

The Performance of Deferred-Acceptance Auctions

PAUL DÜTTING, Stanford University
VASILIS GKATZELIS, Stanford University
TIM ROUGHGARDEN, Stanford University

Deferred-acceptance auctions are auctions for binary single-parameter mechanism design problems whose allocation rule can be implemented using an adaptive reverse greedy algorithm. Milgrom and Segal [2014] recently introduced these auctions and proved that they satisfy a remarkable list of incentive guarantees: in addition to being dominant-strategy incentive-compatible, they are weakly group-strategyproof, can be implemented by ascending-clock auctions, and admit outcome-equivalent full-information pay-as-bid versions. Neither forward greedy mechanisms nor the VCG mechanism generally possess any of these additional incentive properties. The goal of this paper is to initiate the study of deferred-acceptance auctions from an approximation standpoint. We study these auctions through the lens of two canonical welfare-maximization problems, in knapsack auctions and in combinatorial auctions with single-minded bidders.

For knapsack auctions, we prove a separation between deferred-acceptance auctions and arbitrary dominant-strategy incentive-compatible mechanisms. While the more general class can achieve an arbitrarily good approximation in polynomial time, and a constant-factor approximation via forward greedy algorithms, the former class cannot obtain an approximation guarantee sub-logarithmic in the number of items m , even with unbounded computation. We also give a polynomial-time deferred-acceptance auction that achieves an approximation guarantee of $O(\log m)$ for knapsack auctions.

For combinatorial auctions with single-minded bidders, we design novel polynomial-time mechanisms that achieve the best of both worlds: the incentive guarantees of a deferred-acceptance auction, and approximation guarantees close to the best possible.

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

Additional Key Words and Phrases: Single-Parameter Combinatorial Auctions; Deferred-Acceptance Auctions

1. INTRODUCTION

Algorithmic mechanism design studies the conflict between incentives and computation. For a canonical example, consider a *knapsack auction* where there are m identical items and each of n bidders has a valuation (i.e., willingness-to-pay) v_i for s_i of the items. In such an auction, bidder i has no value for fewer than s_i items. The objective is to identify a set A of accepted bidders, with $\sum_{i \in A} s_i \leq m$, that maximizes the welfare $\sum_{i \in A} v_i$. This problem is well understood from a purely algorithmic perspective.

The problem potentially becomes harder when we assume that valuations are *private* — a priori unknown to the seller — and that the bidders are strategic. Is there a protocol — a *mechanism* — that solicits the private valuations from the bidders, computes a near-optimal allocation with respect to the reported valuations, and charges payments to incentivize bidders to report truthfully? For example, when m and all of the s_i 's equal 1, the Vickrey (second-price) auction provides a satisfying solution: it is

Authors' addresses: P. Dütting, V. Gkatzelis and T. Roughgarden, Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA; email:

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC'14, June 8–12, 2014, Stanford, CA, USA. Copyright © 2014 ACM 978-1-4503-2565-3/14/06 ...\$15.00.
<http://dx.doi.org/10.1145/2600057.2602861>

dominant-strategy incentive-compatible (DSIC), meaning that for every bidder it is a dominant strategy to report its true valuation, and assuming all bidders comply, the allocation is optimal — the item goes to the bidder with the highest valuation. The VCG mechanism (see, e.g., [Nisan et al. 2007, Chapter 9]) is the analog of the Vickrey auction for multiple items — it is DSIC and maximizes welfare, but cannot be implemented for knapsack auctions in time polynomial in n and $\log m$ (assuming $P \neq NP$). Is there a good DSIC and polynomial-time mechanism for knapsack auctions? Previous work gives an affirmative answer: Mu’alem and Nisan [2002] showed that combining two greedy algorithms yields a 2-approximate DSIC mechanism, and Briest et al. [2005] gave a DSIC FPTAS.¹

For a second example, consider the problem introduced by Lehmann et al. [2002] of maximizing welfare with single-minded bidders. There are n bidders and m distinct items. Each bidder i wants a specific bundle S_i of items and it has a valuation v_i for it. The objective is to identify a set A of accepted bidders, with $S_i \cap S_j = \emptyset$ for every $i, j \in A$, that maximizes the welfare $\sum_{i \in A} v_i$.

Lehmann et al. [2002] studied polynomial-time approximation mechanisms for this problem. They noted that, under mild complexity-theoretic assumptions, there is no polynomial-time algorithm that has an approximation factor significantly better than $\min\{d, \sqrt{m}\}$, where $d = \max_i |S_i|$ is the maximum bundle size [Hazan et al. 2006; Håstad 1999; Lehmann et al. 2002]. These negative results of course limit the best-case scenario of any approximation *mechanism* that runs in polynomial time, since it is solving an only harder problem (with private valuations). Lehmann et al. [2002] gave two mechanisms, based on appealingly simple greedy heuristics, that are DSIC and achieve approximation factors of d and \sqrt{m} .²

There are many more examples; see [Nisan et al. 2007, Chapters 11-12] for an incomplete list. While many different algorithmic techniques have been used to design DSIC approximation mechanisms, greedy algorithms continue to play a starring role (see [Borodin and Lucier 2010a,b] and the references therein).

Deferred-Acceptance Auctions — Even Better Than DSIC. Milgrom and Segal [2014] recently introduced a remarkable class of mechanisms for binary single-parameter problems they called *deferred-acceptance (DA) auctions*.³ A computer scientist might call them “adaptive backward greedy mechanisms” (cf., [Borodin and Lucier 2010a,b]): unlike common greedy algorithms that greedily accept the most promising candidate, these algorithms greedily reject the least promising candidate.

Milgrom and Segal [2014] were motivated by design challenges posed by the upcoming Federal Communications Commission (FCC) double auction,⁴ and introduced DA auctions for two reasons. The first reason is the computational intractability of the underlying welfare-maximization problem, the same *raison d’être* of algorithmic mechanism design. The second reason is that the incentive properties of DA auctions are superior to those of both forward greedy mechanisms and the VCG mechanism.

But wait, isn’t the DSIC property, possessed by the VCG mechanism and many forward greedy auctions, the “holy grail” of incentive properties in mechanism design? Not necessarily. The VCG mechanism, for example, is almost never deployed in real-world applications with distinct items. One obvious reason for this in settings with at least a modest number of items is that the communication and computational demands of the mechanism — generally exponential in the number of items — are a nonstarter.

¹A *fully polynomial-time approximation scheme (FPTAS)* computes, given an input parameter $\epsilon > 0$, a $(1 - \epsilon)$ -approximate solution in time polynomial in $\frac{1}{\epsilon}$ and in the input size.

²The results in Lehmann et al. [2002] hold even if the bidders’ bundles are private as well.

³The name is due to some similarities to the Gale-Shapley DA algorithm [Gale and Shapley 1962].

⁴See <http://wireless.fcc.gov/auctions/> for details.

This cannot be the whole story, however: the VCG mechanism is almost never used *even in* settings where the number of distinct items is so small that the communication and computational requirements are almost trivial (see, e.g., [Ausubel and Milgrom 2006]). Evidently, the VCG mechanism also suffers from flaws that are not complexity-theoretic; one example is vulnerability to coalitional deviations, even by groups of bidders that have no real stake in the market.

Example 1.1 (Incentive Issues of the VCG Mechanism). Consider two items (A and B) and three single-minded bidders: the first has value 1 for the bundle AB, the second and third have a very small value ϵ for A and B, respectively. If all bidders bid truthfully, then the VCG mechanism gives both items to bidder 1 at a price of 2ϵ ; this is the welfare-maximizing outcome. If bidders 2 and 3 collude and they both report the same false valuation which is greater than 1, however, the VCG mechanism allocates each of these bidders the item that it wants, and for free!

Milgrom and Segal [2014] prove that every DA auction possesses an impressive list of incentive guarantees beyond the basic DSIC guarantee.

- (1) Every DA auction is weakly group-strategyproof: there is no way for a coalition to submit a coordinated false bid that increases the utility of every bidder in the coalition (cf., Example 1.1).⁵
- (2) Every DA auction can be implemented by an ascending-clock auction. Ascending implementations are appreciated in practical applications for a number of reasons: bidders have an opportunity to implicitly share information and coordinate through price discovery; bidders can focus on determining their valuation only for bundles of items that are likely to be relevant, rather than for all bundles as in a direct-revelation mechanism; and, empirically, bidders bid more accurately in ascending auctions than in sealed-bid auctions (see [Cramton 2006] for further discussion).
- (3) The dominant-strategy outcome of a DA auction is the same as the unique Nash equilibrium that survives the iterated deletion of dominated strategies in the full-information game based on the same allocation rule and with the pay-as-bid payment rule. This can be interpreted as a robustness result for the prediction that bidders will play their dominant strategies in a DA auction.

Neither the VCG mechanism nor forward greedy algorithms generally possess any of these three properties.

Example 1.2 (Incentive Issues of Forward Greedy Algorithms). The forward greedy algorithm of Lehmann et al. [2002] ranks bidders by bid, possibly normalized by bundle size, and greedily accepts the highest-ranking bidders that do not cause a conflict with previously accepted ones; it then charges each accepted bidder the smallest bid that would still make it win. The example of Figure 1 shows an instance with four items (the black dots) and four single-minded bidders, whose bundles (the blobs) are labeled with their valuation. If the bidders reported their valuations truthfully, the forward greedy algorithm would accept bids 5 and 1, and the corresponding bidders would be charged 3 and 0. If the two losing bidders were to instead both outbid the highest bid, they would both win, and pay nothing! In general, the impact of group deviations on social welfare can be arbitrarily large.

How Good Are DA Auctions? The goal of this paper is to initiate the study of DA auctions from an approximation standpoint.

⁵This guarantee and its proof are in the spirit of previous work on Moulin mechanisms [Moulin 1999] and acyclic mechanisms [Mehta et al. 2007].




Fig. 1. Single-minded combinatorial auction with four bidders and four items.

Can DA auctions achieve welfare guarantees comparable to those of (forward) greedy mechanisms? More generally, can we attain the best of both worlds — optimal polynomial-time approximation guarantees and strong incentive guarantees — or is there an intrinsic trade-off between the strength of the incentive guarantees and the quality of approximation?

Reverse greedy algorithms have made sporadic cameos in the design of exact as well as approximation algorithms (e.g., [Tarjan 1983; Goemans and Williamson 1997; Charikar 2000; Sakai et al. 2003; Chrobak et al. 2005]) but, as far as we know, DA auctions provide the first reason to study their performance systematically. Indeed, for many problems, reverse greedy algorithms seem like unlikely candidates for good approximation algorithms — they generally don’t even return maximal solutions (see discussion in Section 2).

Our Results. We study the power and limitations of DA auctions through the lens of the two canonical single-parameter mechanism design problems described above. Our first set of results concerns knapsack auctions, where we prove a separation between the welfare approximation guarantee achievable by DA auctions and by arbitrary DSIC mechanisms. Our main lower bound states that no DA auction obtains an approximation sub-logarithmic in the number of items m . This lower bound is information-theoretic and applies to computationally unbounded DA auctions. Recall that forward greedy algorithms lead to a DSIC mechanism with an approximation guarantee of 2, and that general DSIC mechanisms can compute an arbitrarily close constant approximation in polynomial time and an optimal allocation in pseudo-polynomial time. We also give a polynomial-time DA auction with an approximation guarantee of $O(\log m)$.

Second, we consider welfare-maximization problems with single-minded bidders — the original application of greedy algorithms to algorithmic mechanism design. Recall that the forward greedy mechanisms in [Lehmann et al. 2002] achieve d - and \sqrt{m} -approximations of the optimal welfare, where d is the maximum bundle size and m is the number of items. These mechanisms have no incentive guarantees beyond the basic DSIC guarantee (as shown in [Milgrom and Segal 2014]), which proves that they are impossible to simulate in the DA auction framework. Moreover, we show that if the scoring rules in [Lehmann et al. 2002] — rank-by-bid, possibly normalized by a function of the bundle size — are re-deployed in a DA auction, then the approximation guarantees are very far from optimal. In fact, we prove that a class of natural DA auctions fails to match the performance guarantees of [Lehmann et al. 2002]. The only remaining approach to matching these guarantees is to design from scratch good approximation algorithms that can be implemented in the DA auction framework.

Our first DA auction provides an $O(d)$ -approximation when bidders’ desired bundles have size at most d . Conceptually, this auction has two phases (though it can still be implemented in the DA auction framework). The first phase is a sequence of “locally greedy” steps, with each step rejecting all bids but the highest among those whose bundles contain a given item. We prove that the sum of valuations of the remaining bidders is a $1/(d - 1)$ fraction of the optimal welfare, and that every remaining bidder

conflicts with at most two others. Our second phase selects a feasible subset of these bidders while losing an additional factor of 2.

Our second DA auction has an approximation guarantee of $O(\sqrt{m \log m})$, where m is the number of items. Our algorithm computes two solutions, one involving the bidders with small bundles (at most \sqrt{m} items, say), the other involving the bidders with large bundles; ideally, we would then like to take the better of these two. The DA auction framework cannot generally accommodate such “MAX operators,” however. Instead, we identify the highest bidder with a large bundle and use its bid as a welfare target for the rest — this is analogous to the “profit extractors” used for prior-free revenue maximization [Fiat et al. 2002]. We show that this idea can be implemented in the DA auction framework and gives an $O(\sqrt{m \log m})$ approximation.

2. PRELIMINARIES

Binary Single-Parameter Problems. The appropriate abstraction for studying the power and limitations of DA auctions is *downward-closed binary single-parameter problems*. Such a problem is described by a set N of bidders and a set system $\mathcal{I} \subseteq 2^N$. We assume that \mathcal{I} is non-empty and is *downward closed*, meaning that if $T \in \mathcal{I}$ and $T' \subseteq T$, then $T' \in \mathcal{I}$. Each bidder i has a private valuation v_i for “winning,” and the sets of \mathcal{I} are subsets of bidders that can feasibly win simultaneously. In a knapsack auction, for example, \mathcal{I} is given by all subsets of bidders that collectively want at most m items. In this paper, we study the welfare maximization problem: compute the set $A \in \mathcal{I}$ maximizing $SW(A)$, where $SW(A) = \sum_{i \in A} v_i$.

We focus on direct-revelation mechanisms $\mathcal{M} = (f, p)$, which comprise an outcome rule f and a payment rule p . In our setting, given a vector of bids $b = (b_i)_{i \in N} \in \mathbb{R}_+^n$, where b_i denotes the bid reported by bidder i , the outcome rule f computes a feasible solution (a set of \mathcal{I}). On the same input, the payment rule p computes payments $p = (p_i)_{i \in N} \in \mathbb{R}^n$ where p_i denotes the payment of bidder i .

Given a mechanism $\mathcal{M} = (f, p)$ and a bid vector b , bidder i 's utility depends on both its value for winning and the price it is asked to pay. We assume that the bidders have quasi-linear utilities and hence bidder i 's utility equals $u_i^{\mathcal{M}}(b) = v_i - p_i(b)$ if it is one of the accepted bidders (i.e., if $i \in f(b)$) and $u_i^{\mathcal{M}}(b) = 0$ otherwise.

We assume the bidders are strategic in reporting their bids to the mechanism, aiming to maximize their utility. A mechanism \mathcal{M} is *strategyproof* if for every bidder i , for all bid vectors $b_{-i} = (b_j)_{j \neq i}$, and all bids \bar{b}_i of bidder i

$$u_i^{\mathcal{M}}(v_i, b_{-i}) \geq u_i^{\mathcal{M}}(\bar{b}_i, b_{-i}).$$

A mechanism with an allocation rule f achieves an approximation ratio of ρ if

$$\max_v \frac{SW(OPT(v))}{SW(f(v))} \leq \rho,$$

where $OPT(v) = \arg \max_{B \in \mathcal{I}} \{SW(B)\}$ denotes a welfare-maximizing outcome.

Single-Minded Combinatorial Auctions. Let M denote the set of m items and N denote the set of n bidders. The bidders are single-minded: for each bidder $i \in N$ there exists a bundle of items $S_i \subseteq M$ such that bidder i 's valuation is $v_i(S'_i) = v_i$ for all $S'_i \supseteq S_i$ and $v_i(S'_i) = 0$ otherwise. We assume that the sets S_i are publicly known, while the valuations are private. A combinatorial auction with single-minded bidders corresponds to a downward-closed binary single-parameter problem in which \mathcal{I} is all subsets of bidders that desire pairwise disjoint bundles of items. We can assume that for any two bidders i and j , $S_i \neq S_j$. We use $s_i = |S_i|$ to denote the size of bidder i 's bundle, and let $d = \max_i \{s_i\}$ denote the size of the largest bundle across all bidders.

Given a problem instance of this form, it will be useful to think about two different types of graphs. The *bundle graph* G_b is an edge-weighted hypergraph whose vertices correspond to the set of items and whose hyperedges correspond to the n bundles of the single-minded bidders. The *conflict graph* G_c is a vertex-weighted graph whose vertices correspond to the set of bidders, and an edge (i, j) exists if and only if the bundles of bidders i and j are in conflict, i.e., $S_i \cap S_j \neq \emptyset$. The weight of a hyperedge in G_b that corresponds to the bundle of bidder i is the same as the weight of that bidder's vertex in G_c , and they are both equal to v_i . Finally, we let $c_i(G_c)$ denote the degree of vertex i in G_c (the number of conflicts with bidder i 's bundle).

Deferred-Acceptance Auctions. A DA auction is a particular kind of mechanism for a downward-closed binary single-parameter problem. It begins with all bidders being active, it operates in a sequence of stages, and after each stage it rejects some active bidder. This process continues, until the set of active bidders is feasible. At that point it accepts all the remaining bidders and charges each bidder its threshold price, i.e., its smallest winning bid. Losing bidders pay nothing.

Definition 2.1. A DA auction operates in stages $t \geq 1$. In each stage t a set of bidders $A_t \subseteq N$ is active; initially, $A_1 = N$. The DA auction is fully defined by a collection of deterministic scoring rules $\sigma_i^{A_t}(b_i, b_{N \setminus A_t})$ that are non-decreasing in their first argument. Stage t proceeds as follows:

- If A_t is feasible, accept the bidders in A_t and charge each bidder $i \in A_t$ its threshold price $p_i(b_i) = \inf\{b'_i \mid i \in A(b'_i, b_{-i})\}$, where $A(b'_i, b_{-i})$ denotes the set of bidders that would have been accepted if the reported bids were (b'_i, b_{-i}) instead of (b_i, b_{-i}) .
- Otherwise, set $A_{t+1} = A_t \setminus \{i\}$ where bidder $i \in \arg \min_{i \in A_t} \{\sigma_i^{A_t}(b_i, b_{N \setminus A_t})\}$ is an active bidder with the lowest score.⁶

We conclude this section with two examples. The first is an example of a setting in which forward greedy algorithms and DA auctions coincide. The second shows that we cannot implement the forward greedy algorithm of Lehmann et al. [2002] with a DA auction; and also that DA auctions unlike forward greedy algorithms are not guaranteed to produce a maximal solution. Non-maximality is a challenge to designing DA auctions, and for proving good approximation factors for them.

Example 2.2 (Matroid Settings). Suppose the set system \mathcal{I} is a matroid, meaning that it also satisfies the exchange property — whenever $S, T \in \mathcal{I}$ with $|S| < |T|$, there is a bidder $i \in T \setminus S$ with $S \cup \{i\} \in \mathcal{I}$. Then the forward and reverse greedy allocation rules coincide, for any scoring rule that is an increasing function of value only. Indeed, this fact together with the properties established in [Milgrom and Segal 2014] give a novel perspective on previous results stating that the VCG mechanism enjoys additional incentive properties in matroid environments (e.g., [Bikhchandani et al. 2011]).

Example 2.3 (Separation and Non-Maximality). Recall the instance of a single-minded combinatorial auction described in Figure 1. On this input the forward greedy algorithm that uses the scoring rule $\sigma_i^{A_t}(b_i) = b_i/s_i$ accepts bids 5 and 1. The same scoring rule deployed in a DA auction leads to a rejection of all the bids except bid 5 — a non-maximal solution. Note further that, since the input is completely symmetric, every (anonymous) DA auction must start by rejecting the lowest bid and its solution can therefore not coincide with that of this forward greedy algorithm.

⁶Ties can be broken arbitrarily but, for notational simplicity, we assume no two bidders have the same score.

3. KNAPSACK AUCTIONS

We begin this section by showing that, unlike forward greedy algorithms, DA auctions fail to achieve a constant factor approximation of the optimal social welfare in knapsack auctions; in particular, we show that DA auctions cannot implement a simple type of MAX operator. We then present a DA scoring rule that implements an approximate version of this operator, which we leverage in order to design a DA knapsack auction.

3.1. Inapproximability Using DA Auctions

In proving the inapproximability result of this section we identify the following class of problem instances which, as we show, poses a significant obstacle to DA auctions.

Definition 3.1. An *asymmetric set system* comprises two disjoint feasible sets of bidders $G_1, G_2 \in \mathcal{I}$ such that G_2 contains a single bidder, i.e., $|G_2| = 1$.

The fact that G_1 is feasible and that G_2 contains a single bid means that, if the DA auction were to reject the bid of G_2 , then it would be able to accept all the remaining bids of G_1 . The following lemma shows that for any given DA auction we can assign any multiset of bid values to the bidders of G_1 in a way that ensures these bidders may only be rejected in a smallest-bid-first order.

LEMMA 3.2. *Given any DA auction, any asymmetric set system, and any multiset of bid values $b_1 \leq \dots \leq b_{|G_1|}$ for the bidders of group G_1 , there exists an assignment of the bid values to the bidders of G_1 such that:*

- (1) A bidder $i \in G_1$ is never rejected before some other bidder $j \in G_1$ when $b_i > b_j$.
- (2) For any other multiset of bid values $b_1 \leq \dots \leq b_k \leq \bar{b}_{k+1} \leq \dots \leq \bar{b}_{|G_1|}$ that is identical w.r.t. its k smallest bid values, the assignment of these k bid values remains the same.

PROOF. In order to guarantee Property (1), we construct an assignment such that for every $t \leq |G_1|$, if i is the bidder who is assigned the t -th smallest bid value b_t and $j \in A_t$ is any other active bidder who is assigned bid value $b_{t'}$ for $t' > t$, then

$$\sigma_i^{A_t}(b_t, b_{N \setminus A_t}) < \sigma_j^{A_t}(b_{t'}, b_{N \setminus A_t}). \quad (1)$$

For $t = 1$, since $N \setminus A_t$ is empty, Inequality (1) becomes $\sigma_i^N(b_1) < \sigma_j^N(b_{t'})$. In order to satisfy this inequality, we assign b_1 to the bidder i that yields the minimum score *at this value*, i.e., $\sigma_i^N(b_1) < \sigma_j^N(b_1)$ for all other bidders $j \in G_1$. Since every bidder's DA scoring rule needs to be non-decreasing with respect to its bid value, and using the fact that every other bidder $j \in G_1$ will be assigned a value that is greater or equal to b_1 , we conclude that the score $\sigma_j^N(b_{t'})$ of every other bidder $j \in G_1$ at stage $t = 1$ will be at least $\sigma_j^N(b_1)$, so Inequality (1) is satisfied.

For $t > 1$ we assume that Inequality (1) holds up to stage $t-1$, which implies that the first $t-1$ rejected bidders have bid values b_1 to b_{t-1} . Therefore, we know what $b_{N \setminus A_t}$ is and we can use it in assigning the next value. We assign bid value b_t to the bidder $i \in A_t$ that yields the minimum score at this value, i.e.,

$$\sigma_i^{A_t}(b_t, b_{N \setminus A_t}) < \sigma_j^{A_t}(b_t, b_{N \setminus A_t}). \quad (2)$$

Finally, using the fact that $\sigma_j^{A_t}(\cdot)$ is a non-decreasing function of its first argument, Inequality (2) implies that Inequality (1) holds at stage t as well.

To verify that Property (2) is satisfied, note that the assignment of any bid value does not depend on bid values that have not yet been assigned. \square

Using this lemma, we can now prove the following inapproximability result.

THEOREM 3.3. *No DA knapsack auction can guarantee an approximation ratio of $\ln^\tau m$ for a positive constant $\tau < 1$.*

PROOF. Consider the problem instance induced by an asymmetric set system when G_1 contains m bidders of size 1, and group G_2 contains a single bidder of size m . The only maximal solutions for this instance correspond to either accepting all the bids of G_1 and rejecting the bid of G_2 , or accepting the bid of G_2 and rejecting all the bids of G_1 . We prove that no DA scoring rule can always extract a $1/\ln^\tau m$ fraction of $\max\{SW(G_1), SW(G_2)\}$.

In order to assign the bid values (provided below) to the bidders of G_1 we assume that the assignment method described in Lemma 3.2 is used, so the decision of the DA auction at any stage reduces to either rejecting the active bid of G_1 with the smallest bid value, or rejecting the bid of G_2 . Let $c_t = |A_t \setminus G_2|$ denote the number of remaining active bids in G_1 at stage t of the DA auction, e.g., $c_1 = m$. We partition the bidders of G_1 into κ sub-groups, where $\kappa = \lceil 1/(1-\tau) \rceil$ is a constant. In defining which sub-group each bidder belongs to we use the value of c_t at which the bidder would be considered for rejection. Sub-group $\beta \in \{1, 2, \dots, \kappa\}$ includes the bids of G_1 for which $c_t \in \mathcal{C}(\beta)$, where

$$\mathcal{C}(\beta) = \left(\sqrt[\kappa]{m}^{\beta-1}, \sqrt[\kappa]{m}^\beta \right].$$

The highest value bid, i.e., the bid for which $c_t = 1$ belongs to sub-group $\beta = 1$. Therefore, while the DA auction does not reject the bid of G_2 , it will reject bids of G_1 starting from sub-group κ and gradually moving down to sub-group 1.

We restrict m to be a power of κ , thus ensuring that the endpoints of $\mathcal{C}(\beta)$ are integral for every β . Also, for simplicity, the rest of this proof treats the $\sqrt[\kappa]{m}^\beta$ -th harmonic number as if it were equal to $\frac{\beta}{\kappa} \ln m$ for every sub-group β ; this is without loss of generality.⁷ Using this assumption, and according to the definition of the interval $\mathcal{C}(\beta)$, for every sub-group β

$$\sum_{c_t \in \mathcal{C}(\beta)} \frac{1}{c_t} \approx \ln m^{\beta/\kappa} - \ln m^{(\beta-1)/\kappa} = \frac{\beta}{\kappa} \ln m - \frac{\beta-1}{\kappa} \ln m = \frac{\ln m}{\kappa}. \quad (3)$$

Let $L = \frac{1}{\kappa} \ln^{1-\tau} m$ and, for some very small positive constant $\epsilon < 1$, let the value of the bidder of group β that is considered for rejection in round t be

$$\frac{(1-\epsilon^\beta)L^{\kappa-\beta}}{c_t \cdot \ln^\tau m}. \quad (4)$$

Using Equation (3), the total value of the bids in sub-group β is equal to

$$\sum_{c_t \in \mathcal{C}(\beta)} \frac{(1-\epsilon^\beta)L^{\kappa-\beta}}{c_t \cdot \ln^\tau m} \approx \frac{\ln m}{\kappa} \cdot \frac{(1-\epsilon^\beta)L^{\kappa-\beta}}{\ln^\tau m} = (1-\epsilon^\beta)L^{\kappa-\beta+1}, \quad (5)$$

and the total welfare in G_1 , even if we only count the value in group $\beta = 1$ is

$$SW(G_1) \geq (1-\epsilon)L^\kappa = (1-\epsilon) \frac{\ln^{\kappa(1-\tau)} m}{\kappa^\kappa} \geq (1-\epsilon) \frac{\ln m}{\kappa^\kappa}. \quad (6)$$

We now show that, even if we let the value of the bid in G_2 be equal to 1, i.e., $SW(G_2) = 1$, any DA auction that guarantees a $\ln^\tau m$ approximation will have to re-

⁷It is known that $\lim_{k \rightarrow \infty} (H_k - \ln k) \approx 0.577$, so one can express H_k as $(1 \pm \epsilon') \ln k$ where ϵ' becomes arbitrarily small as k goes to infinity. The arguments of the proof remain true as long as m , and hence also $\sqrt[\kappa]{m}^\beta$, is large enough to make ϵ' substantially smaller than ϵ^κ , where ϵ is the constant used in (4).

ject all of the bids in G_1 . In order to show this we make repeated use of the fact that the bid values of G_1 are assigned in a way that satisfies Lemma 3.2. In particular, Property (2) of this lemma implies that, at any stage t , the decisions of the DA auction up to this stage, as well as its upcoming decision regarding whether it will reject the smallest-value active bid b_i of G_1 or the bid of G_2 do not depend on the values of the $c_t - 1$ other active bids. Therefore, if the mechanism chooses to reject the bid of G_2 at round t when facing the instance described above, then it would do the same *even if the remaining $c_t - 1$ active bid values were all equal to b_i* . As a result, an $\ln^\tau m$ approximation mechanism can reject the bid of G_2 at some stage t only if $c_t \cdot b_i$ is at least $\max\{SW(G_1), SW(G_2)\} / \ln^\tau m$.

First, assume that at some round t the mechanism rejects the bid of G_2 in favor of a bid $i \in G_1$ that belongs to sub-group κ . According to (4), the value of bid i is $b_i = (1 - \epsilon^\kappa) / (c_t \cdot \ln^\tau m)$. Therefore, if all the c_t active bids of G_1 also had a value of b_i , then the total welfare extracted by the mechanism would be $c_t \cdot b_i = (1 - \epsilon^\kappa) / \ln^\tau m$. This would not be a $\ln^\tau m$ approximation of $SW(G_2) = 1$ leading to a contradiction.

Similarly, assume that at some round t the mechanism rejects the bid of G_2 in favor of a bid $i \in G_1$ of any other sub-group $\beta < \kappa$. By the time that any bid of sub-group β is considered for deletion though, all the bids of sub-group $\beta + 1$, whose total value according to (5) is $(1 - \epsilon^{\beta+1})L^{\kappa-\beta}$, have already been considered and rejected. This implies that, if the mechanism had instead accepted all of the bids of G_1 , then the extracted welfare would have been at least $(1 - \epsilon^{\beta+1})L^{\kappa-\beta}$, and in order to guarantee the desired approximation factor the mechanism would hence have to extract a welfare of at least $[(1 - \epsilon^{\beta+1})L^{\kappa-\beta}] / \ln^\tau m$. But, similarly to the previous argument, the value of bidder i in sub-group β is $b_i = [(1 - \epsilon^\beta)L^{\kappa-\beta}] / (c_t \cdot \ln^\tau m)$, and if all the other c_t active bidders of G_1 also had a value of b_i the total welfare extracted by the mechanism would be $[(1 - \epsilon^\beta)L^{\kappa-\beta}] / \ln^\tau m$, which would not be $\ln^\tau m$ approximation.

Since the DA auction rejects all the bids of G_1 , the welfare that it extracts is $SW(G_2) = 1$. But, according to (6) the total value of bids in G_1 is $SW(G_1) \geq (1 - \epsilon) \frac{\ln m}{\kappa^\kappa}$. Since κ and ϵ are constants, for large enough values of m the value of $SW(G_1)$ is greater than $\ln^\tau m$, so this mechanism does not guarantee a $\ln^\tau m$ approximation. \square

3.2. Approximate MAX Operator

At the core of the proof of Theorem 3.3 lies the inability of DA auctions to implement a simple MAX operator for asymmetric set systems. Given an asymmetric set system and some DA auction, let c_t denote the number of active bids in G_1 at stage t of the auction that are in conflict with the bid of G_2 . In what follows, we propose a DA scoring rule which implements an approximate MAX operator $AM(G_1, G_2)$ that extracts an $\Omega(1/\log m)$ fraction of $\max\{SW(G_1), SW(G_2)\}$ for any asymmetric set system; we then use this operator in the rest of the paper as a subroutine.

$$AM(G_1, G_2) \text{ is induced by } \sigma_i^{A_t}(b_i) = \begin{cases} b_i & \text{if } i \in G_1 \\ b_i/c_t & \text{if } i \in G_2 \end{cases} \quad (7)$$

THEOREM 3.4. *The $AM(G_1, G_2)$ operator is guaranteed to extract a welfare that is both $\Omega(SW(G_1)/\log m)$, and at least $SW(G_2)$ for any asymmetric set system.*

PROOF. Let $V_1 = SW(G_1)$ and $V_2 = SW(G_2)$ denote the welfare generated by the bids in G_1 and G_2 respectively.

We begin by showing that the value that $AM(G_1, G_2)$ extracts will always be at least V_2 . Clearly, if the one bid of G_2 is never rejected then the value of the accepted bids will be at least V_2 . If, on the other hand, the bid of G_2 is rejected at some stage t then its score is V_2/c_t , and hence there exist at least c_t active bids in G_1 whose score is at least

V_2/c_t (since they were not rejected instead). Therefore, since the score of bids in G_1 is equal to their value, even if the bid of G_2 is rejected, the total value of the remaining bids (all of which are accepted) is at least $c_t \cdot V_2/c_t = V_2$.

We now also show the value that $\text{AM}(G_1, G_2)$ extracts will always be $\Omega(V_1/\log m)$. Note that, as long as the bid of G_2 is not rejected, the score of each bid $i \in G_1$ that was rejected at some stage t is at most V_2/c_t . Therefore, the maximum total value that can ever be rejected by this scoring rule is

$$\sum_{c_t=1}^m \frac{V_2}{c_t} \in O(V_2 \log m).$$

If V_e denotes the welfare extracted by this operator and even if we assume that V_e was extracted from bids of G_1 , we still get that $V_1 = O(V_e + V_2 \log m)$. Since we have already shown that $V_e \geq V_2$, this proves the desired bound:

$$\frac{V_e}{V_1} \in \Omega\left(\frac{V_e}{V_e + V_2 \log m}\right) = \Omega\left(\frac{V_2}{V_2 + V_2 \log m}\right) = \Omega(1/\log m). \quad \square$$

3.3. DA Knapsack Auction

We now propose a DA knapsack auction, which begins by partitioning the set of bidders into two groups, G_1 and G_2 . Group G_1 contains all the bids of size at most $m/2$, and group G_2 contains all the remaining bids. Then, the DA auction applies a different scoring rule on each one of the two groups. The score for each bid i in G_1 is equal to b_i/s_i (rank by density), and the DA auction keeps rejecting the lowest scoring bid in G_1 until the remaining set of active bids in G_1 is feasible. On the other hand, the bids in G_2 are all rejected, except the one with the highest value. Finally, the DA auction combines the remaining active bids $G'_1 \subseteq G_1$ and $G'_2 \subseteq G_2$ using the $\text{AM}(G'_1, G'_2)$ operator.

Algorithm 1: DA Knapsack Auction

- 1 Let G_1 be the set of bids i for which $s_i \leq m/2$ and let $G_2 = N \setminus G_1$.
 - 2 Use scoring rule b_i/s_i on G_1 , rejecting all but a feasible set of bids $G'_1 \subseteq G_1$.
 - 3 Reject all but the highest value bid of G_2 , and place that single bid in $G'_2 \subseteq G_2$.
 - 4 Use the scoring rule of (7) on G'_1 and G'_2 to implement the $\text{AM}(G'_1, G'_2)$ operator.
-

THEOREM 3.5. *The DA knapsack auction achieves a $O(\log m)$ approximation.*

PROOF. We first show that $\max\{SW(G'_1), SW(G'_2)\}$ is a 2-approximation.

If the total size of the bids in G'_1 , i.e., $\sum_{i \in G'_1} s_i$, is at least $m/2$, then $SW(G'_1)$ is at least half of the optimal social welfare. To verify this fact, note that the bids in G'_1 are the ones with the highest value per size density so, even if the optimal solution used up all of the knapsack, it would be with less or equal value per size and it would therefore lead to at most twice the value.

If on the other hand the total size of the bids in G'_1 is less than $m/2$, then $G'_1 = G_1$ since any bid in G_1 would fit in the remaining space in the knapsack. Hence, if the optimal solution does not accept any bids from G_2 , then $SW(G'_1)$ is equal to the optimal social welfare. If the optimal solution accepts a bid $i \in G_2$ (clearly it cannot accept more than one due to their size), then $SW(G'_2)$ is at least as much as the value of i , so it must be the case that $\max\{SW(G'_1), SW(G'_2)\}$ is at least a 1/2 fraction of the optimal social welfare in this case as well.

Finally, according to Theorem 3.4, the $\text{AM}(G'_1, G'_2)$ operator always extracts an $\Omega(1/\log m)$ fraction of $\max\{SW(G'_1), SW(G'_2)\}$, which concludes the proof. \square




Fig. 2. Single-minded combinatorial auction with n bidders and $n + 1$ items.

4. LIMITATIONS OF DA AUCTIONS FOR SINGLE-MINDED BIDDERS

The rest of the paper focuses on auctions involving single-minded bidders; in this section we begin by providing some intuition regarding the limitations that the standard greedy algorithm scoring rules, as well as many of their natural generalizations face when used for the design of DA auctions. The most significant restriction that a DA scoring rule enforces is that, although the score $\sigma_i^{A_t}(b_i, b_{N \setminus A_t})$ of bidder i at stage t may depend on the bundles that the other bidders want, it cannot depend on the values that the active bidders have reported for their bundles.

We first provide an example verifying that casting the forward greedy algorithm of Lehmann et al. [2002] into the DA auction framework does not guarantee efficiency.

PROPOSITION 4.1. *For any $p \geq 0$ the approximation ratio of the DA auction that uses scoring rule b_i/s_i^p is $\Omega(m)$ even if the maximum bundle size is $d = 2$.*

PROOF. Consider the instance of Figure 2 involving n bidders and $m = n + 1$ items. Bidder $i \in \{1, \dots, n\}$ has a valuation of $1 - i\epsilon$ for the bundle consisting of item i and item $i + 1$ (for some arbitrarily small constant $\epsilon > 0$). All the bids have the same size so, starting from bidder n , the DA auction will gradually reject every bidder except bidder 1 for a social welfare of approximately 1. In contrast, the optimal outcome is attained by accepting every bidder i s.t. i is odd, leading to a welfare of approximately $\frac{m}{2}$. \square

A natural way to leverage the structure of the conflict graph would be to use the number of active conflicts of a bid in order to appropriately scale its score. For example, one possible improvement of the b_i/s_i^p scoring rule is the reverse greedy algorithm according to which, the score of bidder i is the ratio $b_i/[c_i(c_i + 1)]$ where $c_i = c_i(G_c)$ is the number of its active conflicts.

This scoring rule, which makes bids with a large number of conflicts look less appealing, was studied by Sakai et al. [2003] who showed that it guarantees an approximation factor of $\Delta(G_c) = \max_i\{c_i(G_c)\}$, the maximum number of conflicts across all bidders. Clearly, this scoring rule performs much better than b_i/s_i^p on problem instances with few conflicts such as the one described in Figure 2 that we used in order to prove Proposition 4.1.

On the other hand, considering the star graph, we see that $\Delta(G_c)$ can be arbitrarily larger than the maximum bundle size d and the number of items m and thus this guarantee is not useful for our purposes. In fact, the following theorem, for which we provide a proof sketch in the appendix, shows that neither this scoring rule nor any other scoring rule from a broader class of natural scoring rules can achieve the performance guarantees that we manage to get in the following sections.

THEOREM 4.2. *No DA auction with scoring rules of the form $\frac{b_i}{f(c_i)g(s_i)}$, where $f(x) \in \Theta(x^\gamma)$ and $g(x) \in \Theta(x^\delta)$ for $\gamma, \delta \in \mathbb{R}$, can guarantee a $O(d)$ or $o(m)$ approximation.*

5. LOCALLY HIGHEST BID: A $O(d)$ APPROXIMATION MECHANISM

In this section, we propose a mechanism that aims to guarantee a good approximation of the optimal social welfare using the maximum bundle size d as a parameter. We first describe and analyze the mechanism. Afterward we show that a slightly modified version can be implemented as a DA auction, and achieves (almost) the same performance guarantee.




Fig. 3. On the left is the bundle graph for the first phase of the LHB mechanism. In this example we have $d = 3$ and the bids appear as triangles containing three items each. The mechanism considers the items in the order $\{a\}$, $\{b, c\}$, $\{d, e\}$, $\{g, h\}$, $\{f\}$ and $\{i, j\}$. Grey ellipses in the left graph contain items that the mechanism considers in tandem. Solid triangles correspond to marked bids, while dashed triangles are rejected bids. On the right is the partition graph G_p induced by the first phase. Bold edges in this graph correspond to the maximum weight matching.

5.1. The Locally Highest Bid Mechanism

The *Locally Highest Bid* (LHB) mechanism follows two different phases; see Figure 3 for an example. The first phase prunes the bundle graph by greedily rejecting all but the locally highest bid.⁸ The second phase translates the resulting hypergraph into a bipartite graph such that matchings in this graph correspond to feasible solutions, and computes a maximum weight matching in this graph.

First Phase. The first phase of the LHB mechanism (see Steps 1-10 of Algorithm 2) begins by considering some arbitrary item u and, among all the bids that contain this item, it rejects all but the bid b_1 with the highest bid value. The mechanism marks bid b_1 to denote that it should not be considered again for rejection during the first phase. We will refer to bids that are neither rejected nor marked as *candidate* bids. The bundle of bid b_1 may contain q other items apart from u , where $q \leq d - 1$ since the size of this bundle is at most d . The next step of our mechanism considers all these q new items in tandem, and it rejects all the candidate bids that contain at least one of these q items, except the bid b_2 with the highest value among them, which is marked instead. Once again, bid b_2 may contain some new items that have not been considered by the mechanism in the past (i.e., different from the $q + 1$ items that have been considered at this point). As long as this is the case, the mechanism considers all these new items introduced by the latest marked bid in tandem, it rejects all but their locally highest bid, and keeps repeating the steps described above. If, on the other hand, either a marked bid does not introduce any new items or the new items it introduces are contained in no candidate bids, the mechanism picks an arbitrary item that has not been considered yet and repeats the same process until all the items have been considered.

Second Phase. The first phase of the LHB mechanism quite naturally induces the following edge-weighted *partition graph* G_p : items that are considered in tandem are represented by a vertex and there is an edge of weight w between two vertices x and y if there is a marked bid of value $b_i = w$ that contains items in the sets of both x and y . The second phase (see Steps 11-12 of Algorithm 2) concludes by computing a maximum weight matching of G_p , and accepting the corresponding set of bids, which is a feasible solution since, by definition of G_p , no two such bids share an item.

⁸Our algorithm can therefore be seen as a reverse greedy variant of an appropriate generalization of an algorithm by Drake and Hougardy [2003] for maximum weight matching, i.e., the special case where $d = 2$.

Algorithm 2: LHB mechanism

- 1 Let all the bids be initially unmarked, and let u be a pointer to an arbitrary item.
 - 2 **if** item u is not contained in any candidate bids **then**
 - 3 | Point u to any other item that has not been pointed to before.
 - 4 Reject all candidate bids that contain item u except the one with the highest value.
 - 5 The bid b that was not rejected contains $q \leq d - 1$ new items.
 - 6 **if** $q > 0$ **then**
 - 7 | Contract the q original items into one item and point u to this item⁹.
 - 8 | Mark bid b and continue with Step 2.
 - 9 **else if** there exists some item that has not been pointed to **then**
 - 10 | Point u to that item and continue with Step 2.
 - 11 Let G_p be the partition graph induced by the first phase of the mechanism.
 - 12 Accept the bids that correspond to the maximum weight matching of G_p .
-

THEOREM 5.1. *The LHB mechanism guarantees a $2(d - 1)$ approximation.*

For the proof of this theorem we use the following auxiliary lemma.

LEMMA 5.2. *The partition graph G_p is a forest of path graphs.*

PROOF. Let b_k denote the k -th bid to be marked during the first phase. In order to prove this lemma, we first focus on the first item that the mechanism considers. Since at most one of the bids that contain this item survives the first phase of the mechanism, the edge corresponding to b_1 is the only one that covers the vertex of this item in G_p . Similarly, the vertex of G_p that corresponds to the q items that bid b_1 introduces may be covered only by b_1 and b_2 . Therefore, among the bids that survive the first phase of the mechanism, bid b_1 will be in conflict only with bid b_2 . Using the same arguments one can verify that for $k \geq 2$, bid b_k may be in conflict only with bids b_{k-1} and b_{k+1} , which implies that G_p will be a forest of path graphs. \square

PROOF OF THEOREM 5.1. The first phase of the LHB mechanism proceeds by rejecting and marking bids until all the bids are either rejected or marked. In order to prove this theorem, whenever some bid is rejected or marked, we assign this bid to the set of items that were being considered (in tandem) when this took place. Therefore, after the completion of the first phase of LHB, every bid has been assigned to one of the sets of up to $d - 1$ items that were being considered in tandem.

Since the set $OPT(v)$ of bids that maximize the social welfare is non-conflicting, it must be the case that at most $d - 1$ such bids have been assigned to the same set of items. To verify this fact note that each such set consists of at most $d - 1$ items, and that the bundle of every bid that is assigned to some set of items has to contain at least one of its items. If d or more bids of $OPT(v)$ were assigned to the same set of items, this would then mean that the bundles of at least two of these bids are sharing one of the items, a contradiction.

Also, note that the mechanism ensures that for every set of items that has rejected bids assigned to it, there is also a marked bid which has been assigned to it, and whose value is greater than or equal to that of any one of those rejected bids (as it is the locally highest one). This implies that, once the first phase is completed, the total value of the marked bids, and hence the total weight of G_p is at least a $1/(d - 1)$ fraction of $SW(OPT(v))$.

⁹The new item is then contained in all the bids that contained at least one of the original q items.

Finally, Lemma 5.2 implies that G_p is a bipartite graph with a maximum degree of 2, and hence the maximum weight matching will extract at least half of the total weight. Since the total weight of the edges of G_p is at least a $1/(d-1)$ fraction of $SW(OPT(v))$, the total value of bids that correspond to the maximum weight matching is at least a $1/[2(d-1)]$ fraction of $SW(OPT(v))$. \square

5.2. Implementation as a DA auction.

We now show how the mechanism described above can actually be implemented as a DA auction with an additional loss of a factor of 2 in the approximation guarantee. For the first phase of this mechanism we show that there exists an appropriate scoring rule that implements exactly the same sequence of steps. For the second phase we provide a simple scoring rule that can be used in order to extract a feasible solution from the marked bids which is a 2-approximation of the maximum weight matching in the partition graph.

First Phase. The DA auction implementation of the first phase of the LHB mechanism imitates every step of the mechanism: whenever it considers some set of items (in tandem), the score of any candidate bid that contains none of these items is set to infinity, while the score of the other candidate bids is equal to their value. This scoring rule is used as long as there exist at least two candidate bids that contain items being considered; once this is not the case anymore the DA auction has rejected all but the locally highest bid. Using the structure of the bundle graph, the DA auction can then identify the new items (if any) introduced by the locally highest bid and the scoring rule changes in order to consider these items instead.

The key observation here is that, at any given point, the score of some bidder i essentially depends only on where its bid is positioned in the bundle graph and on the value b_i of that bid alone. The only step of the LHB mechanism that uses the values of the reported bids is the step that rejects all but the locally highest bidder and the scoring rule described above provides exactly the same outcome. Given this observation, the only remaining subtle point is to clarify how the DA auction keeps track of which bids have been marked. The answer is that these bids are the only non-rejected ones that contain items which have already been considered.

Algorithm 3: DA auction implementation of the first phase of the LHB mechanism

- 1 Let all the bids be initially unmarked, and let u be a pointer to an arbitrary item.
 - 2 **if** item u is not contained in any candidate bids **then**
 - 3 | Point u to any other item that has not been pointed to before.
 - 4 **while** there exist more than one candidate bids containing item u **do**
 - 5 | The score of any candidate bid that does not contain u is equal to infinity.
 - 6 | The score of any candidate bid that contains u is equal to the value of its bid.
 - 7 | Reject the bid with the lowest score value.
 - 8 The bid b that was not rejected contains $q \leq d-1$ new items.
 - 9 **if** $q > 0$ **then**
 - 10 | Contract the q original items into one item and point u to this item.
 - 11 | Mark bid b and continue with Step 2.
 - 12 **else if** there exists some item that has not been pointed to **then**
 - 13 | Point u to that item and continue with Step 2.
-

Second Phase. For the second phase of the mechanism we show how to approximately implement the mechanism using a DA auction. In specific, Lemma 5.2 implies that the conflict graph G_c of the bidders that survived the first phase of the mechanism has a maximum degree $\Delta(G_c)$ of at most two. In this much more convenient setting we can

now use a DA auction according to which, the score of bidder i is the ratio $b_i/[c_i(c_i+1)]$, where b_i is the value of its bid and $c_i = c_i(G_c)$ is the number of other active bids that it is in conflict with. As we discussed in Section 4, this auction guarantees an approximation factor of $\Delta(G_c)$, so it yields a 2-approximation of the maximum weight matching. This additional loss of a factor of 2 implies the following theorem.

THEOREM 5.3. *The DA auction implementation of the LHB mechanism guarantees a $4(d-1)$ approximation.*

6. DIVIDE AND WEIGH: A $O(\sqrt{m \log m})$ APPROXIMATION MECHANISM

In this section we provide a mechanism which aims to guarantee a good approximation of the optimal social welfare using the number of items m as a parameter. In order to achieve an approximation factor better than $O(m)$ which, as we discussed in Section 4, none of the previously considered scoring rules can achieve, we build upon the LHB mechanism of the previous section.

The *Divide and Weigh* (DW) mechanism begins by partitioning the set of bids into two groups G_1, G_2 . Group G_1 contains all the bids of size $s_i \leq \sqrt{m/\log m}$ and group G_2 contains all the remaining bids. Then, our mechanism uses three different DA auctions as subroutines. First, it applies the LHB auction on the bids of group G_1 ; second, it rejects all but the highest bid from the bids of group G_2 ; finally, it combines the corresponding solutions using the approximate MAX operator described in Section 3.2.

Algorithm 4: DW mechanism

- 1 Let G_1 be the set of bids i for which $s_i \leq \sqrt{m/\log m}$ and let $G_2 = N \setminus G_1$.
 - 2 Run the LHB mechanism on G_1 , rejecting all but a feasible set of bids $G'_1 \subseteq G_1$.
 - 3 Reject all but the highest value bid of G_2 , and place that single bid in $G'_2 \subseteq G_2$.
 - 4 Use the scoring rule of (7) on G'_1 and G'_2 to implement the $AM(G'_1, G'_2)$ operator.
-

THEOREM 6.1. *The DW mechanism guarantees a $O(\sqrt{m \log m})$ approximation.*

PROOF. Let V_1, V_2 denote the maximum value that one could extract from a feasible subset of bids of G_1 , and G_2 respectively, and let $V \leq V_1 + V_2$ be the maximum value ($OPT(v)$) that could be extracted from $G_1 \cup G_2$. In what follows, we prove that the value extracted by the DW mechanism is $\Omega(V/\sqrt{m \log m})$.

We first point out that, since the size of every bid in G_1 is at most $\sqrt{m/\log m}$, Theorem 3.4 implies that the total value of the bids in G'_1 , i.e., $SW(G'_1)$ is $\Omega(V_1/\sqrt{m/\log m})$. Also, since the size of every bid in G_2 is more than $\sqrt{m/\log m}$, any feasible subset of bids in G_2 contains less than $\sqrt{m \log m}$ bids; otherwise, this subset would have to contain two bids whose bundles intersect. This implies that, since the bid in G'_2 is the highest value bid of G_2 , then its value is more than $V_2/\sqrt{m \log m}$.

Finally, according to Theorem 3.4, the value extracted by $AM(G'_1, G'_2)$ is guaranteed to be an $\Omega(1/\log m)$ fraction of $SW(G'_1)$, i.e., $\Omega(V_1/\sqrt{m \log m})$, and at least $SW(G'_2)$, i.e., $\Omega(V_2/\sqrt{m \log m})$. Since $V_1 + V_2 \geq V$ this completes the proof. \square

Note that, if we were willing to settle for a randomized mechanism, the algorithm that flips a fair coin and, based on the outcome, either accepts all the bids of G'_1 or accepts the bid of G'_2 , would achieve a $1/2$ approximation of $\max\{SW(G'_1), SW(G'_2)\}$ in expectation, while retaining all the other nice properties that the DA auctions provide. In that case, a $O(\sqrt{m})$ -approximation could be obtained by adjusting the DW mechanism so that G_1 contains all the bids of size at most \sqrt{m} and G_2 contains the rest. The arguments in the proof of Theorem 6.1 would then imply that $\max\{SW(G'_1), SW(G'_2)\}$, which we could approximate within a constant, is an $\Omega(1/\sqrt{m})$ fraction of $OPT(v)$.

7. CONCLUSION

Deferred-acceptance auctions guarantee strong incentive properties beyond the basic dominant-strategy incentive-compatibility (DSIC) guarantee. This paper proposed the research agenda of understanding the power and limitations of these auctions from an approximation perspective. We provided both positive and negative results for welfare-maximization in two canonical problems: knapsack auctions and combinatorial auctions with single-minded bidders.

Looking forward, one natural direction is to consider other binary single-parameter problems for which good DSIC mechanisms are known. More generally, what type of problem structure lends itself to good deferred-acceptance auctions? Are there any approximate “black-box reductions,” which convert (say) a forward greedy mechanism into a deferred-acceptance auction, with some loss in the approximation guarantee?

Also, problems that are not binary single-parameter pose an intriguing challenge. Deferred-acceptance auctions have not been defined for such problems. Even for problems where the definition may seem “obvious,” such as combinatorial auctions with single-minded bidders with unknown (private) desired bundles, the present definition does not guarantee a DSIC mechanism, let alone the stronger incentive properties enjoyed by deferred-acceptance auctions for single-parameter problems. Finally, it is unclear if non-trivial weakly group-strategyproof mechanisms exist for interesting multi-parameter problems.

ACKNOWLEDGMENTS

The authors greatly benefited from discussions with Paul Milgrom and Ilya Segal. This research is supported in part by NSF Awards CCF-1016885 and CCF-1215965, an ONR PECASE Award, and an SNF Postdoctoral Fellowship.

REFERENCES

- L. M. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. In *Combinatorial Auctions, Chapter 1*, pages 17–40, 2006.
- S. Bikhchandani, S. de Vries, S. J. Schummer, and R. V. Vohra. An ascending vickrey auction for selling bases of a matroid. *Operations Research*, 59:400–413, 2011.
- A. Borodin and B. Lucier. Price of anarchy for greedy auctions. In *SODA*, pages 537–553, 2010a.
- A. Borodin and B. Lucier. On the limitations of greedy mechanism design for truthful combinatorial auctions. In *ICALP (1)*, pages 90–101, 2010b.
- P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *STOC*, pages 39–48, 2005.
- M. Charikar. Greedy approximation algorithms for finding dense components in a graph. In *APPROX*, pages 84–95, 2000.
- M. Chrobak, C. Kenyon, and N. E. Young. The reverse greedy algorithm for the metric k -median problem. In *COCOON*, pages 654–660, 2005.
- P. Cramton. Ascending proxy auctions. In *Combinatorial Auctions, Chapter 4*, pages 79–98, 2006.
- D. E. Drake and S. Hougardy. A simple approximation algorithm for the weighted matching problem. *Information Processing Letters*, 85(4):211–213, 2003.
- A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin. Competitive generalized auctions. In *STOC*, pages 72–81, 2002.
- D. Gale and L. S. Shapley. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- M. X. Goemans and D. P. Williamson. The primal-dual method for approximation algorithms and its application to the network design problems. In *Approximation*

- Algorithms for NP-Hard Problems, Chapter 4*, pages 144–191, 1997.
- J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- E. Hazan, S. Safra, and O. Schwartz. On the complexity of approximating k -set packing. *Computational Complexity*, 15(1):20–39, 2006.
- D. Lehmann, L. I. O’Callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM*, 49:577–602, 2002.
- A. Mehta, T. Roughgarden, and M. Sundararajan. Beyond moulin mechanisms. In *EC*, pages 1–10, 2007.
- P. Milgrom and I. Segal. Deferred-acceptance auctions and radio spectrum reallocation. In *EC*, 2014. To appear.
- H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16(2):279–320, 1999.
- A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. In *AAAI/IAAI*, pages 379–384, 2002.
- N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- S. Sakai, M. Togasaki, and K. Yamazaki. A note on greedy algorithms for maximum weighted independent set problem. *Discrete Applied Mathematics*, 126(2–3), 2003.
- R. E. Tarjan. Foundations. In *Data structures and network algorithms, Chapter 1*, 1983.

A. PROOF SKETCH FOR THEOREM 4.2

Lower Bound of $\omega(d)$. We distinguish two cases: $\gamma > 1$ and $\gamma \leq 1$. In both cases, all bids will have the same size, so we disregard the contribution of $g(s_i)$ to the score.

Argument for $\gamma > 1$. Consider the following problem instance with d^2 items which we partition into d rows of d items each. For each row, there are C bidders whose bundle consists of exactly its d items;¹⁰ each one of these dC bidders has a value of 1. Also, there are d more bidders, each of which is interested in a distinct column of d items, and they all have a value of μd^γ . Therefore, the dC “row bidders” initially have $C + d - 1$ conflicts each and a score of $1/f(C + d - 1)$, while the d “column bidders” have dC conflicts and a score of $\mu d^\gamma/f(dC)$.

For any given value of d , let C be arbitrarily large. Since $f(x) \in \Theta(x^\gamma)$, for large enough values of C there exists some constant $\bar{\mu}$ such that $f(dC)/f(C + d - 1) > \bar{\mu} d^\gamma$. If we let $\mu = \bar{\mu}$, this implies that the score of the row bidders will be greater than that of the column bidders. As a result, facing such an instance, the DA scoring rule $b_i/f(c_i)$ would reject all the column bidders before rejecting any row bidder.

Once this happens, the best possible outcome is to accept one bidder from each row, leading to a social welfare of d . If, on the other hand, all the column bidders had been accepted, the social welfare would be $\mu d^{\gamma+1}$, so the approximation ratio is at least d^γ .

Argument for $\gamma \leq 1$. Consider the following problem instance whose conflict graph forms a tree. The root has d children and any other vertex except the leaves has $d - 1$ children and 1 parent (i.e., all the internal vertices have degree d). The weights of the vertices are defined recursively in the following fashion: the weight of the root is 1 and the weight of the k -th of its d children is $(1 - \epsilon)/f(k)$, where $\epsilon > 0$ is some constant which is arbitrarily smaller than the smallest vertex weight. Also the weight of the k -th out of $d - 1$ children of any other internal vertex with weight b_i is $(b_i - \epsilon)/f(k + 1)$.

¹⁰If we wanted to ensure that no two bidders have the exact same bundle, we could just reduce the items per row to $d - 1$ and add one distinct item for each one of these dC bidders.

To verify that this conflict graph can indeed correspond to a set of bids of maximum size d note that, since each bid has at most d conflicts, then we can let each conflict in this graph be due to a distinct item. In other words, each edge of this conflict graph corresponds to a unique item and the bundle of each bid is the set of items corresponding to the edges adjacent to its vertex.

We now show that, at any stage of the DA auction with scoring rule $b_i/f(c_i)$, the score of any vertex that has active children is greater than the score of at least one of its active children. This is initially true for any such vertex since it has some value b_i and exactly d conflicts (a score of $b_i/f(d)$) while one of its children has a value of $(b_i - \epsilon)/f(d)$ (a score less than $b_i/f(d)$ even if this child has no other conflicts). Hence, the scoring rule will begin by rejecting a leaf of the conflict graph. Also, once a leaf is rejected, the property described above remains true. To verify this fact assume that the rejected leaf was child $k \geq 2$ of its parent, then the degree of the parent drops to $k - 1$ and the $(k - 1)$ -th child has a value less than a $1/f(k - 1)$ fraction of the parent's value. If, on the other hand, that leaf was the only child (i.e., $k = 1$), then its parent is now a leaf and the property is still satisfied due to the same arguments used for the initial graph. As a result, the scoring rule will always reject a leaf, until the only bid remaining is that of the root, leading to a social welfare of 1.

We conclude the proof by showing that for $\gamma \leq 1$ there exist values of d such that the optimal social welfare of the problem instance described above grows as a function of m , which can be arbitrarily larger than d .

Since $f(x) \in \Theta(x^\gamma)$ and $\gamma \leq 1$ the series $\sum_{k=1}^{\infty} 1/f(k+1)$ diverges and thus there exists some minimum value q such that $\sum_{k=1}^q 1/f(k+1) > 2$. As a result, for any fixed value of $d \geq q + 1$, the total weight of the $d - 1$ children of some vertex of weight b_i will be $\sum_{k=1}^{d-1} (b_i - \epsilon)/f(k+1) > 2(b_i - \epsilon) > b_i$. Hence, since the weight of the root is 1, the total weight of each level of the tree is at least 1, so the total weight of the tree is at least $\log_{d-1} m$. Clearly, picking the best of either accepting all the odd or all the even depth bids will yield a feasible solution whose social welfare is at least half of that.

Lower bound of $\Omega(m)$. Consider the problem instance that involves m items and $2m$ bidders. The first group of m bidders want a single distinct item each; the second group of m bidders each want $m - 1$ items, but no two of them have the same bundle. The first group bidders have $m - 1$ conflicts each, while those of the second have $2(m - 1)$ conflicts. Finally, the value of the first group bidders is equal to the largest b_i such that

$$\frac{b_i}{f(m-1)g(m-1)} < \frac{1}{f(2(m-1))g(2(m-1))},$$

and the value of the second group bidders is 1.

This ensures that the auction begins by rejecting one of the bidders of the first group, and it keeps rejecting first group bidders until only second group bidders remain. At this stage, all the remaining bids are in conflict with one another so the auction can accept at most one of them. This leads to a welfare of 1, whereas accepting all the first group bids would lead to a total welfare of mb_i . By definition of b_i , we know that

$$b_i + \epsilon \geq \frac{f(m-1)g(m-1)}{f(2(m-1))g(2(m-1))}. \quad (8)$$

Since $f(x) \in \Theta(x^\gamma)$, for any given constant γ we get $f(m-1)/f(2(m-1)) \in \Omega(1)$, and similarly for $g(m-1)/g(2(m-1))$. As a result, Inequality (8) implies that mb_i is $\Omega(m)$.