Derandomization of Auctions

Gagan Aggarwal* Amos Fiat† Andrew V. Goldberg‡ Jason D. Hartline‡
Nicole Immorlica § Madhu Sudan§

February 15, 2005

Abstract

We study the problem of designing seller optimal auctions. Prior to this work, the only previously known auctions that are approximately optimal in worst case employ randomization. Our main result is the existence of deterministic auctions that approximately match the performance guarantees of these randomized auctions. We give a fairly general derandomization technique for turning any randomized mechanism into an asymmetric deterministic one with approximately the same revenue. In doing so, we bypass the impossibility result for symmetric deterministic auctions and show that asymmetry is nearly as powerful as randomization in optimal mechanism design problems. Our general construction involves solving an exponential-sized flow problem and thus is not polynomial-time computable. For a specific auction with good worst-case revenue, we give a polynomial-time construction. To obtain our results, we study coloring problems that are interesting on their own.

1 Introduction

In [9] Goldberg et al. proposed the study of profit maximization in auctions using a worst-case competitive analysis. They focus on the *unlimited supply*, *unit demand*, *single-item* auction problem where an effectively infinite quantity of identical items are for sale to consumers that each desire at most one item. In their worst-case competitive framework, they gave a randomized auction that achieves a constant fraction of the optimal single-price revenue. Further, they prove that such randomization is necessary for *symmetric* auctions, ones whose outcome is not a function of the order of the input bids. Our main result is to show that this result does not hold for asymmetric auctions: there exists an asymmetric deterministic auction that approximates the revenue of the optimal single-price sale in the worst case.

In general, design and analysis of auctions and other mechanisms requires a game-theoretic treatment; in order to understand the performance of an auction, the behavior of the bidders in the auction must be understood. To handle this problem, we adopt the solution concept of *truthful mechanism design*; we only consider mechanisms where each bidder has a dominant strategy of bidding their true value for the good regardless of the actions of any of the other bidders. It is well known that truthful auctions are precisely those auctions that compute an offer price for a bidder

*Stanford University, Palo Alto, CA. gagan@cs.stanford.edu.
†Tel Aviv University, Tel Aviv, Israel. fiat@tau.ac.il.
‡Microsoft Research – Silicon Valley, Mountain View, CA. {goldberg,hartline}@microsoft.com.
§MIT, Cambridge, MA. {nickle,madhu}@mit.edu.

that is not a function of their bid value, but may be a function other bidders' bids. This bidder is then allocated the good if their bid is above the offer price, otherwise the bidder is rejected. The simplest truthful auction that is approximately optimal in worst-case is the *randomized sampling auction* in [9] which randomly partitions the bidders into two sets and uses the optimal sale price for each set as the offer price for all bidders in the opposite set.

Our main result is to show that on any input bid vector, $\mathbf{b}$, with bids $b_i \in [1, h]$, any randomized truthful auction $\mathcal{A}$ that obtains an expected profit of $\mathbf{E}[\mathcal{A}]$ can be converted into a truthful deterministic asymmetric auction with profit $\mathbf{E}[\mathcal{A}]/4 - 2h$. Given any auction that always obtains an expected profit that is within a constant fraction of optimal, such as the randomized sampling auction, this gives a deterministic auction that gets within a constant fraction of the optimal profit less a small additive loss.

Our general derandomization technique involves solving an exponentially large flow problem. We introduce a different technique to derandomize a specific auction of [7] and obtain a polynomial-time deterministic auction with good worst-case revenue guarantees.

We explore the possibility of a general polynomial time derandomization technique through the following toy problem. Given an array of red and blue hats, devise a strategy for guessing the color of each hat while only observing the colors of the other hats. The goal is to design a color-guessing algorithm that simultaneously guesses correctly for about half of the red hats and about half of the blue hats. A simple randomized strategy that solves this problem is to independently flip a coin for each hat and color it red with probability a half and blue with probability a half. Our flow-based derandomization technique applied to this simple coloring algorithm gives an asymmetric deterministic (exponential-time) algorithm for hat-coloring that guesses red for exactly half of the reds and guesses blue for exactly half of the blues (rounded down if the numbers are odd). For this particular toy problem, we also show a rather simple polynomial-time solution for guessing colors that matches the bound of the flow-based derandomization. The flow-based construction generalizes to the case where there are k different colors and we would like to guess correctly for about $1/k$ hats of each color; however, we do not know of any polynomial-time deterministic algorithm that is able to match this bound for any $k \geq 3$.

This coloring problem is related to the auction problem as follows. Consider the case where there are only two types of bidders, those with a high valuation for the item, h; and those with a low valuation for the item, 1. Mapping h to the color red and 1 to the color blue, a solution to the color-guessing problem would offer half the h bids a price of h and half the 1 bids a price of 1 and thus, the profit of such an auction would be at least half of optimal revenue (because either a price of h or a price of 1 was optimal).

The coloring problem is also related to problems in computational learning and coding theory. Relationship to learning is immediate: we cover a hat and try to learn its color from those of other hats. This relationship has been observed before in the context of on-line auctions [3]. In the current paper, we strengthen this relationship by defining guessing auctions, which are very closely related to expert learning. As for coding theory, the coloring problem is similar to the hats problem [4], in particular the hats coloring problem studied in [5, 14]. This relationship motivated our flow-based auction. Fiege [5] independently developed a similar flow-based solution to the hats coloring problem. However, we are unaware of previous applications of the hat problem to truthful mechanism design.

This paper is organized as follows. In Section 3 we formally define the coloring problem and give the flow-based solution. In Section 4 we show how the flow-based technique can be generalized to convert any randomized auction into an asymmetric deterministic auction with approximately the same performance bound. Finally, in Section 5 we give a polynomial-time deterministic auction with good worst-case revenue bounds.

2 Preliminaries

We consider single-round, sealed-bid auctions to sell an unlimited number of copies of a good to any subset of n bidders. As mentioned in the introduction, we adopt the game-theoretic solution concept of truthful mechanism design. A useful simplification of the problem of designing truthful auctions is obtained through the following algorithmic characterization. Related formulations to the one given here have appeared in numerous places in recent literature (e.g., [1, 13, 6, 10]). To the best of our knowledge, the earliest dates back to the 1970s [11].

Definition 2.1 *Given a bid vector of n bids, $\mathbf{b} = (b_1, \ldots, b_n)$, let $\mathbf{b}_{-i}$ denote the vector of bids with b_i replaced with a '?', i.e.,*

$$\mathbf{b}_{-i} = (b_1, \ldots, b_{i-1}, ?, b_{i+1}, \ldots, b_n).$$

Definition 2.2 (Bid-independent Auction, BI_f) *Let f be a function from masked bid vectors (with a '?') to prices (non-negative real numbers). The* deterministic bid-independent auction defined by f, BI_f, *works as follows. For each bidder i:*

1. *Set $t_i = f(\mathbf{b}_{-i})$.*
2. *If $t_i < b_i$, bidder i wins at price t_i*
3. *If $t_i > b_i$, bidder i loses.*
4. *Otherwise, ($t_i = b_i$) the auction can either accept the bid at price t_i or reject it.*

A randomized bid-independent auction is a distribution over deterministic bid-independent auctions.

The proof of the following theorem can be found, for example, in [6].

Theorem 2.1 *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

Given this equivalence, we will use the terminology *bid-independent* and *truthful* interchangeably. We denote the profit of a truthful auction $\mathcal{A}$ on input $\mathbf{b}$ as $\mathcal{A}(\mathbf{b})$. This profit is given by the sum of the prices charged to bidders that are accepted. For a randomized bid-independent auction, $f(\mathbf{b}_{-i})$ and $\mathcal{A}(\mathbf{b})$ are random variables.

It is natural to consider a worst-case competitive analysis of truthful auctions. In the competitive framework of [6] and subsequent papers, the performance of a truthful auction is gaged in comparison to the profit of the *optimal single price sale of at least two units.* There are a number of reasons to choose this metric for comparison; interested readers should see [6] or [8] for a more detailed discussion.

Unfortunately, as we show in Section 4.3, it is not possible for a deterministic auction to always perform well against such an optimal sale. Instead, we assume that all bids are between 1 and h and define OPT as the profit of the *optimal single price* sale. Following [9, 2] we look for auctions that obtain a profit of at least $\text{OPT}/\beta - \gamma h$ for small constants β and γ. We refer to β as the *approximation ratio* and γh as the *additive loss.* Such an approximation framework is tantamount to considering a *promise* problem. If we are promised that $\text{OPT} \gg \gamma h$ then our auction is constant fraction of optimal. This motivates the following formal definition.

Definition 2.3 *We say an auction is* approximately optimal *if its expected profit on any input is at least* $\text{OPT}/\beta - \gamma h$ *for fixed constants* β *and* γ.

3 The Color Guessing Problem

Given n hats each having one of k colors, arranged in a linear order, we would like a scheme for guessing their colors such that when guessing the i'th hat's color we can only look at the colors of the remaining $n-1$ hats. In this "bid-independent" hat-guessing game, we would like our technique, after considering all n of the hats in turn, to have correctly guessed the color of about a $1/k$ fraction from each color class. The following randomized coloring scheme achieves this desired bound in expectation: *for each i, guess each of the colors with probability* $1/k$. We now give a technique that uses the ordering information in place of randomness to achieve the same bound deterministically. It is instructive to view this technique as a derandomization of the simple randomized algorithm proposed above.

Let $\mathbf{c} = (c_1, \ldots, c_n)$ represent the array of colors. Let $\mathbf{c}_{-i}$ represent the array of colors with the i'th color hidden, i.e., $\mathbf{c}_{-i} = (c_1, \ldots, c_{i-1}, ?, c_{i+1}, \ldots, c_n)$.

Consider a flow problem on the following graph. We have a source s and a sink t. For each of the nk^{n-1} possible values of $\mathbf{c}_{-i}$ we have a vertex, $v_{\mathbf{c}_{-i}}$. We place an arc from s to each of these vertices. For each of the k^{n+1} possible values of $(\chi, \mathbf{c})$ (where χ is one of the k colors), we have a vertex $v_{\chi,\mathbf{c}}$. We place an arc between each of these vertices and t. We also add an arc between $v_{\mathbf{c}_{-i}}$ and $v_{c_i,\mathbf{c}}$ signifying that we get $\mathbf{c}$ when we reveal that at position i in $\mathbf{c}_{-i}$ is a hat with color c_i. Notice that the in-degree due to such arcs of a vertex $v_{\chi,\mathbf{c}}$ is precisely the number of hats of color χ in $\mathbf{c}$. The out-degree of a vertex $v_{\mathbf{c}_{-i}}$ is exactly k, one for each possible color of the hat at position i.

Now, imagine the following flow on this graph that represents the randomized color-guessing algorithm above. Between s and each $v_{\mathbf{c}_{-i}}$ place a flow of 1. This corresponds to the randomized algorithm, upon seeing $\mathbf{c}_{-i}$, having a total probability of 1 to spend on guessing colors for the i'th hat. On each of the outgoing arcs from $v_{\mathbf{c}_{-i}}$ we place a flow of $1/k$ corresponding to the probability with which the randomized algorithm picks each color. Now notice that the incoming flow to $v_{\chi,\mathbf{c}}$ is precisely $1/k$ times the number of hats in $\mathbf{c}$ from color class χ. Send all of this flow on the arc from $v_{\chi,\mathbf{c}}$ to t.

Now we set capacities on the arcs. For each χ and $\mathbf{c}$, set the capacity of the arc $(v_{\chi,\mathbf{c}}, t)$ to $\lfloor n_\chi(\mathbf{c})/k \rfloor$ where $n_\chi(\mathbf{c})$ represents the number of hats in $\mathbf{c}$ that are colored χ. On all other arcs place a capacity of 1. The fractional flow set up in the previous paragraph respects these capacities except on the arcs into t. Thus, the minimum cut in this graph separates t from all other vertices, and a maximum flow saturates the capacities of all arcs into t. Furthermore, since all capacities are integral, there is a maximum flow that is integral. Revisiting the analogy between flow and probability, since each of the $v_{\mathbf{c}_{-i}}$ has at most 1 unit of incoming flow, an integral flow places the entirety of this flow on a single outgoing arc corresponding to deterministically guessing a color for the i'th hat in $\mathbf{c}_{-i}$. Thus, such a flow specifies a deterministic "bid-independent" color-guessing algorithm.

Consider the performance of this deterministic color-guessing algorithm on $\mathbf{c}$. Given $n_\chi(\mathbf{c})$ hats with color χ in $\mathbf{c}$, the capacity of the outgoing arc from $v_{\chi,\mathbf{c}}$ to t is $\lfloor n_\chi(\mathbf{c})/k \rfloor$. Since this arc is saturated in an integral maximum flow, it must be that $\lfloor n_\chi(\mathbf{c})/k \rfloor$ of the $n_\chi(\mathbf{c})$ incoming arcs have one unit of flow on them. This corresponds to the deterministic algorithm correctly guessing χ when considering $\mathbf{c}_{-i}$ for $\lfloor n_\chi(\mathbf{c})/k \rfloor$ positions i in $\mathbf{c}$ colored χ. This holds true for all colors χ; thus, this deterministic color-guessing algorithm correctly guesses about a $1/k$ fraction from each color class.

3.1 The $k = 2$ case

We now consider the case where there are $k = 2$ colors, red and blue. The following is an efficient solution for correctly guessing half the reds and half the blues.

Definition 3.1 (2-Color Guessing Algorithm) *On input $\mathbf{c}$, let $M_r(i) = i + \sum$ indices of reds in $\mathbf{c}_{-i}$+ number of reds before position i in $\mathbf{c}_{-i}$. For each hat i, guess i is red if $M_r(i)$ is even and guess blue otherwise.*

Lemma 3.1 *On $\mathbf{c}$ with n_r reds and n_b blues, the above scheme correctly guesses $\lfloor n_r/2 \rfloor$ of the reds and $\lfloor n_b/2 \rfloor$ of the blues.*

Proof: It is easy to see that it correctly guesses about half of the reds. For any hat i that is red, $i + \sum$ indices of reds in $\mathbf{c}_{-i}$ is constant. As the "number of reds before position i in $\mathbf{c}_{-i}$" alternates parity, this algorithm correctly guesses the color of every other red.

Now we show that half the blues are correctly guessed. Define $M_b(i)$ analogously to $M_r(i)$, but with respect to blues. We show that the sum $M_r(i) + M_b(i)$ is constant modulo two (over the choice of i). Thus, our algorithm is equivalent to an algorithm that colors based on the parity of $M_b(i)$, and thus guesses half of the blues correctly. Suppose i is red (a similar argument holds if i is blue), then the sum $M_r(i) + M_b(i)$ is $2i$ plus the sum of the indices of all the reds except for i plus the sum of the indices of the blues plus the rank of number of reds plus the number of blues before position i. The latter part of this sum is simply $i - 1$ and the former is i plus $\sum_j j$. Thus, the total is $2i - 1 + \sum_j j$ which is constant modulo two as desired. □

An obvious question would be to try to extend this algorithm to $k \geq 3$ colors. This problem seems much more difficult even for $k = 3$, and we leave it as an open question.

4 Auction Derandomization

The main goal of our paper is to design deterministic auctions that are approximately optimal. We show that in fact *any* randomized auction has a deterministic counterpart that achieves approximately the same profit. As a corollary of this result, known approximately optimal randomized auctions imply the existence of approximately optimal deterministic auctions. Our proof first reduces any auction to a special type of auction that we define, called a *guessing auction*, and then uses a flow-based construction similar to that in Section 3 to derandomize the guessing auction.

4.1 Guessing Auctions

The flow-based construction for the color-guessing problem in Section 3 works for the case where our goal is to consider $\mathbf{c}_{-i}$ and guess what $\mathbf{c}$ is. For auctions, our performance is not based on guessing a bid value correctly, it is based on correctly guessing below a bid value. The easiest way to resolve this discrepancy is to round the bids to powers of two and consider auctions that only get credit for bidders when they exactly guess at which power of two the bidder's bid value is. We call such an auction a *guessing auction*. Not only are guessing auctions approximately as powerful as standard auctions (in terms of approximating the optimal profit), but it is possible to convert any auction into a guessing auction while only losing a factor of four from the profit.

Definition 4.1 ($\mathcal{G}_\mathcal{A}$) *The guessing auction, $\mathcal{G}_\mathcal{A}$, for an auction $\mathcal{A}$ simulates $\mathcal{A}$ on $\mathbf{b}$. Suppose $\mathcal{A}$ offers bidder i price p_i and let 2^k be the largest power of two less than p_i. Then $\mathcal{G}_\mathcal{A}$ offers bidder i price 2^{k+j} for integer $j \geq 0$ with probability 2^{-j-1}.*

Lemma 4.1 *For any auction $\mathcal{A}$ with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector $\mathbf{b}$, there is a corresponding guessing auction $\mathcal{G}_\mathcal{A}$ whose expected profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{A}]/4$.*

Proof: To see that the guessing auction achieves a profit of $\mathbf{E}[\mathcal{A}]/4$, we show that the expected profit of $\mathcal{G}_\mathcal{A}$ from bidder i, given that bidder i bids above p_i, is at least $2^{k-1} \geq p_i/4$. Suppose $b_i \in [2^{k+j}, 2^{k+j+1})$, then the probability that $\mathcal{G}_\mathcal{A}$ guesses bid i is 2^{-j-1}. The payoff on correctly guessing is 2^{k+j}. Thus, the expected payment of bidder i is 2^{k-1}. Since each bidder's expected payment in $\mathcal{G}_\mathcal{A}$ is a fourth of their payment in $\mathcal{A}$, we have the desired bound. □.

We note that if we are constructing a guessing auction from an auction that only uses prices that are powers of two, then we only lose a factor of two of the profit instead a factor of four.

4.2 The Flow Construction

We now show how to derandomize any guessing auction $\mathcal{G}_\mathcal{A}$.

Lemma 4.2 *Corresponding to any guessing auction $\mathcal{G}_\mathcal{A}$ with expected profit $\mathbf{E}[\mathcal{G}_\mathcal{A}(\mathbf{b})]$ on bid vector $\mathbf{b}$, there is a deterministic auction whose profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{G}_\mathcal{A}(\mathbf{b})] - 2h$ where h is the highest bid value in $\mathbf{b}$.*

Proof: First, round all bid values down to the nearest power of two. We make the analogy between the k colors in the color-guessing problem and the $\log h$ powers of two that are possible bid values and proceed by setting up a flow construction identical to that for the k-color guessing problem except that the fractional flow on an arc from $v_{\mathbf{b}_{-i}}$ to $v_{2^j,\mathbf{b}}$ is the probability that $\mathcal{G}_\mathcal{A}$ on seeing $\mathbf{b}_{-i}$ guesses 2^j. Furthermore the flow from $v_{2^j,\mathbf{b}}$ to t is the expected number of times $\mathcal{G}_\mathcal{A}$ guesses one of the bids at value 2^j correctly. We represent this quantity by $E_j(\mathbf{b})$. We then set the capacities as before such that the capacity on the arc between $v_{2^j,\mathbf{b}}$ and t is $\lfloor E_j(\mathbf{b}) \rfloor$; all other capacities are set to one.

Once again, the above fractional flow implies the existence of an integer-valued flow, and this integer-valued flow corresponds to an auction making a deterministic bid-independent offer upon seeing $\mathbf{b}_{-i}$. The flow out of $v_{2^j,\mathbf{b}}$ is precisely the number of indices i such that the auction upon seeing $\mathbf{b}_{-i}$ correctly guesses 2^j; since this arc is in a minimum cut, it is saturated and the flow out of it is precisely $\lfloor E_j(\mathbf{b}) \rfloor$. Thus, considering a bid vector $\mathbf{b}$ where the expected profit of $\mathcal{G}_\mathcal{A}$ is $\mathbf{E}[\mathcal{G}_\mathcal{A}] = \sum_j 2^j E_j(\mathbf{b})$, the deterministic auction obtains $\sum_j 2^j \lfloor E_j(\mathbf{b}) \rfloor \geq \sum_j \left[2^j E_j(\mathbf{b}) - 2^j\right] \geq \mathbf{E}[\mathcal{G}_\mathcal{A}] - 2h$. □

The following theorem follows directly from Lemmas 4.1 and 4.2.

Theorem 4.3 *Corresponding to any single-round sealed-bid auction $\mathcal{A}$ with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector $\mathbf{b}$, there is a deterministic auction $\mathcal{A}'$ whose expected profit on any input bid vector $\mathbf{b}$ is at least $\mathbf{E}[\mathcal{A}(\mathbf{b})]/4 - 2h$.*

As a corollary, using this derandomization result with the approximately optimal auctions in [9, 6, 7] we obtain deterministic auctions that are approximately optimal. In this construction we assumed that the range of bid values $[1, h]$ is known. This assumption is not necessary. When considering $\mathbf{b}_{-i}$ we can compute 1 and h correctly for all but the minimum and maximum bid value. Assuming the worst, i.e. the auction fails to get any profit from the highest and lowest bid, we only lose an additional additive $h + 1$.

4.3 Additive Loss Term

One discrepancy between the bounds given in this paper and the bounds given in [6, 7] is that our bounds are only interesting if the profit from optimal single price is larger than $2h$. This might not be true for all bid vectors, especially those where h is larger than the profit from the *optimal single-price sale that sells at least two items* (this quantity is denoted $\mathcal{F}^{(2)}$ in [6]). For this case, the bounds obtained in [6] are better because they prove the auctions performance to be a constant fraction of $\mathcal{F}^{(2)}$ without any additive term. We can view these two types of analyzes as the difference between solving a worst-case problem and a promise problem. Given the promise that the optimal single price sale achieves a large profit in comparison to h our auction gets a constant fraction of optimal; otherwise, it may not.

In this section we show that such a promise is necessary for the obtaining a deterministic auction that performs well in the worst case. In particular, we show that there is no deterministic auction that obtains a profit that is a constant fraction of $\mathcal{F}^{(2)}$ on all inputs.

Lemma 4.4 *No deterministic truthful auction obtains a constant fraction of $\mathcal{F}^{(2)}$ on all bid vectors.*

Proof: Assume for contradiction that we are given a deterministic auction BI_f, specified by bid-independent function f, that obtains a profit of $\mathcal{F}^{(2)}/\beta$ on all inputs.

Let $\mathbf{b} = (1, \ldots, 1)$ be the all-ones bid vector. Assume without loss of generality that $f(\mathbf{b}_{-1})$, the price offered the first bidder, is 1. Now, for any $\alpha > \beta$ and $i \in I = \{1, 2, \ldots, \lceil \frac{n+1}{n/\alpha - 1} \rceil\}$, consider $\mathbf{b}^{(i)}$ as the all-ones bid vector except for $b_1 = n\alpha^i$. Let S_i be the set of other bidders (not including bidder 1) that are offered price 1 when the input to the auction is $\mathbf{b}^{(i)}$, i.e., $S_i = \{j > 1 \ : \ f(\mathbf{b}^{(i)}_{-j}) = 1\}$.

Fact 1: $|S_i| \geq n/\alpha - 1$.

This follows directly from the fact that otherwise BI_f's profit would be at most $\mathcal{F}^{(2)}/\alpha$ which would contradict our assumption.

Fact 2: $\bigcap_{i \in I} S_i \neq \emptyset$.

Clearly $|\bigcup_i S_i| \leq n$. But for a contradiction, if the intersection of the S_is is empty then the union of the S_is is of size

$$\begin{aligned} \sum\nolimits_i |S_i| &\geq |I| \, (n/\alpha - 1) \\ &\geq n + 1 > n. \end{aligned}$$

From Fact 2, there exists i, j, and k with $i < j$ and $k \in S_i \cap S_j$. Pick some $h \gg n^n$ (a number bigger than any of the $n\alpha^i$s) and let $p_1 = f(\mathbf{b}'_{-1})$ where $\mathbf{b}'$ is the all-ones input except for $b'_k = h$.

Case 1: $p_1 \leq n\alpha^i$. Then on the input that is all ones except for $b_1 = n\alpha^j$ and $b_k = h$, $\mathcal{F}^{(2)} = 2n\alpha^j$, but auction profit is at most $n\alpha^i + n < 2n\alpha^j/\beta$.

Case 2: $p_1 > n\alpha^i$. Then on the input that is all ones except for $b_1 = n\alpha^i$ and $b_k = h$, $\mathcal{F}^{(2)} = 2n\alpha^i$, but the auction profit is at most n and is therefore not $\mathcal{F}^{(2)}/\beta$. □

4.4 Limited Supply

While we presented these results in terms of the unlimited supply auction problem, they also apply to the limited supply auction. Note that the number of items sold by the derandomized auction is no more than the expected number of items sold by the randomized auction. Thus, if the randomized auction never oversells, neither does its derandomized equivalent.

5 Efficient Deterministic Auction

In this section we describe an efficient competitive deterministic asymmetric auction. This auction uses another coloring problem to derandomize an auction of [7]. We describe the coloring problem first, and the deterministic auction next.

5.1 The Alternating Coloring Problem

In this section we study the following coloring problem. The input to the problem is a bid vector $\mathbf{b}$. We assume that there is a total ordering on the set of bids. In case of ties, we assume that each bidder has a unique, publicly known ID that can be used to break ties. (Alternatively, we can use input positions to break ties.) We say that a function mapping $\mathbf{b}_{-i}$ into $\{-1, 1\}$ (a color of bid i) is *an alternating coloring of* $\mathbf{b}$ if after the components are sorted by their values, the adjacent ones have different colors.

Note that we can reduce the 2-color guessing problem to the alternating coloring problem problem as follows. Interpret color red as -1 and color blue as 1, and break ties using the input positions. Compute an alternating coloring and for each hat interpret the result as the guess for the hat color. Like the coloring algorithm of Section 3.1, the resulting algorithm guesses, modulo rounding, half of the red and half of the blue hat colors correctly.

To describe an efficient deterministic alternating coloring function, we use the notion of *a sign of a permutation*. This is standard notion that is used, for example, in the definition of the determinant of a matrix. A transposition is a permutation that swaps two adjacent elements only. Every permutation can be represented as a composition of transpositions. This representation is not unique. However, for a fixed permutation the parity of the number of transpositions in a representation of the permutation is the same for all representations. The sign of the permutation is defined to be 1 if this parity is even, and -1 if it is odd.

Now we are ready to describe the alternating coloring function ϕ. Let ∞ be a number that is greater than any bid and let $\mathbf{b}'_{-i}$ be $\mathbf{b}_{-i}$ with the i-th component of $\mathbf{b}$ replaced by ∞. Given $\mathbf{b}_{-i}$, ϕ computes $\mathbf{b}'_{-i}$, sorts its components, and outputs the sign of the sorting permutation. Clearly, ϕ can be computed in $O(n \log n)$ time.

Lemma 5.1 *ϕ is an alternating coloring function.*

Proof: Consider two adjacent components of $\mathbf{b}$, i and j, which are adjacent in the sorted order. Any permutation π that sorts $\mathbf{b}'_{-i}$ can be converted into a permutation that sorts $\mathbf{b}'_{-j}$ by first transposing components i and j and then applying π. Thus the values of $\phi(\mathbf{b}_{-i})$ and $\phi(\mathbf{b}_{-j})$ are different and ϕ is a coloring. □

Note that each bidder can compute its own color independently of the bid value, but cannot compute colors of other bidders.

5.2 An Efficient Auction

In this section we describe an auction, *DCORE*, that is a derandomization of a variant of the consensus revenue estimate (CORE) auction of [7]. These auctions use the *cost-sharing auction* of Moulin and Shenker [12] as a building block. Given a target revenue C and a set of bids B, the cost-sharing auction either returns the biggest subset of bidders who can share C equally along with the corresponding price, or an empty set of winners if no set of bidders can share the price. The cost-sharing auction is truthful. Let f_C be the bid-independent function for the cost-sharing auction: bidder i wins at price $f_C(\mathbf{b}_{-i})$ if $f_C(\mathbf{b}_{-i}) \leq b_i$ and loses otherwise. Given $\mathbf{b}_{-i}$, let C_i be the optimal single price revenue from $\mathbf{b}_{-i}$. Let C_i^{-1} and C_i^1 be C_i rounded down to the nearest even and odd integer power of two, respectively. DCORE works as follows:

1. Compute an alternating coloring of $\mathbf{b}$; let c_i be the color of bidder i.
2. Set the offer price for bidder i to $f_{C_i^{c_i}}(\mathbf{b}_{-i})$.

DCORE is bid-independent and therefore truthful. To see intuitively that DCORE should be competitive, recall that the original (randomized) CORE auction chooses to round down to the nearest odd or even integer power of two at random, once for all bidders. It is not hard to see that the expected revenue is the same if we make this decision independently and at random for each bidder, and that DCORE is a derandomization of this variant of CORE. In the following theorem, we use techniques similar to the proof of competitiveness of CORE to prove that DCORE is competitive as well.

Theorem 5.2 *The revenue of DCORE is at least* $\mathrm{OPT}/8 - h$*.*

Proof: If the optimal single price solution has exactly one winner, then the optimal revenue is h and approximating it within an additive h is trivial. Otherwise, for every i we have $\mathrm{OPT}/2 \leq C_i \leq \mathrm{OPT}$, where OPT is the revenue of the optimal single price solution. Since C_i^{-1} and C_i^1 differ by a factor of two, this implies that either all C_i^{-1} values are the same or all C_i^1 values are the same (or both). Let T be equal to C_i^{-1} if this value is constant with respect to i, and C_i^1 otherwise. By definition, $\mathrm{OPT}/4 \leq T \leq \mathrm{OPT}$. Consider the cost-sharing auction on $\mathbf{b}$ with the target value T and note that since $T \leq \mathrm{OPT}$, this auction has revenue T. Suppose this auction has k winners at price p. Note that the winners are the top k bids of the sorted order, and so for $\lfloor \frac{k}{2} \rfloor$ of these, $C_i^{c_i} = T$. Thus, $\lfloor \frac{k}{2} \rfloor$ bidders compute and pay price p, and so the revenue of DCORE is at least $p\lfloor \frac{k}{2} \rfloor \geq \frac{pk}{2} - p$. As p is at most the highest bid, h, and $pk = T \geq \mathrm{OPT}/4$, the result follows. □

6 Conclusions

We have shown the existence of deterministic auctions that are approximately optimal in the worst case. By necessity, these auctions are asymmetric. This gives an affirmative answer to the question left open in [9]. The construction we developed uses a guessing auction as an intermediate and suffers a performance loss because of it. It would be nice to show a more direct derandomization; this is made difficult due to the fact that the natural analogy to the flow problem that takes into account the fact that the auction obtains profit from all bidders above the offered price, is not a totally unimodular linear program and thus is not guaranteed to have an integral optimal solution.

For the color-guessing problem we propose, the $k = 2$ case is really much simpler than $k \geq 3$ since the flow derandomization is related to finding Euler tours between the levels of a hypercube. For the $k = 3$ case where the we are coloring hyperarcs, this formulation is much more complicated.

The flow-based technique allows derandomization of any auction. However, it takes exponential time. We use a different polynomial time technique to derandomize the CORE auction. The existence of a general derandomization technique with a polynomial overhead remains open.

References

[1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[2] A. Blum and J. Hartline. Near-Optimal Online Auctions. In *Proc. 16th Symp. on Discrete Alg.* ACM/SIAM, 2005.

[3] A. Blum, V. Kumar, A. Rudra, and F. Wu. Online Learning in Online Auctions. In *Proc. 14th Symposium on Discrete Algorithms.* ACM/SIAM, 2003.

[4] T. Ebert. *Applications of recursive operators to randomness and complexity.* PhD thesis, University of California at Santa Barbara, 1998.

[5] U. Feige. You Can Leave Your Hat On (if you guess its color). Technical Report MCS04-03, Computer Science and Applied Mathematics, The Weizmann Institute of Science, 2004.

[6] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing.* ACM Press, New York, 2002.

[7] A. Goldberg and J. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Algorithms.* ACM/SIAM, 2003.

[8] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 2002. Submitted for publication. An earlier version available as InterTrust Technical Report STAR-TR-99.09.01.

[9] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.

[10] D. Lehmann, L. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.

[11] J. Mirrlees. An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies*, 38:175–208, April 1971.

[12] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance vs. Efficiency. *Economic Theory*, 18:511–533, 2001.

[13] I. Segal. Optimal Pricing Mechanism with Unknown Demand. *American Economic Review*, 93:509–29, 2003.

[14] P. Winkler. Games People don't Play. In D. Wolfe and T. Rogers, editors, *Puzzler's tribute. A feast for the mind.* A K Peters, 2001.