CSC200: Lecture 16

- Today
  - Review and continue combinatorial auctions
  - Next topic: matching markets (Chapter 10)
- Announcements
  - Assignment 2 due November 25
  - Term test November 27 in tutorial Friday (will indicate scope on Monday if not before)
  - Quiz 4, final quiz of term on December 4

Combinatorial Auctions

- There are much more general types of auctions when selling multiple items under various constraints. Common example: sellers offer many distinct items and buyers want different possible sets of items.
  - buyer's value may depend on the *collection* of items obtained and not just additively; and there may be different collections of interest.
- *Complements:* items whose value increase when combined
  - e.g., a cheap flight to Siena less valuable if you don't have a hotel room
- *Substitutes:* items whose value decrease when combined
  - e.g., you'd like the 10AM flight or the 7AM flight; but not both
- If items are sold separately, knowing how to bid is difficult
  - bidders run an *"exposure" risk:* might win item whose value is unpredictable because unsure of what other items they might win

Flight1

Simultaneous Auctions: Substitutes

Flight1 (7AM, no airmiles, 1 stopover)
Value: $750

Flight2 (10AM, get airmiles, direct)
Value: $950

Bidder can only use ***one*** of the flights:
Value of receiving both flights is $950

- If both flights auctioned *simultaneously*, how should he bid?
- Bid for both? runs the risk of winning both (and would need to hedge against that risk by underbidding, reducing odds of winning either)
- Bid for one? runs the risk of losing the flight he bids for, and he might have won the other had he bid
- If items auctioned *in sequence*, it can mitigate risk a bit; but still difficult to determine how much to bid first time

Simultaneous Auctions: Complements

Flight1

Hotel Room

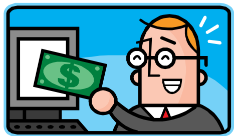

Bidder doesn't want flight without hotel room,
or hotel without flight; but together, value is $1250

- If flight, hotel auctioned *simultaneously*, how should he bid?
- Useless to bid for only one; but if he bids for both, he runs the risk of winning only one (which is worthless in isolation). Requires severe hedging/underbidding to account for this risk.
- If items auctioned *in sequence*, it can mitigate risk only a little bit. If he loses first item, fine. If he wins, will need to bid very aggressively in second (first item a "sunk cost") and could overpay for the pair.

Combinatorial Auction

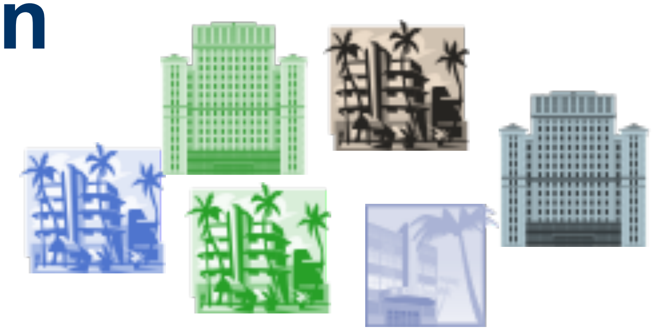

Bidder expresses value for ***combinations*** of items, e.g.:

- Value(flight2, hotel1) = $1250
- Value(flight1, hotel1) = $1050
- Don't want any other package

- *Combinatorial auctions* allow bidders to express *package bids*
  - for any combination of items, bidder can say what he is willing to pay for that combination or package
  - do not pay unless you get exactly that package
  - outcome of auction: assign (at most) one package to each bidder
  - can use 1st-price (pay what you bid) or Vickery-Clarke-Groves (VCG, a generalization of 2nd-price, which we will see later)

Combinatorial Auctions

- Formally:
  - a collection of *goods* *G* for sale
  - *bids* have the form $\{(S_1, v_1), (S_2, v_2), \ldots, (S_k, v_k)\}$, where:
    - each S_i is a subset of *G*, v_i is the price bidder will pay for S_i
    - can assign to any bidder *at most one* subset S_i from his bid
- Goal find an *assignment of goods* to bidders that maximizes the sum of the corresponding prices/valuations
  - i.e., if bidder gets the items that correspond to S_{17} in his bid, he will pay v_{17}; if items correspond to no subset in his bid, he pays nothing
  - sometimes "free disposal" assumed…
- But each item can be assigned to at most one bidder, so some hard choices need to be made by the seller

Combinatorial Auctions: Complexity

- Deciding how to allocate goods to bidders to maximize revenue (or social welfare) is computationally difficult (i.e. set packing problem)
  - formally, it is an *NP-complete problem*, which means that it is widely believed to require *exponential time to solve in the worst-case*
  - informally, it is not known whether you can do much better (in the worst case) than exhaustively searching all ways of assigning items to bidders
    - even *approximately* maximizing the social welfare is NP-hard
    - Recall our early discussion of computational difficulty
  - if I have *n* items and *m* bidders, there are $(m+1)^n$ such assignments (allow for the possibility of an item going to no bidder)

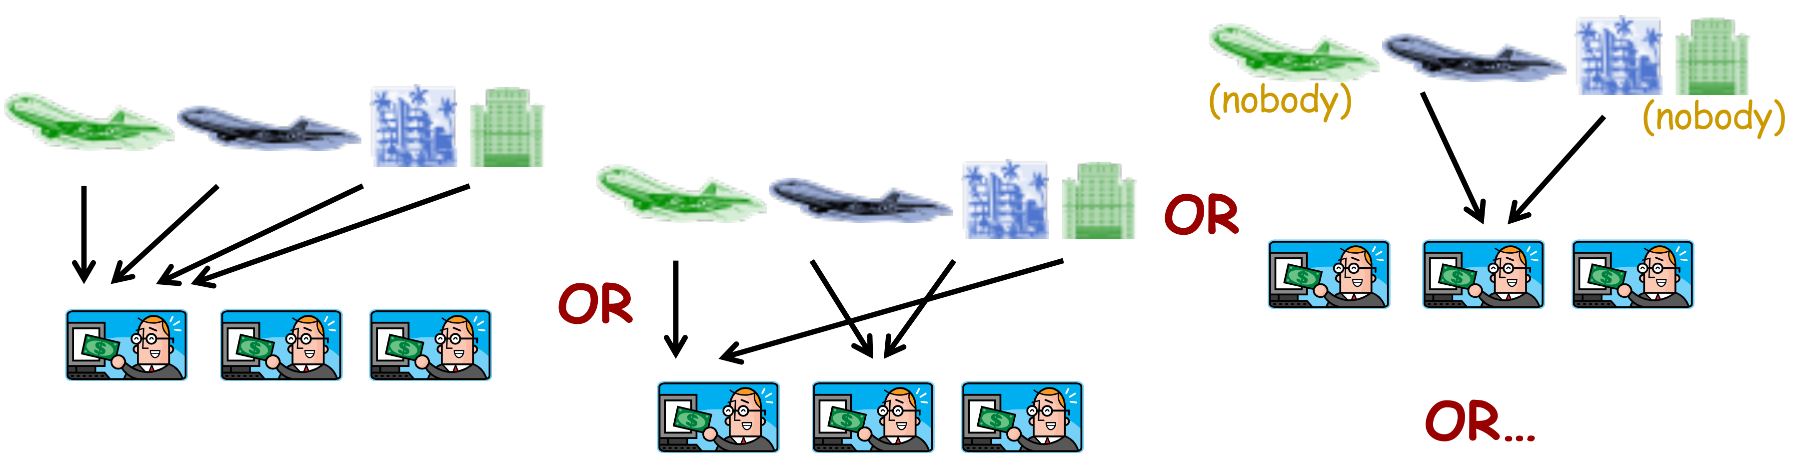

How Many Assignments?

Bidders	Items	Number of Assignments
2	10	59049
4	20	95,367,431,640,625
10	30	1.7×10^{31}
20	100	1.7×10^{132}
100	1000	2.1×10^{2004}

- Of course, you needn't look at assigning arbitrary subsets of items to a bidder: bidder only cares about certain subsets
- Suppose each bidder specifies k different subsets in his bid
  - Most bidders may be interested in "relatively few" combinations of things…
- Need only consider possible selections of a single subset from each bidder: test for feasibility (no overlap) and social welfare
- $(k+1)^m$ such subset assignments ignoring potential overlap (allowing "no subset" for a bidder)

How Many Subset Assignments?

Bidders	Number of Subsets in Bid	Number of Subset Assignments
2	10	121
4	20	194,481
10	30	819,628,286,980,801
20	100	$1.2x\ 10^{40}$
100	1000	1.1×10^{300}
200	10,000	1.0×10^{800}

- This is more manageable, but still not possible to try all.
- Still, CAs used and solved more and more widely used in practice
  - wireless spectrum, airport landing gates, industrial sourcing and procurement, etc…
  - how are they solved?
  - sometimes a good approximation (to optimal) is sufficient

Combinatorial Auctions in Practice

- Despite this, combinatorial auctions are now routinely solved involving hundreds of bidders and hundreds of thousands of items
- Algorithms exploit the fact that most bidders are interested in very few combinations or that these have "structure" (see next slides)
- The valuation functions of the bidders may be very restricted

- Some of these structures do not make the problem easier from a formal perspective: many interesting algorithmic issues remain
- But some of them do make things easier theoretically… and often in practice.

Structure in Combinatorial Auctions

- For example:
- *Single-minded bidders*: every bidder bids on only one subset
  - this is still an NP-hard problem: still 2^m subset assignments
  - even this can't be approximated efficiently according to complexity theory assumptions even ignoring game theory aspects
  - but you can get within a factor of $\sqrt{m}$ of optimizing social welfare using a truthful "greedy allocation mechanism" for single minded bidders whereas there is evidence that there may not be such a deterministic truthful mechanism for multi-minded bidders.
- *Ordered sets* of items: single-minded bidders only bid on "consecutive" subsets of items
  - e.g., only interested in adjacent facilities, plots of land (on a line), consecutive bands of wireless spectrum, adjacent time slots in schedule
  - easier computational problem: solvable in polynomial time
  - Ordered sets can be generalized to items that are arranged in a graph (e.g., an interval or a tree graph) with a special structure.

Bidding Languages

- Bidders may need to specify value for large number of combinations
  - If *n* items, there are 2^n packages that bidder must consider
- But there are usually a lot of packages bidder doesn't care about, and a lot of structure in values
  - e.g., suppose items are strict substitutes: here are the 10 flights I care about, here's how much each is worth, give me only one
  - e.g., independent/additive values: here are the 10 items are I care about, here's how much each is worth, regardless of how many others I get
  - e.g., *k*-of complements: I will pay *$d* for any 10 items from this set of 100
- *Bidding languages* let bidders express values concisely, naturally
  - algorithms can often exploit these languages in practice
  - still not in general theoretically tractable, but often practically so…

Other Issues

- A variety of interesting strategic issues:
  - *envy-free*: find an allocation (and prices to charge) so that no bidder would prefer the bundle of goods allocated to a different bidder
  - *stability*: find payments so no group of bidders could offer to pay more for a set of goods allocated to others, divide it up, and be better off
- *What about pricing:* if people pay what they bid (1st-price), they will obviously hedge their bids; is there an analog of a 2nd-price auction in the combinatorial setting? What would the 2nd highest bid be?

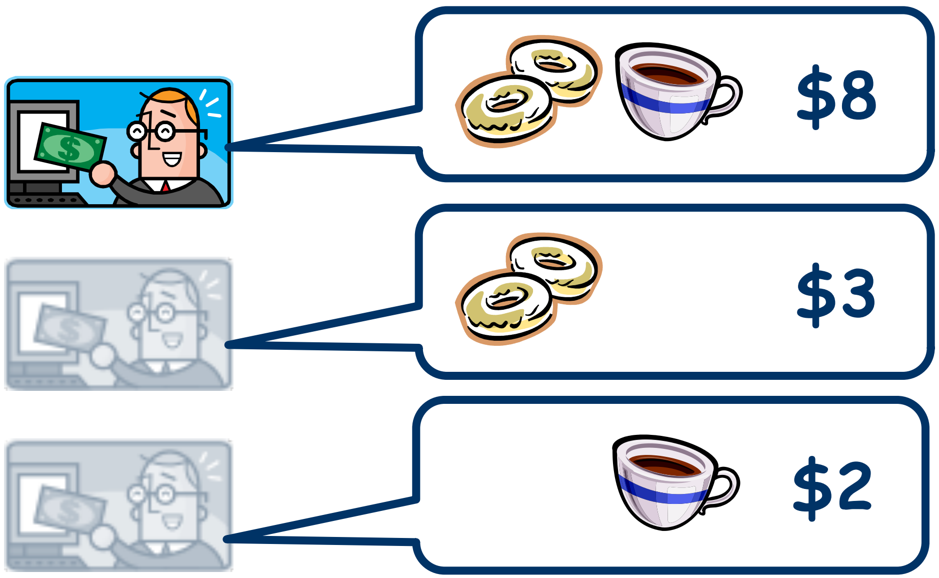

Winner gets bundle of both items. But no other bidder offered a bid on the same bundle: so what price should we charge using "2nd price"?

Another combinatorial auction

- Agents can have very different valuations:
  - *What about pricing:* if people pay what they bid (1st-price), they will obviously hedge their bids; is there an analog of a 2nd-price auction in the combinatorial setting?

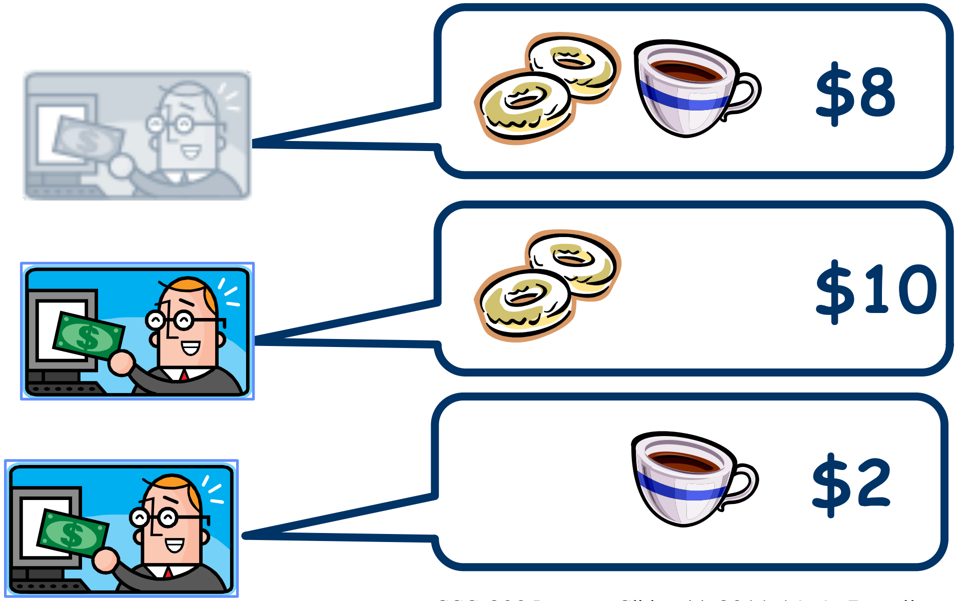

To maximize social welfare we have two winners if agents bid truthfully: second agent gets donuts while third agent gets coffee. So what price should they pay?

The VCG Mechanism

- There is a generalization of the 2nd-price (or Vickrey) auction to CAs
  - The *Vickrey-Clarke-Groves (VCG)* mechanism
- Lets assume (and this is a big assumption) that we know how to allocate items optimally. Now how to charge?
- Roughly speaking, you charge someone based on the “externality” they inflicted on other players by their presence
  - e.g., bidder *X* gets bundle *B*, this (potentially) prevented other bidders *Y, Z, …* from getting (some parts of) bundle *B*
  - figure out social welfare that *Y, Z, …* got in the actual auction
  - then pretend *X* didn’t exist and figure *SW* that *Y, Z, … would have attained* if X hadn’t bid/didn’t exist (can’t be any worse, could be higher)
  - charge X the difference of the two: what he cost *Y, Z,…* by his presence
- In slide 13 winner pays $5 for coffee and donuts. Slide 14 outcome??
- Notice that 2nd-price auction is a special case of VCG with one item
- We’ll see VCG in more detail in Ch.15 (advertising auctions)

Simple auctions

- "Simplicity is the ultimate sophistication" (Leonardo da Vinci

- In an auction, if possible we would like
  - Computational simplicity (eg computationally fast)
  - Strategic simplicity : agents should be able to easily and quickly know what to do
  - Conceptual simplicity : the rules of the auction should be easy to understand and perceived as being fair

Do combinatorial auctions using VCG satisfy these?

Conceptually simple auctions

- If you were a buyer in a combinatorial auction, would you understand what to do?
- Even knowing about (say) the VCG mechanism (but not understanding exactly how the mechanism chooses an optimal allocation) would you feel confidant in bidding?
- Sandholm and Gilpin [2004] argue that mechanisms such as VCG often fail in the real world for various reasons:
  - Buyers may be unwilling to reveal their true values
  - A buyer may be unwilling to participate in an auction where the rules are complex, not fully understood or unintuitive
  - The computational difficulty makes an optimal allocation impossible and VCG in general requires an optimal allocation

Pricing to achieve objectives

- There can be many objectives for a mechanism as we have already seen; social welfare, seller revenue and in our next topic (matching markets) we will consider "market clearing" (which Tyrone spoke about in the context of one buyer and many sellers).
- A number of mechanisms just use item pricing to achieve desired objectives.
- These prices may be offered all at once or dynamically in some manner and then agents choose an allocation based on their valuation function and the item prices.