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

Martin Bichler[1], Paul Milgrom[2], Gregor Schwarz[1]
Technical University of Munich[1]
Stanford University and Auctionomics[2]

March 14, 2020

Abstract

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, 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 introduce a mathematical model of the resulting allocation problem, 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 sizes 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. For the XOR bid language, the missing bids problem and computational hardness incur significant welfare losses compared to the FUEL bid language.

Contents

1 Introduction **3**

1.1 Contributions 4

1.2 Outline 6

2 Bid Languages for Spectrum Auctions **7**

3 FUEL Auction Design **9**

3.1 Product Design 9

3.2 Bid Language 10

3.3 Sell-Side Reserve Bids 12

3.4 Auction Process 12

3.5 Winner Determination Problem 14

4 XOR Bid Language **17**

4.1 Bid Language 17

4.2 Winner Determination Problem 18

5 Experimental Design **19**

5.1 Value Model 19

5.2 FUEL Bid Generation 20

5.3 Deriving XOR from FUEL Bids 22

5.4 XOR Bid Generation 23

6 Results **23**

6.1 FUEL Bid Groups 24

6.2 FUEL Admissible EAs 25

6.3 FUEL vs. XOR 27

6.4 Unrestricted XOR 29

7 Conclusion **30**

1 Introduction

One of the thorniest problems in the theory of resource allocation concerns how to allocate a fixed set of resources efficiently when there is minimal structure on the potential buyers' possible values. The neoclassical economic theory of markets focuses on special cases with convex preferences and constraints, finding that the efficient allocation is supported by prices and that computing the optimum is fast. When there is no such structure and the resources are indivisible, a common recommendation is to use a combinatorial auction with XOR bidding, in which bidders simply enumerate values for all the possible combinations of items. The auctioneer then solves a combinatorial optimization problem to find the allocation that maximizes the total bid. Such auctions have not only attracted the interests of researchers, they 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 can be 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).

We focus on the design of large markets with hundreds of items and bidders and use the recent plans for a private C-band auction as a case for our analysis. In mid-2019, a consortium of companies providing satellite downlink for commercial television in the United States proposed to conduct a private sale of their C-band spectrum rights using a novel combinatorial auction design dubbed FUEL (*Flexible Use and Efficient Licensing*), which was intended to overcome both the communication and computational complexity issues (Milgrom 2019).[1] 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}. The FUEL bid language, described below, was introduced in an attempt to

[1] In February 2020, the US Federal Communications Commission decided against using a private auction, so the FUEL design will not be adopted for the C-band auction.

tame both the communication and computational complexity by requiring bidders to use a particular parameterized valuation model to describe values for all license combinations or for certain subsets of them.[2] This paper explains the rationale for the language and tests and compares the computational tractability of the FUEL and XOR designs assuming a high level of participation in the auction by bidders with similar license values. High participation and similar values are thought to make the optimization more challenging by providing more near-optimal combinations of bids for the software to rule out.

A combinatorial auction for an application like the C-band with an XOR bid language would not only require bidders to omit all but a minute fraction of the relevant combinations, but would suffer another problem as well. Our simulations of the auction 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. However, our computational experiments show 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. Moreover, in our experiments, computation time tends to grow only linearly in the number of package bids.

1.1 Contributions

How does FUEL work and why are its optimizations so fast? The main innovation in FUEL is its bid language, which 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 adjustments within each PEA are general and do not simplify the computations: the simplification comes 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, if the licenses in each PEA were all identical, then 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

[2] The design of parametric bid languages for various auction problems has previously received scholarly attention in Milgrom (2009), Bichler et al. (2011), Eilat and Milgrom (2011), Bichler et al. (2017).

the communication complexity compared to XOR bidding.

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.

For one of the national mobile network operators in the United States, nearly complete nation-wide coverage is important, so the value of a collection of licenses with, say, 40MHz 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 40MHz in every PEA can reflect that value pattern and, if 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.[3]

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 national 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, and a non-national bidder can have a limited number of mutually exclusive bid groups for each Economic Area. There are 176 Economic Areas (EAs) and these partition the set of PEAs, with an average of 2.5 PEAs in a typical EA.

The purpose of these limitations is highlighted by the special case in which there are only local bidders, in which the optimal solution could be computed by separately solving each EA to optimality. Although these problems are also NP-hard, they are small because the EAs

[3]Additional evidence that FUEL may be a good description comes from the package bids submitted in the 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.

are small. The presence of bid groups across EAs by the national bidders makes the problem more complex, but there are only a few of these large national bidders. As every national bidder may win at most one nationwide bid group, the number of winning nationwide bid groups can be upper bounded by the number of national bidders present in the auction. For our experiments we assumed that there are 10 national bidders, which is more than there are in the United States currently. While the details of the search algorithms of commercial-off-the-shelf solvers are not known, we show experimentally that the algorithms find effective cuts and exploit this 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, which explains the only linear increase in runtime that we observe with increasing numbers of bids.

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 1000 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.[4]

The FUEL design may also serve useful as a template for other large combinatorial auctions in which there are several kinds of items with economies of scale and scope among them. In such cases, by exploiting computationally tractable hierarchical valuation structures, some of the most important barriers to large-scale combinatorial auctions may be overcome. Auctions with regional lots as in the procurement of school meals (Kim et al. 2014) and the sale of fishery access rights (Iftekhar and Tisdell 2012) could be appropriate candidates, but so could be TV ad auctions (Goetzendorff et al. 2015).[5]

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

[4] We use a time limit of 30 minutes for each problem instance in our experiments. If problems could not be solved to near-optimality within 30 minutes or if there was a large integrality gap after 30 minutes, they typically could also not be solved to optimality or near-optimality with several hours of computation time.

[5] The bidding language is just one element of a sealed-bid auction design. Two others are the winner determination rule and the payment rule, which together determine bidders' incentives. If a hierarchical language like FUEL language makes optimization tractable, then that would allow selecting the allocation that maximizes the total bid and setting payments using the Vickrey-Clarke-Groves payment rule, which is incentive-compatible (Green and Laffont 1977).

auctions world-wide. Section 5 describes the experimental design. For our computations, we simulate an auction with more bids and participants than are expected for the C-band auction in the United States. In the experimental results described in Section 6, the optimal solution using FUEL could be found in minutes, while optimization based on the fully combinatorial XOR bid language was intractable even for much smaller problems. Section 7 summarizes key insights of the paper.

2 Bid Languages for Spectrum Auctions

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 for 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.

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. To overcome that limitation, the OR* bid language has been proposed, which allows including dummy items in the bid bundles. The dummy items improve bidders' ability to express either-or constraints. For example, suppose that there are two items A and B and a dummy item D. A bidder that submits OR bids for the bundles AD and BD can never win both, which safeguards against winning the package AB. In principle, with an unlimited number of dummies, the OR* bid language can be just as expressive as XOR and that expressiveness can be achieved with many fewer bids for cases in which the additive valuation structure applies well.

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.

Spectrum auctions have so far resisted the design of particular bid languages and, until recently, only XOR bid languages have been used, leading to significant limitations in applications like the Canadian auctions in 2014 with 100 licenses (Kroemer et al. 2017). The even larger numbers of licenses available in spectrum auction in the United States are one reason why combinatorial auctions have not yet adopted there. On the one hand, the missing bids problem in such auctions is huge and the auctioneer cannot confidently solve a large winner determination problem with hundreds of licenses and an XOR bid language. On the other hand, using simple auction formats such as the simultaneous multi-round auction and related clock auctions limits expressiveness of the bids and creates an *exposure problem* for bidders, in which they may win some but not all of the licenses they need for a viable network. Similar issues arise in the UK and other large spectrum markets.

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).[6] 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-Clark-Groves 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 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, however, 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 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.

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

[6] 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.

per PEA. All license blocks can be made available within 36 months of the FCC's final order. Furthermore, in 46 of the 50 most populous PEAs the satellite companies are able to provide 100 MHz (5 blocks) of spectrum even earlier, within 18 months of the FCC's approval of the auction process. With respect to these temporal constraints, licenses for spectrum blocks that are available within 18 months are referred to as *early*, the remaining licenses are called *late*. It is possible to serve a bidder's demand for late blocks with available early blocks, but the reverse is not possible.

3.2 Bid Language

Assuming the C-band auction is executed 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* and incorporates *adjustments*, defining 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 1). 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. Each bidder may be permitted to submit a small number of bid groups.

Bidder 1 LATE SMALL Base price: 177																
		#Licenses														
EA	**PEA**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
60	155		-90	Base	101	174										
60	354			Base												

Figure 1: 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 $177 + $174 + $0 = $351.

In 46 of the PEAs, 5 of the 14 spectrum blocks are available early and bidders may wish to bid more for those blocks. The FUEL bid language addresses this issue by allowing each bidder to designate its bid groups as bids for *early/mixed* spectrum or for *late* spectrum and interpreting the base bids and adjustments differently for each specification.

A late bid group works exactly as described above, with the interpretation that its prices apply to only to late spectrum. The seller, at its sole discretion, can convert early spectrum to late spectrum by delaying rights to use that spectrum.

For an early/mixed bid group, the base bid is interpreted as follows. If the base package contains a PEA providing 5 early and 9 late licenses and a bidder states a demand for k licenses in her base bid, then this is interpreted as demand for $\min\{5, k\}$ early licenses and $\max\{0, k-5\}$ late licenses. Increments in the bid group constitute demand for additional late license blocks, while decrements from the base package first reduce the number of early licenses in the package and, if the reduction takes the number below zero, then apply to reduce the number of late licenses (see Figure 2).

Bidder 2 EARLY SMALL Base price: 2,869																
		#Licenses														
EA	**PEA**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
53	23			-1,879	-1,038	Base	616									
53	114				-131	Base										

Figure 2: Here, bidder 2 designates a bid for an early/mixed bid group. In PEA 23 there are 5 early and 9 late licenses for sale, in PEA 114 all 14 licenses are late. Her base bid of 4 licenses in PEA 23 correspond to 4 early license. Her decrements define adjustments for packages with 3 and 2 early licenses, respectively. If she is allocated 5 licenses (4 early and 1 late) in PEA 23 and 4 (late) licenses in PEA 114, her valuation for this set of licenses is \$2,869 + \$616 + \$0 = \$3,485.

This last restriction — that increments on the base bid must always be served with late licenses while decrements first reduce the number of early licenses present in the base bid before diminishing the number of late licenses — may at first seem surprising. However, this design decision increases the expressiveness of the bid language as Example 1 illustrates.

Example 1. *Consider two PEAs, A and B, which both offer 5 early and 9 late licenses. A bidder would like to purchase 3 early licenses in PEA A and 5 late licenses in PEA B. Certainly, she must submit a early/mixed bid group in order to be able to win early licenses in PEA A. Her only possibility to then specify a demand for late licenses in PEA B is by stating a demand for 0 licenses in her base package for PEA B while at the same time specifying an increment for* 5 *licenses (which must be served with late licenses). It might appear that these bids do not reflect the bidder's demand accurately, as the bidder could theoretically also win 3 early licenses in PEA A but 0 licenses in PEA B. However, the bidder can set the base price of the whole bid group to 0, and specify a suitable markup for 5 licenses in PEA B. By doing so, she can indicate that winning only 3 early licenses in PEA A has no value to her.*

Bid groups are further 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 national 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 will also be restricted by the number of small and large bid groups they are allowed to submit. The exact numbers were to be determined according to the computational feasibility of the underlying allocation problem (see Section 6.1). Moreover, 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 on the same EA.

3.3 Sell-Side Reserve Bids

Reserve prices are commonly used in auctions to set the minimum price at which the auctioneer is willing to sell the products. The FUEL bid language implements reserve prices by placing a bid group on behalf of the auctioneer. In the course of the auction, the reserve bid is treated like any other bid group of a bidder. If the auctioneer's reserve bid is winning when solving the underlying allocation problem, the respective licenses remain unsold.

The auctioneer in the C-band auction sets reserve prices by submitting a single bid group which includes bids on each of the 406 PEAs (see Figure 3). The number of demanded licenses in the base package as well as the base price is set to 0. Through suitable markups the auctioneer then specifies linear reserve price for multiple licenses of the PEA. It is sufficient for the auctioneer to specify reserve prices only for late licenses as every early license can be transformed to a late license. As a consequence, it does not matter whether the auctioneer places the reserve bids as an *early/mixed* or *late* bid group as increments refer to late licenses in both types of bid groups. Note that according to Section 3.2 a bid group is only considered to be *large* if its base package contains licenses exceeding the MHz-pop of two national licenses. The auctioneer, however, is exempt from this rule so that the reserve bid is a valid large bid group even though its base package does not contain any licenses.

3.4 Auction Process

The C-band auction consists of two sealed-bid rounds: an initial *coordination* and a subsequent *main bidding* round. In the coordination round, which is intended as a feature to further reduce the missing bids problem, bidders may voluntarily submit either one large or multiple small bid groups which can be early/mixed or late. In contrast to the main

Auctioneer LATE LARGE Base price: 0																
		#Licenses														
EA	**PEA**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
1	107	Base	38	76	114	152	190	228	266	304	342	380	418	456	494	532
2	77	Base	56	112	168	224	280	336	392	448	504	560	616	672	728	784
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
170	148	Base	30	60	90	120	150	180	210	240	270	300	330	360	390	420

Figure 3: The reserve bid of the auctioneer contains reserve prices. For the sake of brevity, only the reserve bids for licenses in PEAs 107, 77, and 148 are shown. The auctioneer wins 2 licenses in PEA 77 and 1 license in PEA 148, i.e., these licenses remain unsold.

bidding round, none of the bid groups may include any adjustments and bidders may submit at most one small bid group per EA. Furthermore, the base price cannot be chosen freely by the bidder. Instead, the prices are set beforehand by the auctioneer and are based on prices paid for a similar amounts of spectrum in other countries. Bid groups submitted in the coordination round are anonymized and then disclosed to all other bidders.

A bidder might want to participate voluntarily in the coordination round to help others make relevant package bids that combine well with its own planned bids, increasing the chance that its own bids will be winning. A bidder that fails to place bids in the coordination round does not suffer any limits on its bids in the main round.

All bids submitted in both the coordination and main bidding round, as well as the reserve bids will be included in the optimization that determines the allocation. In the main bidding round the bid groups are no longer subject to the constraints formulated above for the coordination round.

As presented in Section 3.5, the allocation problem of the C-band auction can be represented using binary variables. The winning bid groups and also their associated winning adjustments are given by the optimal solution of the corresponding binary program. In accordance with other auctions conducted by the FCC, the determination of winning bids and the price that bidders are obliged to pay are separate calculations. While the binary program is used to determine the allocation of licenses to bidders, a Vickrey-nearest core-selecting rule is applied to compute the prices that the bidders will have to pay.

The solution of the allocation problem gives the number of blocks a bidder wins in the respective PEAs. The final step of the C-band auction determines how specific frequencies within the band are assigned to the winning bidders. If a bidder wins multiple blocks in the same PEA, the bidder will be guaranteed to obtain adjacent frequencies in the band. In case a bidder wins multiple blocks across neighboring PEAs, there will be a limited guarantee to award this bidder similar frequencies across the respective PEAs. Additional preferences concerning the frequency distribution within a band can be articulated by the bidders by placing bids in the subsequent assignment stage.

In the following we will exclusively focus on the most computationally challenging problem: the allocation problem that follows the main bidding round. Incentives and payment rules are not treated in detail here. Rather, we observe that VCG payment rules require the winner determination problem to be solved to optimality. So, among the advantages of the FUEL language is to enable such payment computations.[7]

3.5 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:

$i \in I_0$	Set of bidders including the auctioneer.
$a \in A$	Set of Economic Areas (EAs).
$p \in P$	Set of Partial Economic Areas (PEAs).
$P^E \subseteq P$	Set of PEAs in which early licenses are up for sale.
$g \in G_i^{SE}$	Set of small, early/mixed bid groups of bidder $i \in I_0$.
$g \in G_{ia}^{SE}$	Set of bid groups in G_i^{SE} that include bids on EA $a \in A$.
$g \in G_i^{SL}$	Set of small, late bid groups of bidder $i \in I_0$.
$g \in G_{ia}^{SL}$	Set of bid groups in G_i^{SL} that include bids on EA $a \in A$.
$g \in G_i^{TE}$	Set of large, early/mixed bid groups of bidder $i \in I_0$.
$g \in G_i^{TL}$	Set of large, late bid groups of bidder $i \in I_0$.
$g \in G_i$	Set of all bid groups of bidder $i \in I_0$, i.e., $G_i = G_i^{SE} \cup G_i^{SL} \cup G_i^{TE} \cup G_i^{TL}$.
$p \in P_i^g$	Set of PEAs contained in the base package of bid group $g \in G_i$ of bidder $i \in I_0$.
$k \in K$	Set of possible base bid adjustments, i.e., $K = \{0, \ldots, 14\}$.
$K_i^{gp} \subseteq K$	Set of adjustments in bid group $g \in G_i$ of bidder $i \in I_0$ for PEA $p \in P$.

[7]Core-selecting payment rules, which are widely used in spectrum auctions, also require the allocation problem to be solved to optimality (Day and Milgrom 2008, Goetzendorff et al. 2015).

Parameters:

- e_p Number of early licenses offered in PEA $p \in P$.
- ℓ_p Number of late licenses offered in PEA $p \in P$.
- 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 $i \in I_0$.
- μ_i^{gpk} Markup/discount on the base price of bid group $g \in G_i$ of bidder $i \in I_0$ 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 $i \in I_0$ wins bid group $g \in G_i$.
- $y_i^{gpk} \in \{0,1\}$ Binary variable denoting whether bidder $i \in I_0$ 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 $i \in I_0$ wins multiple small bid groups ($z = 0$) or one large bid group ($z = 1$).

$$\max \sum_{i\in I_0}\sum_{g\in G_i}(x_i^g\,\omega_i^g) + \sum_{i\in I_0}\sum_{g\in G_i}\sum_{p\in P_i^g}\sum_{k\in K_i^{gp}}(y_i^{gpk}\,\mu_i^{gpk}) \tag{1}$$

$$\text{s.t.} \sum_{k\in K_i^{gp}} y_i^{gpk} = x_i^g \qquad \forall i \in I_0, \forall g \in G_i, \forall p \in P_i^g \tag{2}$$

$$\sum_{i\in I_0}\sum_{g\in G_i^{SE}\cup G_i^{TE}}\sum_{k\in K_i^{gp}}\left(y_i^{gpk}\max\Big\{0, \big(\min\{e_p, b_i^{gp}\} - \max\{0, b_i^{gp}-k\}\big)\Big\}\right) \le e_p \quad \forall p \in P^E \tag{3}$$

$$\sum_{i\in I_0}\sum_{g\in G_i}\sum_{k\in K_i^{gp}}(y_i^{gpk}\,k) \le e_p + \ell_p \qquad \forall p \in P \tag{4}$$

$$\sum_{g\in G_i^{TE}\cup G_i^{TL}} x_i^g \le z_i \qquad \forall i \in I_0 \tag{5}$$

$$\sum_{g\in G_i^{SE}\cup G_i^{SL}} x_i^g \le M(1-z_i) \qquad \forall i \in I_0 \tag{6}$$

$$\sum_{g\in G_{ia}^{SE}\cup G_{ia}^{SL}} x_i^g \le 1 \qquad \forall i \in I_0, \forall a \in A \tag{7}$$

$$x_i^g \in \{0,1\} \qquad \forall i \in I_0, \forall g \in G_i \tag{8}$$

$$y_i^{gpk} \in \{0,1\} \qquad \forall i \in I_0, \forall g \in G_i, \forall p \in P_i^g, \forall k \in K_i^{gp} \tag{9}$$

$$z_i \in \{0,1\} \qquad \forall i \in I_0 \tag{10}$$

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 early licenses in PEA $p \in P$. Recall that increments of the base package always refer to late licenses, while decrements first reduce the number of early licenses demanded in the base package.

(4) Supply constraint for late licenses in PEA $p \in P$. Recall that early licenses can be used to serve the demand for late licenses.

(5) Bidder $i \in I_0$ may win at most one large bid group.

(6) 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.

(7) 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 (5).

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 $p_i \geq 0$ and d resources with capacities given by $c \in \mathbb{R}^d_+$ are given. Each item i consumes an amount $w_{ij} \geq 0$ 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*.

Definition 4 (D-FUEL)**.** *There are n' bidders submitting a single bid group i', each specifying a number of desired units $w'_{ij} \geq 0$ for each of m' items. The overall number of units available for sale in each of the $j' \leq m'$ items is $c'_j \geq 0$. The bid price for each bid group is p'_i. 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.

Proposition 1. D-FUEL *is strongly NP-complete. Any instance of* DKP *can be reduced to an instance of* D-FUEL.

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 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. An algorithm A is an fpt-algorithm with respect to a parameter κ, if there is a computable function f and a constant $c \in \mathbb{N}$ such that for every instance I the running time of A on I is at most $f(\kappa(I)) \cdot |I|^c$. If f is also a polynomial, A is referred to as polynomial fpt-algorithm (PFPT) with respect to κ.

For a fixed parameter $\kappa(I) = (c'_1, \cdots, c'_{m'})$ there is a PFPT for the multidimensional knapsack problem (Gurski et al. 2019). In the C-band auction problem for which FUEL was devised, the 14 licenses in each PEA is the small fixed parameter. Given this fixed parameter an algorithm for the multidimensional knapsack problem scales only linearly in the number of the bids and PEAs. Also in our experimental results with a branch-and-cut solver, we find that for a fixed number of licenses per PEA the runtime develops linearly in the number of dimensions (i.e., PEAs) of the problem.

4 XOR Bid Language

We compare the empirical complexity of auctions with the FUEL bid language to auctions with a standard XOR bid language. For this purpose, we will briefly introduce the XOR bid language and the resulting 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 two numbers which indicate the amount of early and late licenses that the bidder would like to purchase in the respective PEA. 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 neither distinguishes between early and late nor 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 k mutually disjoint sets. With k being the parameter, the problem is W[1]-complete, i.e. we cannot expect an algorithm of complexity polynomial in n. With the size of the sets (items in a bid) as parameter, the problem is fixed parameter tractable (Koutis 2005).

While local and national bidders may win at most one bid, the auctioneer is exempt from this rule. In order to represent the auctioneer's reserve bid for a single PEA, only 4 XOR bids need to be generated: one XOR bid each for 1, 2, 4, and 8 licenses.

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 3.5 for sets I, A, P, P^E and Parameters e_p, ℓ_p, ω_i^g.

Sets and indices:

$i \in I$	Set of bidders excluding the auctioneer.
$i \in I_0$	Set of bidders including the auctioneer.
$g \in G_i$	Set of all bids submitted by bidder $i \in I$.

Parameters:

c_i^{gp}	Number of early licenses demanded in bid $g \in G_i$ for PEA $p \in P$.
d_i^{gp}	Number of late licenses demanded in bid $g \in G_i$ for PEA $p \in P$.

Decision variables:

$x_i^g \in \{0,1\}$	Binary variable denoting whether bidder $i \in I_0$ wins bid $g \in G_i$.

$$\max \sum_{i \in I_0} \sum_{g \in G_i} (x_i^g \, \omega_i^g) \tag{1}$$

$$\text{s.t.} \sum_{i \in I_0} \sum_{g \in G_i} (x_i^g \, c_i^{gp}) \leq e_p \qquad \forall p \in P^E \tag{2}$$

$$\sum_{i \in I_0} \sum_{g \in G_i} (x_i^g \, (c_i^{gp} + d_i^{gp})) \leq e_p + \ell_p \qquad \forall p \in P \tag{3}$$

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

$$x_i^g \in \{0,1\} \qquad \forall i \in I_0, \forall g \in G_i \tag{5}$$

Objective & Constraints:

(1) Objective function summing up prices of winning XOR bids.

(2) Supply constraint for early licenses in PEA $p \in P^E$.

(3) Supply constraint for late licenses in PEA $p \in P$. Recall that early licenses can be used to serve the demand for late licenses.

(4) Bidder $i \in I$ may win at most one XOR bid. Note that this constraint does not apply to the auctioneer.

5 Experimental Design

In our experiments we differentiate between *local* and *national* bidders. Local bidders are only active in up to a dozen PEAs. National bidders, on the other hand, are interested in licenses in almost all 406 PEAs. For our tests we assume that there are 10 national bidders and 1,000 local bidders in the auction, which we chose to be larger than (but of the same order as) the actual numbers in any previous spectrum auction.

The bid generator we implemented for our extensive numerical experiments is capable of generating FUEL and XOR bids for local and national bidders. To compare the efficiency of the FUEL and XOR bid language, it is necessary that bidders have the same valuations. For this purpose, our bid generator also provides a feature to derive suitable XOR from FUEL bids. For instances of both the FUEL and XOR bid language, the valuations of bidders are subject to a value model that we present in the following.

5.1 Value Model

A widespread international metric for comparing the prices of spectrum is the license price per MHz-pop. The value model of the FUEL bid generator is based on this convention using the PEA population data provided by the FCC.[8] As the FUEL bid language uses abstract frequency blocks of 20MHz and the population of the smallest PEA is larger than 1,000, the FUEL bid generator divides the MHz-pop of a each license by 20,000 and uses this to scale license values. Let w_p describe the worth of a single license in PEA $p \in P$.

Naturally, 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 worth of a PEA for a particular bidder. In general, local bidders are financially weaker than national bidders so that we choose r_{ip}^{local} and r_{ip}^{national} uniformly at

[8] `https://transition.fcc.gov/bureaus/oet/info/maps/areas/data/FCC_PEA_website.xlsx`

random from the intervals $[1.0, 1.3]$ and $[1.1, 1.4]$, respectively. The auctioneer's idiosyncratic value factor is set to 1.0 for all PEAs, i.e., the auctioneer's reserve prices equal a constant fraction of the MHz-pop in each PEA.

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 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 by the idiosyncratic bidder valuations and shifting it so that is crosses the origin gives

$$d_{ip}(x) = r_{ip}\, w_p \left(\frac{1}{1 + e^{-x+\Delta_i}} - \frac{1}{1 + e^{\Delta_i}} \right),$$

where x is the number of licenses demanded by bidder $i \in I$ in PEA $p \in P$, and $r_{ip}\, w_p$ describes the bidder's idiosyncratic valuations for the respective PEA.

In 46 of the PEAs, 5 out of the 14 offered licenses can be made available 18 months earlier than the remaining 9 licenses. Depending on the current stock of licenses that a bidder owns, receiving early licenses might be essential but may also have little value to the bidder. We therefore introduce another factor t_i, drawn uniformly at random from the interval $[1.0, 1.5]$ for each bidder $i \in I$, that describes by which factor the bidder values early over late licenses. A bidder's valuation for winning x early and y late licenses in PEA $p \in P$ is then given by

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

The valuation function $v_{ip}(x, y)$ is visualized in Figure 4 for exemplary values of w_p, r_{ip}, Δ_i, and t_i.

5.2 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 1 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.

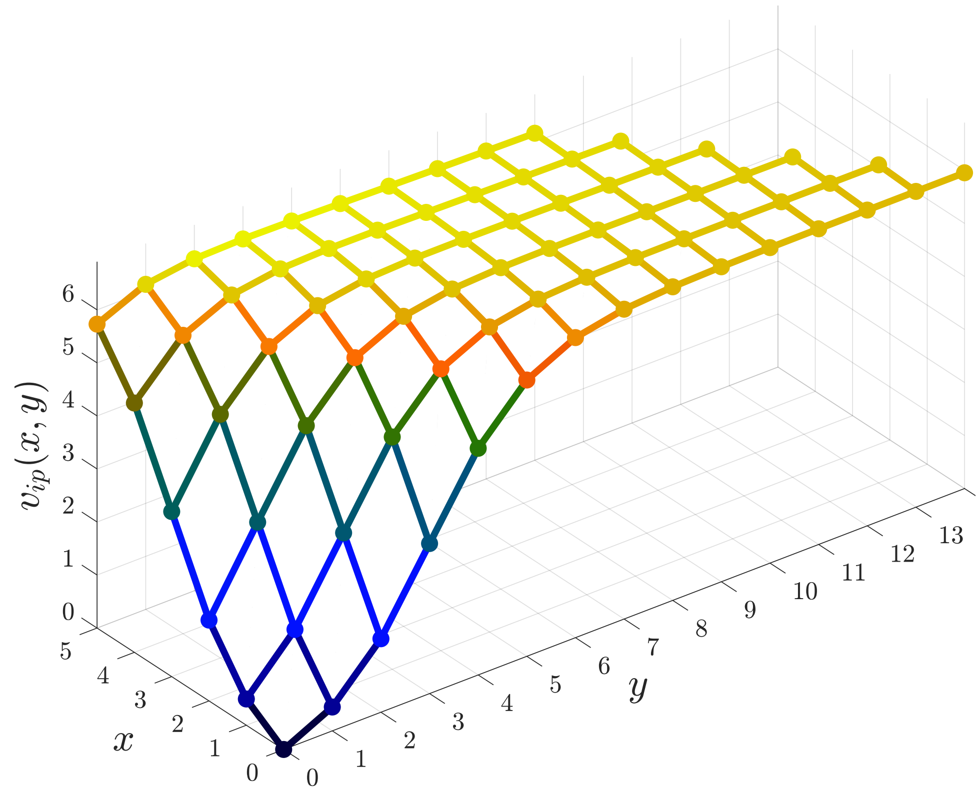

Figure 4: Valuation function $v_{ip}(x, y)$ for a PEA license worth of $w_p = 5$, an idiosyncratic value factor of $r_{ip} = 1.25$, a shift of the point of inflection of $\Delta_i = 3$, and early license value factor of $t_i = 1.1$.

Unlike local bidders, we assume that national 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 national 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.

Of course, every bidder faces the risk of not winning any of her bid groups. National bidders try to limit this risk by placing a so-called *fallback* bid group in addition to their regular bid groups. This dedicated bid group contains bids for all 406 PEAs and specifies a zero-adjustment in every PEA. This gives the auctioneer the option to serve the bidder's demand only in a fraction of the 406 PEAs, which might cause the bidder's winning PEAs to be scattered across the country. Naturally, bidders take this risk into account when determining values for their fallback bids. More specifically they reduce their idiosyncratic valuation factors r_{ip} by 15% for each PEA.

When local and national 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 5.1), 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. Moreover, previous tests suggest that scattered adjustments do not produce significantly different results from consecutive adjustments.

If a FUEL bid group contains PEAs for which early licenses are available, then the bid group is marked early/mixed with a probability of $5/14 \approx 40\%$ as this refers to the fraction of early licenses available in the respective PEAs. Otherwise, the bid group is considered late.

The number of bid groups that local and national 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 6.1.

5.3 Deriving XOR from FUEL Bids

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).

EARLY	**SMALL**	**Base price: 780**		
		#Licenses		
EA	**PEA**	**3**	**4**	**5**
7	44	Base	418	741
7	271	-50	Base	

$\Longrightarrow$

XOR Price: 1,471			
		#Licenses	
EA	**PEA**	**2**	**3**
7	44 Early		Base
7	44 Late	Base	
7	271 Late		Base

Figure 5: A random XOR bid is generated from a FUEL bid by picking a random adjustment combination (highlighted blue). In the XOR bid there are no adjustments and the bidder has to specify for each EA whether she is interested in early or late licenses.

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 national 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 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 5.2. 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).

5.4 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. When generating XOR bids it is assumed that any local bidder's market area contains between 1 and 10 PEAs (potentially belonging to different EAs) which form a highly cohesive component. Each XOR bid of a local bidder contains between 1 and the maximum number of PEAs available in her market area, so that each local bidder's XOR bid encompasses 3.25 PEAs on average. National bidders are active throughout the entire 406 PEAs. Similar to a national 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 national bidder are chosen uniformly at random among the 50% least populous PEAs.

Both local and national bidders demand between 1 and 5 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 case an XOR bid contains a PEA offering early licenses, then the number of early licenses demanded by a bidder is chosen uniformly at random between 1 and the number of licenses demanded by the bidder in this PEA.

6 Results

Here, we report the results of our computational experiments using the Gurobi Optimizer 8.1.1 to solve the winner determination problem up to a tolerance ("MIPGap") of 0.001, i.e., the solution computed by Gurobi differs from the optimal solution by no more than 0.1%. 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.

6.1 FUEL Bid Groups

The original FUEL proposal did not specify the number of bid groups that local and national bidders would be allowed to submit, but proposed to choose those to ensure the computational tractability of the winner determination program. The following experiments successively increase the number of small bid groups that any local bidder submits and the number of large bid groups that any national bidder places in the auction.

Let z_L and z_N denote the number of bid groups that local and national bidders are allowed to submit, respectively. For each configuration of z_L and z_N in Table 2, we generated 25 random instances with the FUEL bid generator. Table 2 summarizes the number of bid groups submitted by all bidders including the auctioneer, 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.

z_L	z_N	Package Bids	Binary Variables	Constr.	Runtime in sec.	TLE	Max. MIPGap	Unsold Licenses
1	1	1,011	27,935	6,613	5	0 of 25	0.0010	1.1
3	3	3,031	63,688	18,871	153	0 of 25	0.0010	2.5
5	5	5,051	99,351	30,121	356	0 of 25	0.0010	1.8
7	7	7,071	135,246	41,374	745	1 of 25	0.0014	1.3
10	10	10,101	188,709	58,193	1,212	3 of 25	0.0030	1.2
15	15	15,151	278,317	86,368	1,640	16 of 25	0.0101	1.3

Table 2: Average values of 25 test instances for different configurations of the number of small (large) bid groups that local (national) bidders are submitting.

Result 1. *If the 1,000 local bidders are allowed to submit 5 small bid groups and the 10 national bidders are allowed to submit 5 large bid groups, we can compute the winning allocation within 356 seconds and a MIPGap of only* 0.001 *on average. If we allow for 10 bid groups per bidder, the average runtime is at 20 minutes and three instances could not be solved to a precision of 0.001 in 30 minutes. The maximum MIPGap of these instances is still at only 0.003.*

The number of small and large bid groups that bidders are allowed to submit has a direct impact on the number of binary variables and constraints present in the winner determination problem. Restricting the number of bid groups in the auction can therefore reduce the runtime substantially. For up to 5 small (large) bid groups per local (national) bidder all test instances can be solved within the time limit of 30 minutes. If local (national) bidders are submitting 10 small (large) bid groups each, the far majority of test instances can still be solved within the time limit.

6.2 FUEL Admissible EAs

Even when restricting bidders to at most five base packages, the binary program for solving the FUEL allocation program is still very large, with roughly 100,000 binary variables and 30,000 constraints. The reason such large binary programs can still be solved within 30 minutes is the hierarchical structure of the binary program. In fact, if no large bid groups were submitted by any national 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 Gurobi optimizer to decompose the problem and apply effective cuts in the branch-and-cut algorithm.

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 1 creates inter-dependencies between small bid groups and therefore effectively prevents the Gurobi 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 1 and at most all PEAs present in the bidder's market area.

The bid generation for national bidders is independent of parameter k as there is no restriction regarding the number of EAs contained in large bid groups that national bidders submit. For our tests, we assume that local bidders submit five small bid groups, while national bidders submit five large bid groups.

Result 2. *If 1,000 local bidders are allowed to submit five bid groups across three EAs (not only within one EA), the winning allocation can be computed within 466 seconds on average. If bid groups that cover PEAs from up to five EAs can be submitted, only 20 out of 25 instances can be solved within the time limit of 30 minutes.*

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 3 lists the number of bid groups submitted by all bidders (including the auctioneer), 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	**Local Bidders**	**Package Bids**	**Binary Variables**	**Constr.**	**Runtime in sec.**	**TLE**	**Max MIPGap**	**Unsold Licenses**
1	1,000	5,051	99,351	30,121	356	0 of 25	0.0010	1.8
3	1,000	5,051	133,197	43,254	466	0 of 25	0.0010	0.9
5	1,000	5,051	169,687	57,126	1,012	5 of 25	0.0025	0.9
7	1,000	5,051	208,852	71,854	1,446	15 of 25	0.0063	1.2

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.

The test results in Table 3 imply that raising the maximum number of EAs in a small bid group has a significant impact on the runtime. If every local bidder is assumed to be active in five EAs, there are already 5 out of 25 test instances that cannot be solved within the time limit of 30 minutes. 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 local bidders now have a larger market area in which they are active so that they tend to be interested in spectrum of more PEAs.

Result 3. *FUEL's hierarchical structure for bids by national and local bidders and the limits on number of EAs in local bidders' bid groups both contribute to promoting fast solutions of large problems.*

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 1 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 1) and at most all PEAs present in the bidder's market area. The results of this test are shown in Table 4.

k	Local Bidders	Package Bids	Binary Variables	Constr.	Runtime in sec.	TLE	Max MIPGap	Unsold Licenses
1	1,000	5,051	99,351	30,121	360	0 of 25	0.0010	1.8
2	500	2,551	95,746	29,729	538	0 of 25	0.0010	3.2
3	333	1,716	95,308	29,870	530	0 of 25	0.0010	4.1
4	250	1,301	95,099	29,921	723	0 of 25	0.0010	4.4
5	200	1,051	95,055	29,982	774	1 of 25	0.0015	5.7

Table 4: 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. For $k = 5$ there exists one test instance that cannot be solved within the time limit of 30 minutes.

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 5 distinct bidders so that there are 1,000 local bidders competing in the auction for the treatment $k = 1$.

6.3 FUEL vs. XOR

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 Section 5.3. 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 national bidders submit five bid groups as our tests in Section 6.1 suggest that such instances can be solved within the time limit of 30 minutes. Similar to our test in Section 6.2, we restrict the maximal number of EAs contained in a local bidder's small bid group by k. In order to keep the XOR allocation problem tractable, we restrict local and national 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.

Result 4. *If bidders are only allowed to submit the same number of bids in the XOR as in the FUEL auction, more than* 5.2% *of all licenses remain unsold, the welfare loss compared* *to FUEL is almost* 10%*, and there are already 2 out of 25 test instances that are intractable. Even when bidders are allowed to submit three times as many XOR as FUEL bids, still more than* 2.8% *of the licenses remain unsold and the welfare loss is* 6.4%*. If local bidders are allowed to include bids for 2 EAs in their small bid groups, then the problem becomes intractable when bidders submit only 2 XOR bids.*

Table 5 shows the auction type, the maximum number of EAs contained in a small bid group (denoted k), the maximum number of bids that local and national bidders may submit (denoted z_L and z_N, respectively), the total number of submitted bids including the reserve bids of the auctioneer (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 0.001, the maximal MIPGap, the number of testcases 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	k	z_L	z_N	$\sum z$	**Bin. Vars.**	**Constr.**	**Runtime in sec.**	**TLE**	**Max MIP-Gap**	**Unsold Licenses**	**Eff.**
FUEL	1	5	5	5,051	99,351	30,121	356	0 of 25	0.0010	1.8	1.000
XOR	1	1	1	2,634	2,634	1,462	32	0 of 25	0.0010	806.3	0.756
XOR	1	5	5	6,605	6,605	1,462	277	2 of 25	0.0022	299.3	0.903
XOR	1	10	10	10,358	10,358	1,462	654	8 of 25	0.0115	202.6	0.925
XOR	1	15	15	13,565	13,565	1,462	770	9 of 25	0.0052	164.0	0.936
FUEL	2	5	5	5,051	115,307	36,346	446	0 of 25	0.0010	1.5	1.000
XOR	2	1	1	2,634	2,634	1,462	27	0 of 25	0.0010	816.6	0.775
XOR	2	2	2	3,644	3,644	1,462	498	5 of 25	0.0032	553.4	0.846
XOR	2	5	5	6,673	6,673	1,462	1,417	19 of 25	0.0121	328.2	0.901
XOR	2	10	10	11,718	11,718	1,462	1,602	21 of 25	0.0104	233.1	0.927

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

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 worth of a license in a PEA is a constant factor of the population living in a PEA according to our value model (see Section 5.1), 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.

6.4 Unrestricted XOR

Deriving XOR bids from previously generated FUEL bids (as done in the tests of Section 6.3) implies that the XOR bids encompass the same EA restrictions as the original FUEL bids. In particular, this means that the XOR bids which are derived from a local bidder's small bid group contain bids for at most k distinct EAs. A fully combinatorial XOR bid language, however, does not pose such restrictions on the bids but allows the auction participants to bid on licenses for any subset of PEAs. Due to the additional interdependencies between bids containing PEAs of different EAs, the corresponding winner determination problem becomes more complex.

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 computational tractable for the C-band auction. Similar to our previous tests, we assume that there are 10 national and 1,000 local bidders.

Result 5. *Even if both local and national bidders are restricted to submit no more than two XOR bids without restrictions on the EAs, 5 out of 25 instances could not be solved within the time limit of 30 minutes.*

For different configurations of the maximum number of XOR bids that local and national bidders may submit in the auction, we generate 25 random test instances. Table 6 shows the maximum number of XOR bids that local and national bidders may submit in the respective treatment (denoted z_L and z_N), the average total number of XOR bids present in the auction including the reserve bids of the auctioneer (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	2,634	2,634	1,462	34	0 of 25	0.0010	748.5
2	2	3,644	3,644	1,462	699	5 of 25	0.0037	444.9
3	3	4,654	4,654	1,462	1,800	25 of 25	0.0158	309.4
10	10	11,193	11,193	1,462	1,800	25 of 25	0.0123	146.4
50	50	45,545	45,545	1,462	1,800	25 of 25	0.0217	55.0

Table 6: Average values of 25 test instances for different restrictions on the number of XOR bids that both local and national bidders are allowed to submit.

Even if all bidders are restricted to submit no more than 3 XOR bids, none of the 25 instances can be solved within the time limit, even though this number of XOR bids is far too small to give a reasonable account of national bidders' preferences. Such limited bids also result in many licenses remaining unsold. When bidders are restricted to a single base bid, an average of 1.1 licenses remain unsold when using the FUEL bid language (see Table 2), while 748.5 licenses cannot be allocated when applying the XOR bid language, causing a significant revenue loss for the auctioneer.

7 Conclusion

The design of large auctions with hundreds of items is a challenge. In this paper we investigate the empirical 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 100,000 binary variables and 30,000 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 result from predominantly the hierarchical structure created by FUEL, which allows the Gurobi optimizer to decompose the binary program effectively. Furthermore, our experiments demonstrate that limiting the number of bid groups that local and national bidders may submit and restricting the maximal number of economic areas that local bidders can bid on in a small bid group effectively reduces the empirical hardness of the allocation problem.

In contrast to FUEL, we show that a fully enumerative XOR bid language quickly becomes computationally intractable. More importantly, 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 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 must be weighed against its powerful 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.

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, for which this would not have been considered feasible only recently.

Acknowledgements

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

References

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

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

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

Bichler, M., Schneider, S., Guler, K., and 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.

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

Cramton, P. (2017). Electricity market design. *Oxford Review of Economic Policy*, 33(4):589–612.

Day, R. and Milgrom, P. (2008). Core-selecting package auctions. *international Journal of game Theory*, 36(3-4):393–407.

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

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

Eilat, A. and Milgrom, P. (2011). The caf auction: Design proposal. Technical report.

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

Garey, M. R. and Johnson, D. S. (1979). *Computers and intractability*, volume 174. freeman San Francisco.

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

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

Green, J. and Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica: Journal of the Econometric Society*, pages 427–438.

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

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

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

Kim, S. W., Olivares, M., and Weintraub, G. Y. (2014). Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science*, 60(5):1180–1201.

Koutis, I. (2005). A faster parameterized algorithm for set packing. *Information processing letters*, 94(1):7–9.

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

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

Lavi, R. and 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., and Sandholm, T. (2006). The winner determination problem. *Combinatorial auctions*, pages 297–318.

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

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

Milgrom, P. (2019). FUEL for 5G: Flexible use and efficient licensing. Technical report, Auctionomics, C-Band Alliance.

Nisan, N. (2006). Bidding languages. *Combinatorial auctions*, pages 400–420.

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

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

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

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

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