of Finance*, 16(1):8–37, 1961.