

Optimizing Prices in Descending Clock Auctions[†]

TRI-DUNG NGUYEN, University of Southampton, Schools of Mathematics and Management
TUOMAS SANDHOLM, Carnegie Mellon University, Computer Science Department *Corresponding author

A descending (multi-item) clock auction (DCA) is a mechanism for buying items from multiple potential sellers. In the DCA, bidder-specific prices are decremented over the course of the auction. In each round, each bidder might accept or decline his offer price. Accepting means the bidder is willing to sell at that price. Rejecting means the bidder will not sell at that price or a lower price. DCAs have been proposed as the method for procuring spectrum from existing holders in the FCC's imminent incentive auctions so spectrum can be repurposed to higher-value uses. However, the DCA design has lacked a way to determine the prices to offer the bidders in each round. This is a recognized, important, and timely problem.

We present, to our knowledge, the first techniques for this. We develop a percentile-based approach which provides a means to naturally reduce the offer prices to the bidders through the bidding rounds. We also develop an optimization model for setting prices so as to minimize expected payment while stochastically satisfying the feasibility constraint. (The DCA has a final adjustment round that obtains feasibility after feasibility has been lost in the final round of the main DCA.) We prove attractive properties of this, such as symmetry and monotonicity. We developed computational methods for solving the model. (We also develop optimization models with recourse, but they are not computationally practical.) We present experiments both on the homogeneous items case and the case of FCC incentive auctions, where we use real interference constraint data to get a fully faithful model of feasibility. An unexpected paradox about DCAs is that sometimes when the number of rounds allowed increases, the final payment increases. We provide an explanation for this.

Categories and Subject Descriptors: [**Artificial Intelligence and Applied Game Theory; Experimental, Empirical, and Applications**]: Auction Theory, Market Design

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Descending Clock Auction, Incentive Auction, Spectrum Auction

1. INTRODUCTION

A *descending (multi-item) clock auction (DCA)* is a mechanism for buying items from multiple potential sellers. In the DCA, bidder-specific prices are initialized at reserve prices and then decremented over the course of the auction [Milgrom and Segal 2012]. In each round, the auctioneer decrements the offer prices to the bidders, who might accept or decline the offers. Accepting means that the bidder is willing to sell at that price. Rejecting means that the bidder has to exit the auction permanently and cannot sell. This process is repeated until the auctioneer's target number of items to purchase would become infeasible if additional bidders were to reject any new (lower) offers. At that point the auction ends and the current prices are paid.

[†] Patent pending.

Author's addresses: T.D. Nguyen, Schools of Mathematics and Management, University of Southampton; email: T.D.Nguyen@soton.ac.uk; T. Sandholm, Computer Science Department, Carnegie Mellon University; email: sandholm@cs.cmu.edu.

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.

ACM 978-1-4503-2565-3/12/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/2600057.2602869>

We consider the following setting. The auctioneer wants to buy items from a pool, N , of n potential sellers. Each seller $i \in N$ has a specific type of item G_i and decides to sell it or not depending on the offer price. The items from the sellers could be substitutable and complementary. The auctioneer has a target number of items to buy, T , and there is a feasibility function $F : 2^N \rightarrow \{0, 1\}$ that specifies, for each subset of potential sellers, S , whether the items from S can fulfill the target T or not; i.e. $F(S) = 1$ if the combined items from $\{G_i, i \in S\}$ fulfill the target.

A simple example of this is the case where the sellers have identical items and the auctioneer wants to buy a target number T them. In this case, the feasibility is simply $F(S) = 1$ if $|S| \geq T$, and $F(S) = 0$ otherwise.

However, in real-world application such as the FCC spectrum reverse auctions discussed later, the feasibility function can be highly complex. Often it cannot be given in closed form, but rather is stated through constraints as an optimization problem.

In the DCA, the auctioneer sends offer prices to the sellers and checks whether they accept those prices. Bidders who accept the offers are called *active*. If the combined items from the active bidders fulfill the target, then the auctioneer reduces the prices further in the next round and repeats the process. If at some point the items from the active bidders do not fulfill the target, then the auctioneer goes back to the last step and conducts a last-round adjustment to offer higher prices to some declined bidders so that feasibility is obtained. Algorithm 1 describes the general DCA framework.

ALGORITHM 1: Descending Clock Auction (DCA)

Input: A set of sellers $\mathcal{N} = \{1, \dots, n\}$ with goods $\{G_1, \dots, G_n\}$, an auctioneer with a feasibility mapping function $F : 2^N \rightarrow \{0, 1\}$. A target number of rounds allowed m . Initial value function estimates v_i .

Output: A set of feasible sellers $\mathcal{A} \subset \mathcal{N}$, i.e., $F(\mathcal{A}) = 1$, and the corresponding vector of offer prices p that aims to minimize the expected payment that the buyer needs to pay.

1. Initialize the price vector p to the reserve prices. Let the set of active bidders be $\mathcal{A}^{(0)} = \mathcal{N}$;
 - for** round $r = 1 \dots m$ **do**
 - 2.1. Find a vector of prices p to offer the bidders;
 - 2.2. Find the set of rejected offers \mathcal{R} ;
 - if** $F(\mathcal{A}^{(r)} \setminus \mathcal{R}) = 1$ **then**
 - 2.2.1. $\mathcal{A}^{(r+1)} \leftarrow \mathcal{A}^{(r)} \setminus \mathcal{R}$;
 - 2.2.2. Update the distributions of the bidders' values;
 - else**
 - 2.2.3. Enter the readjustment round in Step 3;
 - end**
 - end**
 3. Readjust the prices for bidders in the last round to meet the target;
 4. Pay winning bidders the offered prices;
-

A key challenge in a DCA lies in how to set the prices offered to the bidders. A natural approach is to set the offer price at some fixed percentile of the (buyer's model of the) distribution of that bidder's valuation. For example, the prices could be set so that each bidder has the same probability of accepting her offer. The choice of the percentile would depend on what the auctioneer aims for on the trajectory of the sizes of sets of active bidders through the rounds. An example of the trajectory could be such that the expected number of rejections in each round is distributed evenly throughout the auction. Another example of the trajectory is to set a fixed percentage of rejection in each round, that is, the expected number of rejections is proportional to the size of the remaining set of active bidders. We call this class of methods *percentile based* and we will describe it formally later in Section 3.1.

Those methods have several drawbacks. First, having a fixed percentile means there is no way to distinguish between bidders with high bid values and the rest; hence the final payment will likely be unnecessarily high due to the probabilistic inclusion of high-priced bidders. More importantly, those methods do not have any special treatment for the degree of interaction among the items in the feasibility function.

This paper presents a significant improvement within the DCA by addressing these issues by designing an optimization model for setting the prices. The model is designed to minimize the expected final payment while ensuring feasibility in a stochastic sense. It is flexible in that it can incorporate bidder-specific characteristics with respect to feasibility.

1.1. FCC Incentive Auctions

The current flagship application of DCAs is the Federal Communications Commission (FCC) *incentive auctions*. The FCC has been selling radio spectrum licenses via auctions since 1994 [Cramton 1997; Milgrom 2004]—in recent years via combinatorial auctions [Cramton et al. 2006; Day and Milgrom 2008]. However, there is not enough spectrum left to sell for the new high-value spectrum uses that have arisen. The idea of incentive auctions, therefore, is to buy some of the existing licenses back from their current holders, which frees up spectrum, and then to sell spectrum to higher-value users. The idea of such incentive auctions was introduced in the 2010 National Broadband Plan [FCC 2012]. It is motivated by the fact that the demand and the value of over-the-air broadcast television has been declining while the demand for mobile broadband and wireless services has increased dramatically in recent years. Given the limited spectrum resources, incentive auctions were introduced as a voluntary, market-based means of repurposing spectrum. This is done by creating a market that exchanges the usage rights among the two groups of users: (a) existing TV broadcasters and (b) wireless broadband networks. Three key players in this market are existing spectrum owners, spectrum buyers, and the FCC, which acts as the intermediary.

An incentive auction consists of three stages [FCC 2014] (see also a whitepaper about design choices by Hazlett et al. [2012]):

1. *Reverse auction*: parts of the spectrum currently used by TV broadcasters is bought back.
2. *Repacking*: remaining broadcasters are reallocated to a smaller spectrum band.
3. *Forward auction*: freed spectrums is sold via a (combinatorial) auction for use in wireless broadband networks.

In the reverse auction, we need to find a set of stations to be reallocated to lower-band channels and a separate set of stations to be bought off the air, in order to achieve the following goals: (a) meet some target on the number of contiguous channels freed on the higher spectrum band and (b) minimize total payment by the FCC. The FCC is required to respect the broadcasters' carry-right, which means, in the context of a DCA, a station that rejects the offer still has the right to stay on the air, but possibly on a lower spectrum band. This repacking stage needs to ensure that all the stations that rejected their offers can be feasibly repacked into the allocated band without violating the engineering constraints, that is, interference-free population coverage, as we will detail later in the paper.

Once the reverse auction phase is completed and the remaining stations are repacked, the FCC announces the cleared spectrum which is now available for purchase. Buyers then submit bids on bundles of spectrum frequencies. The FCC solves a winner determination problem to decide which bids to accept or reject. This does not involve the engineering constraints, so it resembles standard combinatorial auctions;

hence existing algorithms like those by Sandholm et al. [2005]; Sandholm [2006, 2013] can be used for this. Therefore, we focus on the reverse auction stage.

The FCC aims to free up a number of high spectrum channels—say between channel 32 and 51. In order to do that, all stations need to be assigned to the lower band channels or give up their carry right in exchange for some payment. The process of deciding which stations to retain, which stations to leave, and at what prices, is done through the reverse auction. Within this stage, the repacking problem needs to be solved in order to check whether the remaining stations can be feasibly reassigned to the targeted lower-band channels. Other groups have recently also tackled the repacking part (e.g., Leyton-Brown [2013]).

Two apparent options for the reverse auction design are a single-round Vickrey-Clarke-Groves (VCG) auction or a DCA. In the VCG, stations submit sealed bids. The auctioneer then solves the winner determination problem (WDP) to determine the winning bids, that is, those stations that the FCC buys back at their bid prices. Stations with rejected bids are repacked into the lower-band channels. Unfortunately, solving this WDP is challenging because it involves thousands of binary variables and millions of interference-avoidance constraints. In fact, the FCC has attempted to solve it, but according to Milgrom and Segal [2013], solving an instance of this problem with state-of-the-art optimization packages takes weeks of compute time without finding the optimal solution. Also, even if the list of winners were available, solving for the VCG payment for each winning bidder would involve solving a combinatorial problem similar to the WDP. It is the same as the WDP except that a small portion of the binary variables are fixed. Furthermore, a small approximation error in solving these problems can lead to significant over-payment [Milgrom and Segal 2012].

The DCA was proposed by Milgrom and Segal [2012], who have shown attractive features of the framework. According to documents that the FCC uses to communicate with its stakeholders regarding incentive auctions (wireless.fcc.gov/incentiveauctions/learn-program), the DCA is the FCC's design of choice.

However, there is a key missing piece in the DCA: how to determine what prices to offer to the different bidders in the different rounds. In this paper we present a solution to this. To our knowledge, this is the first paper on the topic.

The importance of this question is highlighted on the FCC web site (wireless.fcc.gov/incentiveauctions/learn-program/rule-option/reverse-auction.html):

The algorithm that determines how prices are decremented for each station in each round ... is an important part of the auction design. This algorithm will rely on a score for each bid that accounts for factors like the potential interference created by a station and the population served by the station, as well as the money amount of the bid. The use of population in scoring could reflect the fact that the value of a broadcasting license depends in part on its population served. Ranking bids and paying winning bidders in relation to their population served or other indicators of value may reduce the amount that the Commission would have to pay to repurpose broadcast television spectrum. The details of the scoring will need to be examined further.

This problem is also timely. The FCC was planning to run the first incentive auction in 2014, but because the auction design is not ready, in December 2013 the FCC postponed the start of the incentive auction to at least mid-2015 [Wheeler 2013].

1.2. Contributions and Paper Outline

We develop a natural percentile-based approach and an optimization model for setting offer prices in DCAs. To our knowledge, this is the first paper on the topic. Section 2 starts with a general model for price setting in DCAs. Section 2.1 presents computa-

tional methods, complexity, and properties of the optimal prices for various settings. A key purpose of DCAs is price discovery; Section 2.2 discusses updating beliefs about the bidders' value functions. Section 2.3 describes optimizing the final round settlement in deciding the winners. Section 3 provides experiments for homogeneous-item DCAs and for FCC incentive auctions using real data. Section 4 concludes and provides a discussion of future directions.

2. TECHNIQUES FOR OPTIMIZING OFFER PRICES IN THE DESCENDING CLOCK AUCTION

A key component of a DCA is to set the prices to offer to the active bidders. The auctioneer needs to consider the tradeoff between minimizing payment to the accepted bidders and fulfillment of the target (i.e., repacking feasibility in the case of incentive auctions). Furthermore, the pricing affects the speed of the auction in terms of the number of rounds. Therefore, there is another tradeoff. On the one hand, if the offer prices are too high, many rounds are required, and that may be undesirable from the perspective of minimizing logistical effort. On the other hand, if the offer prices are too low, many bidders reject and the auction ends too quickly without properly serving its price-discovery purpose.

In this section, we present a method for setting the offer prices. They should be dependent on (a) the estimated value functions of the bidders, (b) the importance of the items for the target to be fulfilled, and (c) how quickly/slowly the auctioneer wants the auction to run. We provide an optimization model that incorporates these considerations. Throughout this section, we assume the auctioneer's target T is a given scalar number and the feasibility mapping F is given in some specific form. Obviously these inputs are not readily available in many settings such as in incentive auctions: the feasibility function involves a complex repacking problem. Later in Section 3.2 we show how to translate such a complex setting to the forms of T and F described in this section.

2.1. Optimization Model for Price Setting in Each Round

Suppose each bidder has a threshold price v_i below which it declines the offer. The auctioneer does not know these threshold prices. Suppose, however, that the auctioneer has an estimate/belief that the threshold price v_i of bidder i follows some distribution on the support set $[l_i, u_i]$. Assume that the auctioneer knows these distributions.

Let $X_i(p_i)$ be the corresponding Bernoulli random variable that indicates whether bidder i will accept the offer at price p_i . The total payment is $c(\mathbf{p}) = \sum_{i \in \mathcal{A}^{(r)}} X_i(p_i)p_i$,

where $\mathcal{A}^{(r)}$ is the set of remaining active bidders in the current round r . One objective of the auctioneer is to set the prices in such a way that minimizes the expected payment which is equal to

$$E[c(\mathbf{p})] = \sum_{i \in \mathcal{A}^{(r)}} E[X_i(p_i)]p_i. \quad (1)$$

The expected payment is nondecreasing in the offer prices. So, setting low prices would lead to low expected payment. However, doing this would lead to a high chance of reaching infeasibility too soon, so little price discovery could be done. The auctioneer, therefore, needs to balance feasibility with low expected payment. The bidders' values are random to the auctioneer, so, for any fixed set of offer prices, the feasibility can only be expressed in a stochastic sense. It is possible to apply ideas from stochastic programming to model this as a chance constraint. However, this often leads to computational intractability since the feasibility mapping tends to be a highly complex function of the prices. We propose a simpler measure of the expected number of bid-

ders accepting the offers $N(\mathbf{p})$ as this is directly related to the chance of feasibility, i.e., the larger the population of active bidders, the higher the chance of feasibility. We have $E[N(\mathbf{p})] = \sum_{i \in \mathcal{A}^{(r)}} E[X_i(p_i)]$. The problem of minimizing the expected payment

while ensuring the expected number of accepted bidders to meet some target can be modeled as the following optimization problem:

$$\text{minimize } \sum_{i \in \mathcal{A}^{(r)}} E[X_i(p_i)]p_i, \text{ subject to } \sum_{i \in \mathcal{A}^{(r)}} E[X_i(p_i)] \geq T^{(r)}, l_i \leq p_i \leq u_i, \forall i \in \mathcal{A}^{(r)}, \quad (2)$$

where $T^{(r)}$ is the targeted number of active bidders at round r . The auctioneer has the flexibility in choosing this target depending on how quickly or slowly the auctioneer wants to complete the auction. One possibility is to set $T^{(r)} = n^{(r)} - \frac{n^{(r)} - T}{m - r + 1}$, where $n^{(r)}$ is the number of active bidders remaining, T is the final target, and $(m - r + 1)$ is the number of remaining rounds. In this case, $T^{(r)}$ is set such that the size of the sets of active bidders reduces evenly throughout the rounds.

Both the payment and feasibility are expressed in expectation due to the stochasticity of the bidders' values. This means the actual number of accepted bidders might exceed or fall below the target $T^{(r)}$. That potential disparity, and its effect on the later rounds, imply that better prices than those from Model 2 could be obtained via a model that incorporates recourse. We present one such model in Appendix A.2. That model is, however, computationally highly complex to solve and might not be appropriate in large-scale DCA settings if quick rounds are needed.

The price-setting strategy in Model 2 simplifies one important fact about the dynamical nature of DCAs: the choice of offer prices in round r will affect the population of active bidders in subsequent rounds as well as the distribution of those bidders' values. Thus, in principle, the pricing problem in DCAs should be modeled as a dynamic program as presented in Appendix A.1. Solving it, however, would be prohibitively complex. Instead, we simplify this process through a dynamic scheduling of the sizes of the sets of active bidders $T^{(r)}$. Specifically, we schedule the size of the set of active bidders $n^{(r)}$ evenly in the last $(m - r + 1)$ rounds. The simplicity of the models described in this section makes them appropriate in situations when it is critical to offer prices to the bidders in a timely manner.

To simplify the notation in Model 2 a bit, let us denote by $F_i(p_i)$ the cumulative distribution of the valuation of bidder i . It can also be interpreted as the probability that bidder i will accept offer price p_i , i.e., $F_i(p_i) \equiv \delta_i \equiv E[X_i(p_i)]$. Model 2 can now be rewritten as

$$\text{minimize } \sum_{i \in \mathcal{A}^{(r)}} F_i(p_i)p_i, \text{ subject to } \sum_{i \in \mathcal{A}^{(r)}} F_i(p_i) \geq T^{(r)}, l_i \leq p_i \leq u_i, \forall i \in \mathcal{A}^{(r)}. \quad (3)$$

2.1.1. Uniform Distribution on Bidder Values. We first consider the case where the bidders' values are drawn from uniform distributions on $[l_i, u_i]$. The probability that bidder i will accept the offer price v_i is,

$$\delta_i(p_i) = F_i(p_i) = \begin{cases} 0, & \text{if } p_i \leq l_i, \\ \frac{p_i - l_i}{u_i - l_i}, & \text{if } l_i \leq p_i \leq u_i, \\ 1, & \text{if } p_i > u_i. \end{cases} \quad (4)$$

Naturally, we can restrict the price to $l_i \leq p_i \leq u_i$. The expected payment is

$$E[c(\mathbf{p})] = \sum_{i \in \mathcal{A}^{(r)}} F_i(p_i)p_i = \sum_{i \in \mathcal{A}^{(r)}} \frac{(p_i - l_i)p_i}{u_i - l_i}.$$

The constraint on the expected number of stations accepting the offers is

$$E[N(\mathbf{p})] \geq T^{(r)} \Leftrightarrow \sum_{i \in \mathcal{A}^{(r)}} \frac{p_i - l_i}{u_i - l_i} \geq T^{(r)},$$

which is linear in the offer prices \mathbf{p} . So, the problem of minimizing the expected payment while ensuring the probabilistic constraints on feasibility can be modeled as the following quadratic program:

$$\text{minimize } \sum_{i \in \mathcal{A}^{(r)}} \frac{(p_i - l_i)p_i}{u_i - l_i}, \text{ subject to } \sum_{i \in \mathcal{A}^{(r)}} \frac{p_i - l_i}{u_i - l_i} \geq T^{(r)}, l_i \leq p_i \leq u_i, \forall i \in \mathcal{A}^{(r)}. \quad (5)$$

The problem has a strictly convex separable quadratic objective and one joint constraint in addition to lower and upper bound constraints. Thus, the problem has a unique optimal solution. In addition, the expected payment increases with \mathbf{p} while the expected number of rejected stations decreases with \mathbf{p} . Thus, we would expect the constraint to be tight at the optimal solution. Let λ be the Lagrangian multipliers for the joint constraint $\sum_{i=1}^n \frac{p_i - l_i}{u_i - l_i} \geq T^{(r)}$. The Lagrangian dual function is

$$\begin{aligned} \mathcal{L}(\lambda, \mathbf{p}) &= \sum_{i \in \mathcal{A}^{(r)}} \frac{(p_i - l_i)p_i}{u_i - l_i} + \lambda \left(T^{(r)} - \sum_{i \in \mathcal{A}^{(r)}} \frac{p_i - l_i}{u_i - l_i} \right) \\ &= T^{(r)}\lambda + \sum_{i \in \mathcal{A}^{(r)}} \frac{(p_i - l_i)p_i - \lambda(p_i - l_i)}{u_i - l_i}. \end{aligned}$$

The Lagrangian dual problem can be derived as

$$\max_{\lambda \geq 0} \left\{ T\lambda + \sum_{i \in \mathcal{A}^{(r)}} \min_{l_i \leq p_i \leq u_i} \frac{p_i^2 - (\lambda + l_i)p_i + \lambda l_i}{u_i - l_i} \right\}.$$

The problem is convex, so there is no duality gap when we take the Lagrangian relaxation. For each fixed set of Lagrangian multipliers λ , the optimal prices can be derived:

$$p_i^* = \begin{cases} \frac{l_i + \lambda}{2}, & \text{if } l_i \leq \lambda \leq 2u_i - l_i, \\ l_i, & \text{if } \lambda \leq l_i, \\ u_i, & \text{otherwise.} \end{cases} \quad (6)$$

PROPOSITION 2.1. *[Symmetry] For any two bidders with the same valuation distribution, the optimal offer prices must be the same.*

PROOF. Straightforward from Formulation 6. For any two bidders i, j with the same boundaries, i.e., $l_i = l_j$ and $u_i = u_j$, any choice of λ leads to $p_i^* = p_j^*$. \square

The Lagrangian dual function can be calculated efficiently in $O(n)$ for each fixed λ . The Lagrangian dual problem is a piece-wise concave maximization problem with one scalar variable on the non-negative orthant. One could apply a conjugate gradient method to solve this problem. Better still, we present an algorithm that exploits the structure of the problem and is $O(n \log n)$:

PROPOSITION 2.2. *The optimal offer prices can be found—i.e., Model 5 can be solved—in $O(n \log n)$ operations.*

PROOF. The optimal offer price p_i^* is a piece-wise linear function on λ with three pieces that are intersected at two points l_i and $(2u_i - l_i)$. We can order all the $2n$ points $\{l_i, 2u_i - l_i\}, \forall i = 1, \dots, n$ on the vertical axis in $O(n \log n)$. From that we obtain $(2n + 1)$ pieces (some potentially with zero length). For λ that falls within each piece, we have a corresponding linear mapping to the offer prices $p_i, i = 1, \dots, n$. Thus, the Lagrangian function on that piece can be calculated as a corresponding quadratic function on λ . The optimum λ on each piece can therefore be calculated. The global optimal λ can be found by taking the maximum Lagrangian function among all the $(2n + 1)$ pieces. \square

2.1.2. *Heterogeneous Bidders with Respect to Target Feasibility.* The feasibility constraint in Model 5 was for the homogeneous case. Consider now a heterogeneous case where the item from each bidder would affect feasibility differently. For example, in incentive auctions, depending on the current list of rejected stations, the repacking feasibility is highly sensitive to the choice of new rejection due to engineering constraints. Suppose there is a weight vector $\omega = (\omega_1, \dots, \omega_n)$ that represents the contribution from the bidders to the feasibility function F ; in Section 3.2 we will show how to derive these weights for the case of incentive auctions. For now, we consider a simple example where bidders all share the same type of item but each of them will have a different number of items to sell. In that case, ω_i can represent the amount available from bidder i . We assume a setting where each bidder considers its item as a single product and considers selling it or not depending solely on the offer price (to avoid the combinatorial complexity). Now, Model 5 can be modified to

$$\text{minimize } \sum_{i \in \mathcal{A}^{(r)}} \frac{(p_i - l_i)p_i}{u_i - l_i}, \text{ subject to } \sum_{i \in \mathcal{A}^{(r)}} \omega_i \frac{p_i - l_i}{u_i - l_i} \geq T^{(r)}, l_i \leq p_i \leq u_i, \forall i \in \mathcal{A}^{(r)}. \quad (7)$$

Applying the same Lagrangian relaxation method as shown in solving Model 5,

$$p_i^* = \begin{cases} \frac{l_i + \omega_i \lambda}{2}, & \text{if } l_i / \omega_i \leq \lambda \leq (2u_i - l_i) / \omega_i, \\ l_i, & \text{if } \lambda \leq l_i / \omega_i, \\ u_i, & \text{otherwise.} \end{cases} \quad (8)$$

Similar to Model 6, the problem is strictly convex, so it has a unique optimal solution and there is no duality gap when we take the Lagrangian relaxation.

PROPOSITION 2.3.

- a) [Symmetry] *For any two bidders with the same valuation distribution and with the same weights, the optimal offer prices must be the same.*
- b) [Monotonicity] *For any bidders with the same valuation distribution, the optimal offer prices are higher for those with higher weights.*

PROOF. In part (a), for any two bidders with identical value distributions and weights, any choice of λ would lead to the same formulations for the offer prices (as shown in Formulation 8) for the two bidders, which means the optimal offer prices are the same. Part (b) can be derived in the same way with a note that both ω and λ are non-negative. \square

The implication of part (b) is that the auctioneer should offer higher prices to more ‘important’ bidders to keep them active. We will show how this is applied to the case of incentive auction where the stations affect the feasibility differently.

PROPOSITION 2.4. *The optimal offer prices can be found—i.e., Model 7 can be solved—in $O(n \log n)$ operations.*

PROOF. The proof is similar to that of Proposition 2.2. \square

2.1.3. *General Valuation Distributions.* Now we drop the assumption that the bidders' valuation distributions are uniform. In this setting, Model 7 generalizes to a nonlinear program:

$$\text{minimize } \sum_{i \in \mathcal{A}^{(r)}} F_i(p_i)p_i, \text{ subject to } \sum_{i \in \mathcal{A}^{(r)}} \omega_i F_i(p_i) \geq T^{(r)}, l_i \leq p_i \leq u_i, \forall i \in \mathcal{A}^{(r)}. \quad (9)$$

Depending on the valuation distributions, this program might be non-convex and might be hard to solve in general. The attractive part of this model, however, is that it has a separable objective function and only one joint constraint. Using the same Lagrangian relaxation method, we obtain the following Lagrangian dual problem:

$$\max_{\lambda \geq 0} \left\{ T^{(r)}\lambda + \sum_{i \in \mathcal{A}^{(r)}} \min_{l_i \leq p_i \leq u_i} (F_i(p_i)p_i - \lambda \omega_i F_i(p_i)) \right\}.$$

For each fixed Lagrangian multiplier λ , the optimal prices \mathbf{p} of the inner problem can be found by solving n nonlinear sub-problems, each with a single scalar variable. In the case where the distribution of the value function is piecewise linear, the inner problem is a piecewise cubic function where the optimal solution of each piece can be found efficiently.

The problem might be non-convex so the Lagrangian relaxation method might produce some optimality gap. The Lagrangian relaxation method has been shown to perform well empirically, that is, with small optimality gaps, for a number of combinatorial problems (see Fisher [2004] for details about the method and successful applications). The offer prices found, however, might not be optimal. In what follows we present a method that could either be used as a stand-alone method for finding the optimal offer prices or can be used in combination with the Lagrangian relaxation method to enhance the performance. This method takes advantage of the discreteness property of the decision variables, the separability of the objective function, and the single joint constraint. Consider an auction design where the offer price for bidder i must take discrete values in a given set $\{P_{i1}, \dots, P_{ik}\}$. This restriction often holds in practical auction settings because (1) allowing fractional bids can increase hassle (e.g., bookkeeping), and (2) allowing unimportant digits to be expressed opens the door for collusion among bidders. Such collusion has been observed in FCC auctions and the FCC has subsequently practiced, for example, “click-box” bidding where bidders have to select their bids from a small set of discrete values [Cramton and Schwartz 2000; Bajari and Yeo 2009]. The problem becomes

$$\text{minimize } \sum_{i=1}^n F_i(p_i)p_i, \text{ subject to } \sum_{i=1}^n \omega_i F_i(p_i) \geq T, p_i \in \{P_{i1}, \dots, P_{ik}\}, \forall i \in \mathcal{N}. \quad (10)$$

This is a mixed-integer nonlinear program. They are generally very difficult to solve. However, this one has a knapsack-type structure so a dynamic program can be utilized:

PROPOSITION 2.5. *Optimal offer prices for Model 10 can be found in $O(KLn^2)$ operations, where K is the maximum number of discrete price levels and L is the number of precision points in the range $[0, 1]$ used to calculate the cumulative values $F_i(\cdot)$.*

The proof of Proposition 2.5 is shown in Appendix A.3. Here we discussed the problem in the first round where $\mathcal{A}^{(1)} = \mathcal{N}$ and has the largest size. Results for other rounds can

be derived similarly. Also, it is possible to combine the idea from Lagrangian relaxation and dynamic programming to improve the computational performance further. This can be done by using the approximated results from the Lagrangian relaxation method and using it to guide the discretization of the feasible space for p before applying the dynamic program. For example, the discretization around the solution suggested by the Lagrangian relaxation could be more refined than elsewhere.

2.2. Updating Value Function Distributions

The auctioneer utilizes the price discovery feature of the DCA to update the estimated value functions in Step 2.2.2 of Algorithm 1. Suppose that, at the beginning of round r , the auctioneer has a belief that the random threshold price v_i for bidder i follows a distribution on the support set $[l_i, u_i]$ with a cumulative distribution $F_i(\cdot)$. Once the auctioneer has made an offer $p_i^{(r)}$ to i , there are two cases. In the first case, i rejects the offer. The rejected list is updated and i is no longer active (except in the last round when the target is not met). In the second case, i accepts the offer and stays active. The auctioneer needs to update the belief about the threshold price of bidder i based on the fact that the bidder accepted. We need to find the conditional distribution for $(v_i | v_i \leq p_i^{(r)})$, which can be calculated as follows:

$$Prob(v_i \leq a | v_i \leq p_i^{(r)}) = \begin{cases} 0, & \text{if } a \leq l_i, \\ 1, & \text{if } a > p_i^{(r)}, \\ \frac{Prob(v_i \leq a)}{Prob(v_i \leq p_i^{(r)})} \equiv \frac{F_i(a)}{F_i(p_i^{(r)})}, & \text{if } l_i \leq a \leq p_i^{(r)}. \end{cases} \quad (11)$$

The auctioneer then uses this new set of beliefs on the value functions and the new set of active bidders to run a new auction round. If $F_i(p_i)$ is piece-wise linear at iteration r , then the updated cumulative distribution is also piece-wise linear at iteration $(r + 1)$. In the special case of a uniform distribution, the updated distribution is uniform on $[l_i, p_i^{(r)}]$, that is, we simply update the upper bound to be $p_i^{(r)}$.¹

2.3. Final-Round Settlement

In Step 3 of Algorithm 1, the auctioneer selects the winners of the auctions after having undertaken the price discovery through the multiple rounds of the DCA. Let p be a vector of prices that the auctioneer offers to the bidders in the last round. Given the offer prices, the bidders might accept or reject the offers. The auctioneer then updates the upper bounds on the bidders' values: upper bounds for bidders that accepted will be updated to the offer prices and those of the bidders who rejected remain unchanged. After the final round, the auctioneer does not have any further opportunity to do price discover, and has to decide which bidders are the final winners to meet the target. For the simple case where each bidder has only one item, the auctioneer chooses T bids with the smallest updated upper bounds and pays each of those bidders that price. For the weighted case, the auctioneer solves the following knapsack problem to determine a set of winners that meets the target with the least total payment:

$$\text{minimize } \sum_{i \in \mathcal{A}^{(m)}} u_i z_i, \text{ subject to } \sum_{i \in \mathcal{A}^{(m)}} w_i z_i \geq T, z_i \in \{0, 1\}, \forall i \in \mathcal{A}^{(m)}, \quad (12)$$

where z_i indicates whether bid i should win.

¹In general, there could also be other factors to take into account when updating the value function distribution. For one, the auctioneer might want to update the bounds based on how other bidders have responded to offers. This, and DCAs in general, beget interesting questions for future research related to interdependent valuations. Within the scope of this paper, our focus is, however, on how to set offer prices given the beliefs.

3. EXPERIMENTS

In this section we instantiate the methodology in two settings, and present optimization experiments in both. We start with a setting where the items are homogeneous. We then proceed to the reverse auction in incentive auctions, using real FCC data.

3.1. DCA with Homogeneous Items

In this section we study a relatively simple auction with n bidders, each of whom has one item to be sold. The objects are identical from the auctioneer’s point of view, and the auctioneer has a target of buying T objects. In this case, the feasibility function is $F(\mathcal{A}) = 1$ if $|\mathcal{A}| \geq T$, and $F(\mathcal{A}) = 0$ otherwise.

To apply the general DCA framework of Algorithm 1 to this relatively simple setting, we need to adapt Algorithm 1 in Step 2.1 and set the target number of accepting bidders to $T^{(r)} = n^{(r)} - (n^{(r)} - T)/(m - r)$, where $n^{(r)} = |\mathcal{A}^{(r)}|$, and solve Model 9 to find a vector of prices p to offer the bidders. In this step, we assume the auctioneer first aims for a trajectory of the sets of active bidders and then optimizes the prices correspondingly. Another strategy that the auctioneer could adopt is to first try to ‘foresee’ the offer prices in the final round by solving Model 9 for $T^{(r)} = T$, and then shrink these final prices to the upper bound by taking advantage of the multi-round guessing in price discovery. We present details about this strategy in Appendix A.6. We also need to adapt Step 3 of Algorithm 1 by solving Formulation 12. Details about this and other steps are presented in Algorithm 2 in Appendix A.4.

We compare the performance of our method with a natural percentile-based method where the prices at round r are set to

$$p_i^{(r)} = l_i + \alpha_i(u_i^{(r)} - l_i), \quad (13)$$

where $\alpha_i = 1 - \frac{Q}{n_r}$ and $Q = \frac{n_r - T}{(m - r + 1)}$ is the expected number of rejections per round. This essentially aims to distribute the expected number of rejections evenly among m rounds.

This percentile-based pricing scheme needs to be incorporated into a DCA such as in Step 2.1 of Algorithm 1. When we refer to a percentile-based method for the homogeneous setting, we mean Algorithm 2 (in the appendix) with Step 2.1 being replaced by the pricing from Equation 13. Similarly, when we refer to a percentile-based method for incentive auctions in the next subsection, we mean Algorithm 3 (in the appendix) with Step 2.1 replaced by the percentile-based pricing from Equation 13.

One can develop other percentile-based methods such as setting the offer prices at a fixed percentile in the bidders’ value distributions, for example, either always aiming for a fixed 5% of rejection at each round or dynamically having this percentile depend on the current number of active bidders and the number of rounds remaining. We conducted extensive tests using these percentile-based methods. Their performance was almost the same on average as the percentile-based method described above. Therefore, we only present the performance of that percentile-based method in comparison with our optimization-based method.

In the experiments, we let there be $n = 100$ bidders. We check the performance of the algorithms for various choices of the target T . We first generate random bounds for the bidders’ valuations. The upper and lower bounds for bidder i are set to $u_i = (1 + \delta)m_i$ and $l_i = (1 - \delta)m_i$, where m_i is a uniform random variable in $[0, 1]$. Here, δ is a measure of how good the auctioneer’s estimate of the bidders’ values is. We vary δ between 10% and 50%. We then draw $M = 10$ sample valuation vectors with bidder valuations from these ranges, that is, $\xi_i^{(k)} \sim U[l_i, u_i]$ for $k = 1, \dots, M$, and for each bidder $i = 1, \dots, n$. That gives us M auction instances to run on. We report the average of them in the figures.

For convenience in reference, in the rest of the paper, we use the label ‘OPT-SCHED’ to refer to DCAs that make use of the price optimization method suggested by Model 9. These include output from Algorithm 2 for the homogeneous case and Algorithm 3 for incentive auctions (presented later in Section 3.2).

Figures 1-a and 1-b show the performance of the DCA for the setting where the number of rounds allowed $m = 50$, the weights $w_i = 1$ for all bidders i , the target $T = 50$, and $\delta = 20\%$.




Fig. 1. Comparison between OPT-SCHED and the percentile-based algorithm in terms of price discovery.

The horizontal axis is the round number in the DCA. The vertical axis shows the number of active bidders, the total payment the auctioneer would pay if the auction ended at that round, the total value of the active bidders, and the optimal (lowest) payment, OPT , with which the auctioneer could procure the needed items if he knew the bidders’ valuations.

The number of active bidders, the total payment, and the total values decrease during the DCA as we expected. The total payment is always above the total value, and the total value is always higher than the optimal value. Comparing the final payments at the last round, one can see that our optimization-based approach significantly outperforms the percentile-based approach.

Figure 2 shows the final payoff as a function of the number of rounds allowed—for different targets T using the same parameters described earlier.




Fig. 2. Comparison between OPT-SCHED and the percentile algorithm in terms of the final DCA payments for a varying number of rounds allowed.

One can see that our optimization-based approach significantly outperforms the percentile-based approach across the board. Surprisingly, the final payment does not decrease when increasing the number of rounds allowed. This surprising finding does not always occur: we can find settings where the final payments for both the percentile-based method and the optimization-based method do keep decreasing as we allow more

rounds. We can also design a specific percentile-based method that avoids this strange behavior completely, for example, by reducing offer prices of only the bidder with the highest upper bound first. However, that percentile-based method will not be effective in minimizing the final total payment.

We phrase the surprising behavior as a paradox and then proceed to explain it.

PARADOX 1. *Having more refined offer prices can lead to a higher total payment.*

Explanation of the paradox: Throughout the multi-round auction, there are two essential effects on the final payment: (a) Desirable effect (price discovery): The prices offered to the remaining active bidders keep decreasing; this lowers the final payment, and (b) Undesirable effect: Some lower-priced bidders are ‘accidentally’ rejected as their offer prices keep decreasing.

The first, desirable effect occurs in both the percentile-based method and the optimization-based method for price setting (OPT-SCHED). The second, undesirable effect occurs more in the percentile-based method and than in OPT-SCHED because the objective function in Model 10 has already aimed to lower this expected payment and hence high-priced bids are supposed to be rejected before others. This leads to overall better performance of OPT-SCHED compared to the percentile-based method, as shown in Figure 2. To summarize, the reason behind the paradox is that, due to the randomness of the bid values, having more rounds allow a higher chance of rejecting good bids before the final adjustment round. Once the good choices are excluded through more rounds, the less choice we have in solving the final settlement problem shown in Model 12, and thus the payment that the model is trying to minimize is higher.

We demonstrate this paradox through a specific example. Consider the following simple case with $n = 3$ bidders and a target number $T = 2$ items to be procured. Suppose all the bidders are still active at round r where the current offer prices are $p_i^{(r)}$. Consider the following two strategies: In strategy 1, we reduce all the prices in round $(r + 1)$ by the amounts δ_i for bidder i , and in strategy 2, we reduce all the prices in round $(r + 1)$ by smaller amounts $\beta_i < \delta_i$ first, observe the bidders’ responses, and then reduce the price by the additional amounts $(\delta_i - \beta_i)$ in round $(r + 2)$.

Consider the following scenario: Bidder 1 accepts new price $(p_1^{(r)} - \delta_1)$ but Bidders 2 and 3 do not accept $(p_2^{(r)} - \beta_2)$ and $(p_3^{(r)} - \delta_3)$. In this case, the number of acceptances is one for both strategies, so the auctions proceed to the adjustment round. For Strategy 1, the adjustment round involves solving a knapsack problem to choose the two smallest offers among three choices $(p_1^{(r)} - \delta_1, p_2^{(r)}, p_3^{(r)})$. For Strategy 2, the adjustment round involves solving a knapsack problem to choose the two smallest offers among three choices $(p_1^{(r)} - \beta_1, p_2^{(r)}, p_3^{(r)})$. Since $\beta_1 < \delta_1$, the knapsack problem resulting from Strategy 1 has lower (or equal) cost compared to that from Strategy 2.

3.2. DCA in FCC Incentive Auctions

In this section we conduct experiments with a model of the FCC incentive auction that uses real FCC data regarding the feasibility. A key feature of reverse auctions in incentive auctions is that the feasibility function is highly sensitive to the set of rejected bids. This is due to a large set of engineering constraints between the stations that restrict them from being assigned to the same or adjacent channels. This means the inclusion of a set of stations in a current reject list would make the characteristics of remaining stations totally different from each other. We first describe these interference constraints and the feasibility function. We then show how the general DCA framework can be applied to incentive auctions.

3.2.1. *Interference Constraints in Repacking and Feasibility Checking.* The description of the FCC incentive auction DCA setting with the engineering constraints is available in detailed files on the FCC web site, which we used [FCC 2013]. There are $n = 2177$ stations and $m = 49$ channels available (ranging from channel 2 to channel 51, with channel 37 not available). The target is to clear channels 33-51, that is, to repack all the rejected stations to channels 2-32. The feasibility function can be defined as $F(S) = 1$ if $S \in \mathcal{P}(\mathcal{C})$, and $F(S) = 0$ otherwise, where $\mathcal{P}(\mathcal{C})$ is the set of feasible assignments to a list of available channels \mathcal{C} and is defined as

$$\mathcal{P}(\mathcal{C}) = \left\{ \mathbf{z} : \begin{array}{l} z_{ik} \in \{0, 1\}, \forall i \in \mathcal{S} \text{ and } k \in C_i, \sum_{k \in C_i} z_{ik} = 1, \forall i \in \mathcal{S}, \\ z_{ik} + z_{jk} \leq 1, \forall (i, j, k) \in \mathcal{I}_c, z_{ik} + z_{j(k+1)} \leq 1, \forall (i, j, k) \in \mathcal{I}_a \end{array} \right\}. \quad (14)$$

Here z_{ik} is a binary variable that indicates whether station i is assigned to channel k , $C_i \subset \mathcal{C}$, $i \in \mathcal{S}$, is the list of feasible channels to station i , \mathcal{I}_c is the list of triplets (i, j, k) such that stations i and j cannot be assigned to the same channel k , and \mathcal{I}_a is the list of triplets (i, j, k) such that stations (i, j) cannot be assigned to neighboring channels $(k, k + 1)$. We present more details about the repacking problem in Appendix A.7.

3.2.2. *Adapting the DCA to FCC Reverse Auctions.* We adapt the general DCA framework of Algorithm 1 to incentive auctions. First, as in the homogeneous setting, we need to adapt Step 2.1 and set the target number of accepting bidders to $T^{(r)} = n^{(r)} - (n^{(r)} - T)/(m - r)$, where $n^{(r)} = |\mathcal{A}^{(r)}|$, and solve Model 9 to find a vector of prices \mathbf{p} to offer the bidders. The key difference between incentive auctions and the homogeneous setting is in the feasibility function. Specifically, the feasibility function in incentive auctions is a complicated function that is highly sensitive to the set of rejected stations. We need to adapt Step 2.2. of Algorithm 1 to find the set of new rejections and check the feasibility on the updated set of rejected stations. Details about this and other steps are presented as Algorithm 3 in Appendix A.5.²

In order to apply the models developed in Section 2 to finding the offer prices in each round, we need to modify the feasibility constraint. As the bidders' true values are random variables and unknown to the auctioneer, any fixed set of offer prices leads to a stochastic set of rejected bidders and hence the repacking feasibility is also stochastic. However, at the beginning of each round, we could simulate the feasibility problem to draw a curve that shows the probability of having feasible repacking as a function of the number of new stations added. We could then choose a target $T^{(r)}$ so that the chance of feasibility when adding $T^{(r)}$ new stations is at some threshold (say 99% chance of feasibility). Once a target $T^{(r)}$ has been determined, we can then solve Model 5 to obtain the offer price. The choice of the feasibility probability (and hence $T^{(r)}$) would depend on how quickly or slowly the auctioneer wants to run the auction. For example, at the beginning of the auction, the auctioneer might want to have small $T^{(r)}$ for more accurate price discovery but then increase $T^{(r)}$ toward the end of the auction to lower the expected payment faster.

Here we present a simple algorithm for price setting in incentive auctions. Within the scope of this paper, we do not undertake extensive simulation to obtain a cumulative function of the feasibility with respect to the number of new stations, $T^{(r)}$, added in each round. Instead, we use an estimated $T = 1177$ number of stations in the final set of active bidders, that is, to have $U = 1000$ stations feasibly rejected.³ We compare

²In Step 3 of Algorithm 3, it is possible to do an improvement and select a smaller set of winners by solving a winner determination problem constrained by having $\mathcal{R}^{(r)}$ as a subset of rejected bids. That problem, however, has a similar structure to a winner determination problem arising from a single-round, sealed-bid auction and is difficult to solve due to the large set of engineering constraints.

³This choice of U comes from the prior experiments we have on the feasibility problem.

the performance of our optimization model with the percentile-based method where both use the same U .

To deal with the heterogeneity of the stations with respect to the feasibility function, we propose to associate each active station with a weight that is proportional to the possibility of causing interference on other stations, especially those that have already been rejected. One possibility for setting such a weight vector is to set

$$w_i = d_i(\mathcal{N}) + d_i(\mathcal{R}^{(r)}), \forall i \in (\mathcal{N} \setminus \mathcal{R}^{(r)}), \quad (15)$$

where $d_i(\mathcal{N})$ is the number of stations in the entire set of stations \mathcal{N} that i might interfere with, and $d_i(\mathcal{R}^{(r)})$ is the number of stations in the current rejected list that i might interfere on. By setting the weights higher for stations with a higher potential of interference on others, we essentially try to avoid having these stations rejected in the next rounds by offering them higher prices (recall Equation 8 and Proposition 2.3).

3.2.3. Experimental Results on Incentive Auctions. Since no incentive auctions have yet been conducted, we have to use simulated data on the bounds of the bidders' valuations. The bounds for the first experiment are generated using a uniform distribution where the upper and lower bound for bidder i are set to $u_i = (1+\delta)m_i$ and $l_i = (1-\delta)m_i$, where m_i is a uniform random variable. Here, $\delta = 0.2$ is a measure of how good the auctioneer's estimate of the bidders' valuations is. We then draw random sample bid values from these ranges, that is, $\xi_i \sim U[l_i, u_i]$ for each bidder $i = 1, \dots, n$. We draw $M = 10$ valuation vectors. Each vector corresponds to a DCA instance. Each instance has one valuation per bidder.

Figure 3-a compares the performance of the DCA when using our optimization-based price-setting method (Algorithm 4) versus the percentile-based method, for the bids generated uniformly in $[0, 1]$, for a sample DCA instance. The number of rounds allowed is 50.

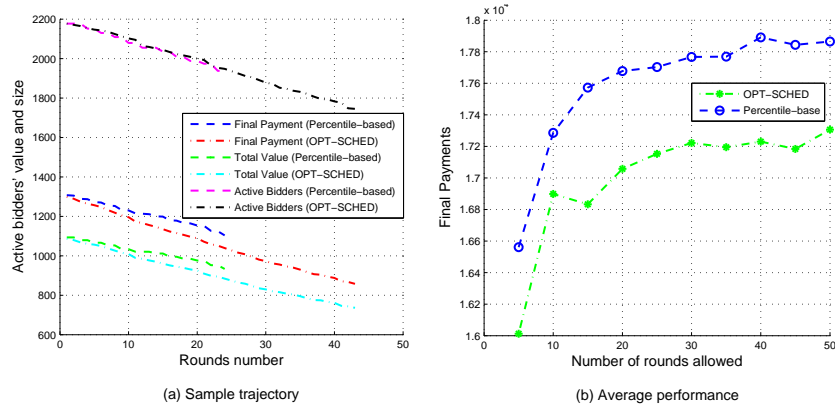


Fig. 3. Performance of our optimization-based pricing strategy versus that of the percentile-based method—using real FCC engineering constraints data and using homogeneous bids.

The horizontal axis of Figure 3-a shows the round number in the DCA. The vertical axis shows the number of active bidders, the total payment the auctioneer would pay if the auction ended at that round, and the total actual values of the active bidders. In each pair of these curves, the dashed curves correspond to the percentile-based method and the dot-dashed curves to our optimization-based method. As can be seen, the number of active bidders, the payments and the total actual values decrease through the rounds. The payment is always above the total actual values as we expected. Comparing the final payments at the last round, one can see that OPT-SCHED result-

s in more stations to be reallocated and also in a lower final payment compared to the percentile-based approach. The lower final payment of OPT-SCHED was partially due to it having more auction rounds before encountering infeasibility: the percentile-based method encountered infeasibility at round 24 while OPT-SCHED encountered it at round 43. The better performance of OPT-SCHED in dealing with infeasibility is due to the added weighted constraints which essentially avoid rejecting stations that are likely to cause interference. Another reason why OPT-SCHED outperformed the percentile-based method is that it has a better way of rejecting high-priced bids. This effect can be seen in each fixed number of rounds. For example, if both algorithms are terminated at round 20, the payment for OPT-SCHED is 1088.7 while that for the percentile-based method is 1153.4.⁴

Figure 3-b shows the average final payoff to bidders versus the number of rounds allowed.⁵ One can see that our optimization-based approach outperforms the percentile-based approach on average for all choices of the number of rounds allowed. The paradox of having the final payment increase as we increase the number of rounds allowed is again observable.

We now present a setting where OPT-SCHED has some issues when high-valued bids interfere with each other, for example, when groups of stations in the same area request high values. We demonstrate this through an experiment where the mean value m_i is set proportional to the population that station i serves. Figure 4 shows the average performance over five randomly generated auction instances.




Fig. 4. Final payments using our optimization-based pricing strategy versus those of the percentile-based method—using real FCC engineering constraints data and using heterogeneous bids.

In this case, OPT-SCHED yields higher payments on average than the percentile-based approach (although it yielded lower payment on two of the five instances). In what follows, we provide some explanation for this. When setting the bids to be proportional to the population served, the bids from the stations in large cliques (e.g., stations in large cities) are often large since the populations served are large. If the number of rejected large bids is greater than the number of channels available, then we will reach infeasibility in repacking. Note that this issue is problematic for general DCAs where many small subsets of bidders can significantly affect the feasibility function. Among the random instances generated, the performance of the percentile-based approach was less affected by this because it uses a very conservative approach of gradually lowering prices for all bidders in the same way. One idea to improve OPT-SCHED is to add clique constraints on these subsets of bidders to enforce that no more than a certain number of stations in each large clique is rejected. Another idea is to

⁴In the incentive auction setting we do not have a number for the optimal payment if the auctioneer knew the bidders' valuations since the WDP is too difficult to solve to optimality according to Milgrom and Segal [2013].

⁵The random variable m_i was uniformly distributed within [9.5, 10.5] for each bidder $i = 1, \dots, n$.

enforce that no price decreases by more than, say, 25% between rounds. We recently conducted initial experiments with these approaches and they help, but so far we did not find settings that outperform the percentile-based approach on average in terms of minimizing payments. This remains an important direction for future research.

4. CONCLUSIONS AND DISCUSSION

A descending (multi-item) clock auction (DCA) is a mechanism for buying items from multiple potential sellers. DCAs have been proposed as the method for procuring spectrum from existing spectrum holders in the FCC's imminent incentive auctions so spectrum can be repurposed to higher-value uses. However, the DCA design has lacked a way to determine the prices to offer the bidders in each round. This is a recognized, important, and timely problem.

We presented, to our knowledge, the first techniques for this. We develop a percentile-based approach which provides a means to naturally reduce the offer prices to the bidders through the bidding rounds. We also develop an optimization model for setting prices so as to minimize expected payment while stochastically satisfying the feasibility constraint. (The DCA has a final adjustment round that obtains feasibility after feasibility has been lost in the final round of the main DCA.) We proved attractive properties of this, such as symmetry and monotonicity. We also developed computational methods for solving the model and very efficient polynomial-time algorithms. (We also developed optimization models with recourse, but they are not computationally practical.)

We presented experiments both on the homogeneous items case and the case of FCC incentive auctions, where we used real FCC interference constraint data to get a fully faithful model of feasibility. The experiments showed that the optimization-based price-setting approach significantly outperforms the natural percentile-based approaches in minimizing the final payment by the auctioneer in DCAs with homogeneous bidders. In incentive auctions, the performance is mixed and the optimization-based price-setting approach is favourable if bids are generated uniformly but this is not the case if subsets of interfering stations request high prices together. Adding clique constraints or enforcing that prices cannot decrease too much between rounds are techniques that help improve the optimization-based approach, but further research on this is needed. An unexpected paradox on the performance of DCAs was that sometimes when the number of rounds allowed increases, the final payment can actually increase. We provided an explanation of this paradox.

There are a number of potential future directions that can be followed up from this research. First, we find the paradox concerning the relationship between the expected payment and the number of rounds allowed quite intriguing. Further research on when this would occur, and on quantification of these relationships, would be interesting. In this project, we have only considered the case where bidders' values are independent. It would be interesting to extend the techniques to settings with interdependent valuations. Another extension would be to incorporate into the model the option of stations agreeing to share a channel or to move to lower bands—something that the FCC is seriously considering. The current DCA framework is designed for settings where bidders can just accept or reject. We are extending the framework to incorporate bidders' multiple choices into DCAs and to find the optimal prices accordingly.

Regarding the implementation of the reverse auctions DCA for the FCC incentive auctions, having extensive simulation for better estimation of the feasibility chance would improve the performance. Other ideas mentioned in the paper, such as using stochastic programming and chance constraints for the feasibility constraints, as well as approximated dynamic programming techniques are worth pursuing for such a high-stakes setting.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library. Electronic appendix is also available at <http://www.cs.cmu.edu/~sandholm/descendingClockAuction.ec14.Appendices.pdf>.

REFERENCES

- Patrick Bajari and Jungwon Yeo. 2009. Auction design and tacit collusion in FCC spectrum auctions. *Information Economics and Policy* 21, 2 (2009), 90–100.
- Peter Cramton. 1997. The FCC spectrum auctions: an early assessment. *Journal of Economics and Management Strategy* 6, 3 (1997), 431–495. Special issue on market design and spectrum auctions.
- Peter Cramton and Jesse Schwartz. 2000. Collusive bidding: Lessons from the FCC spectrum auctions. *Journal of Regulatory Economics* 17 (2000), 229–252.
- Peter Cramton, Yoav Shoham, and Richard Steinberg. 2006. *Combinatorial Auctions*. MIT Press.
- Robert Day and Paul Milgrom. 2008. Core-selecting package auctions. *International Journal of Game Theory* 36, 3 (2008), 393–407.
- FCC. 2012. *NOTICE OF PROPOSED RULEMAKING*. Technical Report. FCC. http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-12-118A1.pdf.
- FCC. 2013. FCC Repacking Constraint Files. (2013). http://data.fcc.gov/download/incentive-auctions/Constraint_Files/.
- FCC. 2014. The Path to a Successful Incentive Auction. (2014). FCC presentation 1/30/2014, http://wireless.fcc.gov/incentiveauctions/learn-program/Incentive_Auction_Jan_30_Present_9am.pdf.
- Marshall L Fisher. 2004. The Lagrangian relaxation method for solving integer programming problems. *Management Science* 50, 12 supplement (2004), 1861–1871.
- Thomas Hazlett, David Porter, and Vernon Smith. 2012. “Incentive Auctions” Economic and Strategic Issues. (2012). Arlington Economics white paper 6/12/2012.
- Kevin Leyton-Brown. 2013. Investigating the Viability of Exact Feasibility Testing. (2013). Working Paper, University of British Columbia / Auctionomics, <http://www.cs.ubc.ca/~kevinlb/talk.php?u=2013-Feasibility-Testing.pdf>.
- Paul Milgrom. 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- Paul Milgrom and Ilya Segal. 2012. Heuristic Auctions and U.S. Spectrum Repurposing. (2012). Presentation, Stanford University, www.stanford.edu/~isegal/fcc.ppt.
- Paul Milgrom and Ilya Segal. 2013. Deferred-Acceptance Heuristic Auctions. (2013). Working Paper, Stanford University, <http://www.milgrom.net/downloads/heuristic.pdf>.
- Tuomas Sandholm. 2006. Optimal Winner Determination Algorithms. In *Combinatorial Auctions*, Peter Cramton, Yoav Shoham, and Richard Steinberg (Eds.). MIT Press, 337–368. Chapter 14.
- Tuomas Sandholm. 2013. Very-Large-Scale Generalized Combinatorial Multi-Attribute Auctions: Lessons from Conducting \$60 Billion of Sourcing. In *Handbook of Market Design*, Zvika Neeman, Alvin Roth, and Nir Vulkan (Eds.). Oxford University Press.
- Tuomas Sandholm, Subhash Suri, Andrew Gilpin, and David Levine. 2005. CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions. *Management Science* 51, 3 (2005), 374–390.
- Tom Wheeler. 2013. The Path to a Successful Incentive Auction. (2013). FCC Chairman’s post on the official FCC blog 12/6/2013.

Online Appendix to: Optimizing Prices in Descending Clock Auctions[†]

TRI-DUNG NGUYEN, University of Southampton, Schools of Mathematics and Management
TUOMAS SANDHOLM, Carnegie Mellon University, Computer Science Department *Corresponding author

APPENDIX

A.1. A Dynamic Programming Model for Optimal Price Setting in DCA

In each round of the descending clock auction, the auctioneer needs to offer each active bidder a price, i.e., to do Step 2.1. of Algorithm 1. Here we show a dynamic programming model that the optimal set of offer prices should solve.

Let $V(m, \mathcal{S}, \mathbf{u}, \mathbf{l})$ be the minimum expected payment that the auctioneer needs to pay to the bidders in a descending clock auction with m rounds, with a set of active bidders \mathcal{S} , with upper bounds \mathbf{u} and lower bounds \mathbf{l} within which the bidders' valuations lie. Let ξ be a realization of the bidders' values. For any offer prices \mathbf{p} in the first round, the state of the auction by the end of that first round will be as follows.

- The number of rounds left will be $(m - 1)$.
- The remaining active bidders will be $\mathcal{S}(\mathbf{p}, \xi)$. This includes bidder i if the offer price p_i is no smaller than the bidder's value ξ_i , i.e., $p_i \geq \xi_i$.
- A new vector of upper bounds $\mathbf{u}(\mathbf{p}, \xi)$ which updates the upper bound of any remaining active bidder i to x_i .
- Unchanged lower bounds \mathbf{l} .

The minimum expected value that the auctioneer needs to pay under the new state of the auction will be $V(m - 1, \mathcal{S}(\mathbf{p}, \xi), \mathbf{u}(\mathbf{p}, \xi), \mathbf{l})$. Thus, the auctioneer's problem in the first round is to choose \mathbf{p} that minimizes the expectation of $V(m - 1, \mathcal{S}(\mathbf{p}, \xi), \mathbf{u}(\mathbf{p}, \xi), \mathbf{l})$. We have the Bellman optimality equation

$$V(m, \mathcal{S}, \mathbf{u}, \mathbf{l}) = \min_{\mathbf{p}} E[V(m - 1, \mathcal{S}(\mathbf{p}, \xi), \mathbf{u}(\mathbf{p}, \xi), \mathbf{l})].$$

Solving this dynamic program would be extremely difficult. In fact, just finding $V(m, \mathcal{S}, \mathbf{u}, \mathbf{l})$ for the case $m = 1$ would be very difficult as shown in a simple case below.

A.2. Optimal Price Setting in the Last Round with Recourse Action

Consider the problem of setting prices in the final round of a descending clock auction. Assume that the actual bid values are uniformly distributed random variables, i.e., $\xi_i \sim U[l_i, u_i]$, where $(l_i, u_i), i = 1, \dots, n$ are known. Let \mathbf{p} be a vector of prices that the auctioneer offers to the bidders. Given the offer prices, the bidders might accept or reject the offers. The auctioneer then updates the best upper bounds on the bids values, that is, upper bounds for accepting bidders will be updated to the offer prices while those of rejected bidders will remain unchanged. The auctioneer chooses T bids with the smallest updated upper bounds and pays each of these bidders those prices. Since the bidders' values are random variables, the acceptance of the bidders for each set of offer prices \mathbf{p} will also be stochastic, so the final payment is stochastic. We consider the problem of finding the optimal offer prices \mathbf{p} such that the expected final payment is minimized. Here expectation is taken over the randomness of the bidders' valuations.

For convenience in notation, we perform a linear transformation on the price vectors \mathbf{p} to \mathbf{x} where $x_i = \frac{p_i - l_i}{u_i - l_i}$, that is, $x_i \in [0, 1]$ can be interpreted as the target chance of acceptance for bidder i . We also have $p_i = l_i + x_i(u_i - l_i)$. Let us denote by $f(\mathbf{x})$ the stochastic payment.

Let us first consider the simple case where $T = 1$ and $n = 2$. Here the payment is $\min(u_1, u_2)$ if both bidders reject the offers, $\min(p_1, p_2)$ if both of them accept, and $p_i, i = \{1, 2\}$ if only bidder i accepts the offer. The probability for each of these four events can be calculated as functions of \mathbf{x} . For example, the chance of rejecting both offers is $(1 - x_1)(1 - x_2)$. Putting all of these together, we have

$$f(\mathbf{x}) = \begin{cases} \min(u_1, u_2), & \text{w.p. } (1 - x_1)(1 - x_2), \\ \min(l_1 + x_1(u_1 - l_1), l_2 + x_2(u_2 - l_2)), & \text{w.p. } x_1x_2, \\ l_1 + x_1(u_1 - l_1), & \text{w.p. } x_1(1 - x_2), \\ l_2 + x_2(u_2 - l_2), & \text{w.p. } (1 - x_1)x_2. \end{cases}$$

The expected payment is

$$E[f(\mathbf{x})] = (1 - x_1)(1 - x_2)\min(u_1, u_2) + x_1x_2\min(l_1 + x_1(u_1 - l_1), l_2 + x_2(u_2 - l_2)) + x_1(1 - x_2)(l_1 + x_1(u_1 - l_1)) + (1 - x_1)x_2(l_2 + x_2(u_2 - l_2)).$$

The problem of determining the optimal offer prices can therefore be formulated as

$$\begin{aligned} \min_{x_1, x_2} \quad & E[f(\mathbf{x})] \\ \text{s.t.} \quad & 0 \leq x_i \leq 1, \forall i = 1, 2, \end{aligned}$$

which is a non-convex quadratic optimization problem. If we extend the problem to the case $n > 2$, the problem becomes a polynomial optimization problem as follows;

$$\begin{aligned} \min_{\mathbf{x}} \quad & \sum_{S \subset \mathcal{N}, S \neq \emptyset} \left[\prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i) \min_{i \in S} \{l_i + x_i(u_i - l_i)\} \right] + \prod_{i \in \mathcal{N}} (1 - x_i) \min_{i \in \mathcal{N}} u_i \\ \text{s.t.} \quad & 0 \leq x_i \leq 1, \forall i = 1, \dots, n, \end{aligned}$$

which is very difficult to solve. Notice that we have considered only the simple case of $T = 1$ and also considered finding the optimal decision in the last round only.

A.3. Proof of Proposition 2.5

PROOF. For each $m \in \{0, 1, \dots, n\}$ and for each budget $B \geq 0$, let us define

$$\begin{aligned} V(m, B) = \min_{\mathbf{p}} \quad & \sum_{i=1}^m F_i(p_i)p_i, \\ \text{s.t.} \quad & \sum_{i=1}^m \omega_i F_i(p_i) \geq B, \\ & p_i \in \{P_{i1}, \dots, P_{i,k}\}, \forall i = 1, \dots, m. \end{aligned}$$

Then we have

$$\begin{aligned} V(m, B) = \min_{p_m} \quad & F_m(p_m)p_m + V(m - 1, B - F_m(p_m)), \\ \text{s.t.} \quad & p_m \in \{P_{m1}, \dots, P_{m,k}\}, \end{aligned} \tag{16}$$

where $V(0, B) = 0, \forall B$. Suppose $F_j(p_j)$ receives one of $(L + 1)$ values in the set $\{0, 1/L, \dots, 1\}$. Then we can calculate $V(1, B)$ for all $B \in \{0, 1/L, \dots, 1\}$. If we knew $V(m - 1, B), \forall B \in \{0, 1/L, \dots, m - 1\}$, then we can plug this in into Formulation 16

and obtain $V(m, B)$ by taking K calculations for $F_m(p_m)p_m + V(m-1, B - F_m(p_m))$ for each $p_m \in \{P_{m,1}, \dots, P_{m,k}\}$ and then choose the minimum, i.e., $2K$ operations in total. To obtain $V(m, B)$ for all possible $B \in \{0, 1/L, \dots, m\}$, we would need to repeat this Lm times, which means the total operations incurred for each m is $2KLm$. Summing this for all $m \in \{1, \dots, n\}$ would require $KLn(n-1)$ operations. Thus the complexity of the algorithm is $O(KLn^2)$. \square

A.4. Descending Clock Auctions using Optimized Price Setting for the Homogeneous Setting

ALGORITHM 2: A DCA Framework using Optimal Price Setting for the Homogeneous Setting

Input: A set of sellers $\mathcal{N} = \{1, \dots, n\}$ with goods $\{G_1, \dots, G_n\}$, an auctioneer with a target T .

A target number of rounds allowed m . Initial valuation estimates v_i .

Output: A set of feasible sellers $\mathcal{A} \subset \mathcal{N}$, i.e., $|\mathcal{A}| = T$, and the corresponding offer price vector p that aims to minimize the expected payment.

1. Let the set of active bidders be $\mathcal{A}^{(r)} = \mathcal{N}$;

for round $r = 1 \dots m$ **do**

2.1. Set the target number of accepting bidders $T^{(r)} = n^{(r)} - (n^{(r)} - T)/(m - r)$ where $n^{(r)} = |\mathcal{A}^{(r)}|$ and solve Model 9 to find a vector of prices p to offer the bidders;

2.2. Find the set of rejected offers \mathcal{R} ;

if $|\mathcal{A}^{(r)} \setminus \mathcal{R}| \geq T$ **then**

2.2.1. $\mathcal{A}^{(r+1)} \leftarrow \mathcal{A}^{(r)} \setminus \mathcal{R}$;

2.2.2. Update the distributions of the bidders' valuations using Formulation 11;

else

2.2.3. Enter the adjustment round in Step 3;

end

end

3. Adjust the prices for bidders in the last round to meet the target by solving Formulation 12;

4. Pay winning bidders the offer prices;

A.5. Descending Clock Auctions using Optimized Price Setting for Incentive Auctions

ALGORITHM 3: A DCA Framework using Optimal Price Setting for Incentive Auctions

Input: A set of stations $\mathcal{N} = \{1, \dots, n\}$, an auctioneer with a feasibility function

$F : 2^{\mathcal{N}} \rightarrow \{0, 1\}$. A target number of rounds allowed m . Initial valuation function estimates v_i .

Output: A set of feasible stations to reject $\mathcal{R} \subset \mathcal{N}$, i.e. $F(\mathcal{R}) = 1$, and the corresponding offer price vector p that aims to minimize the expected payment on the remaining stations.

1. Set the initial prices p at the reserves. Let the set of rejected bidders be $\mathcal{R}^{(r)} = \emptyset$;

for round $r = 1 \dots m$ **do**

2.1. Set the target number of accepting bidders $T^{(r)} = n^{(r)} - (n^{(r)} - T)/(m - r)$ where $n^{(r)} = |\mathcal{A}^{(r)}|$ and solve Model 9 to find a vector of prices p to offer the bidders;

2.2. Find the set of rejected offers \mathcal{R} ;

if $F(\mathcal{R}^{(r)} \cup \mathcal{R}) = 1$, i.e. via solving 17, **then**

2.2.1. $\mathcal{R}^{(r+1)} \leftarrow \mathcal{R}^{(r)} \cup \mathcal{R}$;

2.2.2. Update the distributions of the bidders' valuations using Formulation 11;

else

2.2.3. Enter the final Step 3;

end

end

3. Set all remaining bidders $\mathcal{N} \setminus \mathcal{R}^{(r)}$ as winners and pay them their offer prices;

A.6. Approximation Method for the Multi-Round Case

First, assuming that the auctioneer has only one round left. Then the optimal prices to offer to the bidders will be the solution of Model 9 for the continuous case (or 10 for the discrete case). Now, given that the auctioneer has multiple rounds to do price discovery, he would not offer these ‘optimal prices’ right away. Instead, a set of higher prices will be offered first to learn more about the bidders’ valuations and to update the bounds.

A simple way that the auctioneer can do this is to discretize the prices into m equal intervals between the upper bounds and p^* and offer these to the bidders sequentially until feasibility does not hold.

A better way is to do this dynamically as shown in Algorithm 4. Here, after solving Model 9 (or 10) in Step 1, the auctioneer can offer a guess $p_i = \frac{u_i + (m-1)p_i^*}{m}$ to bidder i and see how the bidder responds. This price is obtained under the expectation that the offer price in the next m rounds will be distributed evenly within the range $[p_i^*, u_i]$. Notice, however, that once the auctioneer has offered the prices to the bidders and received their responses to form the new state of the auction, the auctioneer now has better information and can repeat Step 2.1 of Algorithm 1 to find the new set of offer prices, that is, to run Algorithm 4 again with the updated information. Formally:

ALGORITHM 4: Finding Offer Prices in Round m

Input: Current round r , a current set of active bidders $\mathcal{A}^{(r)}$, most up-to-date valuation estimates v_i .

Output: An offer price vector \mathbf{p} .

1. Solve Model 9 to obtain the optimal offer prices \mathbf{p}^* as if this were the last round;
2. Divide the range $[p_i^*, u_i]$ into m equal intervals and set the actual offer prices

$$p_i = \frac{u_i + (m-1)p_i^*}{m};$$

A.7. Interference Constraints in Repacking and Feasibility Checking

There are two csv data files on engineering constraints available on the FCC web site [FCC 2013]:

- A domain file called “Domain-2013July15.csv”, of size 306KB, that specifies the feasible channels for each station.
- An interference file called “Interference-Paired-2013July15.csv”, of size 6219KB, that specifies the interference constraints that the repacking must meet. This includes:
 - Pairs of (station, station) that must not be assigned to the same channel (among a given list of channels).
 - Pairs of (station, station) that must not be assigned to adjacent channels (among a given list of channels).

The average number of feasible channels that each station can be allocated to is 44.15 (out of 49 channels) with most of the channels being freely allocated to any available channels. However, some stations only have a few feasible channels (that is, there are stations with only three possible channel assignments). There are 2.9×10^6 constraints requiring pairs of stations that are not to be allocated in the same or adjacent channels. Although this is smaller than $2mn^2 = 493 \times 10^6$ in the worst case, i.e., when interference matrices are fully dense, it is still a large number.

Let \mathcal{S} be a set of stations that needs to be repacked into a list of channels in set \mathcal{C} . We use i, j as indices for stations and use k as indices for channels. Let $\mathcal{C}_i \subset \mathcal{C}$, $i \in \mathcal{S}$, be the list of feasible channels to station i . Let \mathcal{I}_c be the list of triplets (i, j, k) such that

stations i and j cannot be assigned to the same channel k . Let \mathcal{I}_a be the list of triplets (i, j, k) such that stations (i, j) cannot be assigned to channel $(k, k + 1)$ respectively. Data for \mathcal{C}_i , \mathcal{I}_c and \mathcal{I}_a are available from the domain file and the interference-paired file on the FCC web site [FCC 2013].

From a given list of channels \mathcal{C} , we say the set \mathcal{S} of stations is feasible with respect to \mathcal{C} if the stations can be packed into the channels without violating any of the constraints. Let $\mathcal{P}(\mathcal{C})$ denote the set of all subsets of stations that can be feasibly packaged into channels in \mathcal{C} .

Let z_{ik} be a binary variable that indicates whether station i is assigned to channel k . We say z is an assignment to the repacking problem. For z to be feasible, we need the following: (a) all the indicator variables z_{ik} are binary, (b) each station is assigned to exactly one channel, and (c) no pairs of stations that might interfere with each other can be assigned to the same or adjacent channels. The set of feasible assignments $\mathcal{P}(\mathcal{C})$ is therefore defined as

$$\mathcal{P}(\mathcal{C}) = \left\{ z : \begin{array}{l} z_{ik} \in \{0, 1\}, \forall i \in \mathcal{S} \text{ and } k \in \mathcal{C}_i, \sum_{k \in \mathcal{C}_i} z_{ik} = 1, \forall i \in \mathcal{S}, \\ z_{ik} + z_{jk} \leq 1, \forall (i, j, k) \in \mathcal{I}_c, z_{ik} + z_{j(k+1)} \leq 1, \forall (i, j, k) \in \mathcal{I}_a \end{array} \right\}. \quad (17)$$

There are a large number—up to 2.9×10^6 —of constraints requiring pairs of stations not to be allocated in the same or adjacent channels. This makes checking the assignment feasibility very challenging for the full problem when all 2177 stations are rejected. In our experiments, however, the largest number of stations being rejected among all the instances tested is less than 1000 and hence CPLEX can still handle the feasibility problem. The feasibility problem does not involve an objective function and hence is much easier to solve than the winner determination problem in a VCG setting.

In the experiments, we used CPLEX to solve the repacking feasibility problem. We could also use a satisfiability (SAT) formulation for this purpose as has been done by Leyton-Brown [2013]. Our choice of CPLEX here was for the convenience of implementation and due to some special network structural properties of the repacking problem that CPLEX could exploit. However, a discussion on the comparison between the performance of SAT and CPLEX is out of the scope of this manuscript since our focus is on the price setting and not on computational method for solving the feasibility problem.