Combinatorial Auctions

Olli Ahonen

1 Introduction

Auctions are a central part of trade in many societies. In an auction, the seller, or the auctioneer, gives a number of people or organizations the option to buy an item or items currently in the auctioneer's possession. In a traditional auction, the auctioneer allows the bidders to express how much they are willing to pay for each item. The bidder who is willing to pay the highest price for the item wins the item. This is essentially how almost all auctions today work. In the future, traditional auctions may be, in part, replaced by combinatorial auctions (CAs).

Combinatorial auctions only make sense if the auctioneer has many items for sale. The essential difference to the traditional auction model is that the bidders are allowed to bid on packages, or combinations, of items. Each bidder can express his willingness to buy different packages of items and the acceptable price for each package. In the general case, the packages of items can be complex combinations of items and mutual dependencies expressed by logical connectives such as AND, OR, and XOR. In real-world auctions, the bundling of items has been restricted to keep the auction feasible for the auctioneer as well as the bidders. Based on the submitted bids, the auctioneer decides which bidder gets which package or packages they bidded for. The prices are fixed by the bids.

Single-item single-bid auctions do not allow bidders to express their valuation of different combinations of the auctioned items. The bidders preferences regarding packages of items can be complex. CAs are particularly suited for auctioning a set of items that includes complements. Items are complements if a set of items has higher value than the sum of values of the individual items. For instance, a pair of shoes has higher value, or utility, than two times the value of a single shoe. Allowing the bidders to fully express their preferences often leads to higher overall economic efficiency, i.e., the items are allocated to those who value them most. Also, auctioneer revenues are larger. [1]

2 Examples of Use

2.1 Airspace Resources

The country-wide air transport system is a massively complex system in large countries, such as the USA. A suitable auction is a good candidate system for the efficient and economical allocation of the scarce airport resources at given times. The use of a CA for the allocation has been proposed and studied, but not yet implemented. CAs might work well here since competing airlines have different valuations for different packages of landing and takeoff slots, for example. In the US, the regulator Federal Aviation Administration is considering and evaluating CAs for allocating resources for LaGuardia airport in New York.

2.2 Truckload Transportation

Combinatorial auctions are being used to buy and sell freight transportation services. In this case, the auctioneer is the shipper, the one who wants goods transported, and the bidders are the carriers of the goods. The auction is a reverse combinatorial auction: The bidders report the minimum prices for which they will take on a specific package transportation contract, and the auctioneer tries to minimize the total cost of all contracts. On average, the following figures are observed in the auctions. There are 120 bidders of which 64 are assigned business in the auction. The total reduction in transportation costs is 6%, and the duration of the entire procurement process 3 months. The truckload transportation industry has benefitted from using CAs as more accurate and comprehensive interaction and more collaboration between shippers (auctioneers) and carriers (bidders). [1]

2.3 Bus Routes

Another reverse auction is the market of London bus routes. The market includes about 800 routes and 3.5 million passengers per day. The London Regional Transport (LRT) is the authority auctioning the routes to individual bus companies. LRT has used some form of CA for the bus route allocation since the mid-1980s. The service contracts are usually for five years. They are so called 'gross cost' contracts, where the winning bidder's compensation is the amount it bid, and all passenger-generated revenues go directly to LRT. The use of CAs for this purpose has been a success. There may still be room for improvement, since the parameter space of the design of the auction is huge and the full evaluation of each combination of options is very laborous. [1]

2.4 Industrial Procurement

Industrial procurement is potentially one of the largest application domains for CAs. Here, the auction is again a reverse auction. In the business-to-business domain, software vendors and procurement managers already show interest on CAs. A number of applications have been reported. Unfortunately, public documentation and analysis of this application is rare—possibly due to efforts to protect trade secrets. At Mars Inc., for instance, the use of CAs in procurement has lead to increased suppliers' margins and cost savings to Mars. The payback time of Mars' investment into the CA system was less than a year. The auctions carried out have never taken more time than the corresponding bilateral traditional negatiations. [1]

3 Winner Determination Problem

In combinatorial auctions, the problem that is of most interest to computer science is the winner determination problem (WDP). This is simply the task that the auctioneer performs after receiving all bids: The auctioneer has to decide which bids win and which lose.

If the auctioneer is a private company, maximal revenue (or in the case of a reverse auction minimal total cost) may be the goal. A public organization may wish to maximize the number of satisfied bidders, i.e., allocate the goods in such a way that a maximal number of bidders have at least one winning or realized bid. Revenue maximization may be a secondary goal. The WDP is an NP-complete optimization problem.

4 Statistical Mechanics Model

T. Galla *et al.* [2] have employed a statistical physics approach to the WDP. In their model, there are N bidders and M items to be auctioned. Not all items need to be sold. Each bidder i submits a single bid. A bid is defined to be a set of items A_i for which the bidder is willing to pay a price of ν_i. A more complex CA where bidders submit lists of bids nested by logical ORs or XORs can be reduced to the one-bidder-one-bid model. This may require enlargening the item and bidder sets. The following assumes that there is exactly one bid per bidder. Variables x_i are defined to describe the allocation of items. Variable x_i is 1 if the bid of bidder i wins, and 0 if it loses.

The winner determination problem is now formulated as follows: Find a configuration $\mathbf{x} = (x_1, ..., x_N) \in \{0, 1\}^N$ which maximizes the auctioneer's revenue R, or alternatively the number of satisfied bidders N_s, or maximize both, in either of the two possible orders. Every bid cannot always be satisfied, because each item can only be sold once, and the item sets in the bids can overlap. That is, since any pair of bids can be wholly or partially on the same items, maximizing R or N_s is computationally hard.

5 Greedy Winner Determination

Before further discussing the statistical mechanics approach, let us mention two different methods for solving the WDP in this and the next section. Let there be a given combinatorial auction where bidders can submit non-overlapping bids connected by logical ORs. In this case, the following greedy algorithm approximately solves the WDP in a polynomial number of steps in M and the number of bids, l, for any given auctioneer's parameter value c [1].

The algorithm: Given an integer c and a set of bids $\{(A_j, \nu_j)\}_{j=1}^{l}$:

- Let $\mathcal{P}_c$ be the set of possible allocations of items to the bidders where at most c bids win.
- Let $\mathbf{x}_{\mathcal{P}}$ be the optimal allocation within $\mathcal{P}_c$.
- Let $\mathcal{B}_c$ be the subset of bids (A_j, ν_j) such that $|A_j| \leq \sqrt{M/c}$.
- Compute a greedy allocation $\mathbf{x}_{\mathcal{B}}$ with respect to bids in $\mathcal{B}$ by sequentially selecting highest price bids that do not overlap with already selected bids.
- Choose the better of the allocations $\mathbf{x}_{\mathcal{P}}$ and $\mathbf{x}_{\mathcal{B}}$.

6 Simulated Annealing

Simulated annealing is a general-purpose algorithm for approximate solving of optimization problems. Applied to the combinatorial auction WDP, simulated annealing starts with a random configuration $\mathbf{x}$ indicating the winning bids. To increase R or N_s, the algorithm randomly switches the state of the bids. Changes in $\mathbf{x}$ increasing R or N_s are always allowed, and those decreasing them are allowed with a probability controlled by a global *temperature*. As the temperature is lowered, the system settles down in a state of local optimum which, in general, often yields results at least nearly globally optimal. [3]

7 Factor Graph and Conflict Graph

A combinatorial auction can be codified into a factor graph, an example of which is shown in Fig. 1. The factor graph shows the item sets and bidders of all bids. The bidders are represented as circles and the items as squares.

There is an edge between a bidder and an item, if the bidder's bid includes the item.

The conflict graph (CG) describes which bids contain same items. An example of a CG is shown on the right. A CG contains only bids (or in this case bidders, since each bidder submits only one bid). This CG is for the same auction as the FG on the left. We see that there is an edge connecting two bids, if they bid for at least one common item.

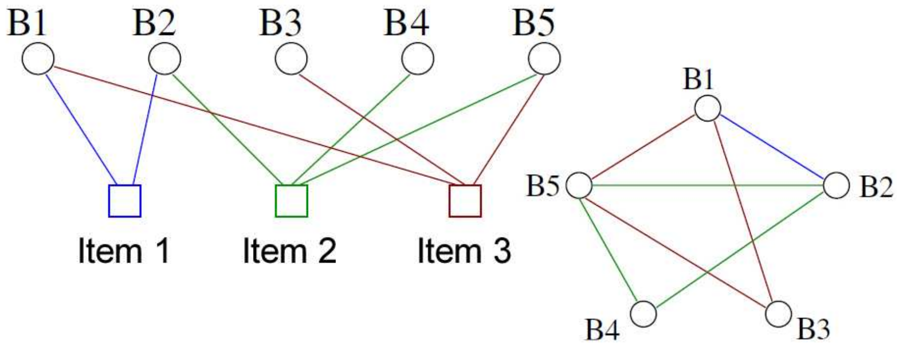

Figure 1: Example of a factor graph (left) and the corresponding conflict graph (right) of a combinatorial auction. (From Ref. [2])

8 Similarity to the Vertex Cover Problem

The problem of allocating the items according to the competing bids is closely related to the vertex cover problem. Finding an allocation that respects the constraint that no item be sold more than once is equivalent to finding a vertex cover for the conflict graph. A covered bid loses and a non-covered bid wins. When a connected vertex is covered in the conflict graph, at least one conflict is removed.

The winner determination problem then becomes a *weighted* vertex cover problem: Find such a vertex cover for the conflict graph that the auctioneer's revenue is maximized. The different vertices are weighted according to the prices in the bids. If one is only interested in maximizing the number of satisfied bidders, one can look for a minimum vertex cover for the conflict graph. The minimum vertex cover contains the minimum number of nodes that solves the problem, thus the number of losing bids is minimized, and that of the winning bids is maximized.

9 Cavity Approach

In Ref. [2], a cavity approach is used to solve the winner determination problem. This approach is suitable for really solving the problem for a given instance of a combinatorial auction. In the algorithm, cavity biases $u_{a\to i}$ and cavity fields $h_{i\to a}$ are associated with the links of the factor graph. This algorithm can also be applied to the conflict graph, with some modifications. The cavity bias $u_{a\to i}$ measures the likelihood that item a has already been assigned to someone else than bidder i. The cavity field $h_{i\to a}$ then measures the likelihood that bidder i would win if his bid did not contain item a. From this one can construct self-consistent equations for solving the biases and the fields, for details see Ref. [2]. The equations are essentially belief propagation.

10 Determining the Winners

Determining the winning bids is done as follows. The self-consistent belief propagation equations are solved by iteration. However, the equations may not converge, which poses a problem for the method. If convergence occurs, it results in local fields H_i. In the case where the auctioneer seeks maximal revenue, the field H_i equals the difference between the price in the bid and the price at which bidder i would win. Thus, if the field is positive for i, the bid of bidder i bid wins in all optimal solutions. If the field is negative, bidder i has not bidded enough and bid i loses. If the field is exactly zero, bidder i wins in some optimal assignments and loses in others.

Once the field H_i is known, one can fix the bids with the highest field value as winning, and then start the belief propagation equations solving procedure over for the reduced problem that contains only the bids that were not assigned to win. Based on this iteration the field H_i is again used to extract winner bids, and so on. By this procedure, the WDP is solved in $N \log N$ steps if a finite fraction of winning bids is assigned in each iteration.

11 Typical Behaviour of the WDP

To obtain estimates of the combinatorial auctions' typical behaviour, one averages the self-consistent equations for the biases $u_{a\to i}$ and the fields $h_{i\to a}$ over random ensembles. This leads to self-consistent equations for histograms of the biases and fields which are then solved. The bidders are assumed to select each item independently with probability z/M. Thus the average number of items wanted by a bidder is z. The following sections discuss two different pricing behaviours.

11.1 Constant Prices

In the first case, the price for any set of items is constant. That is, all bidders that bid for at least one item, offer the same price, 1. Thus $\nu_i = 1$ for all non-empty bids. Bids for zero items have $\nu_i = 0$.

Let us first discuss the inset of Fig. 2. In the solid, dashed, and dotted line, M/N is 0.5, 1, and 1.5, respectively. The lines were obtained by the averaging of the belief propagation equations. The markers are simulated-annealing results of optimal revenue estimates for single auction instances. As one can see, the markers agree well with the belief propagation results. Note that the revenue per item is maximal for an intermediate z. At small z, many items are not part of any bid, and the revenue stays small. At large z, the bids start to develop more and more conflicts, which decreases the total revenue.

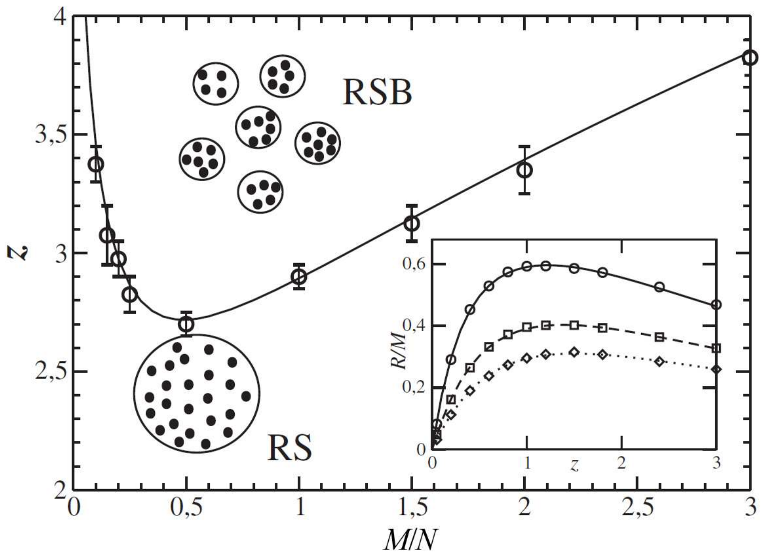

Figure 2: Phase diagram for constant-price combinatorial auctions. In the lower region, replica symmetry (RS) holds, and in the upper region the symmetry is broken. Inset: Maximal auctioneer's revenue per item for $M/N = 0.5, 1, 1.5$, from top to bottom. (From Ref. [2])

The larger figure is a phase diagram. The line marks a transition between two phases: Below the line, replica symmetry holds, whereas above the line, the larger number of conflicts between bids induces replica symmetry breaking (RSB). The onset of RSB is belived to coincide with the onset of computational hardness in this problem. In the RSB phase, the auction solutions are clustered into disconnected sets. Inside the clusters, the solution configurations are connected by a number of steps that stays finite or grows slower than the system size even if the system size goes to infinity. Between the clusters, on the other hand, the configurations are connected by paths with a thermodynamically extensive number of steps.

In the lower region, all maximum-revenue configurations belong to a single cluster, i.e., are near each other in the configuration space. Methods relying on replica symmetry are *believed* to give exact results for the maximal auctioneer's revenue. However, solution methods valid in the replica symmetric phase cannot be trusted in the replica symmetry broken phase. Replica symmetry breaking usually corresponds to non-convergence of the belief propagation equations. This is confirmed for the combinatorial auction model by the symbols in Fig. 2, which mark the points where belief propagation stops to converge.

11.2 Linear Prices

In the second case, the price of each bid is proportional to the number of items in the bid, i.e., the price per item is constant. This probably resembles a real auction more than the constant-price behaviour. The number of items is fixed the same as the number of bidders. For linear prices, selling every item would be optimal for the auctioneer.

Results for linear prices are shown in Fig. 3. The increasing lines show the auctioneer's revenue per item, and the decreasing lines the number of satisfied bidders, N_s. Different line styles correspond to different optimization goals and preferences, see the legend for explanation. The vertical lines and the symbols mark the point of replica-symmetry breakdown. For large z, the results are approximate.

The revenue curves increase with z for small z, as for constant prices. Frustration increases with z and replica symmetry is eventually broken. The increasing frustration can be seen in the curves for the number of satisfied bidders. As z increases, it eventually reaches M, and all items can be sold to one of the bidders, at which point R/M is maximal.

We see that the revenue R/M decreases substantially if we optimize only for the number of satisfied bidders. Optimizing for revenue after this somewhat increases revenue. The second optimization, of course, has no effect on N_s. Changing the order of maximization has a smaller effect on N_s than on the revenue. It makes very little difference if the additional optimization of

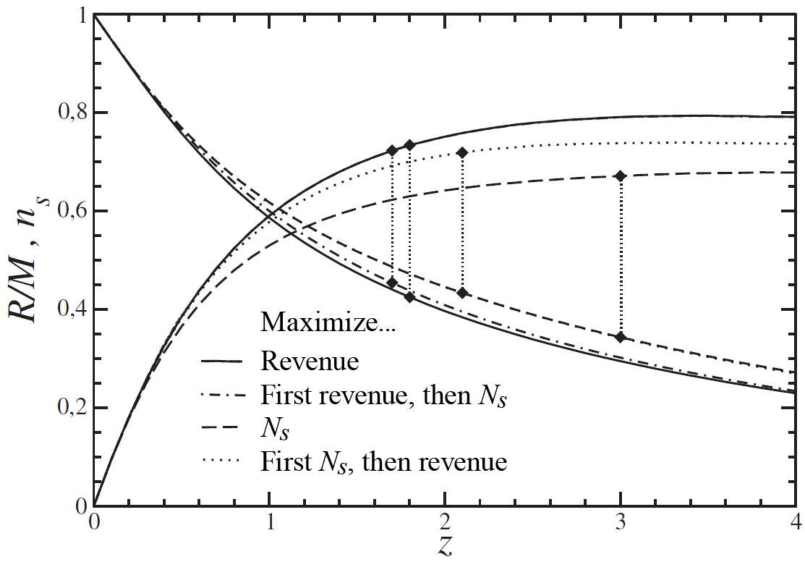

Figure 3: Auctioneer's revenue per item (increasing lines) and number of satisfied bidders (decreasing lines) for different optimization preferences. Vertical lines mark the point of replica-symmetry breakdown. (From Ref. [2], legend added)

N_s is performed after optimizing revenue or not.

12 Summary

Combinatorial auctions have been used and will be used as the means to efficiently buy or sell a large number of products or service contracts. In case studies, benefits from using a combinatorial auction have been observed for both the auctioneer and the bidders. Deciding which bids should win and which should lose is a computationally hard problem. This is due to two facts: Many bids can bid for same items, and each item can only be sold at most once. The auctioneer usually wants to maximize revenue (or minimize costs) or maximize the number of satisfied bidders.

T. Galla *et al.* have shown, using a statistical mechanics model, that there is a phase boundary between computationally easy and hard regimes which corresponds to replica symmetry breaking. In the future, an interesting topic for research would be to see how survey propagation algorithms perform on the combinatorial auction WDP.

References

[1] Cramton, Shoham, Steinberg (Eds.), *Combinatorial Auctions* (The MIT Press, 2006)

[2] T. Galla *et al.*, Phys. Rev. Lett. **97**, 128701 (2006)

[3] R. Carr, *Simulated Annealing*, Mathworld, http://mathworld.wolfram.com/SimulatedAnnealing.html, retrieved Oct 16, 2007