

2011

Theoretical and Experimental Insights into Decentralized Combinatorial Auctions

Georg Ziegler

Technische Universität München, zieglerg@in.tum.de

Tobias Scheffel

Technische Universität München, scheffel@in.tum.de

Follow this and additional works at: <http://aisel.aisnet.org/wi2011>

Recommended Citation

Ziegler, Georg and Scheffel, Tobias, "Theoretical and Experimental Insights into Decentralized Combinatorial Auctions" (2011).
Wirtschaftsinformatik Proceedings 2011. 52.
<http://aisel.aisnet.org/wi2011/52>

This material is brought to you by the Wirtschaftsinformatik at AIS Electronic Library (AISeL). It has been accepted for inclusion in Wirtschaftsinformatik Proceedings 2011 by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact elibrary@aisnet.org.

Theoretical and Experimental Insights into Decentralized Combinatorial Auctions

Georg Ziegler
Technische Universität München
Boltzmannstr. 3
85748 Garching, Germany
zieglerg@in.tum.de

Tobias Scheffel
Technische Universität München
Boltzmannstr. 3
85748 Garching, Germany
scheffel@in.tum.de

ABSTRACT

Combinatorial Auctions (CAs) are promising to increase social welfare by enabling bidders to express their valuation on any combination of items. A major issue of many CAs is the requirement to optimally solve the \mathcal{NP} -hard Combinatorial Allocation Problem. To release a centralized auctioneer from that computational burden he can shift it to the bidders. One of the few discussed decentralized auctions is PAUSE, in which bidders suggest new allocations to the auctioneer. In our theoretical analysis we examine the bidders' bid complexity and determine a worst case bound concerning efficiency, if bidders follow a profit maximizing strategy. Based on these results we conduct computational experiments with different bidding and computation strategies, and analyze their impact on efficiency, auctioneer's revenue and auction runtime. Surprisingly, even if agents deviate from the optimal bid price calculation, PAUSE still achieves high levels of efficiency and auctioneer's revenue compared to the Combinatorial Clock auction.

Keywords

combinatorial auctions, bidding agents, computational experiments

1. INTRODUCTION

The Internet allows for the exchange of complex preference profiles and laid the foundation for the design of new market mechanisms. The promise of these mechanisms is to increase economic welfare by allowing market participants to reveal more comprehensive information about cost structures or utility functions. In the last decade, a growing literature in management science and information systems is devoted to the design of such smart markets [?, ?], with combinatorial auctions (CAs) emerging as a pivotal example [?]. Allocation of spectrum licenses for wireless communication services [?], transportation [?] and industrial procurement [?, ?] are not nearly all domains in which there is an increasing requirement and usage for multi-item auction mechanisms al-

lowing complex bids. However, the auctioneer of a CA faces the \mathcal{NP} -hard Combinatorial Allocation Problem (CAP) [?], for which an optimal solution is untraceable in larger instances, although the availability of computing power grows.

Decentralizing the CAP is the approach of the *Progressive Adaptive User Selection Environment* (PAUSE) auction proposed by [?]. In PAUSE bidders submit not only their own bids, the desired packages of items and the price, but have to propose a new allocation including their new bids and existing bids, being better than the current provisional allocation. Checking bid validity and publishing accepted bids remains the auctioneer's only tasks. Another simplification for the auctioneer is that there is no need for a price calculation mechanism in the iterative process like in most other iterative CAs.

There is only little work in the literature on decentralized auctions, therefore, we study PAUSE theoretically and experimentally. Our theoretical analysis shows the growing complexity for the bidders in PAUSE and gives a worst case bound concerning efficiency, if bidders follow a certain strategy. The determination of a lower bound in CAs has to our knowledge not been done and published yet, but it reveals important insights what can go wrong concerning bidder behavior, value models and auction rules. In this context we analyze PAUSE with computational experiments.

[?, ?] developed some sophisticated bidding strategies for distributed auctions, however, in our experiments we focus on more simple strategies in which bidders reveal as little as possible about their valuations. Further, we use another value model with more items, in which the advantage of not having to calculate the CAP optimally is more decisive. To compare and benchmark we run computational experiments with the Combinatorial Clock auction (CC), which is a centralized CA and known for its sparse need of solving the CAP [?].

2. THE PAUSE AUCTION

PAUSE especially concentrates on achieving the following properties: It should permit bidders to submit any combinatorial bid they choose (*fully combinatorial*) and allow losing bidders to clearly see why they lost (transparent). Furthermore it should allow the auctioneer to determine the winner easily for auctions of any size and achieve high auctioneer payoffs. The basic idea of PAUSE is to place the computational burden of evaluating synergies on the bidder claiming

those synergies. This leaves the auctioneer to simply check that a bid is valid. He no longer faces the CAP. PAUSE is a multi-round, multi-stage CA decentralizing the CAP [?]. A PAUSE auction with m items has m stages.

Stage 1 consists of a *Simultaneous Ascending Auction* (SAA) [?] on all items. During this stage bidders can only place individual bids on items - no package bidding is allowed. The stage ends when bidding ends and the auctioneer determines the provisional allocation by simply choosing the best bid on every item.

In each round of a successive stage $h = 2, 3, \dots, m$ a bidder is required to submit a composite bid (denoted by X^{CB}), which covers all items and includes only disjoint package bids each of maximum cardinality of h . Bidders are allowed to use bids that other agents have placed in previous rounds. The bid price $p(X^{CB})$ of a composite bid is the sum of its package bid prices. For each new package bid in a composite bid, the bidder has to outbid the currently winning composite bid by the minimum increment ϵ . After each round the auctioneer declares the highest composite bid as the provisional allocation and registers the highest submitted package bids in the database. A stage ends when bidding finishes. At the end of each stage h , all agents know the best bid for every subset of size h or less so far.

3. THEORETICAL RESULTS

For our theoretical analysis we assume bidders follow a *straightforward* strategy, by bidding on the package which yields the highest possible payoff at current prices. Since there is no known equilibrium bidding strategy in PAUSE this assumption is justified by the typical use in game theoretical analysis and as it seems natural since bidders reveal as little information as possible keeping the chance for high profits. We assume further that the straightforward bidders do not consider a combination of their package bids, since they are able to bid on those combinations in a single package bid in later stages, thus avoiding a possible exposure problem, which would leave a bidder winning a package of items at prices he is not willing pay.

Let $\mathcal{K} = \{1, \dots, m\}$ denote the set of items and $\mathcal{I} = \{1, \dots, n\}$ the set of bidders indexed by i . In general, bidders have different valuations for packages $S \subseteq \mathcal{K}$. Let $v_i(S) > 0$ indicate the valuation of bidder i for package $S \subseteq \mathcal{K}$. Each bidder i has a demand set $\mathcal{D}_{i,h} := \{S : v_i(S) \geq p_j(S), i \neq j \wedge |S| \leq h\}$, i.e. it contains all packages S for which bidder i has a higher valuation than the price of the current highest bid from another bidder j ($p_j(S)$) and the cardinality of S must not be greater than h . If bidders want to determine the ask-price for a package S , they have to calculate the price ($p(X^{CS}(S))$) of a set of complement disjoint bids, not overlapping with S and covering all items in $\mathcal{K} \setminus S$.

[?] designed PAUSE under the premises of an OR-bidding language, meaning a bidder can win more than just one of his bids, and super-additive valuation functions. We adopted these assumptions in our analysis of the *Bid Determination Problem* (BDP) and the worst case efficiency bound.

Definition 1. The Bid Determination Problem: To maxi-

mize bidder i 's current payoff $\pi_i \in \mathbb{R}_0^+$, he has to bid on the package(s) S determined by:

$$\max_{S \in \mathcal{D}_{i,h}} (v_i(S) - p(X^{CB}) + p(X^{CS}(S)) - \epsilon) \geq \pi_i$$

The inequation ensures that bidder i bids on package(s) S only, if the prospective payoff will not be less than his current payoff. The optimal determination of $p(X^{CS}(S))$ is \mathcal{NP} -hard, as it is a CAP on the complementary set, which has to be calculated for every package $S \in \mathcal{D}_{i,h}$ to determine the straightforward bid.

The following example in Table 1 shows valuations of bidders in $\mathcal{I} = \{1, 2\}$ for the items in $\mathcal{K} = \{1, 2\}$ and sketches the PAUSE auction process with straightforward bidders. PAUSE does not achieve the efficient allocation indicated by the asterisks, but terminates with 51.5% efficiency.

	1	2	1, 2	$p(X^{CB})$	π_1	π_2
v_1	100*	0	103			
v_2	0	100*	103			
Stage1	1 ₁	1 ₂	0	2	99	99
Stage2	0	0	3 ₁	3	100	0
	0	0	4 ₂	4	0	99
	...					
	Termination		103 ₁	103	0	0

Table 1: Bidders' valuations and auction process - an example of low efficiency in PAUSE

THEOREM 1. *PAUSE terminates with an allocation that is at least $1/m$ efficient, if all bidders follow the straightforward strategy and have superadditive valuations.*

Proof: The proof leans towards the example in Table 1. Given the premises stated in the theorem, inefficiencies can only occur in PAUSE, if the auction terminates allocating big packages, although disjoint subsets of them would support the efficient allocation.

Lets assume stage 1 terminates with bids

$$p_i(i) = \max_{i \neq j} v_j(i) + \epsilon \quad \forall i \in \mathcal{I} \quad (1)$$

W.l.o.g. these bids can be considered to support the efficient allocation. The current auctioneers revenue $\Pi_{h=1}$ would be $\sum_i p_i(i)$.

In order to terminate with another allocation we demand no improvement on any of these individual bids. That means once any of these bids $p_i(i) \notin X^{CB} \Rightarrow \exists S \in \mathcal{K}$ which applies to

$$\begin{aligned} v_i(S) - (p(X^{CB})) &> v_i(i) - (p(X^{CB}) - p(X^{CS}(i))) \\ \wedge |S| &\leq h \end{aligned} \quad (2)$$

i.e. bidder i has a better alternative than bidding on the individual item i once his provisional payoff drops to zero.

If $v_i(S)$ is part of the final allocation, we want $p(X^{CS}(S))$ to be as small as possible considering the worst case. Thus we determine $S = \mathcal{K}$. That means as long as $h < m$ every

bid $p_i(i)$ for all i is part of the composite bid, which further means that no new bids are submitted before stage m . In stage m the following must apply:

$$\exists i \in \mathcal{I} \text{ with } v_i(S) - (\Pi_m) > \pi_i \quad (3)$$

Since in this case bidder i bids on the package S , all other bidders $j \in \mathcal{I} \setminus \{i\}$ have a current payoff $\pi_j = 0$ and thus also the following inequation must hold:

$$v_j(S) - (\Pi_m + \epsilon) > v_j(j) - (p_j(j) + p(X^{CS}(j))) \forall j \neq i \quad (4)$$

Efficiency is then calculated by

$$E(X^{CB}) = \frac{\max_i v_i(S)}{\sum_i v_i(i)} \quad (5)$$

To determine the worst case efficiency we need to minimize the numerator and maximize the denominator. Thus we can determine w.l.o.g. $v(S) = v_i(S)$ and $v(i) = v_i(j) \forall i$.

Since the most strict condition on $v_i(S)$ is

$$v_i(S) > \sum_j p_j(j) + v_i(i) - b_i(i) + \epsilon \quad (6)$$

the worst case efficiency results in:

$$\begin{aligned} \min_v E(X^{CB}) &= \min_v \frac{v_i(S)}{\sum_i v_i(i)} \\ &\stackrel{\epsilon=1}{=} \frac{m+v(i)+1}{m \cdot v(i)} \\ &\stackrel{v(i) \rightarrow \infty}{=} \frac{1}{m} \end{aligned} \quad (7)$$

Note assuming a bid increment $\epsilon = 1$ the equations 7 only apply if the valuation $v(i)$ is sufficiently large, i.e. depending on m this valuation must be greater than 2 or 1 respectively. \square

While such situations which lead to $1/m$ efficiency can be considered degenerated cases that will not happen too often in practice, it is very likely to achieve high efficiency on average with more realistic value models.

4. EXPERIMENTAL DESIGN

To analyze the impact of our theoretical results on the outcome of the PAUSE auction in realistic settings, we conduct computational simulations, which consists of three main components. A value model, which defines valuations of all packages for each bidder, auction formats, which define the rules, and bidding agents, who follow certain strategies.

4.1 Value Model

We use a 3×6 *Real Estate* value model that is based on the *Proximity in Space* model from the Combinatorial

Auction Test Suite (CATS) in [?]. Our model contains two different bidder types, one big bidder, interested in all items, and five smaller bidders. Each small bidder is interested in a randomly determined preferred item, all horizontal and vertical neighbors and their respective neighbors. This means small bidders are interested in 6 to 11 items with local proximity to their preferred item. An example is shown in Table 2, in which the preferred item of a small bidder is Q, and all gray shaded items in the proximity of the preferred item have a positive valuation. For

A	B	C	D	E	F
G	H	I	J	K	L
M	N	O	P	Q*	R

Table 2: The value model with the preferred item Q of a small bidder. All his positive valued items are shaded.

each bidder we draw the baseline item valuation $v_i(k)$ from a uniform distribution separately. Complementarities occur upon vertical and horizontal adjacent items based on a logistic function to determine package valuations: $v_i(S) = \sum_{C \in P} \left(\left(1 + \frac{a}{100(1+e^{b-|C|})} \right) * \sum_{k \in C} v_i(k) \right)$, with P being the partition of S containing maximal connected packages C . This complementarity structure takes the lack of economies of scale with small packages and a saturation effect with larger packages into account. For our experiments we choose $a = 320$ and $b = 10$ for the big bidder and $a = 160$ and $b = 4$ for all small bidders, and draw the baseline valuations for the big bidder on the range $[3, 9]$ and for the small bidders on the range $[3, 20]$.

4.2 Auction Formats

We analyze two different auction formats in our economic environment. The PAUSE auction, as described in Section 2 with a minimum increment of 3 and the CC auction [?].

The CC auction is also a multi-round auction, in which bidders are able to place new bids in every round according to new calculated linear ask prices. The price for a package is simply the sum of item-prices. Bidders use an OR-bidding language. Prices for all items are initially zero. In every round bidders identify a package of items, or several packages, which they offer to buy at current prices. If two or more bidders demand an item then its price is increased by the minimum bid increment of 1 in the next round. This process iterates. In a simple scenario in which supply equals demand, the auction terminates and the items are allocated according to the current round bids. If at some point there is excess supply for at least one item and no item is over-demanded, the auctioneer determines the winners to find an allocation of items that maximizes his revenue by considering all submitted bids. If the solution displaces a bidder, who was active in the last round, the prices of items in the corresponding bids rise by the bid increment and the auction continues. The auction ends when no prices are increased and bidders finally pay their bid prices for winning packages.

4.3 Bidding Agents

In PAUSE we use two different bidding strategies and two different approaches to determine the bid price. As introduced in Section 3 we implement the straightforward (*BR*) bidding strategy, and a *Greedy* bidding strategy that allows the agents to reduce their demand set to one package calculated by $\max(v_i(S)/|S|, \forall S \in \mathcal{D}_{h,i})$ in every stage. As shown by our theoretical analysis the optimal calculation of the corresponding complement set $X^{CS}(S)$ is \mathcal{NP} -hard, therefore we explore two different types of calculating it, an optimal (*oCS*) and a heuristic (*hCS*) approach. We propose the following heuristic, with $k(X^{CS})$ denoting the set of items covered by the bids in X^{CS} :

- 1) $X^{CS} := \emptyset$
- 2) while $k(X^{CS}) \neq \mathcal{K} \setminus S$

$$X^{CS} = X^{CS} \cup \arg \max_{T \subseteq \mathcal{K} \setminus (S \cup k(X^{CS}))} p_i(T)$$

We start with an empty complement set X^{CS} , determine all active bids not overlapping the current considered package S , choose the bid with the highest price and add it to our complement set X^{CS} . Then we determine the next bid, not overlapping S and $k(X^{CS})$ with the highest bid price. We repeat until our complement set covers all items of $\mathcal{K} \setminus S$.

For our experiments with the CC auction we use the straightforward bidder and a *heuristic* bidder (*5of20*) bidding on 5 of his 20 best packages in every round, more details to this in [?]. Additionally we implemented a *preselect* bidder (*pres10*) who determines his 10 most valuable packages before the auction starts, and bids in each round on all of them applying to $v_i(S) \geq p(S)$.

5. RESULTS

We run 50 simulations for every of the 4 bidding agents in PAUSE and for the 3 different bidding strategies in CC. All experiments run on an Intel Core2Duo processor with 2.67 GHz, 4 GB of RAM, Windows Vista and the open source IP solver "lp_solve".

A primary measure for the quality of an auction mechanism is the allocative efficiency (X^* denotes the best allocation and $v_i(X)$ bidder i 's valuation for the allocation X):

$$E(X) = \frac{\sum_{i \in \mathcal{I}} v_i(X)}{\sum_{i \in \mathcal{I}} v_i(X^*)}$$

A further measure is the auctioneer's revenue share:

$$R(X) = \frac{\sum_{i \in \mathcal{I}} p_i(X)}{\sum_{i \in \mathcal{I}} v_i(X^*)}$$

Bidders revenue share is:

$$B(X) = E(X) - R(X)$$

As expected by our theoretical analysis, straightforward bidding in PAUSE with more items and higher competition leads to a better efficiency than the lower bound. We find that *BR_{oCS}* agents achieve in many cases a solution near the efficient one (Figure 1) and a high auctioneer's revenue (Table 3). In PAUSE all considered agents are able to find

a highly efficient solution, even *Greedy* agents, who generate only ~60% final bids compared to *BR* agents. Surprisingly, calculating the complement set X^{CS} with our heuristic (*hCS*) leads only to a small deviation in all measures (except the runtime) from the results with agents calculating X^{CS} optimally.

RESULT 1. *Determining the complement set X^{CS} sub-optimally has only a small impact on the auction outcome.*




Figure 1: Auction efficiency with different bidding strategies and auction formats

In contrast to PAUSE, the CC auction mostly ends in allocations with lower efficiency and auctioneers' revenue. We suspect mainly the high number of unsold items (Table 3) to lead to such inefficiencies, together with the bigger size of winning packages ($\varnothing 6.5$ with *BR* agents vs. $\varnothing 5.03$ with *BR_{oCS}* agents vs. $\varnothing 5.45$ in efficient solutions) and the lower number of final bids. To analyze the pure impact of unsold items we ran additional simulations with CC auctions, in which we enforce the agents to bid in the first round on all items they are interested in and found, that the efficiency increases to 89.93% on average with *BR* agents.

RESULT 2. *An auction mechanism forcing agents to bid also on smaller packages, guides them in solving their coordination problem.*

RESULT 3. *CC needs fewer rounds to clear than PAUSE auctions.*

This results from the only moderate increasing of the allowed package size and from the package increment vs. the linear item increment in CC auctions.

Concerning bidders' calculation complexity shows Figure 2 that with an increasing number of items a small *BR_{oCS}* agent is interested in, the required calculation time in the auction increases exponentially. We omit the result of the big bidder, who needs around six hours (particularly $\varnothing 21.451$

	PAUSE				CC		
	BR_{oCS}	BR_{hCS}	$Greedy_{oCS}$	$Greedy_{hCS}$	BR	$5of20$	$pres10$
∅ Efficiency in %	97.71	97.52	90.54	91.01	81.81	91.70	90.95
∅ Auctioneers' revenue in %	88.02	88.44	73.62	73.54	76.22	87.96	88.68
∅ Bidders' revenue in %	9.69	9.08	16.92	17.48	5.59	3.74	2.27
∅ Rounds	126.98	127.74	101.48	101.32	43.14	47.02	44.88
∅ Unsold items	0.00	0.00	0.00	0.00	3.96	1.66	1.50
∅ Auction runtime in sec.	22714.29	2166.12	26.85	25.37	44.81	45.34	11.10
∅ Number of final bids	54.95	55.07	33.33	33.14	35.49	145.35	52.63
∅ Size of winning packages	5.03	5.60	2.87	2.75	6.50	5.34	10.58

Table 3: Summary of simulation results

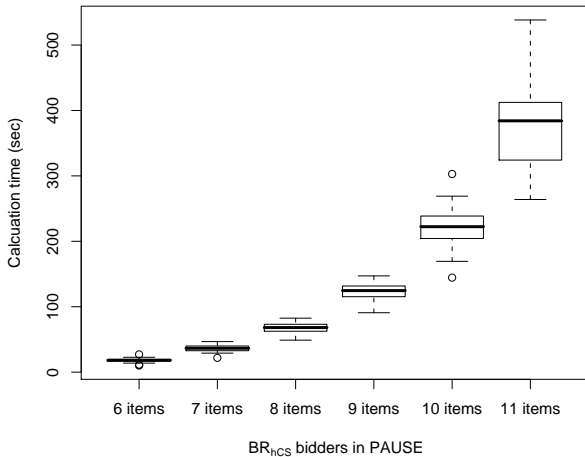


Figure 2: Bidders' required time over the auction dependent on the number of items they are interested in

sec.) per auction. The boxplot further exhibits, with 7 items or more of interest, a single BR_{oCS} agent in PAUSE requires more calculation time than the complete CC auction process.

RESULT 4. *Determining the straightforward bid in PAUSE drastically increase the bidders' complexity.*

Comparing the BR bidders in PAUSE with the $pres10$ bidders in CC or the $Greedy$ bidders in PAUSE with the BR bidders in CC we find the following result.

RESULT 5. *With a similar number of active bids, PAUSE leads to higher efficiency.*

PAUSE collects package bids of every size due to the restrictions of the package size in every stage. This helps to find allocations with high revenue, while in CC more bigger sized package bids are collected which often overlap with each other and so prohibit a "good" allocation.

6. CONCLUSION

We provide a deeper theoretical insight in the decentralized PAUSE auction and present experimental results of two different auction mechanisms. We analyzed effects of the straightforward bidding strategy in PAUSE. First we discover following this strategy leads to a growing bid determination complexity, as bidders are not allowed to submit new package bids without embedding them in a new allocation. Secondly if all bidders follow the straightforward strategy, we determine a worst case bound of $1/m$ efficiency.

Since our theoretical analysis promises better efficiency and auctioneer's revenue by the use of more realistic value models, we conducted computational experiments to verify this prediction. We used an agent-based system to compare different bidding strategies and auction mechanisms and find straightforward bidding with optimal bid price determination in PAUSE leads to very high efficiency and auctioneer revenue. Surprisingly, deviating from the optimal bid price determination does not have a significant impact on the auction outcomes, while the auction runtime is reduced drastically. The comparison to the CC auction exhibits that PAUSE is a better guide solving the bidders' coordination problem since it collects different sizes of package bids.

PAUSE shows some desirable properties, however, before taking it to the field it needs further research concerning bidder behavior and auction rules.

7. ACKNOWLEDGMENTS

Financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/3-1) is gratefully acknowledged.

8. REFERENCES

- [1] M. Bichler, A. Davenport, G. Hohner, and J. Kalagnanam. Industrial procurement auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, 2006.
- [2] M. Bichler, P. Shabalín, and A. Pikovsky. A computational analysis of linear-price iterative combinatorial auctions. *Information Systems Research*, 20(1):33–59, 2009.
- [3] Y. Caplice, C. and Sheffi. Combinatorial auctions for truckload transportation. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.
- [4] P. Cramton. *Auctioning the Digital Dividend*. Karlsruhe Institute of Technology, 2009.
- [5] P. Cramton, Y. Shoham, and R. Steinberg, editors.

Combinatorial Auctions. MIT Press, Cambridge, MA, 2006.

- [6] P. Cramton, Y. Shoham, and R. Steinberg. Introduction to combinatorial auctions. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press, Cambridge, MA, 2006.
- [7] J. Gallien and L. Wein. A smart market for industrial procurement with capacity constraints. *Management Science*, 51:76–91, 2005.
- [8] F. Kelly and R. Steinberg. A combinatorial auction with multiple winners for universal service. *Management Science*, 46(4):586–596, 2000.
- [9] K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auction algorithms. In *ACM Conference on Electronic Commerce*, pages 66–76, 2000.
- [10] K. McCabe, S. Rassenti, and V. Smith. Smart computer-assisted markets. *Science*, 254:534–538, 1991.
- [11] B. Mendoza and J. M. Vidal. Bidding algorithms for a distributed combinatorial auction. In *Proceedings of the on Autonomous Agents and Multiagent Systems*, 2007.
- [12] B. Mendoza and J. M. Vidal. Approximate bidding algorithms for a distributed combinatorial auction (short paper). In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 2008.
- [13] D. Porter, S. Rassenti, A. Roopnarine, and V. Smith. Combinatorial auction design. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 100:11153–11157, 2003.
- [14] M. H. Rothkopf, A. Pekec, and R. M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44:1131–1147, 1998.
- [15] L. D. C. M. M. P. H. R. J. J. Sandholm, T. and D. Begg. Changing the game in strategic sourcing at procter & gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.