MANAGEMENT SCIENCE
Vol. 00, No. 0, Xxxxx 0000, pp. 000–000
ISSN 0025-1909 | EISSN 1526-5501 | 00 | 0000 | 0001

INFORMS
DOI 10.1287/xxxx.0000.0000

Taming the Communication and Computation Complexity of Combinatorial Auctions: The FUEL Bid Language

Martin Bichler
Department of Informatics, 85748 Garching, Technical University of Munich, Germany, bichler@in.tum.de

Paul Milgrom
Department of Economics, Stanford University, Stanford, CA 94305, milgrom@stanford.edu

Gregor Schwarz
Department of Informatics, 85748 Garching, Technical University of Munich, Germany, gregor.schwarz@tum.de

Combinatorial auctions have found widespread application for allocating multiple items in the presence of complex bidder preferences. The enumerative XOR bid language is the *de facto* standard bid language for spectrum auctions and other applications, despite the difficulties, in larger auctions, of enumerating all the relevant packages or solving the resulting NP-hard winner determination problem. We introduce the FUEL bid language, which was proposed for radio spectrum auctions to ease both communications and computations compared to XOR-based auctions. We model the resulting allocation problem as an integer program, discuss computational complexity, and conduct an extensive set of computational experiments, showing that the winner determination problem of the FUEL bid language can be solved reliably for large realistic-sized problem instances in less than half an hour on average. In contrast, auctions with an XOR bid language quickly become intractable even for much smaller problem sizes. We compare a sealed-bid FUEL auction to a sealed-bid auction with an XOR bid language and to a simultaneous clock auction. While the sealed-bid auction with an XOR bid language incurs significant welfare losses due to the missing bids problem and computational hardness, the simultaneous clock auction leads to a substantially lower efficiency than FUEL due to the exposure problem.

Key words: Market design, spectrum auction, bid languages

1. Introduction

One of the thorniest problems in the theory of resource allocation concerns how to allocate resources efficiently in large scale problems with minimal structure on the potential buyers' possible values. Some kinds of structure in such problems allow huge simplifications. Classical economic theory shows that with convex consumer preferences and convex sets of feasible allocations, price messages can decentralize and decompose society's overall resource allocation problem. This means that once the market-clearing prices are known, the problem can be solved by solving the separate problems for each consumer and firm. When there is no such structure and the resources are indivisible, the general problem is much harder. Analysts do not attempt to work out a universal resource allocation, but instead limit their attention to problems of a manageable size, commonly recommending the use of a combinatorial auction with XOR bidding, in which bidders simply enumerate values for all the possible combinations of items. The auctioneer uses the enumeration to solve a combinatorial optimization problem, which finds the allocation that maximizes the total bids. Such auctions have not only attracted the interests of researchers, but have also been used for public and private sector auctions, with auctions of radio spectrum being prominent examples (Bichler and Goeree 2017).

As the number of items to be allocated becomes large, however, a full XOR-based approach to auctioning quickly becomes impractical for two reasons. The first is related to communication complexity (Nisan and Segal 2006). For example, in combinatorial spectrum auctions in Canada using XOR bidding, there have sometimes been more than 100 spectrum licenses for sale, leading to more than 2^{100} different packages — far too many for any bidder to enumerate (Kroemer et al. 2017). The second reason is that computations at this scale easily become impractical. To address that problem, the auctioneer in Canada limited the number of XOR bids that each bidder is permitted to submit to 2,000, treating the many missing packages as if they have received bids of zero. In lab experiments comparing an XOR design to alternatives, the resulting *missing bids problem* from XOR bidding leads to substantial efficiency losses, even with many fewer than 100 licenses (Bichler et al. 2014).

One alternative to a full XOR-based auction is the Simultaneous Multi-Round Auction (SMRA) and related clock auction formats. While these types of auctions mitigate the communication complexity of XOR-based bid languages, they limit the expressiveness of bids and may cause bidders to end up winning only a subset of their desired items at prices exceeding their valuation, creating the well known *exposure problem*. Combinatorial clock auctions avoid the exposure risk for bidders by interpreting bids in each round as package bids (Porter et al. 2003, Ausubel et al. 2006). However, straightforward bidding in the clock phase does not generally lead to an efficient allocation and strategic bidding can lead to further allocative inefficiencies (Bichler et al. 2013).

The *Flexible Use and Efficient Licensing* (FUEL) bid language, as described below, aims to tame the communication and computational complexity of an XOR auction design, but avoids most of the exposure problems of clock auction formats. It enables larger combinatorial auctions by allowing to decompose computations: FUEL leverages the fact that there are general economies of scale within product categories, but limited economies of scope for local bidders across such categories. This reduces the interdependencies among bids and the resulting problem decomposition leads to significant speedups in the computations.

The design of parametric bid languages is at the core of market design. Defining the format to which bidders must adhere when expressing their preferences is fundamental to any auction design. Simplifications and limitations are almost always required. Leveraging domain knowledge is often key for the development of bid languages that are restrictive enough to reduce the communication and computation complexity of combinatorial auctions while still allowing bidders to state their preferences accurately. Devising domain-specific parametric bid languages has received some attention in the literature (Milgrom 2009, Bichler et al. 2011, Eilat and Milgrom 2011, Bichler et al. 2017), but it is rarely an issue for the design of spectrum auctions in practice. In combinatorial spectrum auctions world-wide, the XOR bid language is the de-facto standard.

We explain the rationale for the FUEL language in three parts, related to its enabling of decomposition to simplify computations, its applicability as a good approximation for some significant combinatorial problems, and test results demonstrating its computational tractability and efficiency to both an XOR-based approach and a simultaneous clock auction format.

For our computational experiments, we assume a high level of participation in the auction by bidders with similar valuations as this is thought to make the optimization more challenging by providing more near-optimal combinations of bids for the software to rule out. Our simulations show that even with a vastly reduced bid set, accurate computations with XOR bids require significantly more computation time than FUEL. The optimization problem coded using the FUEL bid language utilizes many binary variables, and just as for the XOR auction, computing the optimal solution is NP-hard (Arora and Barak 2009). Informally, large problems in this class quickly become intractable and it is widely believed that there are no polynomial-time algorithms for NP-hard problems, although that has not been proven. However, our computational experiments indicate that, in practice, even in auctions with more participating bidders than are expected for an actual spectrum auction, optimal solutions for the FUEL auction can be computed on a desktop computer in mere minutes using commercial off-the-shelf optimization software with minimal customization. When comparing FUEL to a clock auction design, we observe that due to the exposure problem the investigated clock auction format incurs significant efficiency losses in settings with large complementarities. If bidders do not exhibit any synergies, the efficiency of the FUEL and clock auction design are similar.

The FUEL bid language was originally proposed in mid-2019 for a private auction by a consortium of companies providing satellite downlink for commercial television in the United States. The C-band offering consists of spectrum licenses at 3.7–3.98 GHz. The spectrum licenses were to be offered in 406 geographical areas — the *Partial Economic Areas* (PEAs) with 14 licenses to use 20 MHz of bandwidth in each, so the number of possible combinations that any bidder might win in the proposed auction was 15^{406}. In February 2020, the US Federal Communications Commission rejected the private auction proposal and chose instead to conduct a public auction using its usual standard simultaneous clock auction format (Federal Communications Commission 2020a,b).

Combinatorial auctions on this scale have received little attention in the literature so far. For our computational experiments, we use the C-band licenses and compare FUEL to a standard XOR bid language as well as to the FCC's simultaneous clock auction design. While the FUEL bid language was constructed to express good approximations for spectrum values, it is likely to be useful as well for other large combinatorial auctions, especially in procurement, in which managers regularly have to buy large quantities of multiple, related items, respecting economies of scale within products and economies of scope among products or delivery locations (Bichler et al. 2006). As an example, consider the procurement of office paper for a multi-national company with hundreds of subsidiaries worldwide that have a demand of several tons of paper per year. There are suppliers who are only active in some countries or continents and there are international suppliers who bid on parts or all of the demand in most or even all regions. Similarly, buyers of sale of fishery access rights in different regions may enjoy economies of scale in each region and of scope across regions (Iftekhar and Tisdell 2012) and buyers of TV ads (Goetzendorff et al. 2015) who typically seek "reach and repetition" may have similar scale and scope concerns. If bidder values in these and other auctions can be similarly well approximated by the FUEL language, our computational tests for the C-band auction hint at similar performance over a wide set of applications.

1.1. Contributions

An important question about any parametric bid language is whether it is sufficiently expressive to approximate actual bidder values. The FUEL language was based in part on an understanding of the common way broadband networks are engineered. A company that needs additional bandwidth for services in a particular PEA can often provide that either by having more frequencies available or by constructing additional cell sites to densify its network in that PEA. If the costs of densification, which tend to be additive across PEAs, are what drive value adjustments for deviations from a base business plan, then the FUEL language may provide a good framework for bidders to describe their actual spectrum values.

How does the FUEL language work and why are its optimizations so fast? FUEL allows each bidder to build collections of package bids called *bid groups*. Each bid group is built from a single all-or-nothing package bid — the *base bid*, which consists of a *base package* and a corresponding *base price*. A bid group is created from the base bid specifying additions to or subtractions from the base price in case licenses are added or subtracted in any PEA. The value expressions are simplified because any price adjustments are summed over PEAs to get the bid for any adjusted package. For example, if the assigned package has different numbers of licenses than the base package for five PEAs, then the implied bid for that adjusted package is the base price plus the sum of the five positive or negative adjustments. In principle, a bid group including all of the 14 possible adjustments for each of the 406 PEAs — 5,685 numbers in all — would specify prices for every one of the 15^{406} packages. In this sense, FUEL tames the communication complexity compared to XOR bidding.

FUEL's bid group structure also tames the computational complexity. For any fixed combination of winning bid groups, computations are fast because the overall problem is then decomposed into many small problems — one for each of the 406 PEAs. If the number of combinations of winning bid groups is small or well structured, then the overall optimum can be computed quickly.

For the nationwide mobile network operators in the United States, nearly complete nationwide coverage is important for network engineering, so the value of a collection of licenses with, say, 40 MHz of spectrum rights everywhere (or at least in all the major cities) may be much greater than that of a less comprehensive collection. Using FUEL, a base bid for 40 MHz in every PEA can reflect that value pattern. If the bid-group adjustments do not include decrements in major cities, then packages that do not include those areas have a value of zero. Less extreme adjustments can also be included in the FUEL bid groups and additional bid groups by the same bidder can incorporate other value patterns.[1] The commonly noted exposure problem for nationwide bidders in clock auctions is avoided by FUEL because a nationwide bidder either wins a package defined by one of its FUEL bid groups or wins nothing at all.

With many bidders or many potentially winning bid groups per bidder, there could still be very many combinations of potentially winning bid groups, leading to potentially long computation times. To tame that possibility, FUEL includes three kinds of limits on the bid groups submitted by each bidder. Each FUEL bid group must be either a *nationwide* or a *local* group. To qualify as a nationwide group, the sum of the populations covered by each license in the base package bid must be at least twice the sum of the populations of all the PEAs. For example, if a base package consists of two licenses in every PEA, then the bid group qualifies as a nationwide group. A nationwide bidder can have only a limited number of mutually exclusive bid groups. Each local bid group can include only licenses from PEAs within a single Economic Area (EA), and a local bidder can have a limited number of mutually exclusive bid groups for each EA. There are 176 EAs which partition the set of PEAs, with an average of 2.5 PEAs in a typical EA.

The function of these limitations is highlighted by the special case in which there are only local bidders. In this specific case, the optimal solution can be computed by separately solving each

[1] Additional evidence that FUEL may be a good description comes from the package bids submitted in the sealed-bid phase of the Canadian spectrum auction in 2014. Those bids can be described very well (R^2 close to 1 for most bidders) with a linear regression model. Thus, these bids could be described with high precision using a single FUEL bid group.

EA to optimality. Although these problems are also NP-hard, they are small because the EAs are small. Problems with only local bid groups can be solved in seconds in our experiments. The presence of bid groups across EAs by the nationwide bidders makes the problem significantly more complex, but there are only a few of these large nationwide bidders. In particular, the nationwide bidders cannot be expected to enumerate all possible package values in an XOR bid language, but a limited number of FUEL bid groups covers a huge number of XOR bids. For our experiments, we assume that there are 10 nationwide bidders, which is more than there are in the United States currently. We show experimentally that the state-of-the-art integer programming solvers can exploit the problem structure very well, as long as the local bidders are restricted to bid groups within one or a few EAs.

For our complexity analysis, we reduce from the multidimensional knapsack problem to show NP-hardness of the FUEL allocation problem. However, this problem is fixed parameter tractable (Downey and Fellows 2012), which explains the only linear increase in runtime that we observe when increasing only the numbers of bids but keeping other parameters such as the number of PEAs and their supply constant. With the FUEL restrictions on bid groups in place, we are able, in practice, to solve large problems with more than 400 licenses and more than 1,000 bidders using a state-of-the-art branch-and-cut algorithm in a few minutes of runtime. Similar problems, we will show, are intractable when coded using the XOR bid language and beyond what one would expect to solve to optimality. In contrast, the efficiency of simultaneous clock auctions suffers from the exposure problem as we show in another set of experiments.

1.2. Outline

In Section 2, we discuss related literature on bid languages. In Section 3, we provide a complete description of the proposed FUEL bid language, introduce a binary program to formulate the FUEL winner determination problem, and show that the problem is strongly NP-hard. In Section 4, we introduce the XOR bid language as it is widely used in spectrum auctions worldwide. Section 5 presents the clock auction format that was adopted for the upcoming C-band auction. Section 6 describes the experimental design. The results of our extensive numerical tests are presented in Section 7 and the key insights of the paper are summarized in Section 8.

2. Bid Languages and Spectrum Auction Formats

Spectrum auctions worldwide have raised hundreds of billions of dollars and become a model for market-based approaches in the public and private sectors (Milgrom 2004, Bichler and Goeree 2017) and multiple researchers have addressed the computational challenges for those auctions, both for allocation and pricing (Kelly and Steinberg 2000, Pekec and Rothkopf 2003, Day and Raghavan 2007, Goeree and Holt 2010, Day and Cramton 2012). The number of package bids submitted and the language used to express those can both affect the computational hardness of these problems.

2.1. Bid Languages

Generally, a bid in an auction expresses a bidder's willingness to pay money for various outcomes and depends both on the bidder's private preferences and its bidding strategy. A *bid language* defines the format used to communicate the bids. For combinatorial auctions, some common bid languages are built from elements including *bundles* (also known as *packages*), which are subsets of the item set, *atomic bids*, which associate a price with a bundle, and *logical rules*, which govern which bids can win simultaneously. The two most popular and intuitive bid languages of this kind are *exclusive-OR* (XOR) and *additive-OR* (OR).

Definition 1 (XOR bid language). *The bid language exclusive-OR* (XOR) *allows bidders to submit multiple atomic bids with the restriction that at most one of each bidder's atomic bids can win. (This means that the bidder either gets* all *items contained in the bundle listed in* exactly one *of her atomic bids or she gets nothing.) In this language, if no bid is submitted for some bundle, a bid of zero is imputed for it.*

In principle, any valuation function can be expressed as a collection of atomic XOR bids, simply by listing each bundle and its associated price. Such a language is said to be *fully expressive.* In the XOR language, however, that expressiveness is achieved by making an exponential number of atomic bids. For example, in a spectrum auction with 100 distinct licenses for sale, there are 2^{100} packages that must be enumerated to achieve this expressiveness. In practice, only a tiny fraction of bundles receive positive bids, and the fraction of all allocations that the auction algorithm can explore is obtained by multiplying the individual bidders' fractions. This is known as the *missing bids problem.* Laboratory experiments have shown that in realistic settings with many fewer items and packages, that problem can lead to substantial efficiency losses compared with the simultaneous multi-round auction, where bids can be submitted only on individual items (Bichler et al. 2014).

One simple way to reduce the number of missing bids is to use an alternative language, such as the OR language.

Definition 2 (OR bid language). *The bid language additive-OR* (OR) *allows bidders to submit multiple atomic bids with the understanding that any non-intersecting combination of atomic bids can win. (This means that the bidder either gets* all *items contained in each of the bundles listed in* some non-intersecting set *of her atomic bids or she gets nothing.)*

The OR bid language can express values for more different combinations in a compact way, but it can represent only valuations that have limited patterns of substitution (Boutilier and Hoos 2001, Nisan 2006): if two disjoint packages are substitutes for a bidder, an OR bid might win both of them rather than just one.

Limitations of general languages such as XOR and OR have encouraged the development of parametrized or compact, domain-specific bid languages (Goetzendorff et al. 2015). These leverage domain knowledge about values and can sometimes remedy the combinatorial explosion. Examples include volume discount auctions for multi-unit and multi-item procurement markets (Bichler et al. 2011), bid languages for TV ad sales (Goetzendorff et al. 2015), and bid languages for electricity markets in the United States (Papavasiliou et al. 2017, Cramton 2017). The structure of values in these domains can often be usefully exploited. For example, bidding languages for procurement may implement the kinds of discount policies that are widely used to reflect the economies of scale and scope in production. Bidding languages in electricity markets also leverage agreed-upon specifications of cost functions in energy production, such as distinguishing between ramp-up costs and marginal costs in a bid.

2.2. Spectrum Auction Formats

Preference elicitation in combinatorial auctions has long been an issue in the literature (Conen and Sandholm 2001, Sandholm and Boutilier 2006). Parkes (2006) identifies the communication complexity and the bidders' cost for determining exact values for many bundles as the two key problems of preference elicitation in combinatorial auctions. While compact domain-specific bid languages can tame the communication complexity, they still require bidders to determine exact values for a large number of bundles. Iterative combinatorial auctions are sometimes seen as a way to improve preference elicitation as bidders do not have to provide their entire valuations in one step. However, with several hundreds of objects and many auction rounds, bidding can become very challenging for bidders or bidding teams in high-stakes auctions. Note that the number of rounds in iterative multi-object auctions can easily become very large unless the bid increments are very high, which can lead to inefficiencies. There is limited experimental research on larger ascending combinatorial auctions, but already in auctions with only 18 objects, bidders preselect packages which leads to significant inefficiencies (Scheffel et al. 2012).

The simultaneous multi-round auction and related simultaneous clock auction formats allow OR bids on single objects only. The bidding process in such simultaneous clock auctions consists of a series of subsequent rounds in which the auctioneer announces a price for each individual item and the bidders respond with the quantities that they demand for each item at the current price. If the aggregated demand of the bidders exceeds an item's supply, the auctioneer raises the item's price in the next clock round. Even though bidders cannot set the prices themselves, their demand response to clock prices can be interpreted as OR bids. In particular, each OR bid consists of the number of units demanded for a specific item and is associated with the item's respective clock price. As the bidders' bids for different items can be accepted independently of one another, simultaneous clock auctions implement an OR bid language. However, this may cause bidders to end up winning only a subset of their desired items at prices exceeding their valuation. This exposure problem causes bidders to bid strategically.

In recent years, combinatorial clock auctions have been adopted for spectrum sales in many countries around the world (Ausubel et al. 2006, Bichler and Goeree 2017). Combinatorial clock auctions consist of two phases: the clock rounds and a subsequent supplementary round. In contrast to simultaneous clock auctions, bidders submit all-or-nothing package bids during the clock phase, mitigating the exposure problem. A standard XOR bid language is used for the subsequent sealed-bid supplementary phase which allows bidders to specify bids for additional packages. In the end, the auctioneer solves a winner determination problem considering all bids submitted in both phases and accepting at most one bid per bidder. Despite its popularity for spectrum sales around the world, the CCA has some drawbacks. First, there can be strategic problems due to the two-stage nature of the auction (Bichler et al. 2013, Levin and Skrzypacz 2016). Second, in order for the winner determination problem to stay tractable, bidders are often restricted in the number of bids they are permitted to submit in the supplementary round (Ausubel and Baranov 2014). The CCA shares the missing bids problem with first-price auctions that have been run in countries such as France or Norway using the XOR bid language (Bichler and Paulsen 2018). The resulting missing bids problem is particularly apparent when the number of products for sale grows large which is the case for the upcoming C-band auction.

FUEL strikes a balance, seeking both to allow bidders to express relevant preferences for many packages using a small number of parameters and also leading to tractable optimization problems when exact solutions are not available and incentive issues are paramount (Nisan and Ronen 2001).[2] The carefully designed, parsimonious bid language of FUEL allows fast, large-scale optimization with currently available integer programming techniques. This makes it possible, for example, to use the Vickrey-Clarke-Groves (VCG) mechanism to implement the efficient allocation while providing incentives for truthful bidding.

3. FUEL Auction Design

We introduce the FUEL bid language as proposed for the C-band auction in the United States. This specific application allows us to discuss a real-world case and generate realistic instances considering all real-world constraints. As indicated earlier, the FUEL design is not limited to the C-band auction but is applicable for a variety of markets with many items and multiple units each.

3.1. Product Design

Similar to previous auctions designed by the Federal Communications Commission (FCC), the market area for the C-band auction is geographically subdivided into smaller entities, so-called *Economic Areas* (EAs). As some local market participants are expected to be only interested

[2] Even when deterministic approximation mechanisms for general valuations are unknown, black-box mechanisms using randomized mechanisms may still be available (Lavi and Swamy 2011) to mitigate the computational problem.

in spectrum for some part of an Economic Area, each EA is split again into *Partial Economic Areas* (PEAs), with the number of PEAs in an EA ranging from 1 to 12. In total, there are 170 EAs and 406 PEAs across the contiguous United States.

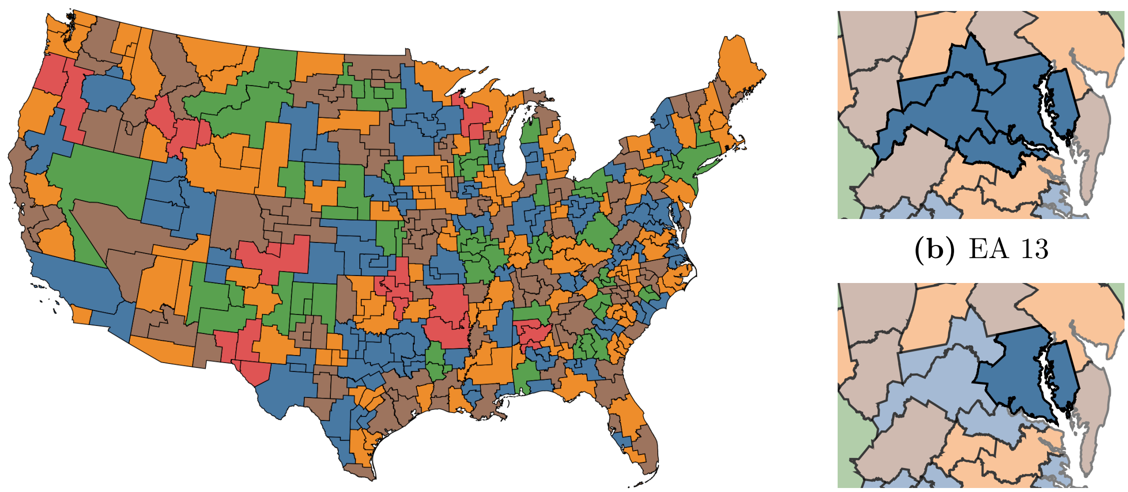

(a) 170 EAs and 406 PEAs throughout the contiguous US

(b) EA 13

(c) PEA 5

Figure 1 Map (a) shows all EAs and PEAs in the contiguous United States. The EAs are highlighted with different colors. Neighboring regions of the same color represent PEAs belonging to the same EA. For illustrative purposes maps (b) and (c) highlight the EA and PEA containing Washington, D.C.

In each PEA, 280 MHz of spectrum is sold in the C-band auction. The spectrum in a PEA is split into 14 homogeneous blocks, each containing 20 MHz of the 280 MHz available per PEA.[3]

3.2. Bid Language

Assuming the C-band auction is organized with 406 PEAs and 14 spectrum blocks per PEA, the number of potential distinct packages equals 15^{406}: far too many to enumerate. The FUEL bid language circumvents this problem by using *bid groups*. Each bid group is based on a single package bid, called the *base bid*, consisting of a *base package* and a *base price*. Bid groups also incorporate *adjustments* that define the price that applies to a package that *increments* or *decrements* the number of licenses to be purchased in a PEA. Each increment is associated with a markup to the base price and each decrement is associated with a discount (see Figure 2). Adjustments are intended to provide a natural and intuitive way for bidders to specify their demand for spectrum and at the same time avoid the missing-bids problem.

[3] In 46 of the 50 most populous PEAs the satellite companies were able to free up 100 MHz (5 blocks) of spectrum earlier than the remaining spectrum blocks. The original proposal of the FUEL bid language differentiated between so-called *early* and *late* spectrum and allowed bidders to submit their bids either in the early/mixed or late category. As the differentiation between early and late spectrum complicates the bid language but does not lead to substantially different results, we will treat all 14 available license blocks to be homogeneous for the remainder of this article.

Bidder 1 SMALL Base price: 200																
		#Licenses														
EA	**PEA**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
60	155		-90	Base	100	170										
60	354			Base												

Figure 2 Bidder 1 submits a bid group whose base package contains 2 licenses in PEA 155 and 2 licenses in PEA 354. She also defines adjustments (in $) for 1, 3, and 4 licenses in PEA 155. If the auctioneer accepts her bid group and assigns her 4 licenses in PEA 155 and 2 licenses in PEA 354 (highlighted blue), her bid for this set of licenses is $\$200 + \$170 + \$0 = \$370$.

Bid groups are classified with respect to the MHz-pop of their base package. The MHz-pop of a set of licenses for the same PEA is given by the product of the frequency bandwidth in MHz and the population of the respective PEA. Summing up the MHz-pop of all PEAs present in the base package gives the MHz-pop of the base package. If the MHz-pop of a base package is no less than the MHz-pop equivalent of two nationwide licenses (i.e. two licenses in all 406 PEAs), then the corresponding bid group is considered to be a *nationwide* bid group and is labeled *large*, otherwise it is a *local* bid group and is classified *small*. While large bid groups may include any combination of PEAs, small bid groups may only contain PEAs from the same EA. Bidders may either win a single large bid group or multiple small bid groups, but never large and small bid groups at the same time. In addition to that, small bid groups can never become winning simultaneously if they contain bids in the same EA.

More formally, if there exist 170 Economic Areas and a bidder submits n_a small bid groups for EA $a \in A$ as well as m large bid groups, the following Boolean expression defines the combination of bid groups that might be accepted simultaneously by the auctioneer. The i-th bid group of the bidder containing bids for EA a is denoted by s_a^i and the j-th large bid group of the bidder is denoted ℓ^j. The logical OR and XOR operators are denoted $\vee$ and $\oplus$, respectively.

$$\left[\left(s_1^1 \oplus s_1^2 \oplus \ldots \oplus s_1^{n_1}\right) \vee \ldots \vee \left(s_{170}^1 \oplus s_{170}^2 \oplus \ldots \oplus s_{170}^{n_{170}}\right)\right] \oplus \left[\ell^1 \oplus \ell^2 \oplus \ldots \oplus \ell^m\right]$$

Depending on the domain, bidders might have to be restricted in the number of small and large bid groups they are allowed to submit in order for the underlying allocation problem to be computationally tractable. Note that in contrast to OR-of-XOR languages as described in Nisan (2006), FUEL has bid groups which allow for adjustments rather than atomic bids on packages of items. We will analyze restrictions for the C-band auction in Section 7.1.

3.3. Winner Determination Problem

The allocation or winner determination problem of the main bidding round can be represented as a binary program, for which the solution identifies the set of winning base bids and associated adjustments. To formalize the winner determination problem of the FUEL bid language, we will introduce additional notation and then express the rules of the FUEL bid language through constraints in the binary program.

Sets and indices:

$i \in I$	Set of bidders.
$a \in A$	Set of Economic Areas (EAs).
$p \in P$	Set of Partial Economic Areas (PEAs).
$g \in G_i^S$	Set of small bid groups of bidder $i \in I$.
$g \in G_{ia}^S$	Set of small bid groups of bidder $i \in I$ that include bids on EA $a \in A$.
$g \in G_i^L$	Set of large bid groups of bidder $i \in I$.
$g \in G_i$	Set of all bid groups submitted by bidder $i \in I$, i.e., $G_i = G_i^S \cup G_i^L$.
$p \in P_i^g$	Set of PEAs contained in the base package of bid group $g \in G_i$ of bidder $i \in I$.
$k \in K$	Set of possible base bid adjustments, i.e., $K = \{0, \ldots, 14\}$.
$K_i^{gp} \subseteq K$	Set of base bid adjustments in bid group $g \in G_i$ of bidder $i \in I$ for PEA $p \in P$.

Parameters:

L_p	Number of licenses offered in PEA $p \in P$. Note that $L_p = 14$ for all $p \in P$.
b_i^{gp}	Number of licenses demanded in the base bid of bid group $g \in G_i$ for PEA p.
ω_i^g	Base price for bid group $g \in G_i$ of bidder $i \in I$.
μ_i^{gpk}	Markup/discount on the base price of bid group $g \in G_i$ of bidder $i \in I$ for a total of k licenses in PEA $p \in P$. The parameter is 0 when k equals b_i^{gp}. Otherwise, if k specifies an increment or decrement, the number is positive or negative, respectively.
M	The maximum number of small bid groups that any bidder submits.

Decision variables:

$x_i^g \in \{0,1\}$	Binary variable denoting whether bidder $i \in I$ wins bid group $g \in G_i$.
$y_i^{gpk} \in \{0,1\}$	Binary variable denoting whether bidder $i \in I$ wins in total $k \in K_i^{gp}$ licenses in PEA $p \in P_i^g$ as stated in bid group $g \in G_i$.
$z_i \in \{0,1\}$	Binary variable denoting whether bidder $i \in I$ wins multiple small bid groups ($z = 0$) or one large bid group ($z = 1$).

$$\max \sum_{i \in I} \sum_{g \in G_i} (x_i^g \, \omega_i^g) + \sum_{i \in I} \sum_{g \in G_i} \sum_{p \in P_i^g} \sum_{k \in K_i^{gp}} (y_i^{gpk} \, \mu_i^{gpk}) \tag{1}$$

$$\text{s.t.} \quad \sum_{k \in K_i^{gp}} y_i^{gpk} = x_i^g \qquad \forall i \in I, \forall g \in G_i, \forall p \in P_i^g \tag{2}$$

$$\sum_{i \in I} \sum_{g \in G_i} \sum_{k \in K_i^{gp}} (y_i^{gpk} \, k) \leq L_p \qquad \forall p \in P \tag{3}$$

$$\sum_{g \in G_i^L} x_i^g \leq z_i \qquad \forall i \in I \tag{4}$$

$$\sum_{g \in G_i^S} x_i^g \leq M(1 - z_i) \qquad \forall i \in I \tag{5}$$

$$\sum_{g \in G_{ia}^S} x_i^g \le 1 \qquad \forall i \in I, \forall a \in A \tag{6}$$

$$x_i^g \in \{0,1\} \qquad \forall i \in I, \forall g \in G_i \tag{7}$$

$$y_i^{gpk} \in \{0,1\} \qquad \forall i \in I, \forall g \in G_i, \forall p \in P_i^g, \forall k \in K_i^{gp} \tag{8}$$

$$z_i \in \{0,1\} \qquad \forall i \in I \tag{9}$$

Objective & Constraints:

(1) The objective is the sum of base prices of winning bid groups and the respective base price markups/discounts of the winning adjustments.
(2) In case a bidder wins a bid group $g \in G_i$, she must win exactly one adjustment in each PEA being part of the bid group's base package.
(3) Supply constraint for the number of licenses in PEA $p \in P$.
(4) Bidder $i \in I_0$ may win at most one large bid group.
(5) If Bidder $i \in I_0$ wins a large bid group ($z = 1$), she cannot win any small bid groups. In case she does not win a large bid group ($z = 0$), the coefficient M ensures that she can win all her small bid groups simultaneously.
(6) At most one bid group of a bidder may become winning per EA. The constraint is explicitly formulated only for small bid groups because it is implicitly given for large bid groups by constraint (4).

The winner determination problem of the FUEL bid language is related to the d-dimensional knapsack problem (DKP).

DEFINITION 3 (DKP). *A set of n items with profits $p_i \ge 0$ and $1 \le j \le d$ resources with capacities given by $c \in \mathbb{R}_+^d$ are given. Each item i consumes an amount $w_{ij} \ge 0$ of resource j. The decision version of the d-dimensional knapsack problem asks whether there exists a selection of items with total profit larger than r such that the chosen items do not exceed the resource capacities c_j.*

Let us briefly define the decision version of the FUEL winner determination problem which we denote D-FUEL. We refer to the PEAs in the C-band auction as *items*, and the available licenses within a PEA as *units*. We focus on the case where bidders do not make use of the optional adjustments and each bidder specifies a single bid group only.

DEFINITION 4 (D-FUEL). *There are n' bidders submitting a single bid group i', each specifying a number of desired units $w'_{ij} \ge 0$ for each of d' items. The overall number of units available for sale in each of the $1 \le j' \le d'$ items is $c'_j \ge 0$. The bid price for each bid group is $p'_i \ge 0$. Is there an allocation of bids, such that the total sum of accepted bids exceeds the revenue r'?*

Now, given an instance I of DKP, there is a 1-to-1 mapping of a variables x in DKP to x' to construct an instance I' of D-FUEL. Any instance of DKP can be reduced to an instance of D-FUEL.

PROPOSITION 1. D-FUEL *is strongly NP-complete.*

Since DKP is strongly NP-complete (Fréville 2004, Varnamkhasti 2012), it follows that D-FUEL must also be strongly NP-complete. If a problem is strongly NP-complete, then it remains NP-complete even if all of its numerical parameters (e.g., object sizes and knapsack sizes) are bounded by a polynomial in the length of the input. Importantly, any strongly NP-hard optimization problem cannot have a fully polynomial-time approximation scheme (or FPTAS) unless $P = NP$ (Garey and Johnson 1979). It has also been shown that there is no efficient polynomial-time approximation scheme (EPTAS) for DKP (Kulik and Shachnai 2010).

Recent work by Gurski et al. (2019) has looked at knapsack problems from the point of view of fixed-parameter tractability (fpt). Parametrized complexity studies the parameters of a problem on which the runtime depends (Downey and Fellows 2012). This allows the classification of NP-hard problems on a finer scale than in the traditional way, where the complexity of a problem is only measured as a function of the number of bits in the input. By separating the problem into two parts (the input and the parameters), one hopes to find an algorithm that has good runtime as a function of one part (the input), while allowing for arbitrarily bad runtime as a function of the other part (the parameters). For example, the number of licenses for sale in the PEAs in our setting doesn't change (the parameters), but the number of bids (input) varies. Fixed-parameter tractability generalizes polynomial time computability by admitting algorithms whose runtime is exponential, but only regarding the parameters.

An algorithm A is an fpt-algorithm with respect to a parameter κ, if there is a computable function f and a constant $t \in \mathbb{N}$ such that for every instance I the runtime of A on I is at most $f(\kappa(I)) \cdot |I|^t$. If f is also a polynomial, A is referred to as polynomial fpt (PFPT) algorithm with respect to κ. Actually, for a fixed parameter $\kappa(I) = (c_1, \cdots, c_d)$, the capacities in DKP, the problem is PFPT (Gurski et al. 2019). The problem is even linear in n, the number of items, in the worst case usind a specific dynamic program. Similarly, in D-FUEL, the 14 licenses ($c'_j = 14$) in each of the $d' = 406$ PEAs describe the fixed parameter set, but we are interested in scaleability in the number of bids. Although, we use a general purpose branch-and-cut solver considering the various additional variables and constraints and not a dynamic program, also in our experimental results we find that for a fixed number of licenses per PEA the runtime develops linearly on average in the number of bids. Apparently, the branch-and-cut solver can exploit the problem structure well.

4. Sealed-Bid XOR Auction Design

We compare the empirical complexity of auctions with the FUEL bid language to auctions with a standard XOR bid language. For this purpose, we briefly introduce the XOR bid language and the corresponding winner determination problem.

4.1. Bid Language

Similar to FUEL bid groups an XOR bid consists of a set of PEAs for which the bidder would like to acquire licenses. For each of these PEAs, the bidder specifies the number of licenses that the bidder would like to purchase. Every XOR bid is also associated with a price which expresses the bidder's valuation for the set of licenses specified in the XOR bid. Bidders may submit multiple XOR bids but at most one of them is accepted. In contrast to the FUEL bid language, the XOR bid language does not distinguish between small and large bids. It is well-known that the winner determination problem with an XOR bid language is strongly NP-hard and can be modeled as a weighted set packing problem (Lehmann et al. 2006). An instance of the problem asks whether a given collection C of n sets (or bids) contains z mutually disjoint sets. It has been shown early on that unless $P = NP$ there is no PTAS for the winner determination problem (Sandholm 2002), and we can only hope for a constant factor approximation (Lehmann et al. 2002). We can also not expect parametrized algorithms. With z being the parameter, the problem is $W[1]$-complete, i.e., we cannot expect an algorithm of complexity polynomial in n, and proportional to $f(z)$ for any function f (Downey and Fellows 2012). These problems are believed to be fixed-parameter intractable.

4.2. Winner Determination Problem

For the binary formulation of the winner determination problem of the XOR bid language, we reuse the notation introduced in Section 3.3 and modify it in the following way. The set G_i no longer refers to the FUEL bid groups but to the XOR bids submitted by bidder $i \in I$. We further introduce variables d_i^{gp} that denote how many license blocks bidder i demands in her XOR bid g for PEA p.

$$\max \sum_{i \in I} \sum_{g \in G_i} (x_i^g \, \omega_i^g) \tag{1}$$

$$\text{s.t.} \quad \sum_{i \in I} \sum_{g \in G_i} (x_i^g \, d_i^{gp}) \leq L_p \qquad \forall p \in P \tag{2}$$

$$\sum_{g \in G_i} x_i^g \leq 1 \qquad \forall i \in I \tag{3}$$

$$x_i^g \in \{0, 1\} \qquad \forall i \in I, \forall g \in G_i \tag{4}$$

Objective & Constraints:

(1) Objective function summing up prices of winning XOR bids.
(2) Supply constraint for the number of licenses in PEA $p \in P$.
(3) Each bidder $i \in I$ may win at most one XOR bid.

5. Simultaneous Clock Auction

In August 2020 the FCC adopted a simultaneous clock auction format for the upcoming C-band auction (Federal Communications Commission 2020b). In the following, we will briefly summarize the most important rules of the clock phase of the adopted C-band auction defined in the FCC's technical guide (Federal Communications Commission 2020a).

5.1. Product Design

Similar to the FUEL product design from Section 3.1, the adopted C-band auction differentiates 406 products, each representing licenses in one of the 406 PEAs. There are 14 homogeneous 20 MHz license blocks offered for each product. The minimum opening bid for each product is set with respect to the MHz-pop of the respective license block. In order to guarantee a speedy auction, the minimum opening bids are chosen by the FCC to approximate the relative value of the licenses. Taking the feedback of the competing telecommunication companies into account, the FCC decided that for PEAs 1–50, 51–100, and 101–406 the minimum opening bid amounts are based on \$0.03 per MHz-pop, \$0.006 per MHz-pop, and \$0.003 per MHz-pop, respectively, subject to a minimum of \$1,000 (Federal Communications Commission 2020b).

5.2. Auction Process

In the first round of the auction, the bidders simply state their demand for each product at the opening prices which equal the minimum opening bids. In any subsequent round, the auctioneer announces two prices for each product, the *start-of-round price* and *clock price*. At the beginning of a clock round the price of a product equals the start-of-round price. While processing the bids, a product's price can be raised up to its clock price.

At certain prices, bidders may want to adapt their demand. They can do so by submitting bids for the respective products. Each bid is associated with a single product, a price, and the number of licenses that the bidder demands from this price onward. The bid price must lie in between the start-of-round and clock price of the current clock round and can also be stated in terms of a *price point*, which defines the percentage position of the bid price between the start-of-round and clock price. A bidder may submit up to five bids per product. The price points of bids for the same product must be distinct and the trajectory of the number of requested licenses must be monotonically increasing or decreasing with respect to the associated price points.

Example 1. Assume a bidder states a demand for 5 blocks for a certain product in the first clock round. In the second clock round, the start-of-round price for the product is \$1,000, the clock price is \$2,000. The bidder reduces her demand to 4 blocks at the 20% price point (i.e., at price \$1200) and to 2 blocks at the 60% price point (i.e., at price \$1600). She is not able to submit a third bid for 3 blocks at the 80% price point (i.e. at \$1800) as her demand trajectory would no longer be monotonic.

If a bidder would like to maintain her current demand for a product throughout the clock round, she submits a bid at clock prices associated with the number of demanded license blocks. In case she does not desire any spectrum in a PEA, she submits a bid stating a demand for 0 blocks at the start-of-round prices.

After all bidders submitted their bids for the clock round, the auctioneer starts to process the bids in the following order. First, all bids that maintain a bidder's demand for a product are applied. Afterwards, the remaining bids of all bidders are inserted into a queue in increasing order with respect to their price points. Ties are broken according to a bid-specific pseudo-random number. The auctioneer processes the bids in the queue one after another and applies each bid to its maximum possible extent, meaning that a bid may either be fully applied, partially applied, or rejected. Bids to reduce demand are only processed to the extent that no excess supply is created or further increased; bids to increase demand are only applied as long as the respective bidder still has enough eligibility (see Section 5.3). Bids that can be fully applied are removed from the queue. In case a bid can only be processed partially, it remains in the queue as it might be possible to process it later on. Whenever a bid can be processed partially or entirely, the queue is re-tested from the beginning in order to check whether any higher prioritized bids can now be applied. This process continues until all bids are processed and no bid in the queue can further be applied. Bids remaining in the queue after this process are discarded at the end of the round.

Example 2. Assume there are two bidders A and B competing for 14 license blocks within the same PEA. Both bidders initially demand 8 blocks. For the current clock round, bidder A reduces her demand to 7 blocks at the 20% price point, while bidder B reduces her demand to 6 licenses at the 50% price point. When the clock round closes, the auctioneer first applies the bid of bidder A as it is associated with a smaller price point. The bid of bidder B can only be processed partially (i.e., the second bidder is allocated 7 license blocks) as fully applying her bid would lead to oversupply. If the auction ended after this round, the second bidder is forced to purchase 7 licenses at the 50% price point even though she wanted to reduce her demand to 6 licenses.

After processing the bids, the auctioneer determines the *posted prices* of the current clock round. If a product is overdemanded after processing the bids from round t, its posted price equals its clock price of round t. In case the auctioneer processed some bids in round t that reduced the aggregate demand for a product so that it equals the supply, then the posted price is set to the price of the bid that reduced the aggregate demand to the supply. In any other case, the product's posted price equals the start-of-round price of round t. The start-of-round prices of round $t+1$ equal the posted prices of round t, while the clock prices of round $t+1$ are set between 10% (the bid increment) higher than the posted prices of round t.[4]

[4] The *bid increment percentage* is set within the range of 5% to 20% inclusively and may change between the auction rounds. The default value is 10% (Federal Communications Commission 2020a).

The auction terminates when there are no more overdemanded products. Each bidder is allocated the products according to the bid processing of the last clock round. The final clock prices of all products are given by the final clock round's posted prices.

5.3. Eligibility Rule

During the clock phase of the C-band auction, bidders are subject to an eligibility rule. For this purpose each product is associated with a number of *bidding units* which is based on the MHz-pop of the respective license block — more precisely, it is approximately 1/20 th of the product's minimum opening bid. A bidder's eligibility in the first round is determined by the bidding units associated with an *upfront payment* that each bidder makes before the start of the auction.[5]

At the beginning of an auction round a bidder may submit bids whose associated bidding units can be as high as 120% of the bidder's eligibility.[6] While processing the bids, however, the auctioneer accepts bids at most up to the bidder's eligibility. Once an auction round closes, a bidder's *processed demand* is given by the number of licenses that are assigned to the bidder at the end of the auction round. The *processed activity* describes the number of bidding units associated with the processed demand. If a bidder's processed activity falls below 95% of its current activity, then the bidder's next round eligibility is set to the ratio of the bidder's processed activity divided by 0.95.[7] Otherwise, the bidder's eligibility in the next round equals the one of the current round.

6. Experimental Design

In the C-band auction of the FCC, there are expected to participate two types of bidders: up to 10 large telecommunication providers who strive for nationwide coverage and a few hundred smaller competitors who are only interested in spectrum for a small subregion. In our tests we model the 10 *national* bidders explicitly, assuming that they are interested in license blocks in almost all 406 PEAs. *Local* bidders, on the other hand, are modeled implicitly. We over-approximate the number of local bidders to 1,000 but assume that they are only active in a single economic area. As small FUEL bid groups may only contain PEAs from at most one economic area, a bidder who is interested in licenses from three economic areas has to submit three separate FUEL bid groups. Thus, from a computational point of view, such a bidder can be modeled by three individual bidders submitting one bid group each. The advantage of modeling local bidders implicitly is that we do not have to fix the market area sizes of local bidders before running our simulation.

[5] For each PEA the minimum opening bid, the number of bidding units, and the upfront payment are listed at `https://www.fcc.gov/file/19175/download`.

[6] The *activity upper limit* is set within the range of 100% to 140% inclusively and may change between auction rounds. The default value is 120% (Federal Communications Commission 2020a).

[7] The *activity requirement percentage* is set within the range of 90% to 100% inclusively and may change between the auction rounds. The default value is 95% (Federal Communications Commission 2020a).

To construct synthetic bids for our numerical simulation of the market, we have to make suitable design decisions regarding the following two aspects: (1) What are the assumed valuations of bidders for the different products and (2) for which products are bidders actually submitting bids in the auction? In Section 6.1, we address the first aspect by proposing a value model that we use for modeling the bidders' preferences for both FUEL and XOR bids. As the FUEL and XOR bid language pose different restrictions on which PEAs can be included simultaneously in a bid, we address the second aspect separately for the two bid languages in Sections 6.2 and 6.3, respectively.

6.1. Value Model

A widespread international metric for comparing the prices of spectrum is the license price per MHz-pop. The FCC follows this convention for defining minimum opening bids for each product in the upcoming C-band auction (see Section 5.1). We use this approximation for our value model and define w_p to be the minimum opening bid for PEA $p \in P$.

The bidders' valuations for licenses in a particular PEA differ depending on their financial strength and their current possession of frequencies. To generate idiosyncratic bidder valuations, we therefore introduce value factors r_{ip} for each bidder $i \in I$ and PEA $p \in P$ which scale the minimum opening bid w_p of a PEA for a particular bidder. In general, local bidders are financially weaker than nationwide bidders so that we choose value factors r_{ip} for local and nationwide bidders uniformly at random from the intervals $[1.0, 1.3]$ and $[1.1, 1.4]$, respectively.

To provide a functional $5G$ network, bidders need spectrum bandwidth of at least 40 MHz which corresponds to 2 license blocks in the C-band-auction. Therefore, bidders have only little interest in being allocated less than 2 licenses. On the other hand, a bidder's marginal valuation for more than 5 licenses is very small. As a consequence, the assumed valuation of a bidder is represented best by a sigmoid function whose point of inflection Δ_i is a bidder specific value chosen uniformly at random from the interval $[2, 4]$. Scaling this sigmoid function with the idiosyncratic bidder valuations above gives a bidder's valuation function for a particular PEA:

$$v_{ip}(x) = \begin{cases} 0 & \text{if } x = 0, \\ w_p \, r_{ip} \left(\dfrac{L}{1 + e^{-x+\Delta_i}} \right) & \text{otherwise,} \end{cases}$$

where x is the number of licenses demanded by bidder $i \in I$ in PEA p, w_p is the minimum opening bid in PEA p, r_{ip} is the idiosyncratic value factor of bidder i for PEA p, Δ_i defines the point of inflection, and $L = 14$ is a constant defining the number of licenses available in each PEA. Figure 3 displays three different valuation functions for exemplary values of w_p, r_{ip} and Δ_i.

While we carefully select the parameters of our value model to resemble the actual preferences of bidders in the C-band auction as accurately as possible, the model is certainly still a simplification

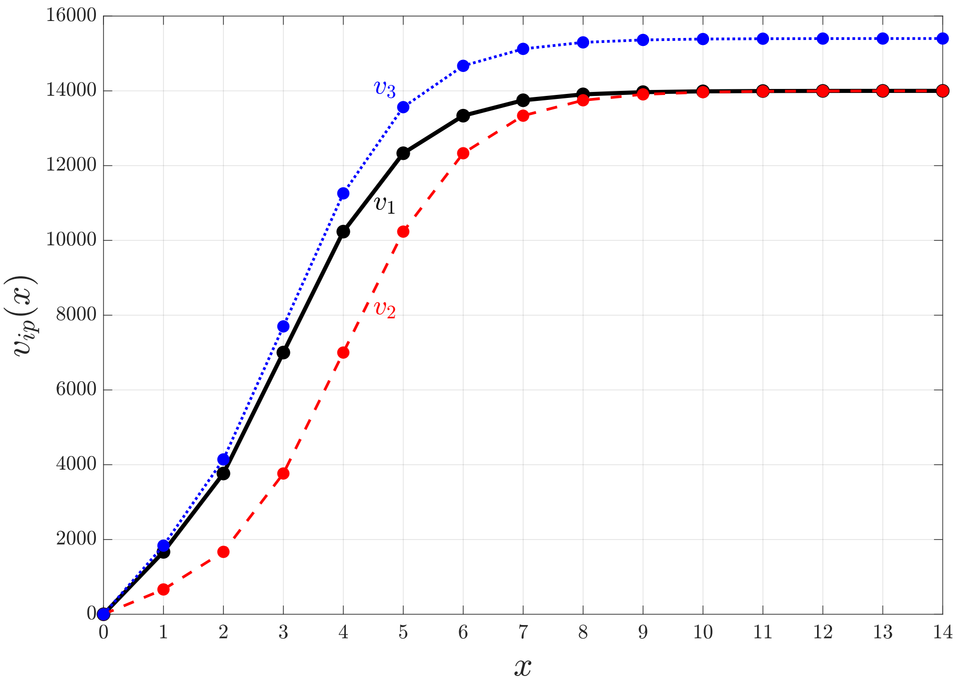

Figure 3 The three valuation functions v_1, v_2, v_3 represent the valuations of three different bidders for the same PEA p with $w_p = 1{,}000$. The bidder specific parameters r_{ip} and Δ_i are chosen as follows: v_1: $r_{1p} = 1.0$ and $\Delta_1 = 3$; v_2: $r_{2p} = 1.0$ and $\Delta_2 = 4$; v_3: $r_{3p} = 1.1$ and $\Delta_3 = 3$.

of the bidders' valuations in practice. For other practical applications with a substantially smaller number of items, less complementarities in the bidders' preference relations, and different assumptions on the bidders' valuations, other auction designs such as the simultaneous clock auction might be preferable.

6.2. FUEL Bid Generation

We assume that every local bidder is active in only a single EA and submits small bid groups containing licenses for that particular EA. The number of PEAs included in the bidder's small bid group is chosen uniformly at random between one and the number of PEAs available in the respective EA (at most 12). In practice, some local bidders may be active in several PEAs belonging to different EAs, but for the purposes of estimating runtime, these bidders can equivalently be represented as multiple independent local bidders as discussed at the beginning of Section 6.

Unlike local bidders, we assume that nationwide bidders are active throughout the whole United States: each of their bid groups covers at least 380 of the 406 PEAs. The PEAs not contained in a bid group are chosen uniformly at random among the 50% least populous PEAs so that nationwide bidders are always able to provide service in all of the most densely populated areas whenever one of their bid groups is accepted by the auctioneer.

When local and nationwide bidders place a bid on licenses in a PEA, they have to state a base bid and may additionally specify some adjustments. According to the sigmoid value model (see

Section 6.1), a bidder's largest marginal gain for a license is at Δ_i, the point of inflection of the bidder's valuation function v_{ip}. Therefore, the base package either contains $\lfloor \Delta_i \rfloor$ or $\lceil \Delta_i \rceil$ licenses in each PEA. Furthermore, it is assumed that a bidder specifies between 0 and 4 adjustments for each PEA in the bid group. The exact number is again chosen uniformly at random. The selected adjustments always constitute a consecutive interval around the base bid as it is assumed that this models the bidders' valuations most accurately.

The number of bid groups that local and nationwide bidders are allowed to submit is a parameter that was still undefined for the C-band auction at the time the design was proposed. It was to be chosen to ensure the computational tractability of the winner determination problem. We will address this question in Section 7.1.

6.3. XOR Bid Generation

In contrast to the FUEL bid language, bids in the XOR bid language are no longer subject to any EA restrictions. In particular, this means that bidders can submit bids for any subset of PEAs even though they belong to different EAs (see Figure 4).

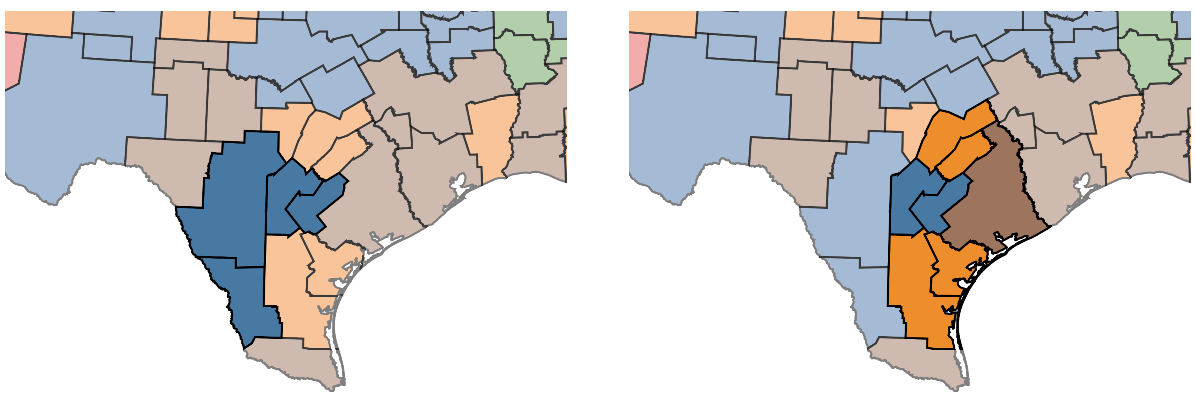

(a) Small FUEL bid group **(b)** XOR bid

Figure 4 Visualization of the PEAs included in the bids of two bidders, where the first bidder submits a FUEL bid group and the second one an XOR bid. EAs are highlighted with different colors. Neighboring regions of the same color represent PEAs belonging to the same EA. While the PEAs in a small FUEL bid group must all belong to the same EA (i.e., all PEAs in the FUEL bid group are of the same color); XOR bids can contain PEAs from different EAs (i.e., bids can contain PEAs of different colors).

When generating XOR bids it is assumed that any local bidder's market area contains between one and seven PEAs (potentially belonging to different EAs) which form a highly cohesive component. Each XOR bid of a local bidder contains between one and the maximum number of PEAs available in her market area. On average every local bidder's XOR bid encompasses 2.5 PEAs which roughly equals the average number of PEAs per EA which is $406/170 \approx 2.4$. Nationwide bidders are active throughout the entire 406 PEAs. Similar to a nationwide bidder's FUEL bid groups, each XOR bid contains at least 380 of the 406 PEAs. The PEAs not contained in an XOR bid of a nationwide bidder are chosen uniformly at random among the 50% least populous PEAs. As explained in our value model, bidders need at least two license blocks for setting up a viable 5G network but are assumed to have little interest in being allocated more than five licenses, we assume that both nationwide and local bidders will submit XOR bids that state a demand between two and five licenses in each PEA. The exact number of licenses is chosen uniformly at random from this interval for each XOR bid and each PEA.

In order to compare different auctions, we need to derive XOR bids from FUEL bid groups, FUEL bid groups from XOR bids, and clock bids from FUEL bid groups. This is described in detail in the Appendices A to C.

7. Results

In this section we report the results of our computational experiments using the Gurobi Optimizer 9.1.0 to solve the winner determination problem up to a tolerance ("MIPGap") of 10^{-4}, i.e., the solution computed by Gurobi differs from the optimal solution by no more than 0.01%. The time limit is set to 30 minutes for all test instances. Our test computer contains two Intel(R) Xeon(R) CPU E5-2620 @ 2.00GHz and 64GB of RAM. All test instances are available upon request.

7.1. FUEL Bid Groups

The original FUEL proposal did not specify the number of bid groups that local and nationwide bidders would be allowed to submit but proposed to choose those numbers to ensure the computational tractability of the winner determination program.

Let z_L and z_N denote the number of bid groups that local and nationwide bidders are allowed to submit, respectively. For each configuration of z_L and z_N in Table 1, we generated 25 random instances with our FUEL bid generator. Table 1 summarizes the number of bid groups submitted by all bidders, the average number of binary variables and constraints in the winner determination problem, the average runtime, the number of test instances that exceed the time limit of 30 minutes (TLE = time limit exceeded), the maximum MIPGap of all 25 test instances, and the average number of licenses that remain unsold out of 5,684 (406×14) licenses. Test instances that exceed the time limit of 30 minutes are weighted with 1,800 seconds when computing the average runtime.

Result 1 *If the 1,000 local bidders are restricted to submit at most seven small FUEL bid groups and the 10 national bidders are limited to place at most seven large FUEL bid groups, the winning allocation can be determined within 144 seconds and a MIPGap of only* 0.0001 *on average. If nationwide bidders were absent in the auction, then the FUEL winner determination problem can*

z_L	z_N	Bid Groups	Binary Variables	Constr.	Runtime in sec.	TLE	Max. MIPGap	Unsold Licenses
1	1	1,010	17,864	6,033	3	0 of 25	0.00010	53.4
3	3	3,030	53,579	18,282	39	0 of 25	0.00010	14.4
5	5	5,050	89,515	29,569	144	0 of 25	0.00010	6.6
7	7	7,070	124,992	40,769	236	0 of 25	0.00010	4.0
10	10	10,100	178,897	57,672	678	5 of 25	0.00044	2.0
15	15	15,150	268,366	85,808	935	6 of 25	0.00051	1.1
15	0	15,000	91,246	26,828	11	0 of 25	0.00010	11.6
30	0	30,000	182,680	52,280	41	0 of 25	0.00010	3.3
50	0	50,000	304,204	86,145	96	0 of 25	0.00010	2.5

Table 1 **Average values of 25 test instances for different configurations of the number of small and large FUEL bid groups that local and nationwide bidders submit, respectively.**

be solved in less than 100 *seconds even when local bidders are allowed to submit up to* 50 *small FUEL bid groups each.*

According to our test results in Table 1, the number of small and large FUEL bid groups that bidders are allowed to submit has a direct impact on the number of binary variables and constraints present in the FUEL winner determination problem. Restricting the number of bid groups in the auction is, therefore, a viable method for auctioneers to tame the computational complexity.

If nationwide bidders are absent, then all bid groups only contain bids for a single EA. Thus, the FUEL winner determination problem can be solved separately for each of the 170 EAs, leading to substantially reduced runtimes even when local bidders are allowed to submit 50 bid groups. Once nationwide bidders are added to the auction, their large bid groups create interdependencies between different EAs which increases the complexity of the allocation problem substantially.

7.2. FUEL Admissible EAs

Even when restricting bidders to at most seven bid groups, the binary program for solving the FUEL allocation program is still very large, with roughly 125,000 binary variables and 40,000 constraints. The reason such large binary programs can still be solved within 30 minutes is the hierarchical structure of EAs and PEAs. In fact, if no large bid groups were submitted by any nationwide bidder, the allocation program could be solved independently in each of the 170 EAs as every small bid group contains PEAs from only a single EA. This hierarchical structure allows the optimizer to decompose the problem and apply effective cuts in the branch-and-cut algorithm. Branch-and-cut involves running a branch-and-bound algorithm and using cutting planes to tighten the linear programming relaxations (Mitchell 2002). Such algorithms are the basis for most of the mixed-integer programming solvers available today.

In the following test, we analyze the degree to which the hierarchical structure impacts the runtimes. For this, we successively increase the maximal number of EAs contained in a small bid group. Raising this number above one creates inter-dependencies between small bid groups and therefore effectively prevents the optimizer from decomposing the full binary problem into smaller subproblems. The benefit of this relaxation for local bidders is that they are given the ability to express synergies between licenses of PEAs belonging to different EAs in their small bid groups.

The bid generation of the FUEL bid generator is adapted as follows. Let k denote the parameter that defines the maximum number of EAs for which bidders may state demand within the same bid group. In a first step, we determine the market area for every local bidder which is given by a set of k EAs that form a highly cohesive component. Within this market area, a bidder submits small bid groups, each containing bids for at least one and at most all PEAs present in the bidder's market area. The bid generation for nationwide bidders is independent of parameter k as there is no restriction regarding the number of EAs contained in large bid groups that nationwide bidders submit. For the following tests, we assume that local bidders submit seven small bid groups and nationwide bidders submit seven large bid groups as our tests in Section 7.1 indicate that for $k = 1$ such instances can be solved within the time limit of 30 minutes.

Result 2 *If* 1,000 *local bidders are allowed to submit seven bid groups across two EAs (not only within one EA), there already exists one test instance that cannot be solved within the time limit. If the admissible number of EAs per small bid group is increased to 15, only one out of 25 test instances can be solved within the time limit.*

For each configuration of k (the maximum number of EAs contained in any local bidder's small bid group), we generate 25 random test instances. Table 2 lists the number of bid groups submitted by all bidders, the average number of binary variables and constraints in the winner determination problem, the average runtime, the number of test instances that exceed the time limit of 30 minutes, the maximum MIPGap, and the average number of licenses that remain unsold of all 5,684 (406×14) licenses.

k	Bid Groups	Binary Variables	Constr.	Runtime in sec.	TLE	Max MIPGap	Unsold Licenses
1	7,070	124,992	40,769	236	0 of 25	0.00010	4.0
2	7,070	147,180	49,131	322	1 of 25	0.00013	2.7
3	7,070	172,622	58,552	452	1 of 25	0.00012	2.8
5	7,070	223,804	77,553	542	1 of 25	0.00012	2.7
7	7,070	278,374	97,584	1,105	3 of 25	0.00075	2.6
10	7,070	354,776	125,916	1,593	15 of 25	0.00108	5.1
15	7,070	473,227	169,964	1,784	24 of 25	0.00423	8.1

Table 2 Average values of 25 test instances for different restrictions on the number of EAs that can be contained in a single small bid group.

The test results in Table 2 imply that raising the maximum admissible number of EAs for small bid groups has a significant impact on the runtime. Moreover, we witness an increase of binary variables and constraints in the binary program despite the number of bid groups staying constant for all configurations of k. This is mainly due to the fact that with a higher value for k local bidders have a larger market area in which they are active so that they tend to be interested in spectrum of more PEAs.

One could argue that the substantial increase in runtime is predominantly due to the fact that the binary program becomes much larger when raising parameter k whereas the impact of the underlying hierarchical structure is negligible. The substantially larger binary program stems from the fact that the number of PEAs in a local bidder's small bid group is chosen uniformly at random between one and the maximal number of PEAs in her market area. As the market area grows when raising k, bidders submit more bids which causes the binary program to grow substantially. In order to keep the number of binary variables and constraints constant across different parameter settings of k, we divide the number of local bidders by k for each treatment. At the same time, we assume that a local bidder's bid group contains at least k (previously one) and at most all PEAs present in the bidder's market area. The results of this test are shown in Table 3.

k	Local Bidders	Bid Groups	Binary Variables	Constr.	Runtime in sec.	TLE	Max MIPGap	Unsold Licenses
1	1,000	7,070	124,992	40,769	236	0 of 25	0.00010	4.0
2	500	3,570	119,955	40,240	491	2 of 25	0.00021	18.5
3	333	2,401	119,466	40,465	549	3 of 25	0.00096	27.3
4	250	1,820	119,008	40,499	734	4 of 25	0.00062	40.5
5	200	1,470	119,063	40,579	767	3 of 25	0.00046	55.2

Table 3 Average values of 25 test instances for different restrictions on the number of EAs that can be contained in a single small bid group.

Even though the number of binary variables and constraints is roughly the same, the runtime grows significantly when raising parameter k. Note that this test also has an interesting economic interpretation. For $k = 5$ there are 200 distinct local bidders whose market area contains 5 economic areas each. For this choice of k, the 200 bidders can express synergies between licenses of different economic areas as their small bid groups may contain bids for up to $k = 5$ EAs. In the treatment where $k = 1$ the same 200 bidders cannot express synergies. Instead, they must formulate individual bid groups for each EA. The FUEL bid generator encodes this by representing each one of the 200 bidders as five distinct bidders so that there are 1,000 local bidders competing in the auction for the treatment $k = 1$.

7.3. Unrestricted XOR

The FUEL bid language poses the restriction that small bid groups may only contain PEAs belonging to the same EA. A standard XOR bid language, however, is fully expressive, which means that bidders must be able to bid on packages of licenses belonging to any subset of PEAs (see Figure 4).

In the following test, we check whether an XOR bid language that only restricts bidders in the maximum number of admissible XOR bids without imposing any further restrictions is computationally tractable for the C-band auction. Similar to our previous tests, we assume that there are 10 nationwide and 1,000 local bidders. We generate the XOR bids according to Section 6.3.

For different configurations of the maximum number of XOR bids that local and nationwide bidders may submit in the auction, we generate 25 random test instances. Table 4 shows the maximum number of XOR bids that local and nationwide bidders may submit in the respective treatment (denoted z_L and z_N), the average total number of XOR bids present in the auction (denoted $\sum z$), the average number of binary variables and constraints in the winner determination problem, the average runtime in seconds, the number of test instances that exceed the time limit of 30 minutes, the maximum MIPGap for any test case in the respective treatment, and the average number licenses that remained unsold out of all 5,684 (406×14) available licenses in the auction.

z_L	z_N	$\sum z$	Binary Variables	Constr.	Runtime in sec.	TLE	Max MIPGap	Unsold Licenses
1	1	1,010	1,010	1,416	11	0 of 25	0.00010	651.1
2	2	2,020	2,020	1,416	200	2 of 25	0.00049	344.2
3	3	3,030	3,030	1,416	630	8 of 25	0.00175	218.8
5	5	5,050	5,050	1,416	977	13 of 25	0.00615	112.8
10	10	9,376	9,376	1,416	1,532	20 of 25	0.00124	43.3
25	25	22,402	22,402	1,416	1,381	18 of 25	0.00256	6.7

Table 4 Average values of 25 test instances for different restrictions on the number of XOR bids that both local and nationwide bidders are allowed to submit.

Result 3 *Even if both local and nationwide bidders are restricted to submit no more than two XOR bids without restrictions on the EAs, two out of 25 instances cannot be solved within the time limit of 30 minutes.*

Our tests indicate that because of the additional inter-dependencies between bids containing PEAs of different EAs, the corresponding winner determination problem becomes more complex. If all bidders are restricted to submit no more than 10 XOR bids, 20 out of the 25 instances cannot be solved within the time limit, even though this number of XOR bids is far too small to give a reasonable account of a nationwide bidder's preferences. Such limited bids also result in many licenses remaining unsold. When bidders are restricted to 5 bid groups, an average of 6.6 licenses remain unsold when using the FUEL bid language (see Table 1), while 112.8 licenses cannot be allocated when applying the XOR bid language, causing a significant revenue loss for the auctioneer. In this subsection, we did not report welfare losses due to the missing bids problem, because we are not able to solve the problem with all possible valuations submitted. However, we study the welfare loss compared to FUEL in the next subsection.

7.4. FUEL vs. XOR

The XOR bid language suffers from the missing bids problem and the fact that the winner determination problem quickly becomes intractable. On the other hand, the FUEL bid language makes some restrictions on the types of package bids that can be expressed. In particular, small bid groups can only be submitted within an EA. In this section, we want to analyze both effects. We analyze two scenarios. In the first one, we derive XOR bids from previously generated FUEL bid groups (see Appendix A) and analyze to what degree the missing bids problem affects the efficiency of the XOR bid language. In the second scenario, we derive FUEL bid groups from previously generated XOR bids (see Appendix B) and analyze the impact of the exposure risk of local bidders in the FUEL bid language on welfare.

7.4.1. XOR Efficiency Losses due to the Missing Bids Problem

To compare the efficiency of a fully combinatorial XOR bid language to a FUEL bid language, we need to ensure that we solve essentially the same problem instances with both the FUEL and XOR bid language. Therefore, we first generate a random FUEL instance and then derive XOR bids from the given FUEL bids as described in Appendix A. As a consequence, bidders in the XOR auction have the same valuations as in the FUEL auction, but they are only able to state a fraction of the potential winning FUEL bid combinations.

For our FUEL instances, we assume that local and nationwide bidders submit seven bid groups as our tests in Section 7.1 suggest that such instances can be solved within the time limit of 30 minutes. In order to keep the XOR allocation problem tractable, we restrict local and nationwide bidders in the maximal number of XOR bids they are allowed to submit and denote these upper bounds by z_L and z_N, respectively.

Table 5 shows the auction type, the maximum number of bids that local and nationwide bidders may submit (denoted z_L and z_N, respectively), the total number of submitted XOR bids (denoted $\sum z$), the average number of binary variables and constraints in the winner determination program, the average runtime in seconds to solve the allocation problem up to a MIPGap of 10^{-4}, the maximal MIPGap, the number of test cases that exceed the time limit of 30 minutes, the average number of unsold licenses out of 5,684 (406×14) licenses, and the average efficiency (denoted Eff.).

Type	z_L	z_N	$\sum z$	**Bin. Vars.**	**Constr.**	**Runtime in sec.**	**TLE**	**Max MIPGap**	**Unsold Licenses**	**Eff.**
FUEL	7	7	7,070	124,992	40,769	236	0 of 25	0.00010	4.0	1.000
XOR	1	1	1,010	1,010	1,416	0	0 of 25	0.00010	764.6	0.822
XOR	3	3	3,029	3,029	1,416	265	3 of 25	0.00015	376.0	0.913
XOR	5	5	5,010	5,010	1,416	478	5 of 25	0.00033	256.8	0.936
XOR	7	7	6,648	6,648	1,416	612	8 of 25	0.00125	204.4	0.945
XOR	10	10	8,718	8,718	1,416	629	8 of 25	0.00367	164.4	0.952
XOR	15	15	12,080	12,080	1,416	683	9 of 25	0.00662	124.9	0.957

Table 5 **Average values of 25 test instances for different limitations on the number of XOR bids that local and nationwide bidders are allowed to submit.**

Result 4 *If bidders are only allowed to submit the same number of bids in the XOR as in the FUEL auction, more than* 3.5% *of all licenses remain unsold, the welfare loss compared to FUEL is* 5.5%*, and there are already 8 out of 25 test instances that are intractable. Even when bidders are allowed to submit more than twice as many XOR as FUEL bids, still* 2.2% *of the licenses remain unsold and the welfare loss is* 4.3%*.*

At first sight, it might be surprising that the efficiency of the XOR bid language is still around 90% even though bidders can only state a fraction of their valuations. This is mainly due to the unequal population distribution among the PEAs. While more than 50% of the population live in the 25 most populous PEAs, the 50% least populous PEAs account for less than 10% of the population. As the value of a license block in a PEA is based on the population living in a PEA according to our value model (see Section 6.1), allocating the licenses in the 50% most populous PEAs already corresponds to serving 90% of the population. Thus, even though a large fraction of licenses remains unsold, the efficiency can still be considerably high when allocating licenses in the most populous PEAs. It is important to note that we can only solve the XOR instances because the XOR bids generated from FUEL reflect the hierarchical structure of EAs and PEAs such that many XOR bids of local bidders do not overlap across EA boundaries. With unrestricted XOR bids, we would suffer from significantly higher MIP gaps as can be seen in Section 7.3.

7.4.2. FUEL Efficiency Losses due to the Exposure Problem When bidders must bid separately for different EA-packages without including them in a single package bid, that introduces an exposure problem: a bidder could win licenses in some EAs without winning the package it needs for a sensible business plan. The next experiments explore the magnitude of the losses this imposes when bidders have values for 10 or 25 XOR packages and zero for all other packages. This raises a question about how to generate FUEL bids when the actual packages are XOR ones. We do this by converting each XOR bid into one or more FUEL bid groups (in those cases where an XOR bid includes PEAs from multiple EAs, see Appendix B).[8] This naive conversion leads to a proliferation of FUEL bid groups and to a consequent failure to solve some FUEL instances to optimality within 30 minutes. However, even when the solution process is terminated after 30 minutes, the maximum MIP gap across the 25 instances is low. Although we cannot compute the XOR winner determination problem with 10 bids per bidder to optimality (see Section 7.3), our data reports assume that the solution with XOR bids is optimal and measure the efficiency loss from the FUEL exposure problem relative to that value. Although this environment is biased against FUEL bid groups, the average efficiency loss in our computations is only around 2% when every bidder submits 25 package bids, even with high synergies. This happens because a certain amount of the package bids of local bidders in our model is within an EA, and for all of these packages bidders can express their synergies appropriately.

As discussed further in Appendix B, when deriving FUEL bid groups from XOR bids, we mitigate the exposure problem by reducing the base price of the corresponding FUEL bid groups by a fixed percentage, which we call the *synergy level*. We consider three different synergy levels in our experiments, namely 30%, 50%, and 70%, using the respective XOR test instances from Section 7.3 as baseline when calculating the efficiency of the derived FUEL instances.

Table 6 lists the synergy level, the number of XOR bids that local and nationwide bidders submit (denoted z_L and z_N, respectively), the average total number of FUEL bid groups derived from these XOR bids, the average number of binary variables and constraints in the FUEL winner determination problem, the average runtime in seconds to solve the FUEL allocation problem up to a MIPGap of 10^{-4}, the maximal MIPGap of any test instance, the number of test instances that exceed the time limit of 30 minutes, the average number of unsold licenses out of 5,684 (406×14) licenses, and the average efficiency (denoted Eff.) with respect to the XOR instances of Section 7.3.

Synergy Level	z_L	z_N	Bid Groups	Bin. Vars.	Constr.	Runtime in sec.	TLE	Max MIPGap	Unsold Licenses	Eff.
30%	10	10	17,371	81,135	66,689	143	1 of 25	0.00058	3.5	0.968
30%	25	25	42,971	202,239	162,204	765	9 of 25	0.00715	0.2	0.983
50%	10	10	17,371	81,135	66,689	370	4 of 25	0.00242	4.9	0.961
50%	25	25	42,971	202,239	162,204	1,057	14 of 25	0.00979	0.2	0.981
70%	10	10	17,371	81,135	66,689	604	8 of 25	0.00369	5.8	0.958
70%	25	25	42,971	202,239	162,204	1,124	15 of 25	0.01030	0.2	0.981

Table 6 Average values of 25 test instances for different synergy levels and number of XOR bids.

[8] One could compute the minimal set of FUEL bid groups required from the XOR bids of a bidder, which would be a topic beyond this analysis.

Result 5 *If local bidders submit* 25 *XOR bids each, the efficiency loss of the FUEL auction compared to the unrestricted XOR bid language is less than* 2% *even if bidders are assumed to have strong synergistic valuations. With* 10 *XOR bids per bidder, the efficiency loss is at around* 4%*. This comparison assumes a naive conversion of XOR bids into FUEL bid groups and assumes that bidders have only* 10 *or* 25 *XOR bids and no value for any other combination, which zeroes out the missing bids problem. Increasing the number of admissible XOR bids per bidder has a positive effect on the efficiency of the derived FUEL instances.*

7.5. FUEL vs. C-band Clock Auction

We also compare the efficiency of the FUEL auction design to the clock auction format, which was the format used for the FCC's C-band auction. As discussed in Appendix C and illustrated in Example 2, bidders might be allocated a set of licenses in the clock auction that does not match any of the packages specified in their FUEL bid groups, that is, licenses for areas that are geographically scattered. In such cases, the bidders' synergy effect for the allocated set of licenses can be expected to be reduced or missing. We reflect this in our tests by reducing a bidder's valuation for such scattered packages by a fixed percentage, which we denote as synergy level in the tests below. More specifically, we first determine the underlying value of the packages using our value model and then reduce that value by either 30%, 50%, or 70% depending on the test scenario. For completeness, we also run our experiments for the 0% scenario in which bidders are assumed to have no synergies and, therefore, do not reduce their values for scattered license packages. For the four different synergy levels, we analyze the efficiency for different limits on the number of admissible FUEL bid groups.

For all our treatments, we use the same FUEL bid groups that we generated for our tests in Section 7.1 and derive the clock bids from them as described in Appendix A. When running the simultaneous clock auction, we apply the default values for the bid increment percentage (10%), activity requirement percentage (95%), and upper activity limit (120%). Table 7 shows the maximum number of FUEL bid groups that local and nationwide bidders may submit (denoted z_L and z_N, respectively), the average number of clock auction rounds, the average number of licenses allocated to bidders whose allocated set of licenses does not match any of their FUEL packages, and the average efficiency for synergy levels of 0%, 30%, 50%, and 70% compared to the respective FUEL treatments from Section 7.1.

Result 6 *If bidders are restricted to a single FUEL bid group and reduce the values for license packages not represented by their FUEL bid group by only* 30%*, the efficiency loss of the clock auction design is already larger than* 10% *and more than* 41% *of the licenses are allocated to bidders who did not desire them. Raising the number of admissible FUEL bid groups above one leads to*

z_L	z_N	Clock Rounds	Undesired Licenses	Efficiency Synergy 0%	Efficiency Synergy 30%	Efficiency Synergy 50%	Efficiency Synergy 70%
1	1	15.6	2,348	1.001	0.891	0.818	0.745
3	3	15.8	3,290	0.997	0.839	0.733	0.628
5	5	16.7	3,965	0.996	0.798	0.667	0.536
7	7	16.4	3,653	0.995	0.812	0.690	0.568
10	10	16.6	3,943	0.995	0.798	0.667	0.536
15	15	16.6	3,763	0.995	0.804	0.678	0.551

Table 7 Average values of 25 test instances for different limitations on the number of admissible FUEL bid groups per bidder, and different synergy levels for scattered license packages.

higher efficiency losses and a larger fraction of licenses being awarded to bidders who did not desire them at the final clock prices.

If bidders are assumed to have no synergies for winning licenses in neighboring regions (i.e., the 0% synergy level scenario), then the efficiency of the clock auction design is close to the FUEL auction design. In fact, as bidders may end up winning license packages in the clock auction which are not included in their FUEL bid groups, the efficiency of the clock auction may even slightly exceed the one of FUEL. However, in spectrum sales, some bidders appear to have strong synergies between license blocks in neighboring geographic areas. The clock auction does not allow bidders to express such synergies as only item-level bids are permitted. Our tests show that in settings with large complementarities this restriction leads to substantial welfare losses compared to the FUEL auction design where package bids are permitted.

When increasing the number of admissible FUEL bid groups, bidders may try to bid on packages associated with different FUEL bid groups during different auction rounds. As the set of PEAs contained in different FUEL bid groups may differ, a bidder's risk to win a set of licenses not aligned with any FUEL package grows. This explains the efficiency losses that we witness in our tests when raising the number of permitted FUEL bid groups per bidder above one.

7.6. Summary

Let us briefly summarize the main findings from our extensive computational experiments. First, with FUEL we can solve problems with 10 nationwide and 1,000 local bidders each submitting seven bid groups within 237 seconds to optimality. These problem sizes are larger than what we can expect in the USA and other countries. In contrast, if local bidders instead submit three (unrestricted) XOR package bids, then 8 out of 25 test instances cannot be solved within the time limit. Obviously, a restriction to three XOR bids would result in a huge missing bids problem with substantial efficiency losses. In contrast, suppose a local bidder submitted a FUEL bid group with 14 adjustments to the base package in each of 10 PEAs. This would describe the valuation for almost 576 billion (15^{10}) package valuations. With only four adjustments for the 10 PEAs, a bidder would already cover close to 10 million package valuations (5^{10}).

In the tests just described, the locals are limited to submit bid groups within an EA. The efficiency loss from these restrictions depends on the synergies assumed. Even under the unrealistic assumption that bidders submitted 25 XOR bids and had no value for any other package, the efficiency loss due to the EA-restriction is only in the order of 2%, even when synergies account for 70% of the package value.

For the simultaneous clock auction, inefficiencies in settings with large complementarities result from bidders not being able to submit package bids. We show that depending on the synergies the efficiency loss in the simultaneous clock auction can be substantial, even if we assume truthful revelation of the payoff-maximizing package in each round. As a result, sealed-bid auctions with an XOR bid language as well as the simultaneous clock auction risk substantial welfare losses in large-scale spectrum sales where bidders exhibit high synergies between the licences of neighboring regions, while FUEL provides a tractable alternative even with many licenses available in each area.

8. Conclusion

The design of combinatorial auctions with hundreds of items is challenging. In large spectrum auctions, the XOR bid language is the standard and has rarely been debated in the auction design. In this paper, we investigate the computational hardness of the FUEL bid language based on the case of the planned C-band auction for the US, which constitutes an important real-world case. Even though the winner determination problem of the FUEL bid language is NP-hard and contains roughly 125,000 binary variables and 40,000 constraints, our experiments indicate that this auction can consistently be solved in less than 30 minutes, and usually much less. We find evidence that the short solution times predominantly result from the hierarchical structure created by FUEL, which allows the optimization algorithms to decompose the binary program effectively, in which combinations of bid groups serve as messages to decompose the problem. This likely explains our empirical finding that limiting the number of bid groups that local and nationwide bidders may submit and restricting the maximal number of economic areas that local bidders can include in a small bid group effectively reduces the empirical hardness of the allocation problem.

In our experiments, in contrast to FUEL, a fully enumerative XOR bid language quickly becomes computationally intractable. As is well-known, bidders would need to specify an exponentially large set of XOR bids to express the same preferences as in a FUEL bid group with adjustments. Although the FUEL bid language is not fully expressive and limits the set of values that can be expressed relative to an XOR bid language, it is based on common spectrum valuation methods and may often be able to express values close to the bidders' actual ones. To the extent that FUEL bids fail to capture actual values, that loss must be weighed against FUEL's mitigation of the missing bids problem that inevitably arises in large auctions using XOR bids. Our experiments show that both the missing bids problem and computation failures using an XOR bid language can lead to significant welfare losses.

When comparing FUEL to the simultaneous clock auction format that was eventually used by the FCC for the C-band auction, the allocative efficiency of the clock auction in our numerical tests is substantially lower. This is predominantly due to the bidders' inability to submit package bids in the clock auction, preventing them to express their assumed strong synergies between license blocks of geographically neighbored areas accurately.

In summary, by allowing bidders to use bid groups with adjustments to their base bids, the FUEL bid language gives bidders an intuitive and compact way to describe their valuations and effectively address the missing bids problem. The hierarchical structure of the bid groups makes it possible to solve very large problem instances exactly on a desktop computer in a matter of minutes. The specifics of the bid language allow for exact solutions in large-scale auctions with several hundred items, which would recently have been considered intractable.

As always, there are some limitations. FUEL assumes a hierarchical structure of the goods, with economies of scale within product categories or regions and economies of scope across them. Our experiments further assume that local bidders (active only in one region or product category) compete against nationwide bidders (active in many regions or product categories). We have argued that many auction applications appear to exhibit such a value structure. Many bidders may find this structure intuitive and, mathematically, it allows decomposition of the problem for each region or product category, conditional on the message of which bid groups are winning.

The FUEL language emphasizes the tradeoff between the structure in the bid language and the ability to compute exact solutions in large auctions. The simultaneous multi-round auction was originally introduced despite its problematic bidding language because it was thought to enable more attention to packages than a sequence of single item auctions. By taming the communication and computational complexity of large-scale combinatorial auctions, FUEL is the first bid language to enable practical combinatorial spectrum auctions in the presence of many products with scale and scope economies, where SMRA based clock auctions have until now been the only practical solution.

Acknowledgments

The authors thank Mitchell Watt and Zi Yang Kang for their valuable comments and suggestions on the manuscript. We also thank Burak Öz for the graphical illustrations. The financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/1-8) is gratefully acknowledged.

References

Arora S, Barak B (2009) *Computational complexity: a modern approach* (Cambridge University Press).

Ausubel LM, Baranov OV (2014) Market design and the evolution of the combinatorial clock auction. *American Economic Review* 104(5):446–51.

Ausubel LM, Cramton P, Milgrom P (2006) The clock-proxy auction: A practical combinatorial auction design. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial Auctions*, chapter 5, 115–138 (MIT Press).

Bichler M, Davenport A, Hohner G, Kalagnanam J (2006) Industrial procurement auctions. *Combinatorial auctions* 593–612.

Bichler M, Goeree J, Mayer S, Shabalin P (2014) Spectrum auction design: Simple auctions for complex sales. *Telecommunications Policy* 38(7):613–622.

Bichler M, Goeree JK (2017) *Handbook of spectrum auction design* (Cambridge University Press).

Bichler M, Hao Z, Adomavicius G (2017) Coalition-based pricing in ascending combinatorial auctions. *Information Systems Research* 28(1):159–179.

Bichler M, Paulsen P (2018) A principal-agent model of bidding firms in multi-unit auctions. *Games and Economic Behavior* 111:20–40.

Bichler M, Schneider S, Guler K, Sayal M (2011) Compact bidding languages and supplier selection for markets with economies of scale and scope. *European Journal of Operational Research* 214(1):67–77.

Bichler M, Shabalin P, Ziegler G (2013) Efficiency with linear prices? a game-theoretical and computational analysis of the combinatorial clock auction. *Information Systems Research* 24(2):394–417.

Boutilier C, Hoos HH (2001) Bidding languages for combinatorial auctions. *International Joint Conference on Artificial Intelligence*, volume 17, 1211–1217.

Conen W, Sandholm T (2001) Preference elicitation in combinatorial auctions. *Proceedings of the 3rd ACM conference on Electronic Commerce*, 256–259.

Cramton P (2017) Electricity market design. *Oxford Review of Economic Policy* 33(4):589–612, ISSN 0266-903X, URL `http://dx.doi.org/10.1093/oxrep/grx041`.

Day RW, Cramton P (2012) Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* 60(3):588–603.

Day RW, Raghavan S (2007) Fair payments for efficient allocations in public sector combinatorial auctions. *Management science* 53(9):1389–1406.

Downey RG, Fellows MR (2012) *Parameterized complexity* (Springer Science & Business Media).

Eilat A, Milgrom P (2011) The CAF auction: Design proposal. Technical report, URL `https://milgrom.people.stanford.edu/sites/g/files/sbiybj4391/f/milgrom_eilat_auction_paper_july_25_2011.pdf`.

Federal Communications Commission (2020a) Auction 107 clock phase technical guide. Technical report, Federal Communications Commission, URL `https://www.fcc.gov/file/19165/download`.

Federal Communications Commission (2020b) Auction of flexible-use service licenses in the 3.7-3.98 GHz band for next-generation wireless services; notice and filing requirements, minimum opening bids, upfront payments, and other procedures for auction 107; et al. Technical report, Federal Communications Commission, URL `https://docs.fcc.gov/public/attachments/FCC-20-110A1.pdf`.

Fréville A (2004) The multidimensional 0–1 knapsack problem: An overview. *European Journal of Operational Research* 155(1):1–21.

Garey MR, Johnson DS (1979) *Computers and intractability*, volume 174 (Freeman).

Goeree JK, Holt CA (2010) Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior* 70(1):146–169.

Goetzendorff A, Bichler M, Shabalin P, Day RW (2015) Compact bid languages and core pricing in large multi-item auctions. *Management Science* 61(7):1684–1703.

Gurski F, Rehs C, Rethmann J (2019) Knapsack problems: A parameterized point of view. *Theoretical Computer Science* 775:93–108.

Iftekhar MS, Tisdell J (2012) Comparison of simultaneous and combinatorial auction designs in fisheries quota market. *Marine Policy* 36(2):446–453.

Kelly F, Steinberg R (2000) A combinatorial auction with multiple winners for universal service. *Management Science* 46(4):586–596.

Kroemer C, Bichler M, Goetzendorff A (2017) About inconsistent bidding and its impact on efficiency in the combinatorial clock auction. Bichler M, Goeree JK, eds., *Handbook of Spectrum Auction Design*, 338 (Cambridge University Press).

Kulik A, Shachnai H (2010) There is no EPTAS for two-dimensional knapsack. *Information Processing Letters* 110(16):707–710.

Lavi R, Swamy C (2011) Truthful and near-optimal mechanism design via linear programming. *Journal of the ACM (JACM)* 58(6):1–24.

Lehmann D, Müller R, Sandholm T (2006) The winner determination problem. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial Auctions*, chapter 12, 297–318 (MIT Press).

Lehmann D, Oćallaghan LI, Shoham Y (2002) Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM (JACM)* 49(5):577–602.

Levin J, Skrzypacz A (2016) Properties of the combinatorial clock auction. *American Economic Review* 106(9):2528–51.

Milgrom P (2004) *Putting auction theory to work* (Cambridge University Press).

Milgrom P (2009) Assignment messages and exchanges. *AEJ-Micro* 1(2):95–113.

Mitchell JE (2002) Branch-and-cut algorithms for combinatorial optimization problems. *Handbook of applied optimization* 1:65–77.

Nisan N (2006) Bidding languages. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial auctions*, chapter 9, 400–420 (MIT Press).

Nisan N, Ronen A (2001) Algorithmic mechanism design. *Games and Economic behavior* 35(1-2):166–196.

Nisan N, Segal I (2006) The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129(1):192–224.

Papavasiliou A, Papalexopoulos AD, Oren S (2017) Market design for a decarbonized electricity market.

Parkes D (2006) Iterative combinatorial auctions. Cramton P, Shoham Y, Steinberg R, eds., *Combinatorial Auctions*, chapter 2, 41–78 (MIT Press).

Pekec A, Rothkopf MH (2003) Combinatorial auction design. *Management Science* 49(11):1485–1503.

Porter D, Rassenti S, Roopnarine A, Smith V (2003) Combinatorial auction design. *Proceedings of the National Academy of Sciences* 100(19):11153–11157.

Sandholm T (2002) Algorithm for optimal winner determination in combinatorial auctions. *Artificial intelligence* 135(1-2):1–54.

Sandholm T, Boutilier C (2006) Preference elicitation in combinatorial auctions. *Combinatorial auctions* 10.

Scheffel T, Ziegler G, Bichler M (2012) 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.

Varnamkhasti MJ (2012) Overview of the algorithms for solving the multidimensional knapsack problems. *Advanced Studies in Biology* 4(1):37–47.

Appendix A: Deriving XOR Bids from FUEL Bid Groups

In an XOR bid language, bidders are unable to specify adjustments. If a bidder wants to state the same information as in a FUEL bid, she has to place one XOR bid for each possible combination of adjustments and adapt the price of the XOR bid according to the chosen markups and discounts (see Figure 5).

SMALL	**Base price: 800**			
		#Licenses		
EA	**PEA**	**3**	**4**	**5**
7	44	Base	400	750
7	271	-50	Base	

$\longrightarrow$

XOR Price: 1,500		
EA	**PEA**	**#Licenses**
7	44	5
7	271	3

Figure 5 A random XOR bid is generated from a FUEL bid by picking a random adjustment combination. In this example the chosen adjustment combination is highlighted blue. The price for the XOR bid is given by the base price of the FUEL bid adjusted by the selected increment and decrement, i.e., $\$800 + \$750 - \$50 = \$1{,}500$**.**

Unfortunately, the number of XOR bids necessary to reproduce a FUEL bid group can become very large. According to our value model, a bidder places (in addition to her base bid) on average two adjustments for each PEA contained in her bid group. If a nationwide bidder submits a bid group containing all 406 PEAs, this means that there would be $3^{406} \approx 5.14 \cdot 10^{193}$ XOR bids necessary to state the same valuations as in a FUEL bid group.

Due to this vast amount of adjustment combinations, it is impossible for bidders to express their assumed valuations to the same degree of accuracy in the XOR bid language as in the FUEL bid language. In fact, in order to guarantee computational tractability of the XOR allocation problem, it is indispensable to restrict the number of XOR bids that any bidder is allowed to submit.

To compare the FUEL and XOR bid language in terms of efficiency and runtime we proceed as follows. For every bidder, we first generate FUEL bid groups according to Section 6.2. For each XOR bid to be generated, we select one of the bidder's FUEL bid groups uniformly at random, pick a random adjustment combination, and finally set the price of the XOR bid to be equal to the implied FUEL bid price, incorporating the markups and discounts for the chosen adjustments (see Figure 5).

Appendix B: Deriving FUEL Bid Groups from XOR Bids

When deriving FUEL bid groups from XOR bids, one has to distinguish three cases: (1) the XOR bid only contains bids for PEAs belonging to the same EA, (2) the licenses specified in the XOR bid exceed the MHz-pop threshold that is necessary for a FUEL bid group to be considered large, and (3) the XOR bid contains PEAs belonging to different EAs but the demanded licenses do not exceed the MHz-pop threshold necessary for a FUEL bid group to be considered large.

In cases (1) and (2) the XOR bid can be transformed straightforwardly to a FUEL bid group as follows: The base package of the FUEL bid group contains the same set of licenses specified in the XOR bid, the base price of the FUEL bid group equals the price of the XOR bid and the FUEL bid group is labeled small or large depending on whether it corresponds to case (1) or (2).

However, XOR bids corresponding to case (3) have to be treated carefully. As the MHz-pop of such an XOR bid does not exceed the MHz-pop threshold, the XOR bid cannot be transformed to a large FUEL bid group. Transforming it to a single small FUEL bid group is also infeasible as small FUEL bid groups may only contain bids for a single EA. Therefore, the only possibility is to express the XOR bid through multiple independent small FUEL bid groups, one for each EA contained in the XOR bid (see Figure 6).

The bidder does not have control over the subset of bid groups that may eventually be accepted by the auctioneer. Thus, the bidder may end up winning a set of licenses that she did not originally specify in one of her XOR bids. This exposure risk has to be taken into account when determining the base prices of the derived bid groups. We reflect such potential efficiency losses by reducing the base price of such small FUEL bid groups by a fixed percentage which we refer to as *synergy level*. As the degree of synergies depends on the domain, we consider several different synergy levels in our tests in Section 7.4.

Appendix C: Deriving Clock Bids from FUEL Bid Groups

To compare the FUEL bid language to the clock auction design adopted for the C-band auction, we first generate suitable FUEL bid groups for each bidder and then derive clock bids from them in each round of the clock auction. A distinctive difference between the FUEL bid language and the clock auction format is

XOR Price: 530			
EA	**PEA**	**#Licenses**	**License Value**
6	41	3	420
6	210	2	70
7	271	2	40

FUEL SMALL Base price: 343			
		#Licenses	
EA	**PEA**	**2**	**3**
6	41		Base
6	210	Base	

FUEL SMALL Base price: 28		
		#Licenses
EA	**PEA**	**2**
7	271	Base

Figure 6 The XOR bid can be expressed by two separate FUEL bid groups. The fourth column of the XOR bid contains the underlying valuations w.r.t. the sigmoid value model. If the synergy level is set to 30%, then the base price of the first FUEL bid group equals $(1-0.3)\cdot(\$420+\$70)=\$343$.

that FUEL allows bidders to submit bids on entire packages of products, while the bids of the clock auction only refer to a single product. Moreover, while a FUEL bid group is entirely accepted or rejected, a bidder may win some but not all of her bids in the clock auction and might be forced to purchase more license blocks for a PEA than specified by her bids (see Example 2).

For simplicity, let us at first assume that each bidder submits only a single FUEL bid group. (We will treat the case where bidders submit multiple FUEL bid groups later on.) Given arbitrary prices for the products, a bidder is able to identify the FUEL adjustment combination that maximizes the bidder's utility at the given prices while requiring fewer bidding units than the bidder's eligibility. Such an adjustment combination for a bid group g of bidder i can be determined by maximizing the following integer program. We reuse the notation from Section 3.3 and add a few more parameters:

c_p Price for a single license block in PEA $p \in P$.
u_p Bidding units associated with a license block in PEA $p \in P$.
e_i Maximum number of bidding units for which bidder $i \in I$ is allowed to submit bids.

$$\max \quad \omega_i^g + \sum_{p\in P_i^g} \sum_{k\in K_i^{gp}} \left(y_i^{gpk}\left(\mu_i^{gpk} - k\, c_p\right)\right) \tag{1}$$

$$\text{s.t.} \quad \sum_{k\in K_i^{gp}} y_i^{gpk} = 1 \qquad \forall p\in P_i^g \tag{2}$$

$$\sum_{p\in P_i^g} \sum_{k\in K_i^{gp}} \left(y_i^{gpk}\, k\, u_p\right) \le e_i \tag{3}$$

$$y_i^{gpk} \in \{0,1\} \qquad \forall p\in P_i^g, \forall k\in K_i^{gp} \tag{4}$$

We assume that all bidders bid straightforwardly throughout the auction, i.e., in every auction round each bidder expresses her utility-maximizing FUEL package by a collection of clock bids. Depending on the bids submitted in a clock round, the price of a product may take on any value between its start-of-round and clock price. A bidder does not have control over the bids placed by her competitors so that she does not know ahead of the bid processing which prices the products may take on. Therefore, the bidder solves the upper integer program not only at the start-of-round and clock prices but also at the 10%, 20%, . . . , 90% price points. In case her demand for a product changes with respect to a previous price point, the bidder places a suitable clock bid. Evaluating the optimization problem at 11 distinct price points allows the bidder to resemble the FUEL bid sufficiently close.

At some point during the clock auction, the prices might get so high that the bidder's utility for any adjustment combination becomes negative. In such a situation, the bidder no longer demands any licenses and therefore specifies a demand of zero licenses for each product. As discussed in Example 2, bids to reduce demand will only be processed by the auctioneer if they do not create or increase oversupply. As a result, a bidder may end up with a set of licenses not resembling any adjustment combination specified in her FUEL bid group. Especially for nationwide bidders, this is very undesirable because they might end up winning a geographically very scattered set of licenses. Of course, such scattered license packages are of less value to a bidder than one of the FUEL packages for which the bidder has synergies. Therefore, we assume for our tests in Section 7.5 that a bidder's valuation for any set of licenses not representing any of her FUEL adjustment combinations is reduced by a fixed percentage.

So far we have focused on the situation where each bidder only submits a single FUEL bid group. Whenever bidders submit multiple bid groups, the utility-maximizing FUEL package may correspond to different bid groups at different price points during a clock round. However, swapping the demand from one FUEL package to another has two major drawbacks for a bidder. First of all, it increases the probability that a bidder ends up with a set of licenses that does not correspond to any FUEL package since FUEL bid groups may contain different sets of PEAs. Secondly, it may create a situation where the bidder's demand for a product oscillates during a clock round. However, according to the rules specified by the FCC, a bidder's demand trajectory for a product within a clock round must be monotonic. Therefore, in situations where bidders submit multiple bid groups, we assume that a bidder first determines the FUEL bid group that maximizes her utility at clock prices, and then only submits clock bids with respect to this single bid group during a clock round.