Combinatorial Auctions - Lecture 13

Olariu E. Florentin

January 12, 2026

Table of contents

Introduction

- Auctions have been a relatively uncommon way to negotiate the exchange of goods and commodities, until the XVII-XVIII centuries when this type of exchange begins to be used for the sale of various items like government surpluses, goods seized by armies, works of art etc.
- Today, the development of Internet makes considerably increases the use of auctions and the number of auctioneers.
- Auctions involve the sale of a large variety of distinct assets. Examples of such assets are
    - business;
    - commodities: Iron, coal, grains etc;
    - licenses for electromagnetic spectrum;
    - airport time slots;
    - railroad segments, delivery routes;
    - cars;
    - network routing.

- *English Auction:*
  - The *auctioneer* starts by announcing the reserve price (the seller's minimum acceptable price) which is usually a low price.
  - The participants can make multiple increasingly higher bids and bid against each other, until either the auction is closed (if no bids were made) or no higher bids are received.
  - The item is sold to the highest *bidder* at a price equal to his bid.
  - This is a fully transparent type of auction ("open outcry") as all the participants know each other.

- *Dutch Auction:*
  - An item is initially priced very high (by the auctioneer) and is lowered until it gets a bid (a participant which is willing to accept it).
  - The first bid made is the winning bid and results in a sale, assuming that the price is above the reserve price. This is in contrast to typical options, where the price rises as bidders compete.
- *Sealed-Bid Auction:* (or Sealed First-Price Auction)
  - All bidders simultaneously submit sealed bids to the auctioneer, so that no bidder knows how much the other auction participants have bid.
  - The highest bidder is usually declared the winner of the bidding process and pays the price he submitted.
  - Is a simple type of auction but is not as transparent as English auction.

- Because of complementarities or substitution effects between different assets, bidders have preferences not just for particular items but for sets of items, sometimes called *bundles*.
- Suppose you must auction off a dining room set consisting of four chairs and a table. Would you wish to auction off the entire set or run five separate auctions for each piece? The answer depends, of course, on what bidders care about.
- If every bidder is interested in the dining room set and nothing less, the first option is preferable. If some bidders are interested in the set but others are interested only in a chair or two it is not obvious what to do.

- If you believe that you can raise more by selling off the chairs separately than the set, the second option is preferable.
- Notice, deciding requires a knowledge of just how much bidders value different parts of the ensemble. For this reason, economic efficiency is enhanced if bidders are allowed to bid directly on combinations of different assets instead of bidding only on individual items.

Auctions where bidders are allowed to submit bids on combinations of items are usually called combinatorial auctions, see for example [Caplice], [Rothkopf], [Fujishima], and [Sandholm].

- However, such auctions were proposed as early as 1976 ([Jackson]) for radio spectrum rights. Rassenti ([Rassenti]), a little later, propose such auctions to allocate airport time slots.
- Srinivasan et al. ([Srinivasan]) have proposed a mechanism for trading financial securities that allows buyers and sellers to offer bundles of financial instruments; their mechanism treats financial securities as divisible.
- Increases in computing power have made combinatorial auctions more attractive to implement.

Combinatorial auctions are auctions where bidders can bid on bundles (packages, combinations) of items. The designer of a combinatorial auction faces a surfeit of choices, some of which we list below:

- Should the collection of bundles on which bids are allowed be restricted? If so, to what?
- Should the auction involve a single round of bidding? If so, how should the bundles be allocated as a function of the bids and what should the payment rules be?
- If the auction is to involve multiple rounds (call such auctions iterative), what information should be revealed to bidders from one round to the next?

- The choice depends on the objectives of the auctioneer. For example, is it to maximize revenue or economic efficiency? Other considerations also matter: speed, practicality, bidders preferences, and the need to discourage collusion and encourage competition among the bidders.
- Nevertheless, no matter how one chooses there are three problems that every auction designer must resolve.
  - The first has to do with *bid expression*.
  - The second is how to *allocate bundles* among bidders so as to optimize some criterion.
  - Third, what are the incentive implications of the solutions offered to the first two.

Bid Expression

- The first and most obvious difficulty faced by an auction that allows bidders to bid on combinations is that each bidder may have to determine a bid for every bundle he is interested in.
- The second problem is how to transmit this bidding function in a succinct way to the auctioneer.
- In theory, a bidder could be interested in every possible combination of items. In practice resource constraints on the part of bidders will limit the number of combinations on which they will submit bids.
- For example, in the auction of spectra, estimating the value of a bundle of spectra requires putting together a business plan. Having decided on which combinations to place a bid, the next step is to communicate this to the auctioneer.

- The difficulty now is to communicate this list, if it is particularly large, in a way that will be computationally useful to the auctioneer.
- One approach, not much explored, is to rely on an "oracle". An *oracle* is a program (black box) that, for example, given a bidder and a subset computes the bid for it. Thus bidders submit oracles rather than bids.
- The auctioneer can simply invoke the relevant oracle at any stage to determine the bid for a particular subset. Another advantage of oracles is that bidders need not be present. Their application does rely on the probity of the auctioneer. Effectiveness of this approach depends on the computational efficiency of the oracle.

- Alternatively, the auctioneer may specify a *bidding language* that all bidders must use to encode their preferences ([Nisan]). In that paper Nisan asks, given a language for expressing bids, what preferences over subsets of objects can be correctly represented by the language.
- What seems clear is that a computationally efficient oracle or language relies on restricting the preferences of bidders, or combinations on which bidders can bid.
- Another way to overcome the complexity of communicating bids and determining the winning bidders is to restrict the collection of bundles on which bidders might bid. Different scenarios along this idea are developed by [Rothkopf].
- Even if this problem is resolved (in a non-trivial way) to the satisfaction of the parties involved, it still leaves open the problem of deciding which collection of bids to accept.

- The problem of identifying which set of bids to accept has usually been dubbed the winner determination problem. The precise formulation will depend on the objectives of the auctioneer.
- Here we focus on the formulation described in [Rothkopf] and [Sandhol
  To distinguish it from other possible formulations we call it the combinatorial auction problem (CAP). (We assume that the auctioneer is a seller and bidders are buyers.)
- (CAP) can be formulated as an integer program. We will survey what is known about the (CAP). It assumes a knowledge of LP and familiarity with basic graph-theoretic terminology.

- To formulate CAP as an integer program, let N be the set of n bidders and M the set of m distinct objects.
- For every subset S of M let $b^j(S)$ be the bid that agent $j \in N$ has announced he is willing to pay for S. (Implicit is the assumption that bidders care only about the combinations they receive and not on what other bidders receive.)
- From the formulation it will be clear that bids with $b^j(S) < 0$ will never be selected. So, without loss of generality, we can assume that $b^j(S) \geqslant 0$.
- Let $y(S, j) = 1$ if the bundle $S \subseteq M$ is allocated to $j \in N$ and zero otherwise.

$$
(\text{CAP}_1)\begin{cases}
\max \displaystyle\sum_{j\in N}\sum_{S\subseteq M} b^j(S)\,y(S,j) & \\
\text{s.t.} \displaystyle\sum_{S\ni i}\sum_{j\in N} y(S,j) \leqslant 1, & \forall\, i \in M \\
\displaystyle\sum_{S\subseteq M} y(S,j) \leqslant 1, & \forall\, j \in N \\
y(S,j) \in \{0,1\}, & \forall\, S \subseteq M, j \in N
\end{cases}
$$

- The first constraint ensures that overlapping sets of goods are never assigned.
- The second ensures that no bidder receives more than one subset. Call this formulation (CAP_1).

Winner Determination CAP_2

- Problem (CAP) as formulated here is an instance of what is known as the set-packing problem (SPP), which is described below.
- When bid functions are super-additive, a more succinct formulation is possible. Let $b(S) = \max_{j \in N} b^j(S)$ and set $x_S = 1$ if the highest bid on the set S is to be accepted and zero otherwise.

Then (CAP) can be formulated as:

$$(\mathrm{CAP}_2) \begin{cases} \max & \sum\limits_{S \subseteq M} b(S) x_S \\ \text{s.t.} & \sum\limits_{S \ni i} x_S \leqslant 1, \quad \forall\, i \in M \\ & x_S \in \{0, 1\}, \quad \forall\, S \subseteq M \end{cases}$$

- Here the constraint $\sum_{S \ni i} x_S \leqslant 1$ ensures that no object in M is assigned to more than one bidder. Call this formulation (CAP_2). It is also an instance of the (SPP). (In the absence of superadditivity, one must impose the additional constraints of (CAP_1) that prevent any bidder from receiving more than one bundle in an optimal solution.)
- There is another possible interpretation of the (CAP). If the bids submitted are the actual values that bidders have for various combinations, then the solution to the CAP is the economically efficient allocation of indivisible objects in an exchange economy.

- We have formulated (CAP_1) under the assumption that there is at most one copy of each object. It is an easy matter to extend the formulation to the case when there are multiple copies of the same object and bidders may want more than one copy of the same unit.
- If the number of units of each type is large, then one could approximate the problem of selecting the winning set of bids using a linear program. The relevant decision variables would be the percentage of each type of good allocated to a bidder.

The formulation for winner determination just given is not flexible enough to encompass some of the variations that have been considered in the literature. Here is a more comprehensive formulation:

$$\begin{cases} \max \sum_{j \in N} \sum_{q \in \Omega_j} b_j(q) y(q,j) & & \\ \text{s.t.} \sum_{j \in N} \sum_{q \in \Omega_j} y(q,j) q_i \leqslant m_i, & \forall\, i \in M & (\text{GCAP}_1) \\ y^j \in P_j^A, & \forall\, j \in N & (\text{GCAP}_2) \\ y \in P^A, & & (\text{GCAP}_3) \\ y^j \in P_j^B, & \forall\, j \in N & (\text{GCAP}_4) \\ y(q,j) \in \{0,1\}, & \forall\, q \in \Omega_j, j \in N & (\text{GCAP}_5) \end{cases}$$

- Here m_i is the number of units of object i available and q is an integral vector whose ith component represents the number of units of object i demanded. $y(q,j)=1$ means that to the agent j is allocated the bundle represented by the vector q.
- The sets $\Omega_j \subseteq \mathbb{N}^M \cap [0,m_1]\times[0,m_2]\times\ldots\times[0,m_n]$ model restrictions on what bidders can bid on. They can be fixed by the auctioneer or he might permit bidders to specify them.
- The constraints (GCAP_1) ensure that no more items are allocated than the available supply.
- The constraints (GCAP_2) are imposed by the auctioneer and enforce capacity constraints on the bidders; for example no bidder is supposed to win more than two items, no bidder is supposed to win more than 40% of the total business etc. Here P_j^A denotes the polyhedron of feasible solutions to these constraints.

- Constraints ($GCAP_3$) permit the auctioneer to restrict the overall allocation. For example, the allocation must be a set of edges that form a path or a tree. Here P_A denotes the polyhedron of feasible solutions to these restrictions.
- Constraints ($GCAP_4$) allow each bidder to restrict the allocations he might win. The feasible solutions satisfying these bidder-imposed restrictions are represented by the polyhedron P_i^B. For example, if he has a sub-additive valuation, he might put $P_i^B = \{y \in \mathbb{R}^{\Omega_j} : \sum_{S \in \Omega_j} y(S, j) \leqslant 1\}$ to ensure that he does not pay more than what he bid.
- Finally, ($GCAP_5$) ensures that we end with an integral allocation.

- We give here a weighted variant. Given a ground set M of elements and a collection V of subsets with non-negative weights, find the largest weight collection of subsets that are pairwise disjoint.
- Let $x_j = 1$ if the jth set in V with weight c_j is selected and $x_j = 0$ otherwise. Define a_{ij} to be 1 if the jth set in V contains element $i \in M$.

Given this, the (SPP) can be formulated as:

$$
(\text{SPP}) \begin{cases} \max & \sum_{j \in V} c_j x_j \\ \text{s.t.} & \sum_{j \in V} a_{ij} x_j \leqslant 1, \quad \forall\, i \in M \\ & x_j \in \{0,1\}, \qquad \forall\, j \in V \end{cases}
$$

- As noted in ([Rothkopf]) and ([Sandholm]) (CAP) is an instance of the (SPP). Just take M to be the set of objects and V the set of all subsets of M.
- Before continuing with a discussion of the (SPP) we mention two of its close relatives. The first is called the set-partitioning problem (SPA) and the second is called the set-covering problem (SCP).
- Both would be relevant if we cast the auction problem in procurement rather than selling terms. The auctions used in the transport industry are of this set-covering type. In that setting, objects are origin-destination pairs, called lanes.

- Bidders submit bids on bundles of lanes that represent how much they must be offered to undertake the deliveries on the specified lanes.
- The auctioneer wishes to choose a collection of bids of lowest cost such that all lanes are served. (In fact, one must specify not only lanes but volume as well, so this problem constitutes an instance of a multi-unit combinatorial auction.)
- While (SPA) and (SCP) are cosmetically similar to the (SPP) they have different computational and structural properties.

Complexity of the (SPP) and (CAP)

- (SPP) is NP-hard, i.e., at least as hard as the hardest problems in NP.
- For the (CAP), this discussion of complexity may have little relevance. Suppose one takes the number of bids as a measure of the size of the input and this number is exponential in $|M|$.
- Any algorithm for (CAP) that is polynomial in the number of bids but exponential in the number of items would, formally, be polynomial in the input size but impractical for $|M|$ large.
- Thus, effective solution procedures for the (CAP) must rely on two things. The first is that the number of distinct bids is not large and is structured in computationally useful ways. The second is that the underlying (SPP) can be solved reasonably quickly.

- The usual way in which instances of the (SPP) can be solved by a polynomial algorithm is when the extreme points of the polyhedron

$$\mathcal{P}(A) = \left\{ x : \sum_{j \in V} a_{ij} x_j \leqslant 1, \forall i \in M, x_j \geqslant 0, \forall j \in V \right\}$$

  are all integral, i.e. 0 - 1. We can solve (SPP) as an LP.
- In most of these cases, because of the special structure of these problems, algorithms more efficient than LP ones exist.
- Nevertheless, the connection to LP is important because it allows one to interpret dual variables as prices for the auctioned objects.
- We list the most important sufficient conditions involving restrictions on the constraint matrix such that $\mathcal{P}(A)$ is integral.

- The most well known of these sufficient conditions is *total unimodularity* (TU). A special case of TU matrices are those with the *consecutive-ones property*: a $0-1$ matrix has the consecutive-ones property (COP) if the non-zero entries in each column occur consecutively (see [Nemhauser]).
- In [Rothkopf] we have the following motivation of the consecutive-ones property in the auction context. Suppose the objects to be auctioned are parcels of land along a shore line.
- The shore line is important as it imposes a linear order on the parcels. In this case it is easy to see that the most interesting combinations (in the bidders' eyes) would be contiguous. If this were true it would have two computational consequences.

- The first is that the number of distinct bids would be limited (to intervals of various length) by a polynomial in the number of objects.
- Second, the constraint matrix A of (CAP_2) would have the COP in the columns.
- If the valuation of each bidder is additive over sets of nonadjacent intervals and super-additive over sets of adjacent intervals, then (CAP_2) models the situation correctly and the problem is polynomially solvable.
- Otherwise one has to use (CAP_1), which adds constraints that violate the consecutive -ones property. It is a consequence of Keil ([Keil]) that this problem becomes NP-hard in general.

- A $0-1$ matrix, B, is *balanced* if it has no square submatrix of odd order with exactly two 1s in each row and column. TU matrices are balanced. If the matrix B is balanced then (see [Schrijver]) the linear program

$$\max\left\{\sum_j c_j x_j \,:\, \sum_j b_{ij} x_j \leqslant 1, x_j \geqslant 0, \forall j\right\}$$

  has an integral optimal solution whenever the c_j's are integral.
- For one instance of balancedness that may be relevant to the (CAP), consider a tree T with a distance function d. For each vertex $v \in T$ let $N(v,r)$ denote the set of all vertices in T that are within distance r of v.

- In other words, the vertices represent parcels of land connected by a road network with no cycles.
- Bidders can bid for subsets of parcels but the subsets are constrained to be of the form $N(v, r)$ for some vertex v and some number r. Now the constraint matrix of the corresponding (SPP) will have one column for each set of the form $N(v, r)$ and one row for each vertex of T.
- This constraint matrix is balanced (see [Nemhauser] for a proof as well as efficient algorithms). In the case when the underlying tree T is a path the constraint matrix reduces to having the consecutive-ones property. If the underlying network were not a tree then the corresponding version of (SPP) becomes NP-hard.

Perfect Matrices

- More generally, if the constraint matrix A can be identified with the vertex-clique adjacency matrix of what is known as a perfect graph, then (SPP) can be solved in polynomial time (see [Grötschel]). The algorithm, while polynomial, is impractical.
- We now describe one instance of perfection that may be relevant to the (CAP). It is related to the example on balancedness.
- Consider a tree T (as before imagine the vertices represent parcels of land connected by a road network with no cycles). Bidders can bid for any connected subset of parcels. Now the constraint matrix of the corresponding (SPP) will have one column for each connected subset of T and one row for each vertex. This constraint matrix is perfect ([Nemhauser]).
- Balanced matrices are perfect ones.

- There are situations where $\mathcal{P}(A)$ is not integral yet the (SPP) can be solved in polynomial time because the constraint matrix of A admits a graph-theoretic interpretation in terms of an easy problem.
- The best-known instance of this is when each column of the matrix A contains at most two 1s. In this case the (SPP) becomes an instance of the maximum-weight matching problem in a graph, which can be solved in polynomial time.
- Each row (object) corresponds to a vertex in a graph. Each column (bid) corresponds to an edge. The identification of columns of A with edges comes from the fact that each column contains two non-zero entries.

- It is well known that $\mathcal{P}(A)$ may contain fractional extreme points. Consider for example a graph that is a cycle on three vertices. A comprehensive discussion of the matching problem can be found in the book by Lovász and Plummer.
- The subclass of (SPP) where each column has at most $K \geqslant 3$ non-zero entries is NP-hard.
- What happens if one restricts the number of 1s in each row rather than column? This subclass of (SPP) with at most two non-zero entries per row of A is NP-hard. These instances correspond to what is called the stable-set problem in graphs. (The instance of CAP produced by the radio spectrum auction in [Jackson] reduces to just such a problem.)

- Another case is when the matrix A has the circular ones property. A $0-1$ matrix has the *circular ones property* if the non-zero entries in each column (row) are consecutive; first and last entries in each column (row) are treated consecutively. (Notice the resemblance to the consecutive-ones property.)
- The constraint matrix can be identified with what is the *vertex-clique adjacency matrix of a circular arc graph.* (Take a circle and a collection of arcs of the circle. To each arc associate a vertex. Two vertices will be adjacent if the corresponding arcs overlap.)
- The (SPP) then becomes the *maximum-weight stable set* problem for a circular arc graph, which can be solved in polynomial time.
- Following the parcels of land on the seashore example, the circular-ones structure makes sense when the land parcels lie on the shores of an island or lake.

- Caplice, C. G., *An optimization based bidding process: A new framework for shipper-carrier relationships*, Ph. D. Thesis, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, 1996.
- Fujishima, Y., Leyton-Brown, K., Shoham, Y., *Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches*, Proc. IJCAI-99, Stockholm, pp. 548 - 553, 1999.
- Grötschel, M., Lovász, L., Schrijver, A. *Geometric Algorithms and Combinatorial Optimization* Springer-Verlag, 1988.
- Jackson, C., *Technology for spectrum markets*, Ph. D. Thesis, Department of Electrical Engineering, Massachusetts Institute of Technology, Cambridge, 1976.

- Karp, R. M., *Reducibility among combinatorial problems*, R.E. Miller, J.W. Thatcher (Eds.), Complexity of Computer Computations, Plenum Press, New York, pp. 85 - 103, 1972.
- Keil, J. M. *On the complexity of scheduling tasks with discrete starting times*, Oper. Res. Lett. 12, pp. 293 - 295, 1992.
- Nemhauser, G. L.,L. A. Wolsey, *Integer and Combinatorial Optimization* John Wiley, 1988.
- Nisan, N., Ronen, A., *Computationally feasible VCG mechanisms*, Proc. ACM Conf. on Electronic Commerce (EC-00), Minneapolis, pp. 242-252, 2000.
- Rassenti, S. J., Smith, V. L., Bulfin R. L., *A combinatorial auction mechanism for airport time slot allocation*, Bell Journal of Economics, vol. 13, 1982.

- M.H. Rothkopf, M. H., Pekec, A., Harstad, A. M. *Computationally manageable combinatorial auctions*, Management Sci. vol. 44, issue 8, pp. 1131 - 1147,1998.
- Sandholm, T., *An algorithm for optimal winner determination in combinatorial auctions*, Proc. of IJCAI-99, Stockholm, pp 542 - 547, 1999.
- Schrijver, A., *Theory of Linear and Integer Programming*, Wiley, New York, 1999.
- Srinivasan, S., Stallert, J., Whinston, A. B. *Portfolio trading and electronic networks. Working paper*, Center for Research in Electronic Commerce, The University of Texas at Austin, Texas, 1998.
- Vries, S. de, Vohra, R. V., *Combinatorial Auctions: A Survey*, INFORMS Journal of Computing, vol. 15, Issue 3, pp. 284-309, 2003.