

Deferred-Acceptance Auctions for Multiple Levels of Service

VASILIS GKATZELIS, Drexel University

EVANGELOS MARKAKIS, Athens University of Economics & Business

TIM ROUGHGARDEN, Stanford University

Deferred-acceptance (DA) auctions are mechanisms that are based on backward-greedy algorithms and possess a number of remarkable incentive properties, including implementation as an obviously-strategyproof ascending auction. All existing work on DA auctions considers only binary single-parameter problems, where each bidder either “wins” or “loses.” This paper generalizes the DA auction framework to non-binary settings, and applies this generalized framework to obtain approximately welfare-maximizing DA auctions for a number of basic mechanism design problems: multi-unit auctions, problems with polymatroid constraints or multiple knapsack constraints, and the problem of scheduling jobs to minimize their total weighted completion time. Our results require the design of novel backward-greedy algorithms with good approximation guarantees.

CCS Concepts: • **Theory of computation** → **Approximation algorithms analysis; Algorithmic game theory and mechanism design;**

Additional Key Words and Phrases: deferred-acceptance auctions; mechanism design; social welfare; approximation algorithms; scheduling

1 INTRODUCTION

The work of Milgrom and Segal [33] recently proposed the framework of *Deferred-Acceptance (DA) Auctions*, a family of single-parameter mechanisms with a number of remarkable incentive properties. In fact, the reverse auction part of the very recently concluded FCC Incentive Auction for reallocating spectrum was a DA auction [14]. DA auctions can be thought of as running an adaptive backward greedy algorithm for deciding the set of accepted bidders. The mechanism maintains a score for each active bidder (which can change from round to round), and rejects in each round the bidder with the lowest score, stopping when the still-active bidders constitute a feasible solution. This contrasts with standard (forward) greedy algorithms, which begin with the empty set and iteratively add the bidder with the highest score (subject to feasibility constraints).

DA auctions are appealing in practice for a number of reasons. First, DA auctions can be implemented as ascending clock auctions (or descending auctions if applied to procurement auctions), which has several desirable consequences. Truthful bidding is a dominant strategy, even in the sense of *obvious strategyproofness* formalized by Li [28]; see the discussion in Section 3. In contrast, a sealed-bid Vickrey auction, for example, is not obviously strategyproof in this sense. This is important for the participation and behavior of non-expert users, who may not fully understand the computations involved in setting the clock prices. Furthermore, every DA auction satisfies

This work is supported by the National Science Foundation, under grants CCF-1216073, CCF-1161813, CCF-1408635, and CCF-1524062. Part of this work took place while the authors were visiting the Simons Institute for the Theory of Computing. 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'17, June 26–30, 2017, Cambridge, Massachusetts, USA.

© 2017 ACM. ISBN 978-1-4503-4527-9/17/06...\$15.00.

DOI: <http://dx.doi.org/10.1145/3033274.3085142>

weak group-strategyproofness (WGSP). That is, no coalition of bidders can collectively submit false bids in a way that makes every bidder of the coalition strictly better off. We refer the reader to [34] for further advantages and motivation behind DA auctions.

Unfortunately, DA auctions do not always achieve a good approximation to the social welfare. Backward greedy algorithms can be inferior to forward greedy algorithms, and the former need not even produce an inclusion-maximal solution. Driven by this concern, Dütting et al. [11] subsequently explored the power and limitations of DA auctions from an approximation algorithms viewpoint (see also related work below). Their main results concern knapsack auctions and combinatorial auctions with single-minded bidders. For example, in the latter setting, novel DA auctions can nearly match the performance of arbitrary truthful and computationally efficient mechanisms.

All work thus far on DA auctions has concerned *binary* problems, i.e., settings where each bidder either “wins” or “loses.” In many scenarios, however, a bidder receives some “level of service” rather than a binary decision. Think, for example, of an air ticket purchase. Depending on the buyer’s willingness to pay, the available choices may include basic economy class, economy class with a more comfortable seat and more leg room, economy class with extra bag allowance, or business class. In a multi-unit auction with non-unit demand bidders, the level of service corresponds to the number of items awarded. In a cloud computing setting, levels of service could correspond to possible completion times of a user’s task. The goal of this paper is to generalize the DA auction framework to non-binary settings, and apply this generalized framework to a number of basic mechanism design problems.

Our results: The first contribution of this paper is to extend the DA auction framework to non-binary settings, where each bidder receives some level of service, subject to feasibility constraints (Definitions 3.1 and 7.1). We consider bidders with downward-sloping valuations, meaning that a bidder’s marginal value for upgrading to the next level of service is non-increasing in the level. The key idea in the generalization is to complement the usual DA auction scoring function with a *clinging function* that specifies the level of service that each bidder has clinched as the auction progresses. Just like in a DA auction, at each stage, a generalized DA auction decides which bidder to remove from the set of active bidders. The difference is that upon removal, the clinging function also determines the level of service received by the bidder who leaves the auction. We show that, provided the clinging functions are monotone, these generalized DA auctions possess the key incentive guarantees of binary DA auctions, including an obviously strategyproof ascending implementation and weak group-strategyproofness.

We demonstrate the usefulness of this framework by designing new DA auctions for four different problems.

- (i) Allocation problems where the service levels of the bidders must obey a polymatroid constraint and where bidders have linear valuations. Here, we show that the VCG mechanism can be implemented in our generalized DA framework (Proposition 4.1).
- (ii) Job scheduling for minimizing the total weighted completion time. Here, we give a novel generalized DA scheduling auction that achieves a $2(1 + \sqrt{2})$ -approximation of the optimal social welfare (Section 5). A priori, it is not obvious that a constant factor approximation can be achieved by any backward greedy algorithm.
- (iii) Knapsack auctions with multiple knapsacks available, where the set of bidders assigned to the same level need to satisfy a knapsack constraint, and bidders have linear valuations. Even in the binary setting with one knapsack, no DA auction can achieve an approximation ratio better than $O(\log m)$, where m is the size of the knapsack. We show that even with multiple knapsacks, there is a DA auction achieving a $O(\log m)$ -approximation of the optimal social welfare (Theorem 6.1).

- (iv) Multi-unit auctions where bidders have general downward-sloping valuations. Even with two bidders and two units, the VCG mechanism cannot be implemented as a generalized DA auction, and neither can any weakly group-strategyproof mechanism (Proposition 7.3). We give a novel DA multi-unit auction that achieves an $O(\log n)$ -approximation to the social welfare, where n is the number of agents (Theorem 7.6).

For settings (ii)–(iv), we are not aware of any existing weakly group-strategyproof mechanisms (DA or otherwise) with non-trivial approximation guarantees.

Related Work. Milgrom and Segal [33] initiated the study of procurement DA auctions, motivated by FCC auctions for reallocating spectrum. In a procurement auction it is the bidders who own the resources that the auctioneer wishes to buy. As discussed in their work, the properties that they established, such as group-strategyproofness and implementability as a clock auction, can be easily shown to hold also for the more common type of selling auctions, which is the focus of our work. In a recent version of their paper, Milgrom and Segal [34] analyze the social welfare that can be attained by DA auctions in a procurement setting with near matroid structure. They show that in this setting a DA auction attains near optimal performance, generalizing well-known results in combinatorial optimization which show that forward and backward greedy algorithms are optimal on matroids [see, e.g., 26].

Following up on the work of [33], Dütting et al. [11] explored the power and limitations of DA auctions from an approximation algorithms viewpoint. They provided both positive and negative results for knapsack auctions and auctions with single-minded bidders. On the positive side, for single-minded bidders, it was shown that we can have almost matching approximation results by DA auctions as the existing ones in the literature, hence providing stronger incentive guarantees to the known approximations. In particular, an $O(d)$ -approximation was obtained via a DA auction, when the size of each requested bundle is at most d , matching the known d -approximation of [27], which is obtained by a forward greedy algorithm. A different DA auction was also shown to achieve a $\sqrt{m \log m}$ -approximation, almost matching the known \sqrt{m} -approximation of [27], where m is the number of available items. On the negative side, for knapsack auctions, it was shown that no DA auction can achieve an approximation that is sublogarithmic in m , establishing a separation with the constant factor approximations that are known if we only want strategyproof mechanisms.

There have been several other works that have already appeared regarding the applicability of the DA auctions framework. In the procurement setting, the performance of DA auctions has been further studied by Kim [25], for spectrum and bandwidth reallocation problems. Nguyen and Sandholm [37] provide an experimental evaluation of the performance of various DA heuristics for the actual interference constraints of the upcoming FCC auction. The DA framework has also been extended to double auctions by Dütting et al. [12] and to multi-lateral markets by Blumrosen and Zohar [7]. The application to double auctions is explored further by Marx and Loertscher [29]. These works also point out the advantages of the DA framework in obtaining budget balance for double auctions. In a similar spirit, Ensthaler and Giebe [13] and Jarman and Meisner [20] point out the advantages of the DA auctions in the design of budget-constrained procurement auctions.

Group-strategyproof mechanisms have been studied extensively prior to the DA auctions framework, among others, in the context of cost-sharing mechanisms [10, 18, 19, 21, 22, 30, 35]. This line of work has also analyzed the economic efficiency that can be achieved with these mechanisms. As has been observed, often the stronger incentive properties come at a significant cost in terms of economic efficiency [see, e.g., 39]. Finally, the stronger notion of “obvious strategyproofness” has been recently introduced by [28].

There is an extensive literature on ascending auctions for both single-parameter and multi-parameter settings pre-dating the DA framework. For single-parameter settings, an ascending auction based on the backward greedy algorithm for matroids is described in Bikhchandani et al. [4]. Babaioff et al. [3] provide a procedure converting any algorithm to a dominant-strategy ascending auction (with some loss in the approximation factor). For multi-parameter settings, Ausubel [1] describes an ascending auction for identical items and bidders with decreasing marginal values. Ascending auctions for non-identical items and bidders with gross substitute preferences appear in [2, 17, 24, 31]. For a comprehensive study of the computational and informational aspects of iterative auctions see [5, 6].

2 PRELIMINARIES

We focus on problems where some type of service needs to be assigned to each of n agents in a set N . This service can be offered at different quality levels (including no service at all) but, due to scarcity or cost of this service, there are restrictions in the levels that can be allocated. We represent the levels of service as rational numbers, hence, the feasible service allocation outcomes are defined by a set system $\mathcal{I} \subseteq \mathbb{Q}^N$.¹ We assume that \mathcal{I} is non-empty and *downward closed*, meaning that if $T \in \mathcal{I}$ and $T' \subseteq T$, then $T' \in \mathcal{I}$. Note that the binary setting, to which previous work on DA auctions was restricted, corresponds to the special case where $\mathcal{I} \subseteq \{0, 1\}^N$.

We assume that each bidder has a private valuation function such that $v_i(\ell)$ is the value of bidder i for obtaining the level of service ℓ . Although all of our mechanisms can be implemented as ascending auctions, our presentation focuses on direct-revelation mechanisms in the form $\mathcal{M} = (f, p)$, which consist of an outcome rule f and a payment rule p . In our setting, given a vector of bids $\mathbf{b} = (b_i)_{i \in N}$, where b_i denotes the bid reported by bidder i (coming either from a single-parameter or multi-parameter domain, depending on the form of the function $v_i(\cdot)$), the outcome rule f computes a feasible solution, i.e., a set of \mathcal{I} , so that every bidder is assigned a level of service. On the same input, the payment rule p computes payments $\mathbf{p} = (p_i(\mathbf{b}))_{i \in N} \in \mathbb{R}^n$ where $p_i(\mathbf{b})$ denotes the payment of bidder i .

Given a mechanism $\mathcal{M} = (f, p)$, and a bid vector \mathbf{b} , let $\ell_i = f_i(\mathbf{b})$ be the level of service assigned to bidder i by the mechanism. We assume that the bidders have quasi-linear utilities and hence bidder i 's utility equals $u_i^{\mathcal{M}}(\mathbf{b}) = v_i(\ell_i) - p_i(\mathbf{b})$. Since bidders can be strategic in reporting their bids to the mechanism, towards maximizing their own utility, we aim for mechanisms that provide incentive guarantees. We say that a mechanism \mathcal{M} is *strategyproof* or *incentive compatible* if for any bidder i , for any bid vector $\mathbf{b}_{-i} = (b_j)_{j \neq i}$, and any bid \bar{b}_i of bidder i

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

We are interested in an even stronger form of resistance to manipulation, namely resistance against coalitions of bidders. Given a subset of bidders $S \subseteq N$, and a bidding vector \mathbf{b} , we denote by \mathbf{b}_S , the vector containing only the bids of S , and by \mathbf{b}_{-S} the bids of the remaining bidders. We say that a mechanism \mathcal{M} is *weakly group-strategyproof* (WGSP) if for any coalition $S \subseteq N$, for any vector $\mathbf{b}_{-S} = (b_j)_{j \notin S}$, and for any vector $\bar{\mathbf{b}}_S = (\bar{b}_j)_{j \in S}$ of the bidders in S , it holds that

$$u_i^{\mathcal{M}}(\mathbf{v}_S, \mathbf{b}_{-S}) \geq u_i^{\mathcal{M}}(\bar{\mathbf{b}}_S, \mathbf{b}_{-S}) \text{ for some } i \in S.$$

Hence, there is no coalition that can make all its members strictly better off by deviating from the truth.²

¹In our exposition, we do not need to impose a priori an upper bound on the maximum possible level of service.

²There is an even stronger form of group-strategyproofness, requiring that there is no coalitional deviation where some members are better off and the rest are not worse off, but DA auctions do not usually guarantee this property.

Apart from the incentive issues, one of the main objectives we try to optimize in mechanism design is the *social welfare*. If $T = (\ell_1, \dots, \ell_n) \in \mathcal{I}$ is a feasible solution, the generated social welfare is $SW(T) = \sum_i v_i(\ell_i)$. A mechanism \mathcal{M} with an allocation rule f achieves an approximation ratio of ρ if

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

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

3 GENERALIZED SINGLE-PARAMETER DEFERRED-ACCEPTANCE AUCTIONS

We begin by proposing a generalization of the DA auctions framework of [34] that captures non-binary single-parameter settings. We assume that each bidder can receive one of multiple levels of service, and for now, that her value for level $\ell \in \mathbb{Q}$ equals to $v_i \cdot \ell$. This linearity assumption is common in many applications, including sponsored search advertising and job scheduling. Section 7 generalizes the framework to the multi-parameter setting of downward-sloping valuations.

We define a generalized DA auction by combining a scoring function with a provisional allocation function, which we will refer to as the *clinging function*. Clinging functions are the key innovation of our generalized framework. The auction begins with all bidders being active, it operates in a sequence of stages, and after each stage it finalizes the level of service of some active bidder. This process continues, until the level of service for every bidder has been finalized. Prices are then set using Myerson's Lemma, to enforce truthfulness (see below).

Definition 3.1. A DA auction operates in discrete stages $t \geq 1$. We denote by $A_t \subseteq N$ the set of currently active bidders in the beginning of each stage t ; initially, $A_1 = N$, and $A_{t+1} \subset A_t$, for every $t \geq 1$. The DA auction is fully defined by two collections of functions:

- The **scoring functions** $\sigma_i^{A_t}(b_i, \mathbf{b}_{N \setminus A_t})$, that are non-decreasing in their first argument.
- The **clinging functions** $g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$, which are non-increasing w.r.t. the set of active bidders, i.e.: $g_i^{A_{t+1}}(\mathbf{b}_{N \setminus A_{t+1}}) \geq g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$.

At each stage t , if $A_t \neq \emptyset$, then the level of service of some active bidder $i \in \arg \min_{i \in A_t} \{\sigma_i^{A_t}(b_i, \mathbf{b}_{N \setminus A_t})\}$ is finalized, possibly with the use of some tie-breaking rule. That is, a bidder i with the lowest score stops being active, we set $A_{t+1} = A_t \setminus \{i\}$, and her level of service is finalized at level $g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$. When we reach $A_t = \emptyset$, then the auction terminates and the payment of each bidder is determined by Myerson's lemma (see Equation (1)).

The scoring and clinging functions generally change throughout the stages. These functions can depend on the bids of non-active bidders, $\mathbf{b}_{N \setminus A_t}$, and on the set of active bidders, A_t , though not the actual bids of active bidders, for incentive reasons.

For a given bid vector $\mathbf{b} = (b_i)_{i \in N} \in \mathbb{R}_+^n$, the DA auction implies an outcome rule f , which defines the level of service $f_i(\mathbf{b})$ of each bidder i . If this outcome rule always outputs a feasible solution (a set of \mathcal{I}), then we say that the DA auction is *feasible*. The following lemma shows that the allocation rule $f_i(\mathbf{b})$ implied by a DA auction is always monotone w.r.t. the bid b_i of bidder i . The proof of the lemma is deferred to the full version of the paper.

LEMMA 3.2. *Any generalized DA auction yields a monotone allocation rule, i.e., for any bid vector \mathbf{b} , and any $b'_i < b_i$, we have $f_i(b'_i, \mathbf{b}_{-i}) \leq f_i(\mathbf{b})$.*

Next we derive payments according to Myerson's lemma [36], given the monotone allocation function $f(b_i, \mathbf{b}_{-i})$, which are appropriate summations of threshold values for each level of service. Formally, suppose that under (b_i, \mathbf{b}_{-i}) , bidder i receives level of service i_ℓ . Let $L_i = \{i_1, i_2, \dots, i_\ell\}$

with $i_1 < i_2 < \dots < i_\ell$, be the levels of service that are attainable, as we let the bid of i increase from 0 to b_i , keeping \mathbf{b}_{-i} fixed. The monotonicity of the levels we encounter is ensured by the monotonicity of the allocation rule $f_i(\cdot)$. For each $i_j \in L_i$, let z_j be the following threshold value:

$$z_j(\mathbf{b}_{-i}) = \inf_{b'_i \geq 0} \{b'_i : f_i(b'_i, \mathbf{b}_{-i}) \geq i_j\}.$$

Then, the payment of bidder i comes from the following formula:

$$p_i(\mathbf{b}) = \sum_{j=1}^{\ell} z_j(\mathbf{b}_{-i}) \cdot (i_j - i_{j-1}), \quad (1)$$

where for $j = 1$, we set $i_0 = 0$.

We next show that the mechanism induced by such an allocation rule and corresponding payment rule is weakly group-strategyproof.

THEOREM 3.3. *Every generalized DA auction is weakly group strategy-proof.*

PROOF. Assume there exists a coalition of bidders $K \subseteq N$ that can coordinate in misreporting their values in a way that strictly improves the utility of every $i \in K$. We show that this assumption leads to a contradiction.

Let v_i denote the true value of each bidder i and v'_i be the value reported by $i \in K$. Also, let $\alpha \in K$ be the bidder of the coalition who would be finalized first (at stage t_α) according to the true values, and let $\beta \in K$ be the bidder of the coalition who is finalized first (at stage t_β) after the deviation.

First, we show that $t_\beta > t_\alpha$. Aiming for a contradiction assume that, after the deviation, bidder β is finalized at stage $t_\beta \leq t_\alpha$, and this deviation strictly increases her utility. According to the definition of the generalized DA auctions, the sequence of bidders finalized during the first $t_\beta - 1$ stages under the deviation is identical to the corresponding sequence of bidders finalized prior to the deviation. To verify this fact, note that the scores of these bidders are not affected by the deviation, since all the bidders in K remain active during the first $t_\beta - 1$ stages (recall that the scores at a stage t maybe affected by the bids of $N \setminus A_t$ and not by active bidders). Hence, both in the original profile as well as in the deviation, the level of service that has been clinched by bidder β by stage t_β , and the total payment for this level is the same. Since we have assumed $t_\beta \leq t_\alpha$, bidder β was receiving a weakly higher level of service prior to the deviation. Moreover, the cost for any additional levels of service compared to the level she receives under the deviation, was at most v_β . Hence, overall the utility of β cannot have increased with the deviation, leading to a contradiction.

Therefore, there exists some bidder $i \notin K$ who is rejected at stage t_α after the deviation, and was finalized at some stage $t_i > t_\alpha$ prior to the deviation. Since bidder α was finalized before bidder i prior to the deviation, and using the fact that both of these bidders' scores are the same before and after the deviation, we infer that $v'_\alpha > v_\alpha$. Also, if ℓ_α and ℓ'_α is the level of service that α receives before and after the deviation, this means that $\ell'_\alpha > \ell_\alpha$ (otherwise the utility of α would be the same after the deviation) and the threshold price of any level higher than ℓ_α is at least v_α (since α would be rejected at stage t_α if she did not strictly increase her bid). Then, according to Equation 1, the additional payment of α , for the additional $\ell'_\alpha - \ell_\alpha$ levels of service is at least

$$\sum_{j=\ell_\alpha+1}^{\ell'_\alpha} v_\alpha \cdot (i_j - i_{j-1}) \geq (\ell'_\alpha - \ell_\alpha)v_\alpha.$$

Note that the payment of α for the first ℓ_α levels is the same as before since $t_\beta > t_\alpha$. Hence, since the value of α for these additional levels of service is $(\ell'_\alpha - \ell_\alpha)v_\alpha$, this means that bidder α could not have strictly increased her utility upon deviating. \square

Apart from group-strategy-proofness, the DA framework also guarantees several other important properties that are desirable for practical purposes. In particular, [34] showed that a (binary) DA auction can be implemented by an ascending clock auction. This means that we can think of the auction as running in a sequence of periods, and offering an ascending sequence of prices to every bidder. Our generalization can also be implemented as such an ascending auction. As a result, it is obvious for the bidders to verify that truth-telling is their optimal strategy, even if they are not aware of how the clock prices are set. This is in contrast, e.g., to the VCG mechanism, where it is not at all clear a priori why a bidder should behave truthfully. This distinction is formalized under the notion of *obvious strategy-proofness* in [28].

The following two propositions summarize the fact that such properties carry over to the setting of generalized DA auctions as well, which further justifies their appeal. The proof of the propositions as well as a discussion regarding clock auctions can be found in the full version of the paper.

PROPOSITION 3.4. *Every generalized DA auction, can be implemented as an ascending clock auction.*

We henceforth refer to the class of ascending clock auctions derived from generalized DA auctions according to Proposition 3.4 as *generalized DA ascending auctions*.

PROPOSITION 3.5. *Every generalized DA ascending auction is obviously strategyproof.*

4 WARM-UP: POLYMATROID CONSTRAINTS

We begin our study of generalized DA auctions by considering settings where the set of feasible outcomes is defined by a polymatroid constraint (defined below).

For instance, consider a multi-unit auction with k identical copies of the same good, and additive bidders, so that bidder i 's value for acquiring ℓ_i units of the good (a level of service ℓ_i) is $\ell_i \cdot v_i$. The constraint in this case is that $\sum_i \ell_i \leq k$, i.e., the number of units to be allocated is at most k . For another example, consider the keyword sponsored search auctions, where the sellers are competing for a sequence of $q < n$ advertising slots, and each slot j has a click-through rate r_j . If bidder i 's value for a click is v_i , then its value for slot j is $r_j \cdot v_i$, and we say that i receives a level of service of $\ell_i = r_j$. This time, the constraint is that the outcome $(\ell_1, \ell_2, \dots, \ell_n)$ needs to correspond to a matching of bidders to slots, with $\ell_i = 0$ for any bidder who is not matched to a slot. For more motivating examples see also [4].

More generally, in this section we consider problems, where the set of feasible outcomes is defined via a given submodular function $h : 2^N \rightarrow \mathbb{R}_+$, as follows:

$$P_h = \left\{ \ell \in \mathbb{N}^n \mid \sum_{i \in S} \ell_i \leq h(S) \ \forall S \subseteq N \right\}.$$

In the first setting provided above, i.e., the multi-unit auction with k available units, the polymatroid constraint is defined by the constant (submodular) function h , with $h(S) = k$ for every $S \subseteq N$. In the second setting, if we assume that $r_1 \geq r_2 \geq \dots \geq r_q$, the polymatroid constraint is defined by $h(S) = \sum_{j=1}^{|S|} r_j$. This implies that any single bidder i gets $\ell_i \leq r_1$, any two bidders i, i' get $\ell_i + \ell_{i'} \leq r_1 + r_2$, and so on.

As a first application of our framework, we show that, in fact, there is a simple generalized DA auction that achieves the optimal social welfare in any problem instance involving polymatroid constraints. Given the submodular function h of the polymatroid constraint, the scoring and clinching functions of this auction are as follows:

- The polymatroid auction scoring function is $\sigma_i^{A_t}(b_i, \mathbf{b}_{N \setminus A_t}) = b_i$.
- The polymatroid auction clinching function is $g_i^{A_t}(b_{N \setminus A_t}) = h(A_t) - h(A_t \setminus \{i\})$.

Note that this DA auction has a very simple scoring function, which is not adaptive. In particular, at every stage, the bidder who is finalized by this DA auction is the one with the smallest bid, among the ones that are still active. What is more interesting is the clinching function, according to which, at each stage t each bidder $i \in A_t$ has clinched a level of service equal to her marginal contribution to the value of $h(A_t)$. Since h is submodular, this marginal contribution weakly increases as A_t shrinks, so this is a valid clinching function.

For a concrete example of how this auction works, we consider the two settings described above. In the multi-unit setting, this auction would keep finalizing the lower valued bidders without letting them clinch any item, until only the highest bidder is active, at which point this bidder clinches all k units. For each one of these units, the price is equal to the second highest bid. In the keyword sponsored search setting, during the first $n - q$ stages, the $n - q$ lowest bidders are finalized without clinching any slots. Then, at stage $n - q + 1$ every active bidder clinches a level of service of r_q , which is the click through rate of the worst slot, q , and the lowest value bidder $i \in A_{n-q+1}$ is finalized at $\ell_i = r_q$, i.e., it is assigned to that slot. Note that, after stage $n - q + 1$, every active bidder is guaranteed a level of service of at least r_q , and the price that the top q bidders pay for clinching this level of service is $b_{q+1}r_q$, i.e., r_q times the $(q + 1)$ -th highest bid. As the auction moves on, more bidders get finalized, and the price that they pay for the additional levels of service that they clinch weakly increases over time.

To verify that this auction always yields the maximum social welfare, note that its outcome is exactly the same as the one that would arise if we instead used the following forward greedy algorithm. First, give the highest bidder, i , the highest level of service possible, i.e., $h(\{i\})$. Then, give the second highest bidder the highest level possible, given the existing assignment to i , i.e., i' gets $h(\{i, i'\}) - h(\{i\})$, and so on. This greedy algorithm is known to be optimal in polymatroid settings [15].

In fact, the auction presented above is exactly what the ascending auction of [4] reduces to if the valuations of the bidders are additive, and what the ascending auction of [16] reduces to if we remove the budget constraints. Furthermore, the allocation and the payments of this auction is the same as that of the VCG auction. Since these auctions lie within the generalized DA framework, all of the incentive guarantees of the previous section (Theorem 3.3–Proposition 3.5) apply.

PROPOSITION 4.1. *When the set of feasible allocations is defined by a polymatroid constraint, the VCG mechanism is a generalized DA auction.*

In contrast Proposition 4.1, we show in Section 7 that the VCG mechanism is not generally a DA auction when bidders have downward-sloping valuations.

5 SCHEDULING CONSTRAINTS

To exhibit the strength of our generalization and the diversity of settings that it applies to, we now consider a well-studied non-binary problem of job scheduling.

Given a set of m identical machines and a set N of n jobs, each of which needs processing time p_i on any one of the machines, a *schedule* is an assignment of each job to a machine. Furthermore, for each machine j , the schedule defines which one of its assigned jobs the machine will process at any given time. As a result, if a schedule x assigns the jobs of bidders 1, 2, and 3 on the same machine, to be processed in that order with no idle time, the *completion times* of the jobs will be $c_1(x) = p_1$, $c_2(x) = p_1 + p_2$, and $c_3(x) = p_1 + p_2 + p_3$ respectively. Depending on its urgency, or importance, each job i also has a weight v_i , and the objective of this scheduling problem is to compute a schedule x that minimizes the weighted completion time $\sum_{i \in N} v_i c_i(x)$.

This problem has received considerable attention (see [9] for a survey of relevant results) and it is known to be NP-hard [8]. But, the following simple greedy algorithm achieves a $(1 + \sqrt{2})/2 \approx 1.207$ -approximation for this problem [23, 40]: order the jobs in non-increasing order of their v_i/p_i ratio and schedule them, in that order, on the first available machine.

In this section, rather than assuming that the scheduler knows the weight v_i of each job, we consider the harder problem where this information is private to agent i , who owns this job. The completion time c_i of job i then corresponds to the delay that user i suffers, and the weight v_i corresponds to the value of agent i for improving the completion time of her job by a unit of time. Every agent is guaranteed completion by some deadline d (defined later on) but, depending on its value, v_i , the agent may be willing to pay in order to improve its completion time, i.e., to get a higher level of service. Our goal is to design a DA auction that elicits bids b_i corresponding to the private values v_i and outputs a schedule x , aiming to minimize the social cost, i.e., $\sum_i b_i c_i(x)$. As we have shown, every DA auction can be implemented as an ascending auction. But, what would an ascending auction for job scheduling look like?

Ascending Job Scheduling Auctions. In an ascending auction for job scheduling, as the prices increase, the bidders who remain active should be guaranteed outcomes of increasingly high quality, i.e., a lower completion time. This means that rather than scheduling the jobs in a bottom-up fashion, starting from lower completion times and moving toward higher ones, like all known list-processing algorithms do, an ascending auction schedules jobs top-to-bottom. In particular, a DA ascending auction uses a clinching function to “promise” increasingly better completion times to the active bidders and to decide their actual completion time when they are finalized. But, a clinching decision is irreversible, and it cannot depend on the bids of jobs in A_t , i.e., the ones still competing for lower completion times. Therefore, a generalized DA auction needs to ensure that, *no matter the values of the active jobs* to be finalized later on, there exists a *feasible* way of scheduling them that *respects the clinching promises* and yields a *good approximation factor*. This introduces a novel and non-trivial trade-off for the designer: on one hand, the designer wants to allow the jobs to clinch low completion times, aiming for a good approximation factor but, on the other hand, if the promised completion times are too low, this may lead to an infeasible schedule.

Aiming for a good approximation, we design a *backward* greedy algorithm whose social cost is within a constant factor of the greedy algorithm described above. In fact, we ensure that our DA auction assigns to every job a completion time within a constant factor of its completion time in the forward greedy schedule. The most natural way to achieve that would be to let our DA auction use the same scoring function, i.e., to finalize bids in increasing order of their b_i/p_i ratio. But, how can we ensure that the completion time of the job i with the worst ratio, the one finalized first, is within a constant of its completion time in the greedy schedule?

Note that the clinching function needs to be oblivious to the values of the active bids and, depending on what these values are, the completion time of job i in the forward greedy outcome may vary significantly. For a concrete example, consider the case when, apart from the job i with the smallest value over size ratio, A_t also comprises two sets of jobs X and Y , where X contains m jobs of size 1 and Y contains $m - 1$ jobs of size m . If the jobs in Y have higher ratios than those in X , the completion time of job i in the forward greedy outcome would be $m + 1$. On the other hand, if the jobs in X have higher ratios than those in Y , the completion time of job i in the forward greedy outcome would be 2! But, we cannot finalize i with a constant completion time for any m , as this would quickly lead to feasibility issues, since the average load per machine is not constant.

5.1 Deferred-Acceptance Scheduling Auction

In this section, we provide a DA auction that determines a feasible schedule of the tasks and achieves a $2(1 + \sqrt{2}) \approx 4.83$ -approximation. At each stage t , our auction will be leveraging two parameters that depend on A_t but not on the values of the bidders in it: a) $M_t = \frac{1}{m} \sum_{i \in A_t} p_i$, i.e., the value of the fractional makespan for the jobs in A_t , and b) $\bar{p}_t = \max_{i \in A_t} \{p_i\}$, i.e., the maximum processing time among the jobs in A_t . Our auction guarantees to every bidder a completion time by the deadline $d = M_1 + 2\bar{p}_1$, and depending on their bids and processing times, some jobs get the opportunity to pay for improved completion times. These improvements of the completion time, beyond the initial deadline, are captured by the clinching function, which also decides what the completion time of a job will be when it is finalized.

Scoring function. As long as $\bar{p}_t \leq M_t/2$, our DA scheduling auction uses a scoring function based on the greedy algorithm ratio, i.e., $\sigma_i(b_i) = b_i/p_i$. But, when there exist jobs whose processing time is higher than $M_t/2$, the auction finalizes them first in a largest job first fashion. Formally,

$$\sigma_i^{A_t}(b_i) = \begin{cases} b_i/p_i, & \text{if } \bar{p}_t < M_t/2 \text{ or } p_i \neq \bar{p}_t \\ 0, & \text{if } \bar{p}_t \geq M_t/2 \text{ and } p_i = \bar{p}_t. \end{cases}$$

Clinching function. For each machine j , we denote by $E_j(t)$ the minimum starting time among all the tasks scheduled on machine j prior to stage t , and let $E_{\max}(t) = \max_j \{E_j(t)\}$. For $t = 1$, we set $E_j(1) = d$ for every $j \in [m]$. The clinching function is defined as:

$$g_i^{A_t}(\mathbf{b}_{N \setminus A_t}) = M_1 + 2\bar{p}_1 - \min \{M_t + 2\bar{p}_t, E_{\max}(t)\}.$$

In more detail, at each stage t , this clinching function defines how much earlier than the deadline ($M_1 + 2\bar{p}_1$) job i is guaranteed to terminate (recall that we want the clinching function to be nonincreasing w.r.t. A_t , hence we take the difference between the deadline and the completion time). In other words, at stage t , every job is guaranteed to complete no later than $\min \{M_t + 2\bar{p}_t, E_{\max}(t)\}$. The job i finalized at stage t is the one with minimum score $\sigma_i^{A_t}(b_i)$, and it is scheduled on some machine $j \in \arg \max_j \{E_j(t)\}$. If this machine is not busy at time $M_t + 2\bar{p}_t$, then job i is scheduled so that its completion time is $c_i(x) = M_t + 2\bar{p}_t$. If, on the other hand, every machine is busy at that time, then the job is scheduled so that its completion time is exactly $E_{\max}(t) = E_{j(t)}(t)$, i.e., right before the previously scheduled tasks on the machine.

THEOREM 5.1. *The DA scheduling auction yields a $2(1 + \sqrt{2}) \approx 4.83$ approximation.*

What is more demanding, is that we also need to verify that the DA auction always yields a feasible schedule. That is, we need to ensure that there exists a schedule where all the jobs can be processed by the completion times they have clinched.

To do this, we will define first some appropriate quantities for every stage t , that will guide us in the analysis. Let $\text{GOAL}(t) = M_t + 2\bar{p}_t$, and for every processor $j \in [m]$, let $\delta_j(t) = \max\{0, \text{GOAL}(t) - E_j(t)\}$. Hence, δ_j is the difference between the goal function and the time at which j starts processing the jobs currently assigned to her. For technical convenience, we set this to 0 when $E_j(t)$ exceeds $\text{GOAL}(t)$. The rest of the analysis focuses on arguing about the behavior of the vector $(\delta_j(\cdot))_{j \in [m]}$ over time, with respect to the features:

$$\text{MAX}(t) = \max_{j \in [m]} \delta_j(t), \text{MIN}(t) = \min_{j \in [m]} \delta_j(t), \text{AVG}(t) = \frac{1}{m} \sum_{j \in [m]} \delta_j(t)$$

For an illustration of these quantities, see Figure 1. We can think of $\text{GOAL}(t)$ as the updated deadline at stage t , which keeps decreasing between successive stages. I.e., recall that $\text{GOAL}(1) =$

$d = M_1 + 2\bar{p}_1$, and it is easy to see that $\text{GOAL}(t + 1) \leq \text{GOAL}(t)$. Our proof examines the area below the line at time $\text{GOAL}(t)$ to argue about the existence of a feasible schedule.




Fig. 1. An illustration of the relevant quantities at a stage t .

To proceed, we prove the following important lemma, which will be useful in verifying that our DA auction always yields a feasible schedule. A rough intuition behind the proof is that, whenever, after stage t , $\text{GOAL}(t)$ drops to $\text{GOAL}(t + 1)$ it causes either $\text{MAX}(t) - \text{MIN}(t)$ or $\text{AVG}(t)$ to drop (or it causes an aggregate drop) by enough time so as to maintain that the inequality remains true.

LEMMA 5.2. *For every stage t of the auction we have*

$$\text{MAX}(t) - \text{MIN}(t) + \text{AVG}(t) \leq 2\bar{p}_t. \quad (2)$$

PROOF. The proof is by induction on t . For $t = 1$ this is straightforward to verify, since $E_j(1) = d$ for every $j \in [m]$. Hence, we now assume that Inequality (2) holds for every stage up to t , and our goal is to show that it then remains true at stage $t + 1$. We distinguish three possible cases: a) $\text{MIN}(t) > 0$ and $\text{MIN}(t + 1) > 0$, b) $\text{MIN}(t) > 0$ and $\text{MIN}(t + 1) = 0$, and c) $\text{MIN}(t) = 0$. In the analysis that follows, we will refer to a job as being *tight*, if it is scheduled without leaving any idle time, till the next job. I.e., for a job scheduled on machine j at stage t , tightness means that its completion time is $E_i(t)$ (thus all machines are busy at time $M_t + 2\bar{p}_t$).

- a) If $\text{MIN}(t) > 0$ and $\text{MIN}(t + 1) > 0$, let t' be the last stage before t when $\text{MIN}(t') = 0$ (the existence of such a stage is clearly guaranteed, e.g., $t' = 1$). All the machines are always busy in the time interval between $\text{GOAL}(t')$ and $\text{GOAL}(t + 1)$ (by definition of t' and since $\text{MIN}(t + 1) > 0$). Therefore

$$\text{AVG}(t + 1) = \text{AVG}(t') - 2(\bar{p}_{t'} - \bar{p}_{t+1}). \quad (3)$$

- If $\text{MAX}(t') \geq \bar{p}_{t'}$, then $\text{MAX}(t + 1) - \text{MIN}(t + 1) \leq \text{MAX}(t')$. This is true because all jobs scheduled between t' and $t + 1$ are tight, and their size is no more than $\bar{p}_{t'} \leq \text{MAX}(t')$. Hence, the gap between $\text{MAX}(t)$ and $\text{MIN}(t)$ will not increase when we go to stage $t + 1$. Using Inequality (3), this implies that

$$\begin{aligned} \text{MAX}(t + 1) - \text{MIN}(t + 1) + \text{AVG}(t + 1) &\leq \text{MAX}(t') + \text{AVG}(t') - 2(\bar{p}_{t'} - \bar{p}_{t+1}) \\ &= \text{MAX}(t') - \text{MIN}(t') + \text{AVG}(t') - 2(\bar{p}_{t'} - \bar{p}_{t+1}) \\ &\leq 2\bar{p}_{t'} - 2(\bar{p}_{t'} - \bar{p}_{t+1}) \\ &= 2\bar{p}_{t+1} \end{aligned}$$

The last inequality above is true by the induction hypothesis applied to t' .

- If $\text{MAX}(t') < \bar{p}_{t'}$, then $\text{MAX}(t+1) - \text{MIN}(t+1) \leq \bar{p}_{t'}$. This is true because all jobs scheduled between t' and $t+1$ are tight, and their size is no more than $\bar{p}_{t'} \leq \text{MAX}(t')$. Hence, the gap between $\text{MAX}(t+1)$ and $\text{MIN}(t+1)$ may increase, but not beyond $\bar{p}_{t'}$. Also, since $\text{MAX}(t') < \bar{p}_{t'}$, we also have $\text{AVG}(t') < \bar{p}_{t'}$. Using Inequality (3), this implies that

$$\begin{aligned} \text{MAX}(t+1) - \text{MIN}(t+1) + \text{AVG}(t+1) &\leq \bar{p}_{t'} + \text{AVG}(t') - 2(\bar{p}_{t'} - \bar{p}_{t+1}) \\ &\leq 2\bar{p}_{t'} - 2(\bar{p}_{t'} - \bar{p}_{t+1}) \\ &= 2\bar{p}_{t+1} \end{aligned}$$

b) $\text{MIN}(t) > 0$ and $\text{MIN}(t+1) = 0$. This case is more lengthy and complicated to analyze, yet it is similar in spirit with Case (a) and we omit it from this version.

c) If $\text{MIN}(t) = 0$, then $\text{MAX}(t) + \text{AVG}(t) \leq 2\bar{p}_t$, and we consider two possible cases:

- If the job scheduled at stage t is not the one with the earliest starting time at stage $t+1$ (i.e., if it does not affect $\text{MAX}(t+1)$), then $\text{MAX}(t+1) \leq \text{MAX}(t) - 2(\bar{p}_t - \bar{p}_{t+1}) - p_t/m$. Combined with the fact that $\text{AVG}(t+1) \leq \text{AVG}(t) + p_t/m$, we get that

$$\begin{aligned} \text{MAX}(t+1) - \text{MIN}(t+1) + \text{AVG}(t+1) &\leq \text{MAX}(t+1) + \text{AVG}(t+1) \\ &\leq \text{MAX}(t) - 2(\bar{p}_t - \bar{p}_{t+1}) + \text{AVG}(t) \\ &\leq 2\bar{p}_t - 2(\bar{p}_t - \bar{p}_{t+1}) \\ &= 2\bar{p}_{t+1} \end{aligned}$$

- If the job scheduled at stage t is the one with the earliest starting time at stage $t+1$ (i.e., if it defines $\text{MAX}(t+1)$), then $\text{MAX}(t+1) \leq \bar{p}_t - 2(\bar{p}_t - \bar{p}_{t+1}) - p_t/m$. Since $\text{AVG}(t+1) \leq \text{MAX}(t+1)$ we get

$$\begin{aligned} \text{MAX}(t+1) - \text{MIN}(t+1) + \text{AVG}(t+1) &\leq \text{MAX}(t+1) + \text{AVG}(t+1) \\ &\leq 2\text{MAX}(t+1) \\ &\leq 2(\bar{p}_t - 2(\bar{p}_t - \bar{p}_{t+1}) - p_t/m) \\ &\leq 2\bar{p}_{t+1} \end{aligned}$$

□

We note that the guarantee provided in the previous lemma is tight. This can be verified by considering for example an instance with two machines in which the jobs have processing times (listed by order of being scheduled): $1, 1/2, 1, 3/4, 1, 7/8, \dots, 1, (2^k - 1)/2^k$.

COROLLARY 5.3. *The DA scheduling auction is always feasible.*

PROOF. Assume that at some stage t there exists a job that cannot fit in any machine, i.e., $\max_j \{E_j(t)\} < \bar{p}_t$. Let t' be the last stage before t when $\text{MIN}(t') = 0$. Given the result of Lemma 5.2, at stage t' we have $\text{MAX}(t') + \text{AVG}(t') \leq 2\bar{p}_{t'}$, and hence $\text{AVG}(t') \leq (\text{MAX}(t') + \text{AVG}(t'))/2 \leq \bar{p}_{t'}$. Since all subsequently scheduled jobs are tight, if P is the total processing time of the jobs scheduled from t' to t , the average busy time below $\text{GOAL}(t')$ increased by P/m , and the drop from $\text{GOAL}(t')$ to $\text{GOAL}(t)$ is exactly $P/m + 2(\bar{p}_{t'} - \bar{p}_t)$. As a result,

$$\text{AVG}(t) \leq \text{AVG}(t') - 2(\bar{p}_{t'} - \bar{p}_t) \leq \bar{p}_{t'} - 2(\bar{p}_{t'} - \bar{p}_t) \leq 2\bar{p}_t - \bar{p}_{t'} \leq \bar{p}_t.$$

But, this means that $\max_j \{E_j(t)\} \geq \text{GOAL}(t) - \bar{p}_t \geq \bar{p}_t$, which contradicts our assumption that $\max_j \{E_j(t)\} < \bar{p}_t$. □

6 MULTIPLE KNAPSACK CONSTRAINTS

As a third application of the generalized DA auctions framework, we also study settings where the set of feasible solutions is defined via multiple knapsack constraints though, due to space constraints, most of this section is deferred to the full version of the paper.

There are multiple motivating examples for knapsack constraints: in FCC auctions there are different bands (VHF, UHF) each with its own capacity constraints, and with a different value. Similarly, the economy, business, and first class seats in an airplane correspond to three different knapsack constraints, each with a different value. As a final example, consider cross country moving trucks that are scheduled to move property or products from a source to a destination. Each truck has a capacity constraint and the earlier the arrival date the higher the value.

Formally, in this section we assume that each possible level of service ℓ has a capacity of m “slots”, and we let s_i be the “size” of each bidder i , i.e., the number of slots that the bidder needs to occupy, if assigned level ℓ . The presence of multiple knapsack constraints then imply that an outcome is feasible if the total size of the set of bidders $N_\ell \subseteq N$ finalized at a level of service ℓ is at most m , i.e., $\sum_{i \in N_\ell} s_i \leq m$ for every possible level of service ℓ .

Dütting et al. [11] showed that, even in the binary setting, if the set of bidders that can be simultaneously accepted is defined by a single knapsack constraint of size m , then no DA auction can achieve an approximation better than $O(\log m)$. The following theorem provides a matching positive result: a generalized DA auction that achieves the optimal logarithmic approximation.

THEOREM 6.1. *There exists a multiple knapsack DA auction achieving a $O(\log m)$ approximation for an arbitrary number of knapsacks.*

Just like in the previous section, the DA auction that we propose manages to strike a balance between allowing the bidders to clinch higher levels of service as quickly as possible, yet ensuring that these “promises” will not eventually violate any knapsack constraints.

7 MULTI-PARAMETER DEFERRED-ACCEPTANCE AUCTIONS

This section further generalizes DA auctions to a multi-parameter setting: bidders with downward-sloping (a.k.a. submodular) valuations in the level of service awarded. For notational simplicity, we assume that the possible levels of service are represented by a set $L = \{1, 2, \dots, k\}$ and that each bidder i reports her marginal value for each additional level, as $b_i(1) = v_i(1)$, $b_i(2) = v_i(2) - v_i(1)$, up to $b_i(k) = v_i(k) - v_i(k-1)$ with $b_i(1) \geq b_i(2) \geq \dots \geq b_i(k)$. Hence, $b_i(j)$ is the added value for the bidder if she jumps from level of service $j-1$ to j .

Just like the generalized DA auction framework that we defined in Section 3, the multi-parameter framework that we propose here also uses a collection of scoring functions and clinching functions for each bidder. The important difference is that the agents now report a sequence of bids that correspond to their marginal valuations for the different levels of service. Given these marginal bids for each bidder i , at each stage t , the multi-parameter DA auction computes the score of each active bidder using her bid corresponding to the next level of service that i is considered for clinching. That is, if $g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$ is the level of service that i has clinched by stage t , which we henceforth denote by just g_{it} , her score is computed based on her marginal value for clinching one more level of service, i.e, based on $b_i(g_{it} + 1)$.

Definition 7.1. A multi-parameter DA auction operates in stages $t \geq 1$. In each stage t a set of bidders $A_t \subseteq N$ is active; initially, $A_1 = N$, and $A_{t+1} \subset A_t$, for every $t \geq 1$. Just as in the single-parameter case, the DA auction is fully defined by two collections of functions:

- The **scoring functions** $\sigma_i^{A_t}(\cdot, \mathbf{b}_{N \setminus A_t})$ that are non-decreasing in their first argument.

- The **clinking functions** $g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$, that are non-decreasing w.r.t. the set of active bidders, i.e., if for every $t \geq 1$, $g_i^{A_{t+1}}(\mathbf{b}_{N \setminus A_{t+1}}) \geq g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$.

At each stage t , if $A_t \neq \emptyset$, the score of bidder $i \in A_t$ is computed as:

$$\sigma_i^{A_t}(b_i(g_{it} + 1), \mathbf{b}_{N \setminus A_t}),$$

i.e., the score is a function of the bidder's marginal value for receiving a level increase, given the level that she has already clinched. The bidder with the smallest score is finalized at level $g_{it} = g_i^{A_t}(\mathbf{b}_{N \setminus A_t})$ and she is removed from the set of active bidders ($A_{t+1} = A_t \setminus \{i\}$).

All the incentive guarantees for single-parameter DA auctions carry over to the present setting. The proof of these properties follows essentially the same arguments as in Theorem 3.3 and Propositions 3.4 and 3.5. These properties are summarized below.

PROPOSITION 7.2. *Every multi-parameter DA auction is weakly group-strategyproof, and has an equivalent clock auction implementation that is obviously strategyproof.*

7.1 Multiunit Auctions with Decreasing Marginal Values

To demonstrate an application of multi-parameter DA auctions, we consider the setting of multi-unit auctions with submodular bidders. Multi-unit auctions are being deployed in practice in a wide range of applications, and more recently they have also been used by various online brokers (for more, see e.g., [32, 38]). Such auctions are defined by a collection of m identical units of some good that need to be allocated among n agents. The level of service in this setting corresponds to the number of units that each agent receives, and the value of an agent for a bundle of items is a submodular function of the bundle's size. Hence the valuation function here is fully described by a nonincreasing vector of marginal values, for bundles up to size m .

Ascending Multiunit Auctions. There is a well-known ascending multi-unit auction, the “clinking auction” of Ausubel [1]. Given an ascending price per unit, the bidders in this auction respond by reporting the number of units that they would be interested in acquiring at the current price, i.e., the number of their marginal valuations that are at least as high as the price. As the price increases, the number of items that the bidders request drops, and a bidder clinches a unit at the point where, even if he were to leave the auction, the total demand of the other bidders would leave that unit unallocated. This auction is known to be strategy-proof and in fact implements the VCG mechanism. However it is also known to be susceptible to demand reduction in group deviations, and hence is not weakly group-strategyproof.

For a simple example, consider an instance involving two units and two bidders: bidder a has marginal values $(1, 1)$, while bidder b has marginal values $(0.6, 0.6)$. If we were to run the clinching auction of [1] in this instance, after the price would exceed 0.6, bidder b would leave the auction and bidder a would clinch both units at that price for a utility of $2 - 1.2 = 0.8$. If, on the other hand, both agents reduced their demand by claiming that they have no value for a second item, they would get one unit each for a price of 0! The equivalence of Ausubel's auction with the VCG mechanism implies the following contrast to Proposition 4.1.

OBSERVATION 1. *In the setting of multi-unit auctions with downward-sloping valuations, the VCG mechanism is not a DA auction.*

We next define a DA auction for multi-unit auctions. But before that, we first establish that some loss in social welfare will be unavoidable, even if we only insist on a weakly group-strategyproof mechanism. In particular, we have the following impossibility result, the proof of which is deferred to the full version.

PROPOSITION 7.3. *Even for 2 players and 2 units, no weakly group-strategyproof mechanism can guarantee an approximation factor to the social welfare, that is better than $\sqrt{2}$.*

The DA auction that we propose in the next subsection provides an $O(\log n)$ -approximation. We leave it as an open question whether or not there is a DA (or at least a weakly group-strategyproof) mechanism with a better approximation factor.

7.2 Deferred-Acceptance Multiunit Auction with Decreasing Marginal Valuations

For notational simplicity, assume that $n = 2^\kappa$ for some $\kappa \in \mathbb{N}_+$, and let $\lambda = \lfloor m/(n \log n) \rfloor$. Also, assume that $m \geq n \log n$, and hence $\lambda \geq 1$.³ In order to achieve the logarithmic approximation, we restrict the distribution of units to agents so that it obeys a very particular structure:

- The number of units that each agent may be allocated is exactly $\lambda \cdot 2^r$, for some $r \in \{0, 1, \dots, \kappa\}$. We therefore distribute the units in blocks of exponentially increasing size.
- The distribution of these blocks of units across the agents is such that the number of agents that receive at least $\lambda \cdot 2^r$ units is at most $n/2^r$.

In this restricted setting, instead of providing m possible levels of service to each bidder (the number of units they can receive), the levels of service are $\kappa + 1$ (the number of unit blocks they can receive). Hence, rather than requiring that the agents report all of their m marginal valuations, our auction requires that the agents report only their *average* marginal value for receiving exactly $\lambda \cdot 2^r$ units of the good, for $r \in \{0, 1, \dots, \kappa\}$, i.e., their average marginal value for each one of the $\kappa + 1$ blocks. Therefore, the agents report a bid vector of size $\kappa + 1$ containing weakly decreasing marginal values. For instance, if $n = 4$ and $m = 8$, then $\kappa = 2$ and $\lambda = 1$. If the marginal valuations of an agent are $[8, 7, 6, 4, 4, 3, 2, 1]$, then her average marginal bid vector would be $[8, 7, \text{avg}(6, 4)] = [8, 7, 5]$. Our auction does not assign more than a κ fraction of the units to one agent, so the remaining marginal values are disregarded.

Definition 7.4. The DA multiunit auction receives $\kappa + 1$ marginal bids from each bidder, corresponding to their average marginal values for each block. At every stage the scoring function of every bidder is the identify function, i.e., $\sigma_i^{A_r}(b_i(g_{it} + 1)) = b_i(g_{it} + 1)$. The clinching function is $g_i^{A_r}(\cdot) = \lfloor \frac{n}{|A_r|} \rfloor$, which means that initially (when $n = |A_r|$), every bidder has clinched the first block, which contains $\lambda \cdot 2^0 = \lambda$ units. Then, after the $n/2$ agents with the smallest score are finalized, bidders that remain active clinch an additional block of $\lambda \cdot 2^1 = 2\lambda$ units. Every time the set of active bidders is halved, those that remain active clinch another block, until they are all finalized.

For the analysis of the auction, it is convenient to also describe it as an ascending auction. Similarly to the Ausubel clinching auction, the blocks of units are clinched when the demand of the competitors drops below a specific threshold but, unlike the clinching auction, ours is a DA auction, so it is weakly group-strategyproof. The ascending auction implementation of our DA auction uses a price per unit p , which happens to be the same for every bidder, and is initially set to $p = 0$. The auction begins by assigning λ units to each bidder at a price of $p = 0$. Then, the agents compete for an additional block of λ units. The auction offers an average price p per unit for this block, and p is gradually increased until exactly $n/2$ agents are still willing to pay this amount per unit of the block. The bidders who do not receive these extra λ units are finalized, and they hence receive just the first block of λ units for a price of 0. The $n/2$ bidders that remain active clinch each unit of the second block for a “threshold price” p and they then compete for the third block, which

³If n is not a power of 2, we may remove the bidders with the lowest first marginal value until we reach a power of 2. This does not affect the order of magnitude of the approximation guarantee that we establish. The same goes for the assumption that $m \geq n \log n$.

contains 2λ more units. A third block can be allocated to at most $n/4$ bidders among them. Note that the price for each unit in this block remains equal to p , i.e., the one used for the previous block, so some of the active bidders might not be willing to pay this amount. If the number of bidders willing to pay p for each unit of the third block is not more than $n/4$, then all these bidders clinch this block for the existing price p per unit. Otherwise, p is increased further, until exactly $n/4$ of these bidders remain interested. We repeat this process by gradually doubling the size of the block, and halving the number of bidders that can receive one, until two agents compete for the single remaining block of $n\lambda/2$ units.

Example 7.5. Consider a problem instance with $n = 4$ and $m = 8$, which implies that $\kappa = 2$ and $\lambda = 1$. Assume that the bid vectors of the four agents $\{A, B, C, D\}$ are $[7, 6, 6]$, $[6, 5, 5]$, $[8, 4, 4]$, and $[5, 4, 3]$. Every agent is assigned one unit at a price $p = 0$ (this means that we disregard the first parameter of the bid vector, but we will use it when m is smaller). Then, the agents compete for an additional unit that only two of them can win. The price is increased to $p = 4 + \epsilon$, at which point agents C and D drop out (since their average marginal value for a second unit is 4) and agents A and B clinch one unit each for this price. Finally, the remaining two active agents compete for a single block of two units. The price is increased to $p = 5 + \epsilon$ and agent B drops out as well, so A clinches the block of two units for this price.

Approximation Guarantee. We now prove that the DA auction that we propose above guarantees a logarithmic approximation factor w.r.t. the social welfare. Note that, although our auction is restricted to allocations that bundle units into groups and does not need the whole marginal valuation vector, the approximation guarantee is comparing the outcome with the optimal allocation that uses all this information.

THEOREM 7.6. *The DA auction that we propose guarantees an approximation factor of $O(\log n)$.*

Benchmark. In order to prove the approximation guarantee, we use a different allocation as a benchmark. In particular, let \bar{x} be the social welfare maximizing allocation with the same bidders if the number of units were $m/(2\kappa)$ instead of m . This allocation finds, among the mn marginal values, the $m/(2\kappa)$ highest ones, and it allocates the corresponding number of units to each agent. Since the values of the agents exhibit decreasing marginal valuations for additional units, this implies that the social welfare in \bar{x} is a 2κ -approximation of the social welfare maximizing allocation x^* with m units:

$$\sum_{i \in N} v_i(x^*) \leq 2\kappa \sum_{i \in N} v_i(\bar{x}).$$

We now show that our auction provides each agent with at least half the value that she would receive in \bar{x} . This implies that the allocation x that our auction computes guarantees:

$$\sum_{i \in N} v_i(x^*) \leq 2\kappa \sum_{i \in N} v_i(\bar{x}) \leq 4\kappa \sum_{i \in N} v_i(x) = 4 \log n \sum_{i \in N} v_i(x).$$

We begin with the following lemma, which shows that for any value that the unit price p takes during our auction, there exist at least $m/(2\kappa)$ marginal values that are at least as high.

LEMMA 7.7. *Let p_{\max} be the highest price that our ascending auction reaches. Then, among the mn marginal values of the agents, there exist at least $m/(2\kappa)$ ones whose value is at least p_{\max} .*

PROOF. Let $r > 1$ be the round when the price is increased for the last time, thus reaching the value of p_{\max} . According to the definition of our mechanism, at that point, there are exactly $n/2^r$ agents who are willing to pay this price per unit for an additional block of $\lambda \cdot 2^{r-1}$ units⁴; they have

⁴The fact that the price was increased during round r means that the agents willing to pay the initial price were more, and the price increase brings them down to exactly $n/2^r$.

already clinched a total of $\lambda \cdot 2^{r-1}$ units in previous rounds. Since the average marginal value of these agents for the additional block is at least p_{\max} , this means that the marginal value of *every unit* within the blocks that these agents clinched during previous rounds is also at least p_{\max} . Therefore, the number of marginal values in the agents' bid vectors which are at least p_{\max} , is at least

$$\frac{n}{2^r} \cdot (\lambda \cdot 2^{r-1}) = \frac{n\lambda}{2} = \frac{nm/(\kappa n)}{2} = \frac{m}{2\kappa}$$

□

Using the lemma above we continue by showing that our auction provides every agent with a value at least half as much as the value that the agent would receive in \bar{x} .

LEMMA 7.8. *Let x be the allocation computed by our auction. For every agent $i \in N$, $v_i(x) \geq v_i(\bar{x})/2$.*

PROOF. To prove this lemma we show that our auction allocates to each agent at least half the number of units that the agent receives in \bar{x} . Since the valuations are concave this implies that the agent receives half the value as well.

Assume that there exists some agent i who is rejected during round $r > 1$ by our auction and hence receives exactly $\lambda \cdot 2^{r-1}$ units, while the same agent receives more than $\lambda \cdot 2^r$ in \bar{x} . Let p be the price per unit at the time when this agent was rejected by our auction. The fact that the agent is rejected means that the average marginal value of that user for an additional block of $\lambda \cdot 2^{r-1}$ units is less than p . This implies that the lowest marginal value over all these units, which is the last, is less than p . Therefore, the agent's marginal value for a $(\lambda \cdot 2^r + 1)$ -th unit is less than p . Given Lemma 7.7 though, there exist at least $m/(2\kappa)$ units with marginal values at least $p_{\max} \geq p$, which means that the $(\lambda \cdot 2^r + 1)$ -th unit of that agent cannot be one of the $m/(2\kappa)$ highest value marginals and thus part of the assignment \bar{x} , leading to a contradiction. □

Finally, we conclude by exhibiting that our auction is feasible, i.e., the number of units it allocates does not exceed the number of available units. The proof of the following lemma is omitted.

LEMMA 7.9. *For $n \geq 4$, the DA multi-unit auction is feasible, i.e., it always assigns at most m units.*

REFERENCES

- [1] L. Ausubel. 2004. An Efficient Ascending-bid Auction for Multiple Objects. *American Economic Review* 94 (2004), 1452–1475.
- [2] L. Ausubel. 2006. An Efficient Dynamic Auction for Heterogeneous Commodities. *American Economic Review* 96 (2006), 602–629. Issue 3.
- [3] M. Babaioff, R. Lavi, and E. Pavlov. 2009. Single-value Combinatorial Auctions and Algorithmic Implementation in Undominated Strategies. *J. ACM* 56, 1, Article 4 (2009), 4:1–4:32 pages.
- [4] S. Bikhchandani, S. de Vries, S. J. Schummer, and R. V. Vohra. 2011. An Ascending Vickrey Auction for Selling Bases of a Matroid. *Operations Research* 59 (2011), 400–413. Issue 2.
- [5] L. Blumrosen and N. Nisan. 2009. On the Computational Power of Demand Queries. *SIAM J. Comput.* 39, 4 (2009), 1372–1391.
- [6] L. Blumrosen and N. Nisan. 2010. Informational limitations of ascending combinatorial auctions. *Journal of Economic Theory* 145, 3 (2010), 1203–1223.
- [7] L. Blumrosen and O. Zohar. 2015. Multilateral Deferred-Acceptance Mechanisms. In *Proceedings of the 11th Conference on Web and Internet Economics*. 173–186.
- [8] P. Brucker. 2004. *Scheduling algorithms (4. ed.)*. Springer.
- [9] C. Chekuri and S. Khanna. 2004. Approximation Algorithms for Minimizing AverageWeighted Completion Time. In *Handbook of Scheduling - Algorithms, Models, and Performance Analysis*.
- [10] N. R. Devanur, M. Mihail, and V. V. Vazirani. 2005. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems* 39 (2005), 11–22.
- [11] P. Dütting, V. Gkatzelis, and T. Roughgarden. 2014. The performance of deferred-acceptance auctions. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*. 187–204.

- [12] P. Dütting, T. Roughgarden, and I. Talgam-Cohen. 2014. Modularity and greed in double auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*. 241–258.
- [13] L. Ensthaler and T. Giebe. 2015. A Dynamic Auction for Multi-object Procurement under a Hard Budget Constraint. *Research Policy* 143 (2015), 179–189.
- [14] FCC. 2017. Broadcast Incentive Auction. <https://www.fcc.gov/about-fcc/fcc-initiatives/incentive-auctions>. (2017).
- [15] S. Fujishige. 1980. Lexicographically Optimal Base of a Polymatroid with respect to a Weight Vector. *Mathematics of Operations Research* 5 (1980), 186–196. Issue 2.
- [16] G. Goel, V.S. Mirrokni, and R. Paes Leme. 2015. Polyhedral Clinching Auctions and the AdWords Polytope. *J. ACM* 62, 3 (2015), 18:1–18:27.
- [17] F. Gul and E. Stacchetti. 2000. The English Auction with Differentiated Commodities. *Journal of Economic Theory* 92 (2000), 66–95. Issue 1.
- [18] N. Immorlica, M. Mahdian, and V. S. Mirrokni. 2008. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms* 4, 2 (2008).
- [19] N. Immorlica and E. Pountourakis. 2012. On Budget-Balanced Group-Strategyproof Cost-Sharing Mechanisms. In *Proceedings of the 8th Workshop on Internet and Network Economics*. 244–255.
- [20] F. Jarman and V. Meisner. 2015. Ex-post Optimal Knapsack Procurement. In *SSRN Working Paper No. 2548543*.
- [21] R. Juarez. 2013. Group strategyproof cost sharing: The role of indifferences. *Games and Economic Behavior* 82, C (2013), 218–239.
- [22] R. Juarez. 2013. Optimal Group Strategyproof Cost Sharing. In *Working paper*. Available from: <http://www2.hawaii.edu/~rubenj/optimalgsp.pdf>. Last accessed: May 8, 2017.
- [23] T. Kawaguchi and S. Kyan. 1986. Worst Case Bound of an LRF Schedule for the Mean Weighted Flow-time Problem. *SIAM J. Comput.* 15, 4 (Nov. 1986), 1119–1129.
- [24] A. Kelso and V. Crawford. 1982. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica* 50 (1982), 1483–1504. Issue 6.
- [25] A. Kim. 2015. Welfare Maximization with Deferred Acceptance Auctions in Reallocation Problems. In *Proceedings of the 23rd Annual Symposium on Algorithms*. 804–815.
- [26] B. Korte and J. Vygen. 2007. *Combinatorial Optimization: Theory and Algorithms*. Springer Verlag, Berlin Heidelberg, Germany.
- [27] D. Lehmann, L. I. O’Callaghan, and Y. Shoham. 2002. Truth Revelation in Approximately Efficient Combinatorial Auctions. *J. ACM* 49 (2002), 577–602. Issue 5.
- [28] S. Li. 2015. Obviously Strategy-Proof Mechanisms. In *SSRN Working Paper No. 2560028*.
- [29] L. M. Marx and S. Loertscher. 2015. Prior-free Bayesian optimal Double-Clock Auctions. In *Working paper*. Available from: <https://faculty.fuqua.duke.edu/~marx/bio/papers/PriorFree.pdf>. Last accessed: May 8, 2017.
- [30] A. Mehta, T. Roughgarden, and M. Sundararajan. 2009. Beyond Moulin mechanisms. *Games and Economic Behavior* 67, 1 (2009), 125–155.
- [31] P. Milgrom. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy* 108 (2000), 245–272. Issue 2.
- [32] P. Milgrom. 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- [33] P. Milgrom and I. Segal. 2014. Deferred-acceptance auctions and radio spectrum reallocation. In *ACM Conference on Economics and Computation, EC ’14, Stanford, CA, USA, June 8-12, 2014*. 185–186.
- [34] P. Milgrom and I. Segal. 2015. Deferred-Acceptance Auctions and Radio Spectrum Reallocation. In *Working paper*. Available from: <http://web.stanford.edu/~isegal/heuristic.pdf>. Last accessed: May 8, 2017.
- [35] H. Moulin. 1999. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare* 16, 2 (1999), 279–320.
- [36] R. Myerson. 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 1 (1981), 58–73.
- [37] T. Nguyen and T. Sandholm. 2014. Optimizing prices in descending clock auctions. In *Proceedings of the 15th ACM Conference on Economics and Computation*. 93–110.
- [38] A. Ockenfels, D. H. Reiley, and A. Sadrieh. 2006. *Economics and Information Systems*. Handbooks in Information Systems, Vol. 1. Elsevier, Chapter 12. Online Auctions, 571–628.
- [39] T. Roughgarden and M. Sundararajan. 2009. Quantifying inefficiency in cost-sharing mechanisms. *Journal of the ACM* 56, 4 (2009).
- [40] U. Schwegelshohn. 2011. An alternative proof of the Kawaguchi-Kyan bound for the Largest-Ratio-First rule. *Operations Research Letters* 39, 4 (2011), 255 – 259.