

FAucS: An FCC Spectrum Auction Simulator for Autonomous Bidding Agents

János A. Csirik, Michael L. Littman, Satinder Singh, and Peter Stone

AT&T Labs – Research

180 Park Ave.

Florham Park, NJ 07932

{janos,mlittman,baveja,pstone}@research.att.com

<http://www.research.att.com/~{janos,mlittman,baveja,pstone}>

April 22, 2001

Abstract

We introduce **FAucS**, a software testbed for studying automated agent bidding strategies in simulated auctions, specifically the United States FCC wireless frequency spectrum auctions. In addition to the complexity of these auctions, which provides ample opportunities for intelligent approaches to bidding, this type of auction has huge commercial importance, each bringing in billions of dollars to governments around the world. We implement straightforward sample agents in **FAucS** and use them to replicate known beneficial bidding strategies in this type of auction. We then discuss potential in-depth studies of autonomous bidding agent behaviors using **FAucS**. The main contribution of this work is the implementation, description, and empirical validation of the **FAucS** testbed. We present it as a challenging and promising AI research domain.

1 Introduction

In recent years, governments around the world have used spectrum auctions as a mechanism for allocating rights to radio spectrum. These auctions have generated extremely large revenues, often with billions of dollars at stake.

The United States' Federal Communications Commission (FCC) has already engaged in more than 30 such auctions, altering the rules from auction to auction in the hopes of maximizing revenue and reducing opportunities for tacit, mutually beneficial agreements among participants [Weber, 1997]. The result is a complex set of rules that require sophisticated strategizing on the part of auction participants. To the best of our knowledge, participants in these auctions have always generated their strategies and bids manually, even though all bids are submitted and acknowledged online.

Recent work in the field of artificial intelligence has introduced the possibility of creating autonomous bidding agents to participate in auctions. In particular, the first trading agent competition (TAC) was held in Boston in July, 2000 [Wellman *et al.*, 2001]. TAC agents acted as simulated travel agents and had to procure goods for their clients in different types of auctions, bidding against other autonomous agents. ATTac-2000, our entry in TAC, finished in first place [Stone *et al.*, 2001]. TAC was very successful at attracting many competitors from around the world by creating an artificial domain that was simple enough to understand quickly, but complex enough to prevent a trivial winning strategy. However, one potential criticism of TAC is that the domain was not realistic enough to generate meaningful, scalable agent strategies.

This paper introduces spectrum auctions as a complex and realistic domain for exploring general autonomous bidding strategies. Our simulator, **FAucS**, is faithful to the FCC regulations for the most recent FCC auction, Auction 35, which involved 422 licenses in 195 markets and ran from December 12th, 2000 to January 26th, 2001. **FAucS** allows autonomous bidding agents and humans to participate simultaneously in mock auctions. We introduce several implemented bidding strategies and use **FAucS** to replicate some of the observed phenomena in past auctions such as strategic demand reduction (SDR) and strategic bidding opportunities such as budget stretching.

The main contribution of this work is the implementation, description, and empirical validation of the **FAucS** testbed. We present it as a challenging and promising AI research domain.

The remainder of the paper is organized as follows. Section 2 introduces the spectrum auction domain and our FCC spectrum auction simulator, **FAucS**. Section 3 discusses the strategies used by our sample bidding agents. Section 4 reports the results of a set of simulations. Section 5 suggests future uses for **FAucS**, Section 6 describes related work, and Section 7 concludes.

2 FCC Spectrum Auction Simulator

The goods available in the FCC spectrum auctions are a set of *licenses*, or blocks of spectrum; each in a *market*, or region of the United States. In Auction 35, licenses were 10 or 15 megahertz in size, and each of the 195 markets had between 1 and 4 licenses available. A total of 422 licenses and more than 80 bidders were involved.

2.1 Overview

To a first approximation, the rules of the auction are straightforward (official rules are presented in FCC document DA 00–2038). All of the FCC spectrum auctions, including Auction 35, use a simultaneous multiple round (SMR) system. In an SMR auction, all goods are available at the same time, and bidding occurs in discrete rounds. After each round, each bidder’s bids are announced publicly. The *provisionally winning bids* are also announced: These are the highest bid received up to that point on each license. (In case of a tie, the first bid submitted to the FCC’s system wins.) The auction ends immediately after the first round with no new activity. Each license is then sold to its provisional winner, for a price equal to the provisionally winning bid.

2.2 Allowable Bids

The amount that can be bid on a particular license is constrained to be an element of a small fixed set. If the license has attracted no bids yet, this set consists only of a *minimum bid price* as pre-announced by the FCC. Once the license has attracted at least one bid, the acceptable bids are those that top the provisionally winning bid by 1 to 9 times the *increment*, which is a value between 10% and 20% of the provisionally winning bid. The increment is calculated by the FCC using an exponential smoothing formula, which tends to make the increment high or low depending on whether the number of bids the license is drawing is high or low. We refer to any bid other than the 1-increment bid as a *jump bid*.

The bid increment rule was added to prevent bidders from communicating with one another via the low-order (i.e. relatively insignificant) digits in the bids.

2.3 Eligibility Constraints

The most important rules that make FCC spectrum auctions different from other auctions concern *eligibility constraints*. The rules are intended to prevent bidders from sitting out during the early rounds of bidding and delaying placing their serious bids until late in the auction. It was feared that such a strategy might be employed and would result in lower FCC revenues and an unnecessarily long auction.

For the purposes of the eligibility rules, each license is assigned a number of *bidding units*. The number of bidding units per license ranges from 2,500 (Minot, ND) to 15,906,000 (New York City), and it correlates with the population of the corresponding market. For each round of the auction, each bidder has a certain *eligibility* and *activity*, both of which are expressed in terms of bidding units.

Activity(b,r): A bidder b 's activity for a round r is defined as the total number of bidding units in the licenses where b has (a) placed a bid in the current round; or (b) held the provisionally winning bid as of the end of the previous round.

Eligibility(b,r): b 's eligibility in round r is a quantity that controls the allowable activity for that round. Eligibility is defined as follows:

$$\text{Eligibility}(b,r) = \text{Min}(\text{Eligibility}(b,r - 1), \text{Activity}(b,r - 1)/.80)$$

The eligibility in the first round can be chosen by the bidders themselves (it is directly proportional to the deposit they put down). In each round, each bidder's activity is constrained to be no higher than its eligibility. If a bidder submits a set of bids such that the activity is less than 80% of the eligibility, its eligibility is reduced¹. Thus, once eligibility is lost, it cannot be recovered.

¹This constant starts at 80% and is increased by the FCC to 90% and then 98% as the auction draws near its end.

2.4 Withdrawals and Waivers

Bidders are also allowed a small number of *bid withdrawals*, but the rules governing doing so are so punitive² that withdrawals play little role in recent FCC spectrum auctions and in our simulations.

Additionally, bidders are allowed up to five *activity waivers*. A bidder using an activity waiver in a certain round ensures that their eligibility will not be reduced in that round, even if its activity is low. These waivers are intended to ensure that equipment failures and the like do not put bidders in impossible situations, since a bidder with no provisionally winning bids could reduce its eligibility to zero by not bidding in a single round of the auction, and thereby shut itself out completely from the rest of the auction³. Activity waivers have potential strategic uses. In real FCC spectrum auctions, bidders tend to save up most of their activity waivers until they are almost ready to drop out.

2.5 FAucS

Our FCC spectrum auction simulator, **FAucS**, models all of the above rules in their entirety. It uses a client-server architecture with the server and the bidding agents (clients) all written in Perl and using TCP sockets to communicate with each other. The software is highly customizable in that it

- Reads in license files defining the goods that are available for sale;
- Can model a wide range of rules used by the FCC in past auctions or considered by the FCC for future auctions via simple comand-line parameters.
- Allows for the automatic inclusion of arbitrary numbers of parameterized agents;
- Allows humans and agents (or only humans) to simultaneously participate in a simulated auction;
- Allows for tracking FCC spectrum auctions as they happen;
- Allows for loading data from a partially completed auction, real or simulated, and finishing it off with the participation of various bidding agents and/or humans;
- Was designed so as to be easily extensible to combinatorial auctions such as those planned in the future by the FCC.

Typical auctions last between 100 and 150 rounds, with Auction 35 lasting 101 rounds. In real life, they take several weeks, or even months to complete. We have completed auctions involving humans using **FAucS** that have completed within a day, and agent-only auctions with more than 60 licenses and about 10 agents complete in 15 to 20 minutes⁴. A good deal of this time is typically agents waiting for each other to finish their bidding, since a round can only advance when all agents have placed their bids.

²A withdrawn provisionally winning bid must still be paid unless another bidder matches that price on the license. Thus, bid withdrawals are a potentially very expensive way of freeing up eligibility.

³The FCC will automatically grant bidders activity waivers (if they have any left) in any round in which they weren't heard from.

⁴Times are on a multiprocessor, 270Mhz SGI workstation.

3 Prototype Autonomous Bidding Agents

In this section, we present sample autonomous bidding agents for use in **FAucS**. These agents are not meant to be realistic in all regards, but are rather intended to demonstrate the full functionality of **FAucS** and the challenging nature of the problem it presents. As such, the agents incorporate several approximations to fully realistic (i.e. human-like) behavior. We point out the elements that represent approximations throughout.

First, we define the utility measure we use to evaluate bidding agents. Second, we present a sub-problem essential to all intelligent bidding strategies. Third, we describe the basic agent (BA), which uses a strategy designed to maximize the bidder's utility without any regard to the constraints of the other bidders. Fourth, we describe the budget-stretching agent (BSA), which actively tries to force other agents to spend more money in some markets in order to reduce their abilities to bid in others.

3.1 Agent Utilities

Merrill Lynch & Co. analyzed the wireless communications industry and estimated that each megahertz of spectrum is worth

- \$2 per person in markets with fewer than 750,000 people (144 markets from McCook, NE to Worcester, MA)
- \$4 per person in markets with between 750,000 and 2,500,000 people (33 markets from Syracuse, NY to Pittsburgh, PA), and
- \$6 per person in markets with more than 2,500,000 people (18 markets from Tampa, FL to New York, NY) [Hall, 2000].

Following this basic insight, we imagine that each bidder has a dollar value for obtaining a license in a given market that is related to the size of that market, and a somewhat lower value for obtaining a second license in the same market. We use the following terms:

Market Value(m,b): the dollar worth (possibly 0) of a single license in market m to bidder, b .

Secondary market value(m,b): the dollar worth (possibly 0) of a second license in m to b .

Supply(m): The number of licenses for sale in market m .

Demand(m): The number of non-zero market and secondary market values in market m among all the bidders.

Budget(b): The total dollar amount available to bidder b . b can never have outstanding bids in excess of this amount.

For the purposes of this paper, no bidder has a non-zero value for a third license in a market. The precise market values vary among the agents within a range but are commonly known: we use a private-value, complete-information model⁵.

⁵The complete-information aspect is not entirely realistic: In real auctions, the bidders only have rough ideas of each others' market values.

In the absence of other constraints, the expected final price of a market in which supply is greater than or equal to demand will be the minimum bid price. The final price of a market with supply of k and demand greater than k is within one bid increment of the $(k + 1)$ st highest market value. Note that multiple licenses in the same market are sold independently; they can and do sell for different prices⁶.

At the end of the auction, the total of the market values and secondary market values of the licenses a bidder won, minus the cost of those goods, is the bidder's *utility*:

$$\text{Utility} = \text{Total value} - \text{Total cost.}$$

Each bidder seeks to maximize its utility, which can be done by winning as much market value as possible at as low a price as possible.

3.2 The Allocation Sub-Problem

In the absence of budget or eligibility constraints, bidders can be quite successful with the strategy of simply bidding in each market in which the ask price is less than the bidder's value. The presence of the additional constraints make the basic bidding strategy somewhat more complex. This section describes a constrained allocation problem and our solution to it. The resulting algorithm is then used by the bidding agents described in later sections.

The allocation sub-problem is this: given a set of licenses, each with a utility, a number of bidding units, and a cost; a maximum total number of bidding units (eligibility); and a maximum total cost (budget), find the subset of licenses with maximum total utility subject to budget and eligibility constraints.

Ignoring the eligibility constraint, the allocation sub-problem is precisely the classic "knapsack" problem. The standard exact dynamic-programming solution to knapsack computes a table with one entry for each possible integer value from 0 to the total utility of all licenses. In the FCC auction domain, however, the total can easily be in the billions of dollars. We developed a refined version, which only creates table entries for totals that are actually encountered during the search for an optimal solution.

To find a subset of licenses that obeys both budget and eligibility constraints, our allocation algorithm calls the knapsack solver to find a set of licenses with maximum total utility subject only to the budget constraint. If the resulting set violates the eligibility constraint, the algorithm repeats the optimization using a reduced budget (one dollar less than the cost of the previous optimal solution). This process terminates eventually with a solution that satisfies both budget and eligibility constraints.

In general, the procedure described above can produce solutions that are arbitrarily worse than optimal. However, we found that, in practice, the first solution produced by the algorithm generally satisfied the eligibility constraint and was therefore an optimal allocation.

REPEAT (once per round) <ol style="list-style-type: none"> 1. Get market prices from server 2. Compute remaining budget and eligibility 3. Compute current values and costs of markets 4. Choose desired licenses within constraints 5. Submit bids to server UNTIL game over
--

Table 1: High-level overview of our basic agent (BA) algorithm.

3.3 Basic Agent

Our sample bidding agent algorithm is summarized in Table 1. The quantities in Steps 2 and 3 are computed as follows.

Remaining eligibility: subtract from the bidder’s current eligibility the bidding units tied up in licenses of which it is provisional winner;

Remaining budget: subtract from the bidder’s total budget the money tied up in licenses of which it is provisional winner;

Current values for markets: use the market value or the secondary market value, depending on the number of licenses in the market in which the bidder is already provisional winner;

Current costs for each market: find the two least expensive licenses in each market.

This basic agent strategy is myopic in that it assumes, at each bidding round, that the bidder will win all licenses of which it is provisional winner, and that it could win any other license for a 1-increment bid. For choosing the desired licenses in each round (Step 4 in Table 1), the bidder uses the allocation algorithm described in Section 3.2

Once the set of desired licenses was determined, the basic agent always bids (Step 5 in Table 1) at the 1-increment price (no jump bids) for those and only those licenses. In a single-good ascending auction, this strategy is optimal [Klemperer, 1999]. However, in SMR auctions, and particularly in the presence of budget constraints, other bidding strategies may be more beneficial. One such strategy is described next.

3.4 Budget-Stretching Agent

Budget constraints “provide incentives to ...try to reduce opponents’ budgets in early sales in order to lower subsequent sale prices.” [Klemperer, 1999]. To demonstrate this effect within the

⁶The licenses represent different frequency bands. Some real bidders exhibit preferences among the bands [Cramton, 1997]. For the purposes of the experiments in this paper, we consider licenses in the same market to be of identical value to the agents.

spectrum auction scenario, we created an agent that aggressively bids in markets in which its opponents are interested in order to raise the prices of licenses in those markets and reduce their competitiveness.

Specifically, BSA acts identically to BA except for in Step 5 of Table 1, at which point it does the following.

1. Submit all bids that BA would submit.
2. Compute the **stretchable markets**—those in which the agent has no market value and in which the total demand of all bidders is equal to the supply. This computation can be done once at the beginning of the auction and stored.
3. In each stretchable market:
 - Let **OPlimit** be the smallest market value for a license in the market among bidders with demand in that market.
 - Compute:
 - **BSAlimit** = OPlimit/1.2
 - **NJmin** = BSAlimit/1.4
 - **NJmax** = BSAlimit/1.2
 - Bid at the largest possible price p such that $p < \text{NJmin}$ OR $(p > \text{MJmax AND } p < \text{BSAlimit})$

Figure 1 illustrates this strategy. Notice that BSA relies on knowing the demands and values of the other bidders: We are assuming a complete-information scenario.




Figure 1: A general picture of the budget-stretching strategy (numbers in parentheses correspond to the example in the text). The agent avoids bidding at prices covered by a dashed line. Subject to that restriction, it bids at the maximum allowable price.

To understand the purpose of budget stretching, consider the following example with two agents and 2 markets, each of which has only 1 license available:

Agent S (*stretcher*): budget \$250M, market 1 value \$300M

Agent V (*victim*): budget \$525M, market 1 value \$300M, market 2 value \$300M.

If both agents use the BA strategy, **V** will win the license in market 1 for (within one bid increment of) \$250M and the license in market 2 for the minimum bid price (since there's no competition). **S** will win neither license.

However, assuming that **V** follows the BA strategy, **S** can win the market 1 license. Early in the auction, it can raise the price of the market 2 license to just below \$300M (the opponent's budget limit, OPlimit). Care must be taken such that the next lowest acceptable bid is still under \$300M

so that **V** can still afford it. **V** will then spend (close to) its limit on that license, leaving **S** to win the license in market 1 for \$225M.

Recall that the bid increment can be as high as 20% of the current price. Thus, **S** must bid on the license in market 2 for no more than $OPLimit/1.2 = \$250M$ (the $BSALimit$). Additionally, if its final bid is only a single increment over the previous price (i.e., not a jump bid), **V** will place the same bid and have a chance of winning the license at the lower price. Therefore, **S**'s best strategy is to plan such that it can reach a price of close to the $BSALimit$ with a final bid that is a jump bid. In particular, if the going price is ever between $BSALimit/1.4$ ($NJmin$ for "no jump minimum") and $BSALimit/1.2$ ($NJmax$), then there will be no jump bids available below the $BSALimit$ level in the next round. Thus, in this example, **S** bids in the early rounds at the highest price possible covered by a solid line in Figure 1. Due to the jump bids, the market 2 license rises in price much faster than does the market 1 license, and, after a few rounds, **V** steps in and wins the license at a price close to its \$300M budget.

In general, budget stretching can be effective in markets such that the agent has no demand and total demand is equal to supply, the stretchable markets (observe that market 2 in the above example is stretchable). Without budget stretching, the agents with demand in that market win the licenses at the minimum bid price. With budget stretching, agents with demand in the market win the licenses at a significantly higher price, close to their market values. Even though the prices are high, the bidders still win the licenses because the budget stretching happens in the early rounds, so they can still all afford to bid for the licenses, as they do according to the BA strategy⁷.

4 Empirical Results

This section describes a set of simulated auctions we ran to illustrate two families of known bidding strategies: strategic demand reduction and budget stretching. The primary purpose of these experiments was to verify the potential usefulness of **FAucS** for studying autonomous agent bidding strategies in a complex, realistic domain. However, the parameter choices are all justified by data from Auction 35. In this sense, the experiments provide a way of quantifying the potential benefits of strategies that were known to be good ideas in principle.

4.1 Test Auction Setup

A fixed set of market values were used in all our runs. We began by defining a set of five "big bidders," intended to emulate the presence of five national companies participating in Auction 35⁸. We also included 4 "small bidders" in each market with smaller market values. These small bidders represent regional companies that are budget-constrained and are only interested in specific markets. They essentially served to raise the minimum final prices to more realistic levels.

⁷The opponent agents could come up with a better response to budget stretching. BSA is designed to exploit weaknesses in BA.

⁸AT&T Wireless, Cingular Wireless, Sprint PCS, Verizon Wireless, and Voicestream.

4.1.1 Market Values

In our experiments, we randomly assigned market values to bidders based on a model that incorporates the Merrill Lynch estimates presented in Section 3.1. This model was created in consultation with the Auction 35 bidding team from AT&T Wireless, who felt that it constituted a reasonable approximation of bidders such as Sprint, Verizon, and AT&T, which possess nationwide networks.

First, the total number of licenses needed by each bidder was determined using a power law distribution (specifically, so that the total number of licenses with non-zero market values for the various bidders forms a geometric progression of ratio 0.8). Given the resulting total number of licenses a bidder needs, we selected its markets with non-zero value by repeatedly drawing (with replacement) from the set of all markets. The only constraint was that no bidder had non-zero demand for more than 2 licenses in a given market.

Second, each bidder was given smoothed and randomly perturbed versions of the Merrill Lynch market values. For markets in which the bidder had a demand of 2, the bidder used this market value, and had a secondary market value of 95% of the first value. For markets in which the bidder had a demand of 1, only the 95% discounted value was used. All other markets were given market values of 0. The small bidders were assigned market values of approximately 75% of the Merrill Lynch market values. Initial eligibilities were set to allow bids on all licenses.

As an example, we show the bidders’ resulting market values for Los Angeles used in our experiments. Bidders 5–8 are the small bidders. There were 3 licenses available in Los Angeles. Thus, for our experiments, demand (4) was greater than supply in this market.

Bidder	Market Value	Secondary Market Value
0	\$1,148 Million	\$1,090 Million
1	\$0	\$0
2	\$1,073 Million	\$0
3	\$0	\$0
4	\$1,120 Million	\$0
5–8	\$862 Million	\$0

Table 2: The bidders’ market values for Los Angeles used in our experiments. Bidders 5–8 are the small bidders. There were 3 licenses available and an overall demand of 4.

In our runs, we included only the largest 62 markets from Auction 35 (this is the subset of the top 100 U.S. markets that are available in Auction 35), as these constitute a large majority of the value in the auction. There were 155 licenses available in these markets.

4.1.2 Baseline Budgets

To pick a baseline budget for each of the five big bidders, we performed the following “sealed bid” analysis. Each bidder placed bids in each market at its market value and secondary market value. Then, in each market, we assigned the k licenses available to the top k bidders at the bid price placed by the $(k + 1)$ st highest bidder. Ties were broken randomly and expected values computed. Each bidder’s baseline budget was computed as the total expected expenditure in this sealed-bid analysis. As a result of this analysis, the big bidders’ baseline budgets ranged from \$1.7–\$3.8 billion, which is roughly in line with their apparent budgets in Auction 35.

Bidder	BA	Run 13	SDR	Δ
0	933M	893M	1401M	50.2%
1	716M	699M	927M	29.4%
2	391M	409M	579M	48.2%
3	390M	328M	788M	102.0%
4	294M	241M	509M	73.2%

Table 3: Utilities of the big bidders in an average of 50 runs with big budgets (175% of baseline); run 13 of 50, which yields utilities closest to the average; and an average of 20 runs with strategic demand reduction based on run 13. The percent improvement (Δ) compares the SDR and BA runs.

4.2 Sample Experiment 1: Strategic Demand Reduction

Strategic demand reduction (SDR) is a tactic by which bidders pay less for the same goods by agreeing to “want” less. When adopted unilaterally, SDR can potentially be disastrous. However, if a group of bidders coordinate their demand reduction so that each wants different goods, they can get what they want at nearly the minimum bid price.

Explicit coordination of the kind required to arrange and enforce SDR among a group of bidders is specifically prohibited by the FCC spectrum auction rules; it is essentially a form of collusion. Nevertheless, there is evidence that this type of behavior has often been arranged via tacit bidding signals in previous auctions [Weber, 1997]. Specifically, bidders have signaled their intentions to reduce demand in particular markets by sending coded messages in the low order digits of bids. The messages have consisted of market identification codes or bidders’ names spelled out on telephone keypads. This type of communication has been ruled out in Auction 35 by the bid increment system described in Section 2. Nonetheless, if the benefit is sufficiently high compared to the cost (or punishment), bidders are likely to find a way to signal their intentions to one another.

To quantify the potential effect of SDR in our simulator, we ran the following study. First, a series of 50 runs was carried out in which bidders, following the BA strategy, had extremely large budgets (175% of their baseline budgets). The average utilities for the big bidders is given in Table 3 in the big-budget column.

Of the 50 runs, we picked a single run that gave utilities closest to the average (both in mean squared error and max-norm). The utilities for the bidders in this run, run 13, are listed in Table 3. In run 13, each bidder ended up purchasing a particular number of licenses in each market. In a sense, the bidders should be reasonably happy with the total value of the licenses they got in run 13, since it is about average. Therefore, they ought to be even happier if they could get those same licenses at a lower price.

Following this train of thought, we constructed a modified set of market values for the five large bidders by zeroing out the value for licenses that were not won in run 13. Using these modified market values in an auction would result in the bidders winning precisely the same number of licenses per market as in run 13. However, because the competition between big bidders is eliminated (demand is always greater than or equal to supply), prices should only go as high as necessary to beat the price set by the small bidders. Thus, we’d expect the utilities to be substantially larger in this case. This expectation is born out in the SDR and Δ columns of Table 3, which give average utilities over 20 runs with the modified market values and the percent improvements compared to the BA column. All improvements from the BA case are statistically significant by a

Bidder	BA	BSA	Δ	Sig.
0	437M	490M	12.2%	yes
1	457M	403M	-11.7%	yes
2	231M	176M	-23.5%	yes
3	306M	385M	25.9%	yes
4	162M	141M	-12.9%	no

Table 4: Comparison of average utilities (65 runs) between runs in which all agents used the basic agent (BA) strategy and runs in which bidder 3 used the budget-stretching (BSA) strategy. Use of the BSA strategy improved bidder 3’s utility substantially.

one-sided t-test ($p \ll .05$).

SDR is a known, clearly beneficial technique in auctions. Although it is known that SDR improves bidder utilities, if the absolute amount of the improvement is small, the potential costs of coordination may not be worthwhile. These experiments serve to quantify the potential benefits available to the agents when using SDR in our model of Auction 35.

For SDR to work, the agents must find a means of signaling to each other and negotiating at low prices. One potential mechanism for signaling in Auction 35 is via jump bids in desired markets. We leave investigation of such communication methods to future research. However, the above study gives an upper bound on the potential benefits of SDR.

4.3 Sample Experiment 2: Budget Stretching

To study budget stretching in our simulation, we ran bidder 3 with the BSA strategy⁹. With our randomly generated market values, bidder 3 had 8 stretchable markets.

We ran the bidders with 75% of their baseline budgets (to make sure that the budgets would have a constraining effect), and compared their utilities in runs in which all bidders used the BA strategy to runs in which bidder 3 used the BSA strategy (65 runs of each). Table 4 gives the mean utility for each bidder under the BA and BSA conditions, along with the increase (positive) or decrease (negative) in utility and whether the difference is statistically significant ($p < .05$ by a two-sided t-test).

Note that these utilities are smaller than those given in Table 3 because the budgets are a great deal more constrained—more licenses are being won by the small bidders in the budget-stretching runs.

Not surprisingly, the biggest change in utility is the increase for bidder 3. Budget stretching was quite successful in this case, improving bidder 3’s score by almost \$80 million. All other bidders were hurt, except bidder 0, who also enjoyed a significant improvement in utility. Although bidder 0 is budget-stretched in some markets, it benefits fortuitously from the fact that some of its competitors no longer have enough capital to compete with it in some of its most valuable markets: it more than makes up for its losses in 15 different markets with significant gains across 21 markets. For comparison, bidder 3 has small losses in 4 markets and gains in 14 (5 of them larger in magnitude than the largest loss) as a result of budget stretching. The gains of bidder 3 and

⁹We selected bidder 3 since it has the largest number of stretchable markets.

bidder 0 both result in large part from a significant saving in a single big market (Los Angeles and San Francisco, respectively).

Bidder 0's benefits derive from a complex set of situations that would be difficult to predict in advance. This fact points out one of the strengths of the simulation approach enabled by **FAucS**—it makes it possible for us to examine the effects of complicated interactions in auction dynamics that would be difficult, if not impossible, to predict by direct analytic methods alone.

5 Potential Uses for **FAucS**

The agents presented in this paper are straightforward agents meant to suggest the potential uses of **FAucS**. We anticipate that **FAucS** will be useful for a wide variety of future studies, both by ourself and by other researchers. Potential directions for future bidding agents include those designed to study

- the development of stable strategies;
- inference under incomplete information;
- strategic uses for activity waivers; and
- communication via bid signaling.

Furthermore, as **FAucS** was developed so as to be easily extensible to combinatorial auctions, it can also be used to study agent bidding strategies in such cases.

One of the original motivations behind **FAucS** was our desire to create a tool to aid human bidders who are participating in multi-unit simultaneous auctions. It can be used to study on-going auctions in depth and to project future scenarios by modeling the other participants as agents.

Another exciting potential use for **FAucS** is as the substrate for a competition such as TAC [Wellman *et al.*, 2001], which would undoubtedly lead to a set of competitive agent strategies.

Finally, **FAucS** can be used to study bidding interactions between autonomous bidding agents and human bidders.

6 Related Work

Although there has been a good deal of research on auction theory, especially from the perspective of auction mechanisms [Klemperer, 1999], studies of autonomous bidding agents and their interactions are relatively few and recent. The Trading Agent Competition [Wellman *et al.*, 2001], briefly described in Section 1, is one example. Our first-place team at TAC [Stone *et al.*, 2001] used adaptive bidding agents to successfully avoid the most contentious, and thus expensive, markets. However, TAC involved simplified bidding scenarios that were not particularly realistic and that in some cases dramatically affected agent strategies [Greenwald and Stone, 2001]. FM97.6 is another auction test-bed, which is based on fishmarket auctions [Rodriguez-Aguilar *et al.*, 2000]. Automatic bidding agents have also been created in this domain [Gimenez-Funes *et al.*, 1998]. Outside of, but related to, the auction scenario, automatic shopping and pricing agents for internet commerce have been studied within a simplified model [Greenwald and Kephart, 1999].

FAucS addresses a much more complex scenario than has been previously studied with autonomous bidding agents: the FCC spectrum auctions. Spectrum auctions have been analyzed retrospectively [Weber, 1997; Cramton, 1997], but little is known about them from a theoretical perspective.

As suggested by the results presented above, perhaps the most challenging aspect of spectrum auctions is the presence of budget constraints. As noted by Benoit and Krishna [2000], “Traditionally, auction theory . . . assumes that bidders have unlimited budgets . . .” They argue that even with just 2 goods for sale, and complete information about the goods’ values and the bidders’ budgets (as assumed by our budget-stretching strategy), budget constraints complicate auctions sufficiently to warrant study.

Indeed, Benoit and Krishna [2000] study such a scenario—extending previous research involving just 2 bidders—and find that sequential auctions are more advantageous to the seller than simultaneous auctions (Cramton [1997] discusses why the FCC chose to use sequential auctions), and that budget constraints can arise endogenously in addition to externally. That is, agents can do better if they choose to be budget constrained, provided that the other agents are aware of this constraint and adjust their strategies accordingly. Pitchick [1996] also studies budget-constrained bidders in sequential auctions, finding that the order of sale can affect price.

Our study reported in this paper significantly extends this budget-constraint research to consider multiple bidders, multiple goods, and private values (bidders value goods differently) for the goods in simultaneous, multiple-round auctions. We retain the constraint that bidders have complete knowledge of each other’s values for the goods.

7 Conclusion and Future Work

This paper introduces **FAucS** as a rich and realistic simulator for studying autonomous bidding agents in a complex, real-world domain. We have illustrated its usefulness via quantitative experiments using prototype agents in one randomly generated, competitive scenario over a series of runs. The chosen scenario is closely modeled after the most recent FCC spectrum auction, Auction 35.

Our on-going research agenda includes more agent development and experimentation in this new domain. For example, although we empirically demonstrated the *potential* benefits of SDR, they cannot be realized until the agents can transmit and understand bidding signals. In related research, we have studied abstract mechanisms for agents to issue and receive such signals in a game-theoretic scenario [Littman and Stone, 2001].

The BSA strategy relies on the other agents all using the BA strategy. Our research agenda includes using **FAucS** to explore potential responses to BSA as well as competitive equilibria in this domain. BSA also relies on the agent having complete information about the market values of all the other agents. Although bidders do have some knowledge of each other’s values in general, we plan to investigate strategies that are less dependent on this assumption.

FAucS currently reflects all of the details of the United States FCC wireless frequency spectrum auction 35. However, it is implemented in such a way as to be easily modifiable to facilitate human and autonomous agent bidding in other FCC spectrum auctions including future combinatorial auctions that will allow bidders to bid on packages of licenses instead of simply single licenses.

Acknowledgements

We thank Mark Bradner and Tim Schaffer of AT&T Wireless for sharing their insights on FCC spectrum auctions.

References

- J. P. Benoit and V. Krishna. Multiple object auctions with budget constrained bidders. *Review of Economic Studies*, March 2000. In press, available at <http://econ.la.psu.edu/~vkrishna/research.html>.
- Peter C. Cramton. The FCC spectrum auctions: An early assessment. *Journal of Economics and Management Strategy*, 6(3):431–495, 1997.
- Eduard Gimenez-Funes, Lluís Godo, Juan A. Rodríguez-Aguiolar, and Pere Garcia-Calves. Designing bidding strategies for trading agents in electronic auctions. In *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 136–143, 1998.
- Amy Greenwald and Jeffrey O. Kephart. Shopbots and pricebots. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 506–511, 1999.
- Amy Greenwald and Peter Stone. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, 5(2):52–60, March/April 2001.
- Jessica Hall. U.S. carriers likely to be choosy in wireless auction. Reuters Newswire, November 2000.
- Paul Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13(3):227–86, July 1999.
- Michael L. Littman and Peter Stone. Leading best-response strategies in repeated games. Research note, 2001.
- Carolyn Pitchick. Budget-constrained sequential auctions with incomplete information. Unpublished manuscript. Under revision November 2000., 1996.
- Juan A. Rodríguez-Aguilar, Francisco J. Martín, Pablo Noriega, Pere Garcia, and Carles Sierra. Towards a test-bed for trading agents in electronic auction markets. *AI Communications*, 2000. In press. Available at <http://sinera.iiia.csic.es/~pablo/pncve.html>.
- Peter Stone, Michael L. Littman, Satinder Singh, and Michael Kearns. ATTac-2000: An adaptive autonomous bidding agent. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001. To appear, available at <http://www.research.att.com/~pstone/papers.html>.
- Robert J. Weber. Making more from less: Strategic demand reduction in the FCC spectrum auctions. *Journal of Economics and Management Strategy*, 6(3):529–548, 1997.
- Michael P. Wellman, Peter R. Wurman, Kevin O’Malley, Roshan Bangera, Shou-de Lin, Daniel Reeves, and William E. Walsh. A trading agent competition. In press. Available at <http://tac.eecs.umich.edu/reports.html>, 2001.