SATS: A Universal Spectrum Auction Test Suite

Michael Weiss
Department of Informatics
University of Zurich
mail@mweiss.ch

Benjamin Lubin
Questrom School of Business
Boston University
blubin@bu.edu

Sven Seuken
Department of Informatics
University of Zurich
seuken@ifi.uzh.ch

ABSTRACT

For the past 17 years, much of the work on combinatorial auctions (CAs) has used the Combinatorial Auction Test Suite (CATS) by Leyton-Brown et al. [24]. However, CATS does not include a good model for spectrum auctions, which have become the most important application of CAs. In this paper, we make four contributions. First, we propose the *Multi-Region Value Model (MRVM)* which captures the difficult to model geographic complementaries of large US and Canadian auctions. Second, we also encode our model as a MIP, making the auction's winner determination problem tractable. Third, we introduce a new *spectrum auction test suite (SATS)*, and release it to the public under an open-source license. SATS includes our new MRVM model, as well as six previously introduced value models from the literature. Fourth, using SATS, we evaluate our MRVM model experimentally: after fitting the model parameters to the bidding data from the 2014 Canadian auction, we show that the MRVM model can represent this auction well.

1. INTRODUCTION

Combinatorial auctions (CAs) allow bidders to submit bids for packages of items. They are one of the must successful applications of market design, with practical uses in many areas. For example, they have been used for the procurement of bus routes [8] and industrial goods [29], as well as for selling TV-ad slots [17]. The most important application of CAs has been to *spectrum auctions* [10].

Since 1994, many governments have used spectrum auctions to sell licenses for radio frequencies (i.e., spectrum). Original designs required bidders to submit bids on individual licenses. By contrast, the *combinatorial clock auction (CCA)* [4], which allows for package bidding, has gained momentum in recent years. Between 2012 and 2014 alone, ten countries have used the CCA, raising approximately \$20 billion in revenues, with the 2014 Canadian 700 MHZ auction being the largest auction, raising more than \$5 billion [3].

1.1 Research on better Auction Designs

Economists, operations researchers, and computer scientists have all contributed to the study of CAs in recent decades. Important contributions have included the study of winner determination algorithms [28], bidding languages [26], payment rules [27], and preference elicitation algorithms [23]. In practice, additional implementation details like reserve prices and activity rules also need to be carefully considered [3]. Together, these aspects of an auction's design determine how well bidders can express their preferences, the extent to which bidders can and need to strategize, the computational complexity of the winner determination and pricing problems, and the auction's efficiency and revenue in equilibrium.

Obtaining better CA designs is difficult for multiple reasons. First, deriving analytic results for CAs is very challenging. Second, insights from small, stylized models often do not translate to practical, real-world problems. Unfortunately, using data from real spectrum auctions also has its limitations. Many governments do not make the bidding data publicly available. Further, even when they do, the data only provides a few individual auction instances, while a researcher typically needs many thousand instances to test the performance of different auction designs.

An alternative approach is to use a *value model*, i.e., an analytic or algorithmic description of a bidder's value function. Given a value model, one can develop an *auction instance generator*, i.e., a piece of software that can produce an arbitrary number of *different* auction instances upon request. When a new instance is requested, the generator then draws some or all *parameters* of the underlying value model from a pre-specified *probability distribution*. Thus, each use of the generator generally produces a different instance, enabling the creation of a *distribution* of auction instances.[1]

Many researchers have used value models/auction generators in their work, for example to evaluate the run-time of core pricing algorithms [7, 14], to evaluate different reserve prices [13], or to evaluate the performance of different payment rules [25].

1.2 CATS and other Spectrum Value Models

For the past 17 years, the majority of this simulation-based work has used the *Combinatorial Auction Test Suite (CATS)* by Leyton-Brown et al. [24], making it the de-facto standard for running CA experiments. CATS contains five main CA value models, and corresponding software for realizing them as auction instance generators. While CATS has provided immense value to the research community, it has two main shortcomings. First, the value models included in CATS cannot model spectrum auctions particularly well.[2] This severely limits the utility of CATS in research on spectrum auctions. Second, CATS does not have a number of features that are often needed in CA simulations (e.g., providing a value for a specific bundle); and its C-based code is hard to extend.

Consequently, researchers have developed new value models specifically for spectrum auctions. This includes models capturing the separation of licenses into different bands [6], as well as mod-

Appears in: *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.*

[1] Note that it is essential for some model parameters to be drawn from a distribution. Without this, each program execution would produce the same auction instance.

[2] Already its authors wrote: "Clearly the problem of realistic test data for spectrum auctions remains an area for future work." [24]

Value Models		Model Description	Example/Motivating Auctions
Stylized Models	1. Base Value Model [6]	Two bands: In band A bidders have very high intra-band synergies for having up to 4 blocks; in band B values are additive. Single-region market with one bidder type.	e.g., European 2.6 GHz auctions
	2. Multi-Band Value Model [6]	Similar to the Base Value Model, but different bands and synergies: 24 licenses divided into 4 bands containing 6 blocks each. One band has higher value.	e.g., multi-band auctions in the UK, Austria, etc.
	3. Global Synergy Value Model [16]	Abstract geographic model: regions arranged on two circles; adjacency modeled as distances on these circles. Non-linear complementarities for packages of adjacent items. Two bidder types: national and regional. No model of bands or frequency.	No specific market
	4. Local Synergy Value Model [30]	Similar to the Global Synergy VM, but the regions are arranged on a grid and adjacency is modeled by Manhattan distance.	No specific market
	5. CATS Regions Generator [24]	Regions are connected on a planar graph; complementarities for adjacent bundles. Bidders have both a common and a private value for each region.	No specific market
Realistic Models	6. Single-Region Value Model [22]	Has three bidder types with distinct preference profiles over bands which capture both inter- and intra- band synergies. Does not model geography/regions.	UK 4G Auction (2013)
	7. Multi-Region Value Model	Highly parameterizable model of both geography and frequency dimensions, with multiple regions and bands. Three bidder types: local, regional and national.	Canadian 4G Auction (2014)

Table 1: Overview of value models from the literature, and our new Multi-Region Value Model (MRVM)

els capturing geographic division, where individual licenses cover only part of a country [16, 30]. However, these models only focus on individual aspects of the spectrum auction design problem, and most of them are highly stylized. It is not surprising that none of the prior models has captured the separation of licenses into different bands *as well as* geographic division, because modeling the interaction between band division and geographic separation is highly complex, and especially so when the goal is to create realistic (as opposed to stylized) auction instances.

1.3 Overview of Contributions

In this paper, our goal is to develop a new value model and software suite for research on combinatorial and spectrum auctions. To this end, we make four main contributions:

1. We present the *Multi-Region Value Model (MRVM)*. This model is specifically designed to capture the complex ways in which *multiple frequency bands* as well as *geographic complementarities* determine the bidders' values in large real-world auctions such as in the US and Canada.

2. We also present a concise mixed integer programming (MIP) formulation for solving the winner determination (WD) problem (i.e., finding the efficient allocation) for our MRVM model. The MIP formulation encodes the values for *all bundles* in a succinct way by only using the model parameters. This lets us find the efficient allocation or compute VCG payments without exponential bundle enumeration. Using our MIP formulation, the median run-time for solving the WD problem for instances of the size of the 2014 Canadian auction (10 bidders, 98 licenses, 14 regions) is 12 seconds.

3. We introduce a new *Spectrum Auction Test Suite (SATS)* and release it to the public (under an open-source license) at: `www.spectrumauctions.org`. SATS includes our own MRVM model as well as six previously introduced value models. We describe how researchers can use SATS via three different interfaces: a simple web interface, a command-line tool, and a Java API that offers full access to the software.

4. We evaluate the MRVM model experimentally. We first fit the model parameters to real-world data from the 2014 Canadian auction. We then use SATS to generate 1,000 different MRVM auction instances and compare various statistics of our simulated auctions with the Canadian auction. The results show a close match between the statistics, demonstrating that our model can capture the Canadian auction well.

2. PRELIMINARIES

We consider a spectrum auction of the type conducted by many governments. A range of spectrum frequencies is called a *band* (e.g., in the 2014 Canadian auction, the "Lower 700 Mhz band" consisted of the frequencies between 698 MHz and 746 MHz). A band is divided into multiple *blocks* (e.g., the D-block in the Canadian auction, ranging from 716-722 MHZ). The government sells *licenses* for each block to mobile network operators (MNOs). In most smaller countries (which we call *single-region markets*), a license gives the MNO the right to use the corresponding block in the whole country. In large countries (which we call *multi-region markets*), the government sub-divides the country into *regions*, and with each license, the MNO only obtains the right to use the corresponding block in a particular region (e.g., in the 2014 Canadian auction, separate licenses were sold in each of 14 regions [11, 20]).

Generally, licenses in lower frequency bands (between 700 and 900 MHz) are more valuable than licenses from bands covering higher frequencies, due to their physical properties (low frequencies carry signals further and have a better ability to penetrate buildings). Obtaining multiple licenses for blocks in the same band can induce *intra-band synergies*. For example, four 5-MHz blocks in the 2.6 GHz band enable peak performance using LTE [6].

A block can be *paired* (i.e., two blocks are sold as a pair: one for uplinking, one for downlinking) or *unpaired*. Due to the technology used by today's smartphones, paired licenses are significantly more valuable than unpaired licenses. Sometimes, blocks within a band are assumed to be *generic*, i.e., to have roughly the same value. This allows a more compact bidding language, because bidders can then place their bids, not for specific licenses, but for a quantity of licenses of a particular *generic band* (see [9] for details).

In multi-region markets, the MNO's value function becomes more complex to model and depends on the type of the buyer. *National* MNOs have large *synergies* if they acquire licenses for all or almost all regions of a country, but their value may fall drastically if they have large holes in their coverage area. In contrast, *regional* or *local* MNOs only have value for specific regions of the country.

3. PREVIOUS VALUE MODELS

Table 1 provides an overview of seven value models for spectrum auctions (six that have been introduced in the literature, and our new MRVM model). For each model, we provide the citation to the original paper, a brief description of the model, and an example/motivating real-world auction that is well captured by the model (if any). All seven value models are included in SATS.

Models 1-5 are relatively stylized, often focusing on a particular aspect of spectrum auctions. They have mostly been used in lab experiments to study bidding behavior of human players, which explains why they do not aim to capture the full complexity of real-world auctions. In contrast, models 6 and 7 are quite realistic, capturing specific real-world auctions. As a researcher, one might use a stylized model if one is only interested in studying a specific aspect of an auction/setting. If one is striving for external validity, then one might instead want to use one of the realistic models.

4. MULTI-REGION VALUE MODEL

In this section, we present our *Multi-Region Value Model (MRVM)* – a new value model that is significantly more realistic than any of the previously introduced models. The MRVM model is motivated by the 2014 Canadian 700MHz auction [20]. Note that our modeling goal was not to capture many different real-world auctions in one model, but rather to capture a single important auction well. Nevertheless, the model also captures key features of other large-scale auctions with geographical division as conducted in the US (including the Forward phase of the Incentive Auction [15]) and Australia [1].[3]

We first conducted an in-depth study of the physics and engineering behind the technology deployed to use the spectrum (e.g., [2] and [5]), identifying a long list of influences on spectrum value, called *value drivers*. Using this domain knowledge, we then modeled the value drivers essential for fidelity, while also targeting understandability and tractability of the model. The resulting model is designed to be rich enough to capture the complexities of large real-world spectrum auctions (like the Canadian one), while at the same time being succinct enough to be formalized as a concise MIP.

At a high level, for large countries like the US and Canada, *geographical division* and *frequency band division* emerged as the most essential value drivers. We aggregated multiple low-level value drivers into a single modeling parameter whenever this was warranted, to keep complexity manageable.[4] We now first build some intuition for how the bidders' value functions are constructed in the MVRM model in three steps (also see Figure 1), before providing the formal definition of the model in the next section.

- **Step 1: Bands/Licenses → Bandwidth:** We model different bands with different valuations, and bands have *intra-band synergies*. The number and kind of licenses an MNO acquires in each region determines his *total bandwidth*.
- **Step 2: Bandwidth → Regional Value:** We transform the total bandwidth the MNO has acquired into a regional value. The more spectrum a bidder buys in a particular region (each of which has a particular number of subscribers), the higher the quality of service (QoS) will be, and thus the higher the value per subscriber. A bidder's monetary value for a region is calculated as a function dependent on his bandwidth, the share of the population he intends to serve, the size of the population of that region, the bandwidth required to provide a good QoS, as well as his maximum value per subscriber.
- **Step 3: Regional Discounting:** The bidder's final value for the whole bundle is calculated as the sum over his regional values, discounted by a bidder-specific term (that depends on which exact regions the bidder covers). National players have large synergies for covering the whole country; regional players have synergies between regions close to their headquarters; local players have no synergies between regions.

[3] We currently include the region map (adjacency graph) for the Canadian regions, and users can also supply their own maps. In the future, we plan to include a synthetic map generator in SATS.

[4] For example, while in practice, multiple factors influence the quality of a band, we capture this in a single parameter.

4.1 Formal Definition of MRVM

We now provide the mathematical definition of the MRVM model. Here, we provide the most natural (but sometimes non-linear) description. For the MIP formulation (provided in Appendix A) we spent considerable effort to linearize all functions. Note that five of our model parameters are not fixed, but are assumed to be drawn from non-degenerate distributions. Whenever we introduce one of those parameters, we denote this by †. See Appendix B for the distributions we used.

World Setup.

We first define the *world*, i.e., everything independent of bidders.

- R denotes a set of *regions*. For each $r \in R$, let $p_r \in \mathbb{N}$ denote the *population* of that region. The elements of R are embedded in a planar *adjacency graph*, where two regions (nodes) are connected if they share a border.
- B denotes a set of *bands*. For each $b \in B$ let $n_b \in \mathbb{N}$ denote the number of *blocks* in b. We let $c_b \in \mathbb{R}_+$ denote the *base capacity* of b.† c_b captures the amount of *bandwidth* an MNO obtains when buying one license (for one block) in band b. This allows us to encode different values for different bands, e.g., due to their frequency, the number of MHz a block covers, and whether the blocks are paired/unpaired.
- A tuple $l = (r, b, j) \in R \times B \times \{1, \ldots, n_b\}$, is called a *license* and L is the set of all licenses. For all licenses $l = (r, b, j)$, let $r(l)$ denote r and $b(l)$ denote b. We use $x \subseteq L$ to denote a *bundle*, i.e., a set of licenses.
- For all bands $b \in B$, the *intra-band synergies* are captured by a function $\mathrm{syn}_b : \mathbb{N} \mapsto \mathbb{R}_{\geq 1}$, where $\mathrm{syn}_b(n)$ denotes the synergy for having n licenses of b in the same region.[5]

† This parameter is assumed to be drawn from a non-degenerate distribution. See Appendix B for the distribution we use.

[5] $\mathrm{syn}_b(n)$ has to be defined in a way that ensures free disposal of licenses. Furthermore, $\mathrm{syn}_b(1)$ should be set to 1 for consistency.

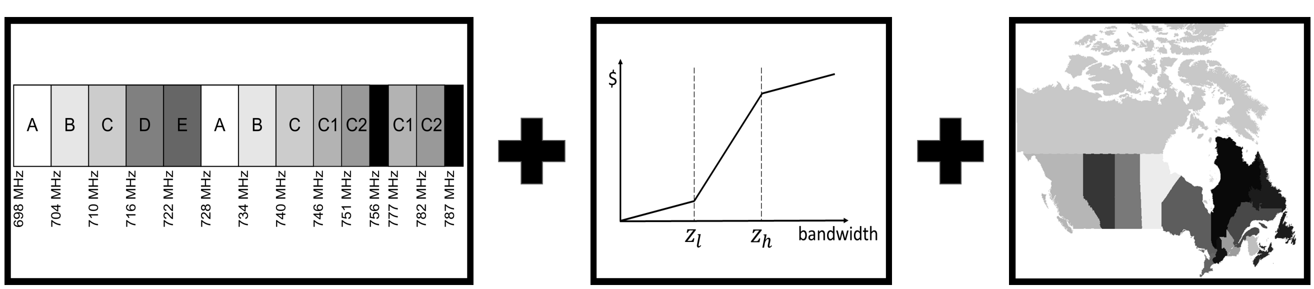

(a) (b) (c)

Figure 1: Flow chart for the Multi-Region Value Model (MRVM). Sources for (a) and (c): [18]

Bidder-specific Parameters.

We let $\mathcal{N}$ denote the set of bidders. Each bidder $i \in \mathcal{N}$ has the following parameters:

- A *bidder type* $t_i \in \{local, regional, national\}$, which is used to differentiate bidders by their business plan.
- A *maximum value* per subscriber $\alpha_i \in \mathbb{R}_{\geq 0}$. α_i specifies the relative bidder strength, distinguishing stronger from weaker (otherwise identical) bidders.†
- A *market share* $\beta_{i,r} \in [0,1]$, i.e., the share of the population that bidder i serves in region r.†
- *Quality function parameters* $z^l_{i,r}$ and $z^h_{i,r}$, indicating the amount of bandwidth below (respectively above) which per-subscriber quality of service is low (respectively high) in region r.

Overall Value Function.

Each bidder $i \in \mathcal{N}$ has a *value function* $v_i : \mathcal{P}(L) \to \mathbb{R}_{\geq 0}$, defining the bidder's value for any bundle $x \subseteq L$ as follows:

$$v_i(x) = \sum_{r \in R} \underbrace{\beta_{i,r} \cdot p_r \cdot \mathrm{sv}_{i,r}\big(\alpha_i, \beta_{i,r}, p_r, \underbrace{\mathrm{c}(r,x)}_{\text{(1) Bandwidth}}\big)}_{\text{(2) Monetary Value per Region}} \cdot \underbrace{\Gamma(i,r,x)}_{\text{(3) Regional Discounting}}$$

As described in the introduction to this section, the value function is constructed in three steps, which we describe next in detail.

Step 1: Bands/Licenses → Bandwidth

We first define an auxiliary function $\mathrm{c}(r,x) : R \times \mathcal{P}(L) \to [0,1]$, which calculates the *bandwidth* in region r which a bidder obtains when purchasing bundle x

$$\mathrm{c}(r,x) = \sum_{b \in B} \mathrm{cap}(b,r,x) \tag{1}$$

where $\mathrm{cap} : B \times R \times \mathcal{P}(L) \to \mathbb{R}_{\geq 0}$ is a function that obtains the capacity for band b in region r for bundle x, defined as:

$$\mathrm{cap}(b,r,x) = c_b \cdot |x_{b,r}| \cdot \mathrm{syn}_b(|x_{b,r}|) \tag{2}$$

where c_b is the base capacity, $|x_{b,r}| = |\{l \in x | r(l) = r \wedge b(l) = b\}|$, and $\mathrm{syn}_b(n)$ is the intra-band synergy for having n licenses in band b.

Step 2: Bandwidth → Regional Value

Next, we transform the total bandwidth held by a bidder in a particular region, $\mathrm{c}(r,x)$, into a dollar-denominated *monetary value per region* (i.e. component (2) of the overall value function) according to:

$$\beta_{i,r} \cdot p_r \cdot \mathrm{sv}_{i,r}\big(\alpha_i, \beta_{i,r}, p_r, \mathrm{c}(r,x)\big) \tag{3}$$

The first two factors are straightforward: The product of the market share $\beta_{i,r}$ and the population in the region p_r is the number of subscribers in region r. This is then multiplied by the value the bidder (MNO) can obtain from a single subscriber to obtain an overall value for the region. We denote this latter quantity the *subscriber value*, $\mathrm{sv}_{i,r}(\cdot)$, which is itself a non-linear function that accounts for the quality of service that can be provided to a subscriber in the region. We illustrate such a function in Figure 1 (b). When the available bandwidth, $\mathrm{c}(r,x)$, is small (for the population being served), then an individual subscriber will obtain poor quality of service. Consequently, the subscriber will have low value, which, when monetized by the MNO, will likewise yield a low monetary value for the MNO. By contrast, if the MNO obtains a large amount of bandwidth then the quality a subscriber obtains will be high, her value will be high, and thus the monetized value to the MNO will be high; as additional bandwidth is obtained past this point, though, the value only rises slowly. The thresholds in this sigmoid shape are governed by the constants $z^l_{i,r}$ and $z^h_{i,r}$. And the maximum dollar value per subscriber that can be obtained is controlled by the constant α_i. Formally, we define the function $\mathrm{sv}_{i,r}(\,\cdot\,, \mathrm{c}(r,x))$ to be piecewise linear with the following control points:

$$\begin{aligned} &(0,\ 0) \\ &(z^l_{i,r} \cdot p_r \cdot \beta_{i,r},\ 0.27 \cdot \alpha_i) \\ &(z^h_{i,r} \cdot p_r \cdot \beta_{i,r},\ 0.73 \cdot \alpha_i) \\ &(\mathrm{c}(r,L),\ \alpha_i) \end{aligned} \tag{4}$$

Step 3: Regional Discounting

In the last step, the bidder's total value is calculated as the sum of his regional values when discounted (by Γ) according to his type:

- Every **local bidder** $i \in \mathcal{N}$ has a set $I_i \subseteq R$ of *regions of interest*.† A local bidder has zero value for all licenses outside of I_i and full value otherwise.
- Every **regional bidder** $i \in \mathcal{N}$ has a *headquarters* $h_i \in R$.† Furthermore, the bidder has a discount factor $\lambda_i \in (0,1]$. This bidder type models bidders that are particularly strong in one specific region but also plan to provide service in nearby regions.
- Every **national bidder** $i \in \mathcal{N}$ has discount factors $\gamma_{i,k} \in [0,1]\ \ \forall k \in \{0,\ldots,k_{\max}\}$, where k indicates the number of regions for which bidder i does not have any license, which we cap at a constant $k_{\max} \in \mathbb{N}$ (i.e., we winsorize the number of missing regions to $k_{\max}$). Such bidders attempt to cover every region, and incur an increasing utility loss for every region (up to $k_{\max}$) in which they do not provide coverage.

Formally, we have

$$\Gamma(i,r,x) = \begin{cases} \mathbf{1}_{r \in I_i}, & \text{if } t_i = \textit{local}, \\ \lambda_i^{\mathrm{distance}(h_i,r)}, & \text{if } t_i = \textit{regional}, \\ \gamma_{i,k(x)}, & \text{if } t_i = \textit{national}, \end{cases} \tag{5}$$

where $\mathrm{distance}(r_1, r_2)$ is the shortest path distance of $r_1, r_2 \in R$ in the adjacency graph of the regions, $k(x) = \min\{k_{\max}, |\overline{R_x}|\}$, and where $R_x \subseteq R$ is the set of regions for which at least one license is in x; thus, $|\overline{R_x}|$ is the number of regions without a license.

4.2 Discussion & Limitations

MRVM includes many value drivers, the most important include: (1) geographical division, (2) frequency band division, (3) intra-band synergies, (4) higher values for licenses in high density regions, (5) the match between a bidder's targeted market share and obtained bandwidth, (6) the distinction between local, regional and national bidders in terms of their different preferences for different regions (i.e., geographical division). To the best of our knowledge, there does not exist another spectrum value model for multi-region markets of similar accurateness.

However, our model also has its limitations. Most importantly, it currently does not include a notion of *spectrum endowment* for the MNOs; thus, it could not model well a situation where an MNO wants to "fill up" holes in a frequency band for which it already has some licenses. Second, the model does not capture *inter-band synergies*, i.e., combinatorial effects on value between different bands (e.g., an MNO may want licenses for low-frequency and high-frequency bands simultaneously). In principle, the model could be extended to include both of these aspects.

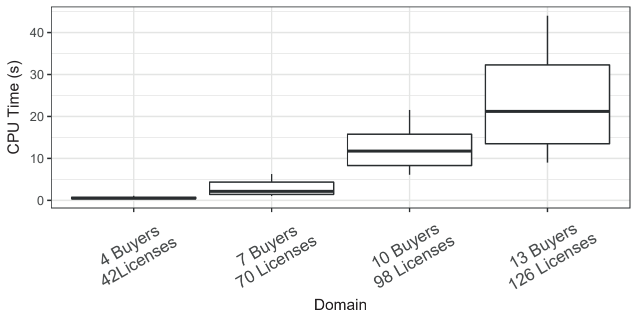

Figure 2: Run-time results for finding the efficient allocation using our MIP formulation for differently-sized problems. The box plot provides the first, second and third quartiles; the whiskers are at the 10th and 90th percentile.

5. MIP FORMULATION

For the formulation of the MRVM model described in the last section, we presented the mathematically most natural (and sometimes non-linear) formulation. However, when designing the model, we were careful to make sure that it can also be represented as a concise mixed integer program (MIP), which we provide in Appendix A; this MIP is also included in the SATS software suite. *But why is this MIP formulation so essential?*

Researchers conducting auction simulations typically need to compute the social-welfare maximizing allocation of an auction; for example, to evaluate the efficiency of different mechansims, or to compute VCG prices. Standard algorithms for solving CAs use bidding languages like XOR or OR*, but bidders are restricted to a few bids (e.g., 500). However, computing the true efficient allocation requires access to the bidders' *full* value functions, and not just a sample of XOR bids. Reporting a bidder's *full* value function in the MRVM/Canadian model would require 2^{98} XOR bids – obviously, even generating that many XOR bids is infeasible. By contrast, our MIP formulation encodes the values for *all bundles* in a succinct way by only using the model parameters (i.e., there is a one-to-one mapping between the model's value function and the MIP objective). This lets us find the efficient allocation without exponential bundle enumeration, something that has been impossible (and was therefore often glossed) in existing work based on CATS.

To evaluate our MIP formulation in terms of computational tractability, we conducted run-time experiments with varying problem sizes. The experiments were run on a PC with a 1.8 GHZ Intel Core i7 CPU and 8 GB of RAM. All mathematical programs were solved using CPLEX 12.6. Figure 2 shows the results. For each of the four domains (listed on the x-axis) we generated 500 auction instances and computed the efficient allocation. On the y-axis, we show the CPU time it took to compute the efficient allocation. The third domain from the left (10 buyers, 98 licenses) corresponds exactly to the 2014 Canadian auction. For the other three domains, we scaled the number of bidders and licenses up/down, to create more/less difficult domains. Clearly, run-time increases with problem size. However, even for large problem sizes, the WD problem can be solved quickly. Specifically, the median run-time to solve auction instances of the same size as the 2014 Canadian auction (10 bidders, 98 licenses) was 12 seconds. We emphasize that this domain has 2^{98} bundles, and our MIP formulation is finding the *fully* efficient allocation by implicitly representing this *entire* space.

6. THE SATS SOFTWARE SUITE

In this section, we describe the *SATS software suite*, which we released to the public in May 2017 under an open-source license at `www.spectrumauctions.org`. SATS includes all value models listed in Table 1: our new MRVM model as well as six models previously introduced in the literature (i.e., the Base and Multi-Band Value Models, the Global and Local Synergy Value Models, CATS Regions, and the Single-Region Value Model).

6.1 Accessing/Using SATS

SATS can be accessed in three different ways, to ensure it can be used for very basic and very sophisticated use cases.

A Simple to Use Web Interface. Via the web interface, users can quickly access the most commonly-used features of SATS. Using the website (see Figure 3) is straightforward: the desired value model can be selected, a few parameters can be modified, then the model is run on the server, and finally the resulting value files can be downloaded in the browser. This simple access method is meant for students, or for researchers who want to get a quick first impression about a value model.

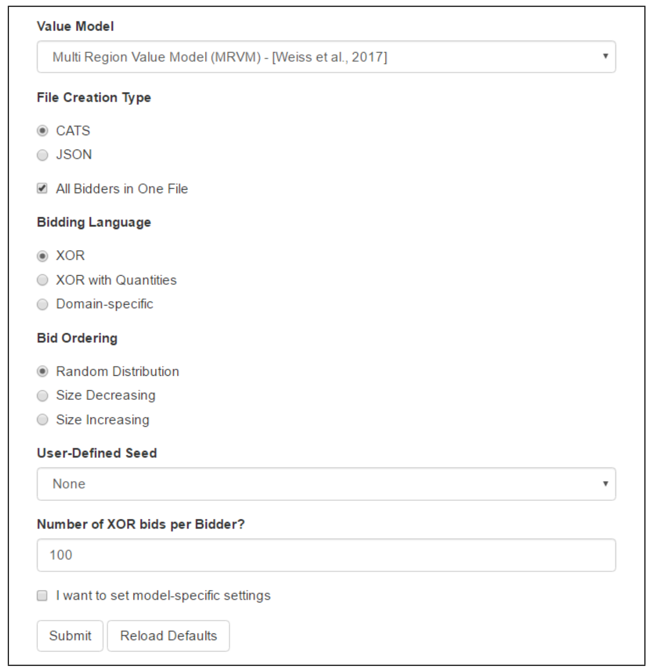

Figure 3: Part of the web interface to generate value files

The Command-Line Tool. The SATS *command-line tool* provides an efficient way to generate a large number of value functions, without writing any code. It is designed to be similar in its use to CATS. Listing 1 shows an example command, illustrating the creation of valuation data using the MRVM's default configuration, with 60 bids per bidder.

```
java -jar sats.jar --model MRVM --bidsPerBidder 60
```

Listing 1: Simple Example of SATS on the command line

The Java API. SATS is written in Java, and by directly accessing its API, one gains access to its full feature set, including the winner determination solver, value queries for user-specified bundles, and the ability to specify all model parameters and distributions.[6] The classes have been designed to enable simple use and easy extension. Listing 2 illustrates this, showing the snippet of code used to

[6]For example, a user could set the number of bands in the MRVM to 1, if he only wanted to study geographic division.

```
// Choose a value model with default settings
DefaultModel<?,?> model = new MultiRegionModel();
// Create a new set of bidders of this model
for(Bidder<?> bidder : model.createNewPopulation()){
        // Choose a way to access the value function
        // We choose XOR-Bids, ordered by bundle size
        XORLanguage<?> xorBids =
            bidder.getValueFunction(SizeOrderedXOR.class);
        // Pass the generated values to your code, e.g.,
        yourSimulator.addXorBids(xorBids.iterator());
}
```

Listing 2: Simple Example of the SATS Java API

create a new set of value functions. The code is easily modified: as shown, the code will produce MRVM instances. Simply replacing *new MultiRegionModel()* with *new BaseValueModel()* causes BVM model instances to be created instead.

SATS was designed for both basic and sophisticated usage. Because we support full API access, new ways of doing large-scale simulations are now possible, as we detail in the next section.

6.2 The Features of SATS

SATS provides a number of useful features:

1. SATS supports two output file formats. First, it supports the standard CATS format, enabling drop-in use of its output in experiments that are already coded for CATS. Second, SATS supports a simple JSON-based native file format, which is easily read by existing libraries, and enables human introspection for debugging purposes.[7]
2. A user can select whether to receive bids in *XOR*, *XOR-Quantity*, or *Domain-Specific* bidding languages. The XOR-Quantity language matches the bids used in most CCA auctions, where bidders place bids not for specific licenses but for a quantity of licenses in a particular band. This increases the number of bids that can be expressed concisely. The domain-specific language directly outputs the drawn model parameters, allowing for the exact specification of the full value function. For models, such as MRVM, that can be encoded into the objective of a winner determination formulation, the domain-specific language concisely provides all the information needed for *fully expressive bidding*.
3. A user can choose in which order XOR or XOR-Quantity bids are generated by the system, including options for (1) random, (2) size-increasing, and (3) size-decreasing.
4. The value models encoded in SATS specify the bidders' *complete preferences*, and users of the suite can, for any given instance, request the value for any given bundle. In multi-round auction simulations, the value for bundles can also be requested "on-the-fly," i.e., during the process of the auction, rather than generating all bids beforehand.[8]
5. The API also supports even more sophisticated use cases, where a user can specify a *world* (domain), and then, during a simulation, can dynamically request new value functions that are consistent with this world. This is essential for applications such as the search for equilibria in large auctions when payoffs are only available as outputs from a simulation.
6. SATS also includes MIP formulations of the winner determination problem for the two realistic value models, SRVM and MRVM.[9] This allows SATS users to compute the efficient allocation over the full value space even for very large SRVM and MRVM auction instances using a MIP solver.[10] This is impossible when using value generators like CATS which use non-compact bidding languages like XOR.

[7]This format is required for the domain-specific bidding languages which cannot be represented in the CATS file format.

[8]Simple XOR file formats can only represent a small fraction of all bundles. Thus, to access all bundles, use of either the domain-specific bidding language or the API is required.

7. EVALUATION

In this section, we show that our MRVM model is capable of generating bidding data that matches large-scale real-world auctions.

Background and Setup. To perform our analysis, we must first fit our model to real-world data. Recall that our generator is *randomized* such that even after being fitted, each run of the generator produces a new distinct auction instance (we obtain this randomization by drawing certain model parameters from distributions). Therefore, we do not seek an exact match to the one Canadian data point – in particular, having a generator that produces *one data point* perfectly is not very useful. Instead, we seek a generator that produces a realistic *distribution* of instances consistent with the Canadian data. As a consequence of this, there will not be a one-to-one match between the bidders in the benchmark data and those in the generated instances – even though we use the same *number of bidders*. As a comparison at the level of individual bidders would thus be ill-defined, we instead perform our evaluation at the level of *bidder types* (local, regional, national). Specifically, we compare summary statistics over each bidder type as observed in the benchmark with the results of our generator by tracking the following metrics:

- **Licenses:** Number of licenses won by each bidder type.
- **Regions:** Regions in which at least one license was won.
- **Bid per MHz-Pop:** Winning bid divided by the MHz bought and the population served. Division by MHz-Pop is a commonly-used way to control for the amount of people a license covers and the amount of bandwidth it provides.

Given this, we can now define the goal of the fitting procedure: we search for model parameters such that the *means of the distributions over the three metrics* produced by our auction generator match the corresponding statistics of the benchmark as well as possible. The result is a (probabilistic) model constructed such that, *on average*, the generated instances match the benchmark, while possessing a useful variance for experimentation purposes.[11]

Note that fitting the model parameters to the data is a highly non-linear problem, and thus we cannot use global search techniques which are otherwise standard to fit simpler models to data. Instead, we use a human-guided hill-climbing algorithm to greedily improve each of the parameters in turn to minimize the distance between the generated instances and the real-world data. More precisely, we search for model parameters such that the mean-squared-error (MSE) between the means of the distributions of the generated instances and the benchmark data is minimized for all three metrics.

Experimental Data. To conduct these experiments it is important to choose the right benchmark. We choose the 2014 Canadian 700

[9]The MIP formulations for the other models will be added soon.

[10]SATS currently supports the CPLEX solver, but adapters for other solvers could straightforwardly be implemented.

[11]Note that an alternative fitting procedure is also possible. We could fix all model parameters deterministically (or equivalently, set all variances to zero) to then exactly fit the model to the one observed Canadian auction instance. The generator would then produce only a single instance, which would be its best match to the benchmark data. One could then analyze the resulting model for possible further insights into the benchmark auction. However, such an econometric analysis is not the goal of this paper and outside the scope of this work.

	Type	SATS	Canadian
Licenses	Local	0.355 (0.013)	0.333 (0.333)
	Regional	4.707 (0.047)	3.250 (1.436)
	National	26.035 (0.247)	27.667 (2.848)
Regions	Local	0.355 (0.013)	0.333 (0.137)
	Regional	4.132 (0.047)	3.250 (1.436)
	National	13.980 (0.005)	13.667 (0.333)
Bid/MHz-Pop	Local	0.142 (0.005)	0.137 (0.137)
	Regional	1.925 (0.016)	1.968 (0.466)
	National	4.162 (0.031)	4.163 (0.832)

Table 2: SATS and Canadian Auction benchmark data for each of our metrics, showing mean and standard error

MHz auction because (1) the auction was of a substantial scale with 10 bidders, 14 regions, and 98 licenses; (2) the auction was complex with good diversity in geography, frequency, and bidder types; (3) the full set of bidding data is publicly available; and (4) the auction used a CCA auction format with a distinct supplemental round [3].

In the CCA, the supplemental round consists of a sealed-bid CA with a payment rule that, while not strategyproof, is chosen because it is thought to induce "very good incentives."[12] We take advantage of this by adopting the reported bids as statements of value and use them to fit our model. Because there is no obvious methodology for correcting for strategic play in any of the real-world CA designs (and because such an econometric exercise would in any case be beyond the scope of this paper), we adopt this approach as a second-best approximation. To mitigate the effects of this approximation as much as possible, we only compare the *efficient allocation* computed at the real-world bids, with those computed at the values our model produces; we do not implement or compare the final payments under the CCA mechanism.

The Canadian auction, like many similar real-world auctions, included side constraints (called *aggregation limits*) in its winner determination algorithm to promote industry competition. Specifically, these limits ensured that no bidder won more than two of the five paired blocks. As a further allocation detail, any bundle that included two paired blocks from the same region had to include at least one A block [11]. For realism, we also included similar constraints in our evaluation, corresponding to our implementation of the Canadian auction.[13]

Results. Table 2 shows a comparison of 1,000 instances generated by the SATS MRVM model, with parameters chosen so as to be consistent with the Canadian 700 MHz auction. Comparing the columns, it is clear that our generator is able to match the three statistics for all three bidder types extremely well.

Limitations. We have focused our evaluation on the Canadian 700 MHz auction, because it is the only auction of this complexity (with multiple frequency bands and geographic division) where full bidding data is available and where the auction format was conducive to obtaining useful bid data (CCA-style combinatorial auction). Thus, we could not validate our model with more auction instances, which we leave to future work should additional datasets become available. That said, we note that the Canadian auction setting is broadly similar to several auctions in the US and Australia.

[12] Cramton [10] argued that the CCA "induces truthful bidding", while Day and Raghavan [14] showed that MRC-selecting payment rules (as used in the supplemental round) minimize the bidders' total potential gains from strategic manipulation. Day and Milgrom [12] have argued that, if finding a beneficial deviation from truthful bidding is very hard, then many bidders may just report truthfully.

[13] Note that we modeled the Canadian auction as having three separate bands, with 3/2/2 blocks each (see Appendix B.1).

8. CONCLUSION

Over the last 15 years, many researchers have used computational experiments to evaluate different designs for CAs in general, and spectrum auctions in particular. Having good spectrum auction value models, as well as software tools, is essential to conduct these kinds of experiments.

In this paper, we have made four main contributions. First, we have introduced the MRVM model, a realistic model capturing frequency division as well as geographic complementarities. Second, we have encoded our model as a concise MIP, making large-scale auction simulations tractable. Third, we have introduced SATS, a new spectrum auction test suite, and released it to the public under an open-source license at `www.spectrumauctions.org`. As of now, SATS includes our new MRVM model as well as six previously introduced models. Fourth, we have evaluated the MRVM model experimentally, showing that it is rich enough to match the 2014 Canadian auction data well on various metrics.

We are planning to successively extend SATS by adding new features, incorporating feedback from the research community. When a community adopts a set of benchmark problems, a number of positive externalities can be realized, enabling significant advances [21]. We believe that SATS can serve in this role, in a manner similar to how CATS has served the community for the past 15 years, enabling research that would not otherwise have been possible. As a first example of this, in Spring 2017, SATS is expected to be used as part of a CS course at Brown University in which students are building autonomous bidding agents for a CA.[14] In the future, we hope to expand upon this experience in order to use SATS as the basis for broader competitions, similar to the trading agent competition (TAC) that has been run for over 15 years.

APPENDIX

A. MRVM WINNER DETERMINATION

The following MIP is implemented in SATS to solve for the efficient allocation of a set of bidders with MRVM valuations *exactly*, using a MIP Solver (e.g., CPLEX):

$$\text{argmax} \quad \sum_{i\in\mathcal{N}_L} v_L(i) + \sum_{i\in\mathcal{N}_R} v_R(i) + \sum_{i\in\mathcal{N}_N} v_N(i) \qquad (6)$$

$$\text{s.t.} \quad \sum_{i\in\mathcal{N}} X_{i,r,b} \leq n_b \; \forall\, r \in R, b \in B \qquad (7)$$

Here $X_{i,r,b} \in \{1, ..., n_b\}$ is the primary decision variable which denotes the number of licenses assigned to bidder i in region $r \in R$ from band $b \in B$, $\mathcal{N}_L \subseteq \mathcal{N}$ is the set of local bidders, $\mathcal{N}_R \subseteq \mathcal{N}$ the set of regional bidders and $\mathcal{N}_N \subseteq \mathcal{N}$ the set of national bidders, $\mathcal{N} = \mathcal{N}_L \cup \mathcal{N}_R \cup \mathcal{N}_N$, and $v_L(i)$, $v_R(i)$ and $v_N(i)$ are variables holding the value of local, regional or national bidders respectively, and defined as follows:

$$v_L(i) = \sum_{r\in I_i} \Omega_{i,r} \quad \forall i \in \mathcal{N}_L \qquad (8)$$

$$v_R(i) = \sum_{r\in R} d_{i,r} \cdot \Omega_{i,r} \quad \forall i \in \mathcal{N}_R \qquad (9)$$

$$v_N(i) = \sum_{k\in\{0,...,\bar{k}\}} \gamma_{i,k}\Psi_{i,k} \quad \forall i \in \mathcal{N}_N \qquad (10)$$

where $\Omega_{i,r} \in \mathbb{R}_{\geq 0}$ is the undiscounted regional value for bidder i in region r defined below, $I_i \subseteq R$ are the regions of interest for local bidder i, $d_{i,r} = \lambda_i^{\text{distance}(h_i,r)}$ is a discounting constant associated with the distance between region r and bidder i's headquarters h_i.

[14] CS1951k (Alg. Game Theory) taught by Amy Greenwald.

A national bidder i's value is parameterized in terms of the number of regions k (capped at $k_{\max}$) not covered by any license in the bidder's bundle. It is specified as a sum over a discount factor for missing exactly k regions, $\gamma_{i,k}$, multiplied by the undiscounted value for a bundle missing exactly k regions, $\Psi_{i,k}$. Importantly, $\Psi_{i,k}$ is defined to be 0 whenever the current bundle does not have exactly k missing regions.

Undiscounted Regional Values.
For all bidders $i \in \mathcal{N}$, the undiscounted value for a given region r in the current bundle is defined as:

$$\Omega_{i,r} = \beta_{i,r} \cdot p_r \cdot \mathrm{sv}_{i,r}(c_{i,r}) \tag{11}$$

where: $\beta_{i,r}$ is bidder i's market share in region r, p_r is the number of subscribers in region r, and $\mathrm{sv}_{i,r}(\cdot)$ is a piecewise linear function specifying the value bidder i obtains for region r when receiving capacity $c_{i,r}$ in r, defined below. It is defined using the control points listed in (4), and included by appeal to standard MIP formulations for piecewise linear functions. Next, we define regional capacity:

$$c_{i,r} = \sum_{b\in B} \mathrm{cap}_{i,r,b} \tag{12}$$

where $\mathrm{cap}_{i,r,b} \in \mathbb{R}_{\geq 0}$ is the capacity obtained by bidder i in region r for band b, defined as:

$$\mathrm{cap}_{i,r,b} = c_b \cdot X_{i,r,b} \cdot \mathrm{syn}_b(X_{i,r,b}) \tag{13}$$

where c_b is the capacity per unit of band b, and $\mathrm{syn}_b(X_{i,r,b})$ is another piecewise linear function capturing the capacity synergy from having multiple bands in the same region.

National Bidders' Value.
Lastly, we specify $\Psi_{i,k}$, the undiscounted value national bidders obtain from having a bundle that is missing exactly k regions, with $0 \leq k \leq k_{\max}$. We build toward this by introducing several auxiliary variables. To begin, we define $W_{i,r} \in \{0, 1\}$ to be 1 iff bidder i possesses at least one license in region r:

$$W_{i,r} \leq \sum_{b\in B} X_{i,r,b} \qquad \text{No license} \tag{14}$$

$$W_{i,r} \geq \frac{1}{\sum_{b\in B} n_b} \cdot \sum_{b\in B} X_{i,r,b} \qquad \text{At least 1 license} \tag{15}$$

This enables us to define $W_i \in \{0, ...|R|\}$ as the number of regions bidder i covers with at least one license as:

$$W_i = \sum_{r\in R} W_{i,r} \qquad \text{Covered regions} \tag{16}$$

With this, we can define $\hat{W}_{i,k} \in \{0, 1\}$ to be 1 iff bidder i has exactly k missing regions for $k \in \{0, ..., k_{\max}\}$ using the following constraints:

For $0 \leq k < k_{\max}$ uncovered:

$$W_i - (|R| - k) \leq M_{\hat{W}_{i,k}} \cdot (1 - \hat{W}_{i,k}) \ \forall\, k < k_{\max} \tag{17}$$

$$W_i - (|R| - k) \geq -M_{\hat{W}_{i,k}} \cdot (1 - \hat{W}_{i,k}) \ \forall\, k < k_{\max} \tag{18}$$

For $\geq k_{\max}$ uncovered:

$$W_i - (|R| - k_{\max}) \leq M_{\hat{W}_{i,k_{\max}}} \cdot (1 - \hat{W}_{i,k_{\max}}) \tag{19}$$

$$\hat{W}_{i,k_{\max}} \geq \frac{|R| - k_{\max} - W_i + 1}{|R| - k_{\max} + 1} \tag{20}$$

where the "big-M" constant $M_{\hat{W}_{i,k}} = |R|$.

We can now define the bidders' undiscounted value for a bundle missing exactly k regions, using the following:

Value if k uncovered:

$$\Psi_{i,k} \leq \sum_{r\in R} \Omega_{i,r} + M_{\Psi_{i,k}} \cdot (1 - \hat{W}_{i,k}) \tag{21}$$

$$\Psi_{i,k} \geq \sum_{r\in R} \Omega_{i,r} - M_{\Psi_{i,k}} \cdot (1 - \hat{W}_{i,k}) \tag{22}$$

0 otherwise:

$$\Psi_{i,k} \leq M \cdot \hat{W}_{i,k} \tag{23}$$

where $M_{\Psi_{i,k}}$ is the largest non-discounted value bidder i can have.

B. PARAMETERIZATION OF MRVM

The following are the parameters of the MRVM model after being fitted to the 2014 Canadian auction data.

B.1 World Parameters

The region geographic proximity and their populations are set according to the *Tier Two* information from Industry Canada [19]. The frequency bands B are modeled as:

Band Description	Parameters: Block Quantity	Parameters: Capacity c_b
Lower 700 Mhz Paired	3	$\mathcal{U}[3, 4]$
Upper 700 Mhz Paired	2	$\mathcal{U}[1.5, 2.5]$
Unpaired	2	$\mathcal{U}[0.5, 1]$

We set the band synergies as: $syn_b(n) = \begin{cases} 1, & \text{if } n = 1 \\ 1.2, & \text{otherwise} \end{cases}$

B.2 Bidder Parameters

Bidder Type	Parameters: Number of Bidders	Parameters: Maximum Value per Customer α_i	Parameters: Market Share β_i[‡]
Local	3	$\mathcal{U}[200, 400]$	$\mathcal{U}[0.05, 0.15]$
Regional	4	$\mathcal{U}[700, 950]$	$\mathcal{U}[0.1, 0.2]$
National	3	$\mathcal{U}[800, 1400]$	$\mathcal{U}[0.1, 0.2]$

For the fit to the Canadian auction data, we set the thresholds for the function governing the capacity-to-quality mapping in an endogenous way based on the other model parameters as follows (where, for all r, we set $\beta_{i,r} = \beta_i$ for simplicity):[15]

$$z^l_{i,r} = \frac{\max(0, \beta_i - 0.3) \cdot c(r, L)}{p_r \cdot \beta_i} \tag{24}$$

$$z^h_{i,r} = \frac{\min(1, \beta_i + 0.3) \cdot c(r, L)}{p_r \cdot \beta_i} \tag{25}$$

There are several additional type-specific bidder parameters:

Local

- The set of regions of interest, I_i, is drawn uniformly at random from all regions, with $|I_i| \sim \mathcal{U}[3, 7]$.

Regional

- The head region, head_i, is drawn uniformly from all regions.
- The distance discount is $\lambda_i = 2^{-0.9}$.

National

- The discount index cap is $k_{max} = 4$.
- The missing region discount is $\gamma_{i,k} = 1 - 2^{k*0.2}$.

[‡]For simplicity, we set $\beta_{i,r} = \beta_i \ \forall\, r$.

[15]Note that setting these parameters in this endogeneous way is not necessary. Instead, one can also set/fit those parameters directly.

REFERENCES

[1] ACMA. Allocation of the 700MHz (digital divided) and 2.5 GHz bands. http://archive.acma.gov.au/WEB/STANDARD/pc=PC_312315, 2013. Accessed: 2017-1-15.

[2] T. Ali-Yahiya. *Understanding LTE and its Performance*. Springer, New York, 2011.

[3] L. Ausubel and O. Baranov. A practical guide to the combinatorial clock auction. *Economic Journal*, 2017. Forthcoming.

[4] L. M. Ausubel, P. Cramton, and P. Milgrom. The clock-proxy auction: A practical combinatorial auction design. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.

[5] C. Bazelon and G. McHenry. Spectrum value. *Telecommunications Policy*, 37(9):737–747, 2013.

[6] M. Bichler, P. Shabalin, and J. Wolf. Do core-selecting combinatorial clock auctions always lead to high efficiency? an experimental analysis of spectrum auction designs. *Experimental Economics*, 16(4):511–545, 2013.

[7] B. Bünz, S. Seuken, and B. Lubin. A faster core constraint generation algorithm for combinatorial auctions. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 827–834, 2015.

[8] E. Cantillon and M. Pesendorfer. Auctioning bus routes: The london experience. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.

[9] P. Cramton. A review of the 10-40 ghz auction. Technical Report: Office of Communication, UK, 2008.

[10] P. Cramton. Spectrum auction design. *Review of Industrial Organization*, 42(2):161–190, 2013.

[11] P. Cramton. Lessons from the canadian 700 mhz auction, 2014. URL https://ecfsapi.fcc.gov/file/7521096791.pdf. Accessed 2016-08-15.

[12] R. Day and P. Milgrom. Core-selecting package auctions. *International Journal of Game Theory*, 36(3):393–407, 2008.

[13] R. W. Day and P. Cramton. Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research*, 60(3):588–603, May 2012.

[14] R. W. Day and S. Raghavan. Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science*, 53(9):1389–1406, 2007.

[15] FCC. Incentive auction dashboard. https://auctiondata.fcc.gov/public/projects/1000, 2017. Accessed: 2017-1-15.

[16] J. K. Goeree and C. A. Holt. Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior*, 70(1):146–169, 2010.

[17] A. Goetzendorf, M. Bichler, P. Shabalin, and R. W. Day. Compact Bid Languages and Core Pricing in Large Multi-item Auctions. *Management Science*, 61(7):1684–1703, 2015.

[18] Industry Canada. http://www.ic.gc.ca/. Accessed: 2016-08-15.

[19] Industry Canada. Service areas for competitive licensing. http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/h_sf01627.html, 2014. Accessed: 2016-11-15.

[20] Industry Canada. Spectrum management and telecommunications: 700MHz auction (2014). http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/h_sf10598.html, 10 2015. Accessed: 2016-08-15.

[21] W. Ketter and A. Symeonidis. Competitive benchmarking: Lessons learned from the trading agent competition. *AI Magazine*, 33(2):103, 2012.

[22] C. Kroemer, M. Bichler, and A. Goetzendorff. (Un)expected bidder behavior in spectrum auctions: About inconsistent bidding and its impact on efficiency in the combinatorial clock auction. *Group Decision and Negotiation*, 25(1):31–63, 2016.

[23] S. Lahaie, F. Constantin, and D. C. Parkes. More on the power of demand queries in combinatorial auctions: Learning atomic languages and handling incentives. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI)*, 2005.

[24] K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 66–76, 2000.

[25] B. Lubin, B. Bünz, and S. Seuken. New core-selecting payment rules with better fairness and incentive properties. In *Proceedings of the 3rd Conference on Auctions, Market Mechanisms and Applications (AMMA)*, 2015.

[26] N. Nisan. Bidding and allocation in combinatorial auctions. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 1–12, 2000.

[27] D. Parkes. On indirect and direct implementations of core outcomes in combinatorial auctions. Technical report, Harvard University, September 2002.

[28] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial intelligence*, 135(1):1–54, 2002.

[29] T. Sandholm. Very-large-scale generalized combinatorial multi-attribute auctions: Lessons from conducting $60 billion of sourcing. In N. Vulkan, A. E. Roth, and Z. Neeman, editors, *The Handbook of Market Design*. Oxford University Press, 2013.

[30] T. Scheffel, G. Ziegler, and M. Bichler. On the impact of package selection in combinatorial auctions: an experimental study in the context of spectrum auction design. *Experimental Economics*, 15(4):667–692, 2012.