Available online at www.sciencedirect.com

SCIENCE DIRECT®

Games and Economic Behavior 55 (2006) 321–330

GAMES and Economic Behavior

www.elsevier.com/locate/geb

Note

Side constraints and non-price attributes in markets ☆

Tuomas Sandholm [a,*], Subhash Suri [b]

[a] *Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA*
[b] *Department of Computer Science, University of California, Santa Barbara, CA 93106, USA*

Received 26 December 2001

Available online 19 July 2005

Abstract

In most real-world (electronic) marketplaces, there are additional considerations besides maximizing immediate economic value. We present a sound way of taking such considerations into account via side constraints and non-price attributes, and show that side constraints (such as budget, limit on the number of winners, and exclusive-or) have a significant impact on the complexity of market clearing.

JEL classification: D44; C63; L81; L86

1. Introduction

For a long time, auctions and exchanges have been proposed as mechanisms for allocating items (resources, tasks, etc.) in multiagent systems, especially systems consisting of self-interested agents. Some of the market mechanisms that lead to economically efficient outcomes are computationally complex to clear. In particular, there has been a recent surge of interest in algorithms for clearing auctions where bids can be submitted on pack-

☆ This work was funded by, and conducted at, CombineNet, Inc., 311 S. Craig St., Pittsburgh, PA 15213. An early version of this paper appeared at the IJCAI-01 Workshop on Distributed Constraint Reasoning, pp. 55–61, Seattle, WA, August, 2001.

* Corresponding author.
E-mail addresses: sandholm@cs.cmu.edu (T. Sandholm), suri@cs.ucsb.edu (S. Suri).

0899-8256/$ – see front matter
doi:10.1016/j.geb.2005.06.001

ages of items (e.g., Rothkopf et al., 1998; Sandholm, 2002a, 2002b; Fujishima et al., 1999; Lehmann et al., 2002; Sandholm and Suri, 2003; Nisan, 2000; Tennenholtz, 2000). The bulk of that literature has focused on clearing the auction so as to maximize unconstrained economic value. In most real-world marketplaces, especially in business-to-business commerce, there are other considerations besides maximizing immediate economic value that must be taken into account.

In this paper we introduce and analyze two methods for incorporating these additional considerations. First, we study side constraints on the trading outcome. Second, we show how non-price attributes can be soundly integrated into markets with package bidding and side constraints. (Side constraints and attributes could be imposed by any participant in the market: the buyer(s), the seller(s), the marketplace executor, the technology provider, or a regulatory body such as the SEC, etc.) Our main focus is to explore what impact these features have on the computational complexity of market clearing. (Throughout the paper we use the term "market" to mean a centrally cleared auction, reverse auction, or exchange.)

We first discuss side constraints in markets where bids are on individual items. Without side constraints such markets are easy to clear, but we show that under several practical side constraints, the market clearing becomes $\mathcal{NP}$-hard. We then move to markets where bids can be submitted on packages of items, and show the complexity implications of side constraints. Finally, we show how to integrate non-price attributes into markets with package bids and side constraints.

2. Bids on individual items

We will show that certain practical side constraints can make even noncombinatorial auctions hard to clear.

Definition 1 *(WDP).* The seller has m items (one unit each) to sell. Each bidder places a set of bids on individual items. The *winner determination problem (WDP)* is to determine a revenue-maximizing allocation of items to bidders.[1]

In the absence of side constraints, WDP can be solved in polytime by picking the highest bid for each item independently. The rather natural budget constraint[2] below illustrates how sharp the $\mathcal{P}$ vs. $\mathcal{NP}$-complete cutoff is in the space of side constraints. This is especially surprising since a similar constraint, where the *number of items sold to each bidder* is constrained, leads to a WDP that was recently shown to be polytime solvable using *b-matching* (Tennenholtz, 2000).[3]

[1] As usual, the clearing problem is defined with respect to the given bids, which may or may not be truthful.

[2] Budget constraints occur naturally in markets, and they have been studied from the *bidding* perspective in the literature before (Rothkopf, 1977).

[3] Multi-item auctions (with bids on individual items only) with certain types of structural side constraints are also solvable in polytime using b-matching (Penn and Tennenholtz, 2000).

Definition 2 *(BUDGET).* WDP where the amount sold to any bidder does not exceed her budget.

Theorem 2.1. *The* (*decision version of*) *BUDGET is $\mathcal{NP}$-complete, even with integer prices.*[4]

Proof. We reduce PARTITION (Garey and Johnson, 1979) to BUDGET. In the PARTITION problem, we have a set of integers $S = \{x_1, x_2, \ldots, x_n\}$, and the goal is to partition S into two subsets A and B (i.e. $A \cap B = \emptyset$ and $A \cup B = S$) s.t.

$$z = \sum_{i \in A} x_i = \sum_{i \in B} x_i,$$

where $z = (1/2) \sum_{x \in S} x$. We create an instance of BUDGET as follows. Corresponding to each x_i, we create an item i. There are two bidders, say, Andy and Bob; each places the bid of same price, x_i, for item i. The budget for Andy and Bob each is z (half of the total). This instance of BUDGET has a solution with revenue $2z$ if and only if the original partition problem has a valid partition. □

Another practical side constraint is the number of winners.

Definition 3 *(MAX-WINNERS).* WDP where at most k bidders receive items.

Theorem 2.2. *The* (*decision version of*) *MAX-WINNERS is $\mathcal{NP}$-complete, even with integer prices.*

Proof. We reduce SET-COVER (Garey and Johnson, 1979) to MAX-WINNERS. Given an instance of set cover, namely, a ground set $X = \{1, 2, \ldots, m\}$, and a set of subsets $\mathcal{F} = \{S_1, S_2, \ldots, S_n\}$, where $S_i \subseteq X$, we formulate an instance of MAX-WINNERS as follows. We create an item i for each element i in the ground set X. Corresponding to each set S_i, we create a bidder B_i, who places a \$1 bid on each item in the set S_i. We claim that there is a set cover of size k (or less) if and only if the auction has a solution with revenue m and max number of winners k.

We first argue that the sets corresponding to the k winning bidders form a set cover. (That is, if bidder i receives at least one item, then we put the set S_i in the cover.) Since the revenue is m, each item must be awarded to some bidder, and hence it must be covered by the set cover. Conversely, consider a solution to the set cover. For each set S_i in the cover, we make bidder i a winner. Since each item is covered in the set cover, each item is bid upon by at least one bidder in the just constructed winning set, but the item may be claimed by more than one winning bidders. However, since each bid is for the same price, we can arbitrarily award the item to any of the winning bidders claiming this item. This gives a solution to the auction problem with revenue m and number of winning bidders k. □

[4] Recently, the $\mathcal{NP}$-completeness of winner determination in auctions with budget constraints has been independently proven by Lehmann et al. (Lehmann et al., in press). (They use a different reduction.)

In some settings, a bidder may want to submit bids on multiple items, but may want to mutually exclude some of the items. For example, a buyer may want to buy a VCR and a TV, and either of two TVs (but not both) would be acceptable. She could express this by placing bids on each of the three items, and inserting an XOR-constraint between the bids on the TVs.

Definition 4 *(XORS).* WDP with XOR-constraints. Whenever two bids are combined with XOR, at most one of them can win.

Theorem 2.3. *The* (*decision version of*) *XORS is $\mathcal{NP}$-complete, even with integer prices.*

Proof. We reduce INDEPENDENT-SET (Garey and Johnson, 1979) to XORS. Corresponding to each vertex, generate an item and a \$1 bid for that item. Corresponding to each edge, insert an XOR-constraint between the bids. Now, XORS has a solution of \$$k$ iff there is an independent set of size k. □

Remark 1. In each of the three side constraint classes above, the clearing problem remains $\mathcal{NP}$-complete even if there is free disposal, that is, even if the seller is not required to sell everything.

Remark 2. In MAX-WINNERS and XORS, the problem remains $\mathcal{NP}$-complete even if bids can be accepted partially. (This is because a fractional solution can be converted into a solution where bids are accepted all-or-nothing without affecting the number of sets in the set cover in MAX-WINNERS and the number of vertices in the independent set in XORS. Therefore, a polynomial-time algorithm for the fractional case would yield a polynomial algorithm for the all-or-nothing case, contradicting the hardness of the all-or-nothing case.) In BUDGET, on the other hand, allowing partial bids makes the problem solvable in polynomial time—using linear programming.

Remark 3. In XORS, if each bidder places an XOR-constraint between *every* pair of his bids (so that at most one of his bids can be accepted), then the problem becomes the *assignment problem*, which can be solved in polynomial time (Kuhn, 1955).

3. Bids on packages of items

In this section we discuss how the complexity of clearing a market with package bids changes as different types of side constraints are imposed on the outcome. In a *combinatorial auction* (*CA*), bidders may submit bids on packages of items. This allows the bidders to express the fact that the value of a packages of items may differ from the sum of the values of the individual items that constitute the package.

Definition 5 *(CAWDP).* The auctioneer has a set of items, $M = \{1, 2, \ldots, m\}$, to sell, and the buyers submit a set of bids, $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$. A bid is a tuple $B_j = \langle S_j, p_j \rangle$, where $S_j \subseteq M$ is a set of items and p_j, $p_j \geqslant 0$ is a price. The *combinatorial auction* *winner determination problem* (*CAWDP*) is to label the bids as winning or losing so as to maximize the auctioneer's revenue under the constraint that each item can be allocated to at most one bidder:

$$\max \sum_{j=1}^{n} p_j x_j \quad \text{s.t.} \quad \sum_{j \mid i \in S_j} x_j \leqslant 1, \quad i = 1, 2, \ldots, m.$$

If there is no free disposal (auctioneer is not willing to keep any of the items, and bidders are not willing to take extra items), an equality is used in place of the inequality.

In the binary version (BCAWDP), each bid must be accepted fully or not at all. The continuous version (CCAWDP) permits bids to be accepted partially. The binary version is $\mathcal{NP}$-complete (Rothkopf et al., 1998) (and inapproximable (Sandholm, 2002a)) while the continuous version can be solved in polynomial time using linear programming.

A *combinatorial reverse auction* (Sandholm et al., 2002) models situations with one buyer and multiple sellers, where each seller can place bids on self-selected packages of items. Similar to the forward auction, we have the binary and continuous version of the winner determination problem for reverse auction as well.

Finally, in a *multi-unit combinatorial exchange*, both buyers and sellers can submit package bids, and in one bid, a bidder might be selling units of some items and buying units of other items (Sandholm, 2002b; Sandholm and Suri, 2003). Thus, we have the multi-unit combinatorial exchange winner determination problem (MUCEWDP), with its binary and continuous versions.

While most research on combinatorial auctions has focused on the binary version, there exist numerous important continuous combinatorial markets as well. For example, CombineNet, Inc. and Manhattan Associates, Inc. run combinatorial reverse auctions for long-term trucking lanes (the volume of each lane is numerous truck-loads), where the carriers' bids can be accepted fractionally. As another example, `BondConnect` was a combinatorial bond exchange, and there it was possible to accept bids fractionally. In practice, whether the market should be binary or continuous depends on the items being traded. By using the continuous version, better (at least no worse) objective values in the optimization are achieved. If the items are arbitrarily divisible, the continuous version is applicable. If the items are not divisible, only the binary market is applicable.

It turns out that different side constraints introduce sharp cutoffs in the complexity of clearing a market with package bids. Seemingly similar side constraints lead to problems that lie on different sides of the $\mathcal{P}$ vs. $\mathcal{NP}$-complete cutoff. In the first subsection we present side constraints under which the continuous case remains easy and the binary case remains hard. In the next subsection we present side constraints that make both cases hard. In the last subsection we present a side constraint that make both cases easy.

3.1. Side constraints under which the continuous case remains easy, and the binary case remains hard

The following classes of domain-independent side constraints, which we view as practically important and quite general, turn out to be easy for the continuous winner determination problem, and remain hard for the binary case. All of these constraints impose a lower or upper bound on the bids from a certain subset with respect to the bids from another subset. There is a long list of such constraints. Due to lack of space, we just mention two such classes. An interested reader can see the full paper for a more comprehensive catalog (Sandholm and Suri, 2001).

1. *Maximum or minimum trade constraint*. Specifying an upper or lower bound on the total acceptance for a certain set of bids $\mathcal{B}' \subseteq \mathcal{B}$. The acceptance can be measured in any reasonable way, such as revenue (net or gross), absolute amount or percentage of total, number of item units or volume in currency. The bound can be also be imposed relative to *another* set of bids $\mathcal{B}'' \subseteq \mathcal{B}$.
2. *Item trades*. Specifying a lower bound on the amount of trade (absolute or percentage) for a subset of items $M' \subset M$.

Theorem 3.1. *The binary versions of the combinatorial auction* (*BCAWDP*)*, combinatorial reverse auction* (*BCRAWDP*)*, and combinatorial exchange* (*BMUCEWDP*) *are $\mathcal{NP}$-complete for any constraint from the classes presented in this section, even with integer prices.*

Remark. The continuous versions of these problems can be solved in polynomial time using linear programming.

3.2. Side constraints under which the continuous and binary case are hard

The most interesting results of this paper show that some classes of side constraints that are among the most important ones from a practical perspective, make even the continuous case $\mathcal{NP}$-complete to clear.

Theorem 3.2. *If no more than k winners are allowed, then combinatorial auctions, reverse auctions, and exchanges are all $\mathcal{NP}$-complete, both in their binary and continuous versions. These problems remain $\mathcal{NP}$-complete with or without free disposal and even if the prices are integer.*

This theorem follows from the fact that MAX-WINNERS is $\mathcal{NP}$-complete even if bids can be submitted on individual items only (Theorem 2.2).

In a combinatorial auction where the bids are combined with OR, a bidder can only express complementarity, not substitutability. For example, say a bidder has submitted three bids: $\langle\{1\}, \$4\rangle$, $\langle\{2\}, \$5\rangle$, and $\langle\{1, 2\}, \$7\rangle$. Now the auctioneer can allocate items 1 and 2 to the bidder for \$9. To allow bidders to express any valuation $v : 2^M \rightarrow \{\Re_+ \cup 0\}$, it was proposed that bidders can submit XOR-constraints between bids (Sandholm, 2002a). If two bids are combined with an XOR-constraint, only one of them can win. It turns out that in the continuous case, there is an inherent tradeoff between the full expressiveness of XOR-constraints and computational complexity (recall that in the binary case, winner determination is $\mathcal{NP}$-complete even without XOR-constraints):

Theorem 3.3. *With* XOR*-constraints between bids, combinatorial auctions, reverse auctions, and exchanges are all $\mathcal{NP}$-complete, both in their binary and continuous versions. These problems remain $\mathcal{NP}$-complete with or without free disposal, and even if prices are integer.*

This is immediate from the fact that XORS is $\mathcal{NP}$-complete even if bids can be submitted on individual items only (Theorem 2.3). (However, while the XORS problem becomes polynomial-time solvable when for every bidder XOR-constraints are placed between all of the bidder's bids, in the case of package bids this remains $\mathcal{NP}$-complete. In fact, the problem remains $\mathcal{NP}$-complete even if every bidder bids on only one package—this follows trivially from the $\mathcal{NP}$-completeness of the combinatorial auction (or combinatorial reverse auction or combinatorial exchange) winner determination problem because those formulations do not even have a notion of a bidder.)

It follows that winner determination under the other fully expressive bidding languages that have been proposed for combinatorial auctions (which are generalizations of the XORS language)—OR-of-XORs (Sandholm, 2002b) and XOR-of-ORs (Nisan, 2000)—is $\mathcal{NP}$-complete even in the continuous case. (The widely advocated idea of expressing mutual exclusion among bids via dummy items that the bids share (Fujishima et al., 1999; Nisan, 2000), does not lead to a fully expressive bidding language in the continuous case because the dummy items may be partially allocated to different bids.)

The heart of the difficulty with the side constraints of this section is that they would require a bid to be "counted" even if it is accepted only partially. As the theorems of this section entail, such a counting device cannot be encoded into a linear program (of polynomial size) unless $\mathcal{P} = \mathcal{NP}$.

3.3. Side constraint under which the continuous and binary case are easy

Some side constraints restrict the space of feasible allocations so dramatically that the winner determination problem becomes easy even for the binary case. Currently we are not aware of any constraint in this class that would be of great practical interest, but the following artificial constraint serves as an existence proof.

Definition 6 *(EXTREME-EQUALITY).* Each bid and ask has to be accepted to the same extent: $\forall j, x_j = x$.

Theorem 3.4. *Both the binary and continuous versions of the multi-unit combinatorial exchange* (*and thus also the combinatorial auctions and reverse auctions*) *are solvable in polynomial time under EXTREME-EQUALITY* (*with and without free disposal*).

Proof. The continuous case is directly solvable by linear programming. In the binary case, simply try accepting all offers ($x = 1$) and rejecting them ($x = 0$).

In either case, if there is any XOR-constraint, or any two bids share items (actual or dummy), then there is no feasible solution.

4. Incorporating non-price attributes

There are at least two reasons for introducing multi-attribute techniques into markets. First, in a basic auction (or reverse auction or exchange), each item has to be completely specified. In many settings, this is overly restrictive. It would be more desirable to leave some of the parameters of the items unspecified, so that each player could propose in her bids the most suitable parameter combinations for her, such as delivery date, quality (Che, 1993), insurance, etc. (each player could also specify different parameter combinations in different bids). This would avoid the problem of having to enumerate alternative parameter combinations as separate items up front. (Different parameter settings would, in general, not be equally desirable to the recipient of the bids.) Second, a bid from one bidder can be more desirable than the same bid from another bidder (e.g., due to historical data on timeliness and quality of different bidders).

Consider a (combinatorial) auction or reverse auction such as the ones discussed in this paper so far. Let $\vec{a}_j$ be a vector of the additional (non-price) attributes. These may be attributes of the items being traded, attributes that characterize the bidders themselves, or both. Some of the attributes can be specific to one bid (say j) while others might not (such as quality of a certain line of products). The vector can include attributes revealed by the bidder as well as attributes whose values the recipient gets from other sources such as historical performance databases. As a preprocessor to winner determination, we re-weight the bid prices based on the additional attributes. The new price of any bid j is $p'_j = f(p_j, \vec{a}_j)$. The re-weighting function f could be specified by any party, but in most markets it would be set by the bid-taker (seller in an auction, buyer in a reverse auction) to characterize his preferences. Usually this function would be set before the bid-taker receives the bids (or in some cases even after, but this would, in general, affect the bidders' incentives). We then run the winner determination in the (combinatorial) market using prices p'. The following result is thus immediate:

Theorem 4.1. *Consider an auction or a reverse auction. Whether or not p' is used* (*for some of the bids*) *in the objective, and whether or not p' is used* (*for some of the bids*) *in the side constraints, all of the easiness* (*polynomial-time*) *results of this paper still hold* (*under the assumption that f can be evaluated in polynomial time*)*, and the $\mathcal{NP}$-hardness results still hold* (*at least if f is not restricted*).

Unlike in auctions and reverse auctions where multiple attributes can be handled as a preprocessor to winner determination like this, in exchanges multiple attributes cannot be handled as a preprocessor! The reason is that in an exchange there are multiple bid takers (each buyer and each seller is a bid taker in this sense), and they may have different preference functions f_i over attribute vectors. The heart of the difficulty then is that which f_i should be used in the preprocessor depends on the outcome of the clearing, which is not known at preprocessing time.

Multiple attributes can still be handled, but their handling has to be incorporated into the optimization problem (winner determination problem) itself. This can be accomplished as follows. We treat items that have different values of the item attributes as different items. We use a separate decision variable not just for each such item, but for each ⟨item, buyer, seller⟩ tuple. This way each buyer (seller) can condition his bid price on the item attributes and on whom he is buying from (selling to). (Conditioning on whom he is buying from (selling to) is pertinent when bidder attributes have to be taken into account.)

5. Related research on bidding languages

There has been considerable recent work on bidding languages for auctions. Many of them take package bids as the atomic constructs which are then combined with logical operators such as OR and XOR (Fujishima et al., 1999; Nisan, 2000; Sandholm, 2002a, 2002b). Some of the newer languages also allow recursive logical formulae (Hoos and Boutilier, 2001).

Side constraints, such as those in this paper, can be viewed as a *compact combinatorial bidding language* (although they can also be used by the bid taker(s) for other purposes such as what-if analysis).

Recently, certain side constraints in markets have been independently suggested (Davenport and Kalagnanam, 2001; Bichler et al., 2003). They present integer programming formulations of certain winner determination problems with side constraints. Their winner determination problems are $\mathcal{NP}$-complete even without the side constraints, so in their settings, side constraints do not fundamentally change the complexity. This is unlike many of our settings where the side constraints cause the winner determination problem to move to the other side of the $\mathcal{P}$ versus $\mathcal{NP}$-complete cutoff.

6. Conclusions and future research

In most real-world (electronic) marketplaces, there are other considerations besides maximizing immediate economic value. We presented a sound way of taking such considerations into account via side constraints and non-price attributes, and showed that side constraints have a significant impact on the complexity of clearing the market. Future research includes analyzing the complexity entailed by other side constraints. We are also developing search algorithms that perform well on average on $\mathcal{NP}$-complete clearing problems that include side constraints.

References

Bichler, M., Kalagnanam, J., Lee, H.S., 2003. RECO: Representation and evaluation of configurable offers. In: 8th INFORMS Computing Society Conference: Optimization and Computing in the Network Era, Phoenix, AZ.

Che, Y.-K., 1993. Design competition through multi-dimensional auctions. RAND J. Econ. 24, 668–680.

Davenport, A.J., Kalagnanam, J., 2001. Price negotiations for procurement of direct inputs. Technical Report RC 22078. IBM, May.

Fujishima, Y., Leyton-Brown, K., Shoham, Y., 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI), Stockholm, Sweden, August, pp. 548–553.

Garey, M., Johnson, D., 1979. Computers and Intractability. Freeman.

Hoos, H., Boutilier, C., 2001. Bidding languages for combinatorial auctions. In: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI), Seattle, WA, pp. 1211–1217.

Kuhn, H.W., 1955. The Hungarian method for the assignment problem. Naval Res. Logist. Quart. 2, 83–97.

Lehmann, B., Lehmann, D., Nisan, N., Combinatorial auctions with decreasing marginal utilities. Games Econ. Behav. In press. Early version in ACMEC-01.

Lehmann, D., O'Callaghan, L.I., Shoham, Y., 2002. Truth revelation in rapid, approximately efficient combinatorial auctions. J. ACM 49 (5), 577–602. Early version appeared in ACMEC-99.

Nisan, N., 2000. Bidding and allocation in combinatorial auctions. In: Proceedings of the ACM Conference on Electronic Commerce (ACM-EC), Minneapolis, MN, pp. 1–12.

Penn, M., Tennenholtz, M., 2000. Constrained multi-object auctions and *b*-matching. In: Inform. Process. Lett., North-Holland, Amsterdam, July, pp. 29–34.

Rothkopf, M., 1977. Bidding in simultaneous auctions with a constraint on exposure. Operations Res. 25, 620–629.

Rothkopf, M., Pekeč, A., Harstad, R., 1998. Computationally manageable combinatorial auctions. Manage. Sci. 44 (8), 1131–1147.

Sandholm, T., 2002a. Algorithm for optimal winner determination in combinatorial auctions. In: Artif. Intell., Springer, Berlin, pp. 1–54. January. First appeared as an invited talk at the First International Conference on Information and Computation Economies, Charleston, SC, October 25–28, 1998. Extended version appeared as Washington Univ., Dept. of Computer Science, technical report WUCS-99-01, January 28th, 1999. Conference version appeared at the International Joint Conference on Artificial Intelligence (IJCAI), Stockholm, Sweden, 1999, pp. 542–547.

Sandholm, T., 2002b. eMediator: A next generation electronic commerce server. Comput. Intelligence 18 (4), 656–676. Special issue on Agent Technology for Electronic Commerce. Early versions appeared in the Conference on Autonomous Agents (AGENTS-00), 2000, pp. 73–96; AAAI-99 Workshop on AI in Electronic Commerce, Orlando, FL, July 1999, pp. 46–55; and as a Washington University, St. Louis, Dept. of Computer Science technical report WU-CS-99-02, January 1999.

Sandholm, T., Suri, S., 2001. Side constraints and non-price attributes in markets. In IJCAI-2001 Workshop on Distributed Constraint Reasoning, Seattle, WA, pp. 55–61.

Sandholm, T., Suri, S., 2003. BOB: Improved winner determination in combinatorial auctions and generalizations. In: Artif. Intell., Springer, Berlin, pp. 33–58. Early version: Improved Algorithms for Optimal Winner Determination in Combinatorial Auctions and Generalizations. National Conference on Artificial Intelligence (AAAI-00), Austin, TX, July 31–August 2, pp. 90–97.

Sandholm, T., Suri, S., Gilpin, A., Levine, D., 2002. Winner determination in combinatorial auction generalizations. In: International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Bologna, Italy, July, pp. 69–76. Early version appeared at the AGENTS-01 Workshop on Agent-Based Approaches to B2B, Montreal, Canada, May 2001, pp. 35–41.

Tennenholtz, M., 2000. Some tractable combinatorial auctions. In: Proceedings of the National Conference on Artificial Intelligence (AAAI). Austin, TX, August.