

Combinatorial Auction Design for Bandwidth Trading

An Experimental Study¹

Charis Kaskiris
kaskiris@sims.berkeley.edu
School of Information Management & Systems
University of California, Berkeley

Rahul Jain, Ram Rajagopal, Pravin Varaiya
{rjain, ramr, varaiya}@eecs.berkeley.edu
Electrical Engineering and Computer Science
University of California, Berkeley

Abstract

We experimentally investigate a combinatorial double-sided auction mechanism for allocation of bandwidth between buyers and sellers. The purpose of the experiment is to investigate the efficiency of the mechanism proposed in Jain & Varaiya (2004a) as well as the bidder behavior. We have implemented the mechanism in a combinatorial auction experimental platform which was used for this investigation. We performed experiments utilizing a simplified version of the theory using two different network structures, one with full valuations over every link and another with valuations only on specific combinations with restricted supply. Experimental results show that the mechanism gets close but not achieve competitive equilibrium. The market conditions also affect the efficiency of the mechanism. In cases where the supply was restricted and the demand was on packages rather than single links, has disciplined sellers to overbid less and the buyers to underbid less. Experience from both buyers and sellers led to more efficient results. Observed prices were closer to competitive equilibrium prices under the constrained market conditions.

1. Introduction

We study the interaction amongst internet service providers who lease bandwidth from owners of individual links to form desired routes. These bandwidth markets have been a rather polarized issue since their inception in the late nineties. These dynamic markets were setup to deal with the inefficient bilateral negotiations used in the industry. Market dynamics, technical difficulties, and the collapse of Enron have brought these markets to an end (Borthick, 2001). Ferreira, Mindel & McKnight (2003) identified two main factors that hamper the implementation of dynamic and fluid spot markets for bandwidth trading. The first one has to do with the excessive time needed to disseminate new routing information, which is an artifact of the bilateral nature of contracting for routes. The second has to do with balance loading once carriers become multi-connected. The inefficiency of bilateral contracting is due to the exposure effect (Milgrom, 2000) from not being able to form attractive routes from individually leased links. This is a problem identified with markets of complimentary goods/services. The combinatorial auction market mechanism can be used to alleviate the exposure effect. Jain and Varaiya (2004b) propose a double-sided combinatorial auction mechanism for bandwidth trading which we experimentally explore in this paper. Auctions in general have gained prominence as resource allocation problems with the spectrum license auctions.

¹ Nisha Pilai and Jun Shu contributed to the design and implementation of the experimental software used in this experiment. We want to thank NetExchange Inc. and DARPA for financial support and the xLab for providing subject recruiting and facilities.

In the early nineties the Federal Communications Commission was chartered to auction spectrum licenses for wireless communication licenses. This was one of the first times that economics was used as engineering in solving real-world problems.² What was different with the spectrum auctions was the existence of complementarities between the different licenses. These economic environments have been experimentally investigated in the context of airline slot allocation (Rassenti, Smith & Buflin 1982), payloads for NASA's Space Station (Banks, Ledyard & Porter 1989), tracking routes (David et al., 2002), pollution license trading (Ishikida et al., 2001), and spectrum auctions (Ledyard, Porter & Rangel, 1997)(Plott, 1997)(Porter et al., 2003). Combinatorial auction problems have been investigated further within the computer science and economics communities. Vries and Vohra (2003) provide a survey of the current research in both communities.

Dealing with complex economic environments, where complementarities exist, has proven a formidable task for auction theorists. The theoretical properties of different auction formats such as the simultaneous ascending bid auctions and combinatorial auctions were poorly understood. Designers of such systems turned to experimental economics to investigate the properties of such mechanisms. Experimental economics is the application of the laboratory method to test the validity of various economic theories and to test new market mechanisms. Using cash-motivated subjects, economic experiments create real-world incentives to help us better understand why markets and other exchange systems work the way they do (Friedman & Sunder 1994)(Kagel & Roth, 1995). Many of these new auction mechanisms have been introduced with improvements in computational power of computing, especially combinatorial auctions. Combinatorial auctions are being studied by two main areas of research. Economists deal with the economic rationality of self-interested agents; Computer Scientists deal with the computational and informational constraints of such auctions.

2. Combinatorial Auctions

Combinatorial auctions enhance our ability to allocate multiple resources efficiently in complex economic environments due to their generalized bid expression. Their bid expression allows the auction participants to bid on packages of items with related values or costs. They also allow bidders to impose logical constraints that limit the feasible set of auction allocations. They also can handle functional relationships amongst bids or allocations such as budget constraints or aggregation limits. As such large flexibility is provided the task of devising an optimal strategy for bidding becomes a computationally intensive task for bidders and sellers.

Combinatorial auctions allow for more expressive bidding in which participants can submit package bids with logical constraints that limit allowable outcomes. This type of auction can be useful when

² Game theory was also used for the design for assigning physicians to hospitals using two-sided matching (Roth & Sotomayor, 1990)(Roth 1999).

participants' values are complementary or when participants have production and financial constraints. There are several reasons to prefer to have the bidding message space expanded. One of the problems that combinatorial auctions solve is the "exposure" problem, evident in simultaneous ascending auctions (Milgrom, 2000). With individual bidding, a bidder is "exposed" to the risk of winning a few licenses it wants without winning other complimentary licenses it wants. Fearing that, a bidder may not bid aggressively, not participate in the auction, or try to collude (Cramton & Schwarz, 2000). Hence, allowing combinatorial bidding has better efficiency in the presence of rigidities in the demand (i.e. a bidder extracts any valuation only if the whole package is fulfilled).

However, combinatorial auctions are currently rare in practice. The main problems confronted in implementing these auctions are that they have computational uncertainty with regards to winner determination with large numbers of items and participants. The auction is also cognitively complex and can lead participants to pursue perverse bidding strategies. They also lead to inefficiencies in cases where the "threshold" effect is in effect (Bykowsky, Cull & Ledyard, 2000). The threshold effect refers to the case when aggregating smaller bids would have displaced a larger bid, but the incentives to do so are not aligned.

The computational uncertainty in winner determination comes from the fact that winner determination in combinatorial auctions is equivalent to a Set Packing Problem (Vries & Vohra, 2003), which is non-deterministic polynomial time complete or hard problem. Bidding in combinatorial auctions is burdensome, both strategically and cognitively, for all participants including devising optimal strategies. In designing combinatorial auctions, a set of design questions need to be considered. How does the format of the auction withstand the "threshold effect"; does iterative bidding allow for strategy building through learning; what is the appropriate level of information feedback to the bidders; what is the computational cost of the algorithms proposed?

What has been observed in the field and during experiments is that in complex economic environments iterative auctions which permit the participants to observe the competition and learn when and how to bid, produce better results than sealed bid auctions. There are two current frameworks used for iterative procedures. The first one is the use of continuous auctions (Banks, Ledyard & Porter, 1989) during which bidders may see a set of provisional winning bids as well as a set of bids to be combined from a standby list. The standby list consists of non-winning bids and these bids are there to signal willingness to combine bids to outbid larger-package bids. The second one is the use of multiple rounds using sealed-bid formats (Kwansinka et al., 2002), which solves repeated integer programming problems. In general auction systems that provide feedback and allow bidders to revise their bids seem to produce more efficient outcomes (Porter et al., 2003)

In this paper we compare the revenue, efficiency, and bidding properties of the particular combinatorial auction setup in the presence of complementarities among the objects being allocated. Specifically, we conduct laboratory experiments allocating three links with private values and complementarities using the combinatorial auction format under different degrees of complementarity. In the benchmark case, every seller owns all types of links and every buyer has private values over all subsets of links. In alternative cases, sellers own two types of links with one type being owned by all sellers. Buyers have valuations over bundles of links but only over one singleton link.

The remainder of the paper proceeds as follows: section three presents the theoretical framework of the combinatorial seller bid double auction (c-SeBiDA); section four presents the information and valuation structure used in the experiments. Section five gives an overview of the experimental design. Section six reports the results of the experiments. Section seven concludes with a discussion.

3. Theoretical Background

The formal mechanism is described in Jain & Varaiya (2004), on which the exposition below is based. There is a network of links $1, \dots, L$. There are two types of players in this economy, buyers, indexed by i and sellers indexed by j . Each buyer i has a reservation utility of $v_i \in [0, V]$ per trunk along the route R_i and submits a buy bid of b_i per trunk and demands $\delta_i \in [0, D]$ trunks along the route. Seller j operates each trunk at cost $c_j \in [0, C]$ and offers to sell up to $\sigma_j \in [0, S]$ trunks per link l_j at a unit price of a_j . The set of links is denoted by $L_j = \{l_j\}$.

The mechanism receives all the bids and matches buy with sell bids together to generate the maximal surplus. The matches are described by indicator variables, x_i, y_j , each taking a value of 1 if accepted or a value of 0 if not. The mechanism determines the allocation (x^*, y^*) as a solution to the surplus maximization problem, which is a mixed-integer program, MIP:

$$\begin{aligned} & \max_{x,y} \sum_i x_i b_i - \sum_j y_j a_j \\ & s.t. \sum_j y_j 1(l \in L_j) - \sum_i x_i 1(l \in R_j) \geq 0, \forall l \in [1 : L] \\ & x_i \in [0 : \delta_i], \forall i \\ & y_j \in [0 : s_j], \forall j \end{aligned}$$

MIP generates the set of matched bids and asks which maximizes the total surplus. Uniform prices per link are assigned by the highest winning sell bid price.

$$\hat{p}_l = \max\{a_j : y_j^* > 0, l \in L_j\}$$

The payments made by buyers to sellers are determined by the number of trunks at the determined uniform prices. One further consideration is with regards to the total payment made by buyers. It is conceivable that the total amount owed may exceed the bid level b_i , then that bid is rejected as more surplus can be generated by removing it. Therefore the outcome of the auction respects individual rationality.

Strategic analysis of the mechanism is presented in Jain and Varaiya (2004b). Under full information, a Nash equilibrium exists for the combinatorial seller bid double auction, with everyone's strategy being truthful revelation, with the exception of the matched seller with the highest bid per link. The resulting equilibrium allocation is efficient. The proof of the theorem is presented in Jain & Varaiya (2004b) as proof of theorem 2. Under incomplete information, the mechanism is asymptotically Bayesian incentive compatible, provided that the both buyers and sellers have ex post individual rationality constraints. The proof for the theorem is presented in Jain and Varaiya (2004b) as proof of theorem 3.

We experimented using an incomplete information setup, where buyers and sellers had common knowledge of the distribution of buyer valuations and seller costs. We investigated the impact of having

4. Information and Valuation Structure

We propose the use of the economic experimental methodology to the design and understanding of mechanisms for allocating resources in engineering and electronic commerce applications. Even though economic theory has already been applied to engineering problems, most of the models are of theoretical nature. The appropriate use of economic theory in engineering needs to address human participation, and experimental economics provides a way for testing the robustness of such theories. They also provide an environment for formulating new theories and testing improved designs for such systems.

We propose investigating different properties of combinatorial auction settings, using the combinatorial auction platform developed by Kaskiris et al. (2003). The platform allows for single-sided, double-sided, XOR/OR, combinatorial bidding, and short-selling (Jain & Varaiya, 2004a). The particular interest is in the design of a combinatorial auction system that could potentially be used in the allocation of links and trunks in bandwidth trading markets.

Jain and Varaiya (2004b) propose a Combinatorial Seller's Bid Double Auction (c-SeBiDA) mechanism, which maximizes the auctioneer's profit and announces payments based on sellers' bids. They also show that the announced allocations and prices exist which represent a competitive equilibrium, under assumptions on bidding format and valuations. We replicate these conditions and also investigate how robust the mechanism is to differing private valuations. This will necessitate the use of different market environments with different valuations for combinations of goods. In particular, we want to replicate a simple bandwidth trading environment, where the bidding is over links (i.e. goods) and number of trunks

(i.e. quantity of goods) on each link. Obtaining different lengths of paths (i.e. combinations of links) provides different valuations for users.

We followed a similar valuation structure used in Morgan (2003), where users were given valuations over combinations of links and trunks. Subjects were provided with valuation charts over different combinations of goods at different quantities.

4.1 Sellers

Each seller owns a combination of links and trunks on those links. Each seller has a cost of operation of each item-trunk pair is drawn from a uniform discrete distribution between 5 and 15. The cost of each additional trunk within the same link is uniform. Operation costs are only incurred when a link-trunk combination is provided to a seller. There is no cost and no benefit associated with unsold link-trunk combinations. Table 1 presents an example of a seller who is endowed 3 trunks on each one of the links available.

Trunks	Link A	Link B	Link C
1	7	5	6
2	14	10	12
3	21	15	18

Table 1: Example of Seller Valuations

In the alternative setup, all sellers own 3 trunks of one type of link, and two sets of two sellers own 3 units of each remaining type of link.

Sellers can submit multiple bids (asks) with the restriction that bids cannot be combinatorial. Seller bids are loose, meaning that a single bid may be matched with multiple buyer bids. This would be the case when the seller submits a multi-trunk bid (e.g. \$12/trunk for 3 trunks of Link A).

4.2 Buyers

Buyers begin each round without owning any links and trunks. They are however provided with chart of private valuations over the all possible subsets and trunks that they may obtain. In the benchmark setup valuations for each subset of items are generated in the following way:

1. Valuation for each item is generated from a discrete uniform distribution between 10 and 20. For example, item A maybe valued at 12.
2. Valuation for each subset of two items is generated by adding the valuation for the two items. Then a number from a uniform distribution between 0 and 5 is added. For example, item A is valued at 12, item B is valued at 14, and the bundle AB is valued at 29.

3. Valuation for having all three items is generated by the maximum additive valuation between combinations of two items and the valuation of the remaining object. Then a number from a uniform distribution between 0 and 2 is added.
4. Each additional trunk for each combination is valued at -1 of the previous single item trunk; -2 of the previous double item combination; -3 of the previous all item combination.

Table 2 demonstrates an example of a valuation chart based on the procedure described above.

Trunks	A	B	C	AB	BC	AC	ABC
1	20	12	14	37	27	35	52
2	39	23	27	72	52	68	101
3	58	34	40	107	74	101	150

Table 2: Example of Buyer's Valuations

In the alternative setup, each buyer values different sets of links that include only one type of link (i.e. A, AB, AC, ABC). This type of link is never the type of link that all sellers own.

Buyers may bid on combinations of items, but are restricted to have an equal number of trunks (quantity) on each link (item). Buyer bids are not loose and need to be completely satisfied to be matched. All buyer bids are XORed together hence only one of them can match at each round.

The objective of each bidder would be to improve her endowment position through trading. Initially everyone starts with a level of endowment in goods and money. Users are induced to perform well by only rewarding changes from the initial endowment point. In each round of the experiment, subjects accumulate points based on their buyer/seller surplus they generated. At the end of the experiments, \$500 was split in proportion to the total surplus generated. If a participant does nothing s/he will receive nothing at the end of the experiment. Negative balances provide subjects with no payoff (except a show-up fee).

5. Experimental Protocol

The experiment consisted of a 3-hour experimental session which was conducted at the end of July 2004 at the xLab facilities at the University of California Berkeley. Subjects were recruited from the graduate programs in electrical engineering, information management and systems, and economics using e-mail postings. Participants were either required to be familiar with basic networking and/or auction understanding. There were two sessions of four rounds each. The subjects were instructed on how to bid using the web-based interface and also explained on how the system calculates prices and performs matching. Test runs were conducted so that the subjects could get a feel of how the system worked and where information was displayed.

In the first session, subjects had a 50% chance of being a buyer or a seller in each of four rounds. Each round had an equal number of sellers and buyers. Two sessions of four rounds each were conducted. During the first session, subjects participated in four rounds of using the Combinatorial Seller's Bid Double Auctions using the benchmark setup. During the second session, subjects participated in four rounds using the same auction format but with the alternative setup valuations.

To ensure that both sessions used the same procedures, we adopted a written protocol which we used on both sessions. In all sessions, the participants were seated in a large room, each sitting at a desk with a laptop computer. They were read instructions and given an opportunity to ask questions. Throughout the session, participants would only communicate through the submission of bids.

The submission of bids in the system was monitored through the server platform and once everyone submitted the bids were entered into the combinatorial engine which would provide prices for each link and which buyer and seller matched. The market information, namely the prices of the links, would be posted for everyone to see at the conclusion of the round. The participants who matched would also be notified of which link and what quantity they matched on. Before being asked to bid, participants received a handout depending on whether they were a seller or a buyer. At the conclusion of the experiment, the subjects were paid in private with checks according to their performance during the two sessions.

6. Results

We present the results for both sessions in terms of efficiency, revenues, and bidding behavior. We pool results from each round and present the average behavior of subjects during each session. We also present how the bidding behavior of subjects changed over consecutive rounds.

6.1. Efficiency and Revenue

We begin by considering how each auction setup fared in achieving the efficient allocation of the objects. Buyer efficiency is calculated by dividing the value of the objects actually realized by the bidders by the theoretical maximum obtainable. In the benchmark rounds of session 1, the valuations of the buyers were sufficiently high to be satisfied by the sellers selling all their links and trunks. Efficiency was then calculated by comparison between the actual outcomes and the maximum possible valuations of the buyers. Table 3 below shows the associated standard deviation (σ) is given below each of the mean figures for efficiency generated during each round of each session. During session 2, given the restricted supply on two links, the valuations were compared on the highest possible combination of valuations given the supply.

Round	Session 1				Session 2			
	1	2	3	4	1	2	3	4
Mean	56.24	37.65	76.31	81.20	63.19	50.00	55.42	71.19
σ	32.82	11.43	31.93	12.99	33.74	57.74	31.75	33.27

Table 3: Summary of Buyer Percentage Efficiency in Each Round³

For the mechanism to be efficient, it should be able to induce participants with high valuations to be the ones who match. Buyers with the highest valuations for each combination of links and trunks, can be identified from their valuation tables. Given our mechanism, the prices should be the highest bids of the sellers with the highest operating costs. This is because of the distribution functions used in assigning valuations to buyers and sellers. What we observe is that even in the case of limited supply, the mechanism performs at about 60% buyer efficiency. What we also observe is that during session 1, subjects have performed better with each new session. Figure 1 shows the progression over time of the percent efficiency achieved by buyers over the four rounds of each session. What is observed it that as participants gained more experience they tended to perform better.




Figure 1: Buyer Valuation Efficiency Per Round

On the seller side efficiency was calculated as percent of items sold given their initial endowments. Table 4 shows mean percentage efficiency of the mechanism in distributing the sellers links and trunks.

³ In session 2 round 3 a seller sold one more unit that s/he had, for which /she got penalized. The efficiency reading on the buyer side was not adjusted for that problem, since the difference would be minimal.

Round	Session 1				Session 2			
	1	2	3	4	1	2	3	4
Mean	55.56	41.67	72.22	75.00	63.89	66.67	58.33	72.22
σ	17.35	22.05	4.81	0.00	17.35	28.87	22.05	25.46

Table 4: Summary of Seller Percentage Efficiency of Sales in Each Round

What we observe is the mechanism performs at about 65% seller efficiency when there is no restrictions on supply or demand and 61% of seller efficiency when there are restrictions on demand and supply. What we also observe is that during session 1, subjects have performed better with each new session, a result reflected in the buyer efficiency discussion.




Figure 2: Seller Percent Efficiency

Figure 2 shows the percent efficiency per round achieved by the sellers. We observe that overall as sellers gained experience with bidding, they tended to improve the efficiency of their final allocations.




Figure 3: Session 1 Differences Between Actual Prices and Competitive Equilibria Over Different Rounds
 Aside revenues, we also examine whether the market clearing prices during each round were close to the theoretical prices. Figure 3 shows session 1 deviations from the competitive equilibrium assigned to each link versus the actual prices generated by the mechanism.

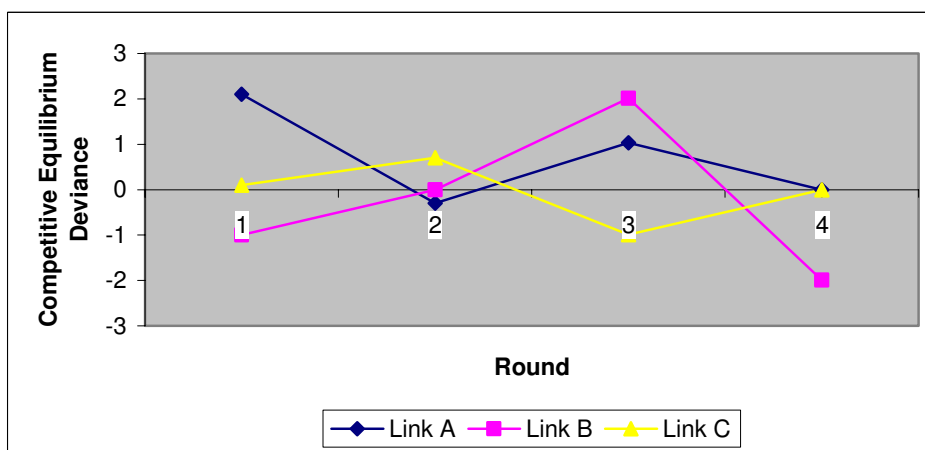


Figure 4: Session 2 Differences Between Actual Prices and Competitive Equilibria Over Different Rounds
 Figure 4 shows the session 2 deviations from the competitive equilibrium assigned to each link versus the actual prices generated by the mechanism. What we observe is that prices did vary from the competitive equilibrium with the session 2 results being closer to the theoretical competitive equilibrium.

6.2. Bidding

Another important aspect of investigating a mechanism is with regards to patterns of underbidding or overbidding. This means that we need to pay closer attention to the bidding of participants when their type of bidder (buyer or seller) changes between rounds. The shading factor is a measure of how much

lower than the actual valuation of a link-trunk combination has been submitted during each round. The shading percentage for each buyer i is given below:

$$d_i = \frac{v_i - b_i}{v_i}$$

The shading factor for each seller j is given below:

$$d_j = \frac{a_j - c_j}{c_j}$$

The average shading factor of both buyers and sellers for each round is shown in Table 5. What we observed was that sellers tended to overbid above their own costs of operation. What we did not observe however is bidding with regards to the expected price of each good. Since operating costs were randomly drawn from a discrete uniform distribution between 5 and 15, then the expected operating cost is 10. Suppliers who had operating costs of 5 would bid marginally higher, possibly reflecting high risk-averse behavior.

Round	Session 1				Session 2			
	1	2	3	4	1	2	3	4
Mean	23.95	28.79	19.30	31.85	27.00	18.00	13.00	23.00
σ	24.42	26.15	8.59	30.06	26.00	9.00	14.00	21.00

Table 5: Aggregate Average Percentage Shading Factor Per Round

Table 6 shows the overbidding behavior of sellers with respects to their true costs. The impact of this bidding behavior is also reflected in the ability of the mechanism to assign the competitive equilibrium prices, since the final uniform price per link is the maximum successful seller ask bid value.

Round	Session 1				Session 2			
	1	2	3	4	1	2	3	4
Mean	27.51	14.59	23.44	32.40	8.75	14.25	11.33	30.75
σ	35.14	9.97	2.83	35.10	11.81	12.01	11.50	24.96

Table 6: Seller Overbidding Percentage Over Cost of Operation

Buyers would underbid in most cases as reflected below. Given the way that prices were determined in the mechanism, specifically, by having the highest accepted seller bid dictate the price, some risk-loving buyers bid with prices exceeding their own valuations in expectation that the actual price paid would be lower. Table 7 shows the underbidding behavior of buyers.

Round	Session 1				Session 2			
	1	2	3	4	1	2	3	4
Mean	20.39	42.98	18.08	31.31	44.50	21.50	14.25	16.00
σ	11.08	30.98	10.08	29.59	24.06	5.32	13.78	14.90

Table 7: Buyer Overbidding Percentage over Cost of Operation

Our results show that the mechanism does not induce truth-revelation as a VCG auction (Morgan 2002), Sellers tend to overbid by 26% and buyers underbid by a 20% on average. Comparisons of this mechanism with the alternative mechanism of simultaneous double auctions of each link or the VCG auction format, could shed some light on the comparative efficiency and strategic interactions of this auction mechanism. A mitigating factor might have to do with the number of bidders in the auction.

6.3. Role of Experience

The role of experience over participating in the auction over time let to lower underbidding by buyers but higher overbidding by sellers. The session 2 market environment has consistently created incentives for smaller shading factors for both buyers and sellers.




Figure 5: Buyer Underbidding Behavior

The Wilcoxon rank-sum test between the two distributions of overbidding by sellers (p -value = 0.69; Figure 6) and for underbidding by buyers (p -value=0.20; Figure 5), rejects the null that the distributions across sessions are the same. Hence, market structure does affect bidding behavior.




Figure 6: Seller Overbidding Behavior

7. Discussion

Experimental economic approaches can be used to aid engineers in the design of mechanisms for allocation of resources which exhibit complimentary value to users. Such environments include bandwidth trading, spectrum auctions, airport slot planning, hospital staff scheduling, utility pricing, etc. In many cases the theoretic properties of different allocation mechanisms are unknown. Similarly, implementation challenges with respect to such systems can be identified through the use of experimental economic methods. We have briefly investigated some basic properties of the mechanism developed by Jain and Varaiya (2004b) and investigated its properties. Investigating these properties however was not straightforward. Initial dry run experimental investigations with undergraduate students revealed that understanding the structure of the experiment and formulating strategies was a formidable task. Much of the bidding behavior was in error and participants would get confused over the interchange of being a buyer or a seller. Learning effects have been critical in understanding and getting a feel for the experimental setup and subsequent bidding behavior. Given the theoretical asymptotic efficiency results, we intent to vary the number of market participants in future investigations and observe the revenue efficiency of the mechanism. We will further examine the efficiency of expanding the bidding space allowing for buyers to submit both XOR and OR bids. A further option is to permit short-selling using multi-

round setups. We also want to investigate the impact on efficiency when we allow the combinatorial auction to be multi-round ascending.

References

- Banks, J., Ledyard, J. and Porter, D. "Allocating Uncertain and Unresponsive Resources: An Experimental Approach." *The RAND Journal of Economics*, 1989, 20(1), pp. 1-25.
- Borthick, S. "Bandwidth Trading: What's in it for you?," *Business Communications Review*, pp. 26-30, Jan. 2001.
- Bykowsky, M.; Cull, R. and Ledyard, J.. "Mutually Destructive Bidding: The FCC Auction Design Problem." *Journal of Regulatory Economics*, 2000, 17(3), pp. 205-28.
- Clearwater, S. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. Singapore, River Edge, N.J.: World Scientific, 1996.
- Cramton, P. and Schwartz, J. "Collusive Bidding: Lessons from the FCC Spectrum Auctions." *Journal of Regulatory Economics*, 2000, 17(3), pp. 229-52.
- Ferreira, P., Mindel, J., McKnight, L. "Why have bandwidth trading markets not matured? Analysis of technological and market issues," *International Journal of Technology, Policy and Management* 2003 - Vol. 3, No.2 pp. 142 - 160
- Friedman, D. and Shyam, S. *Experimental Methods: A Primer for Economists*. New York: Cambridge University Press, 1994.
- Jain, R. and Varaiya, P. "Efficient Bandwidth Allocation Through Auctions," (in submission) 2004.
- Jain, R. and Varaiya, P. "An Efficient Incentive-Compatible Combinatorial Market Mechanism" Allerton, 2004.
- Kagel, J. and Roth, A. *The Handbook of Experimental Economics*. Princeton, N.J.: Princeton University Press, 1995.
- Kahneman, D. and Tversky, A. *Choices, Values, and Frames*. New York: Cambridge University Press, 2000.
- Kaskiris, C., Pillai, N., Rajagopal, R., and Varaiya, P. "The Berkeley Combinatorial Auction Platform," Berkeley: University of California, 2003.
- Klemperer, P. "Auction Theory: A Guide to the Literature." *Journal of Economic Surveys*, 1999, 13(3), pp. 227-86.
- Krishna, V. *Auction Theory*. San Diego: Academic Press, 2002.
- Ledyard, J., Porter, D. and Ragel, A. "Experiments Testing Multi object Allocation Mechanisms." *Journal of Economics and Management Strategy*, 1997, 6(3), pp. 639-75.
- Milgrom, P. *Putting Auction Theory to Work*. Cambridge, UK; New York: Cambridge University Press, 2004.
- Morgan, J. "Combinatorial Auctions in the Information Age: An Experimental Study," M. Baye, *Advances in Applied Microeconomics*. JAI Press, 2002.

- Plott, C.. "Laboratory Experimental Testbeds: Application to the PCS Auction." *Journal of Economics and Management Strategy*, 1997, 6(3), pp. 605-38.
- Porter, D. Rassenti, S., Roopnarinen, Anil and Smith, Vernon. "Combinatorial Auction Design." *Proceedings of the National Academy of Sciences*, 2003, 100(19), pp. 11153-57.
- Porter, David; Torma, David P.; Ledyard, John O.; Swanson, Joseph A. and Olson, Mark. "The First Use of a Combined-Value Auction for Transportation Services." *Interfaces*, 2002, 32(5), pp. 4-12.
- Roth, Alvin E. and Peranson, Elliott. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *The American Economic Review*, 1999, 89(4), pp. 748-80.
- Roth, Alvin E. and Sotomayor, Marilda A. Oliveira. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. New York: Cambridge University Press, 1990.
- Sandholm, Tuomas. "Emediator: A Next Generation Electronic Commerce Server," *Proceedings of the fourth international conference on Autonomous agents*. Barcelona, Spain: ACM Press, 2000, 341 - 48.
- Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *The Journal of Finance*, 1961, 16(1), pp. 8-37.
- Vries, Sven de and Vohra, Rakesh V. "Combinatorial Auctions: A Survey." *INFORMS Journal on Computing*, 2003, 15(3), pp. 284-309.
- Wurman, Peter R. and Wellman, Michael P. "Akba: A Progressive, Anonymous-Price Combinatorial Auction," *Proceedings of the 2nd ACM conference on Electronic commerce*. Minneapolis, Minnesota, United States: ACM Press, 2000, 21-29.
- Wurman, Peter R.; Wellman, Michael P. and Walsh, William E. "The Michigan Internet Auctionbot: A Configurable Auction Server for Human and Software Agents," *Second International Conference on Autonomous Agents*. Minneapolis, Minnesota, United States: ACM Press, 1998, 301 - 08.

Appendix: Experimental Instructions

EXPERIMENTAL INSTRUCTIONS

This is an experiment to compare different mechanisms for allocating tradable resources. We will specifically work in the context of a bandwidth trading market. Funding for this research has been provided through DARPA. The instructions are simple and if you follow them carefully and make good decisions you might earn a considerable amount of money which will be paid to you in cash at the end of the experiment.

We will contact a set of different sessions, each of which will be testing a mechanism under particular market conditions. Each session will consist of 4 independent rounds. During each round there will be an equal number of buyers and sellers. Buyers and sellers are randomly assigned in the beginning of each round and will be denoted on the top of the experimental dossier you will receive at the beginning of each round.

Two kinds of papers will be distributed – information for buyers and information for sellers. The sheets are identified and numbered. The numbers are strictly for data collection purposes. If you have received seller's information, you will function only as a seller in this market. Similarly, if you have received buyers' information, you will function only as a buyer in this market. The information you have received is for your own private use. **DO NOT REVEAL IT TO ANYONE.**

MARKET

The experimental market you are participating in simulates a bandwidth trading market. Three items (links) {A, B, C} each of which has capacity of three trunks (quantities), are traded in the market. Trunks reflect discrete quantities of bandwidth. The links represent communication (telephone) links, and a link's bandwidth is measured in trunks (one trunk can carry 100 telephone calls simultaneously).

SELLERS

Each seller owns a combination of links and trunks on those links. Each seller has a cost of operation of each item-trunk pair drawn from a uniform discrete distribution between 5 and 15 per unit. The cost of each additional trunk within the same link is uniform. Operation costs are **ONLY** incurred when a link-trunk combination is provided to a seller. There is no cost and no benefit associated with unsold link-trunk combinations. Table 1 presents an example of a seller who is endowed 3 trunks on each one of the three links. As a seller you may sell up to the total amount of trunks per links you own. You may submit as many bids as you like as long as the total amount trunks is not exceeded (don't sell something you don't own). Only non-combinatorial bids will be admissible for matching.

Trunks	A	B	C
1	7	5	6
2	14	10	12
3	21	15	18

Table 8: Example of Seller Valuations

BUYERS

Buyers begin each round without owning any links and trunks. They are however provided with a chart of private valuations over the all possible subsets and trunks that they may obtain. Valuations for each subset of items are generated in the following way:

1. Valuation for each item is generated from a discrete uniform distribution between 10 and 20. For example, item A maybe valued at 12.
2. Valuation for each subset of two items is generated by adding the valuation for the two items. Then a number from a uniform distribution between 0 and 5 is added. For example, item A is valued at 12, item B is valued at 14, and the bundle AB is valued at 29.
3. Valuation for having all three items is generated by the maximum additive valuation between combinations of two items and the valuation of the remaining object. Then a number from a uniform distribution between 0 and 2 is added.
4. Each additional trunk for each combination is valued at -1 of the previous single item trunk; -2 of the previous double item combination; -3 of the previous all item combination.
5. At the beginning of each round, each buyer owns no trunks.
6. ONLY one bid will match per round.

Table 2 demonstrates an example of a valuation chart based on the procedure described above.

Trunks	A	B	C	AB	BC	AC	ABC
1	20	12	14	37	26	35	52
2	39	23	27	72	50	68	101
3	58	34	40	107	74	101	150

Table 9: Example of Buyer’s Valuations

Buyers may submit a number of bids, but only one bid can be matched in each round. Sellers are restricted to providing bids on each individual item {A,B,C} and not on combinations of items. Buyers may bid on combinations of items, but are restricted to have an equal number of trunks (quantity) on each link (item).

NETWORK STRUCTURE

During some of the sessions the network structure will be such that the buyer valuation table will also admit zero valuations. Similarly, sellers might also not own a particular type of link. The selection of which links will not be available is randomly selected prior to each session. You will be notified of what type of network you are competing in at the beginning of each session.

BIDDING

Bids are submitted in the Bid Submission part of the interface, which you can access by pressing the [BID SUBMISSION] menu of the Combinatorial Auction Experiment Screen. The BLUE section denotes the area where bids can be submitted. There are three main components to a bid:

- Price per unit: the price per unit of the link/bundle of links you specify.
 - Quantity: the number of trunks you wish to purchase on each link...
 - Links: select which combination of links you wish the price and quantity to be applied to.
- You may add additional bids by clicking on [ADD NEW BIDS]. Do not submit empty bids. If you make a mistake click again on [BID SUBMISSION] and re-enter all your bids.

Specific Instructions for Sellers:

- In order to submit a sell bid (ask) you need to specify your price as a **negative number**. For example, if you want to sell 3 trunks of link1 for \$10/trunk, you specify Price = -10, Quantity = 3 and select A.
- As a seller you can **ONLY** select one link per bid. Combinatorial bids will be dropped.
- Seller bids can potentially all match (they are OR bids).
- You should not make bids for quantities higher than the ones you own (i.e. if full table, then you own 3 trunks on each link). Failing to meet this goal will lead to disqualifying your gains from a particular round.

Specific Instructions for Buyers:

- In order to submit a bid you need to specify your price as a positive number.
- You may also submit a number of bids with different combinations of selected links and/or trunks.
- If you submit multiple bids, **ONLY** one can be matched.

MARKET CLEARING

The combinatorial auction system matches the combination of seller and buyer bids which generates the highest surplus. Uniform prices are generated for each link and each participant pays the combination of prices based on their successful bids. The uniform price generated is the maximum successful seller bid price for each item.

Submitting Bids

You can submit multiple bids by clicking on the [ADD NEW BID] button in the blue row. When you are done formulating bids, click on the [SUBMIT BIDS] button in the RED striped row. **Only do this when you are completely done with bidding.**

IMPORTANT THINGS TO REMEMBER

- All prices submitted as prices per unit of a particular link/bundle of links.
- Sellers must denote prices by negative values.
- Sellers can only submit single link bids (all combinatorial bids will not be admitted)
- Sellers can only sell as many trunks on a link as they own.
- All Seller bids may match.
- Only one Buyer bid can match.
- Do not submit empty bid rows.

RESULTS

The results of the auction can be viewed by clicking on [VIEW RESULTS]. The results will be made available once the experimenter announces that the bids have been processed.

How You Make Money

Each participant in this experiment is given \$5 participation fee. On top of the participation fee, you will also be paid an additional award based on how well you perform. Performance is calculated as the change from your starting position of money and valuations. Sellers start with a set of links and trunks, which provide zero value to them. The performance of a seller is calculated as the difference between money received and operation costs of the successful bids. The performance of a buyer is calculated as the difference between the valuation of the final set of link-trunk combinations and the total costs of the purchase of those combinations. We will apply an exchange rate between experimental dollars and real dollars as to share \$500 for the whole experiment.

Experimental Trading Mechanism

The experiment will proceed as follows. You will be provided with your valuation table. Each participant may bid according to their own type. Once all bids are collected, the experimenter will run the solver and generate the matching results. The results will be disseminated in the form of the prices for each link. Each participant will be informed whether they have acquired/sold different links and at what priced. You will be instructed on where to view the results.

Then we will repeat the same experiment with new random assignment of roles. We will run 4 runs of each session. You will be paid at the end of the whole set of experiments.