

Distributed Software Development

More Auctions

Chris Brooks

Department of Computer Science

University of San Francisco

24-2: Previously on “Let’s Make a Deal”

- Auctions are a nice way to solve resource allocation problems with self-interested agents.
 - Explicit, publically known rules
 - Multilateral - can handle many exchanges simultaneously.
 - Well-understood outcomes.

24-3: Previously on “Let’s Make a Deal”

- English auction: ascending, first-price.
 - Dominant strategy, easy to understand
 - Can require many rounds, may not be appealing to the seller.
- First-price sealed-bid.
 - One-shot, goes to the highest bidder.
 - No dominant strategy.

24-4: Previously on “Let’s Make a Deal”

- Dutch auction: descending, first-price.
 - Equivalent to first-price sealed-bid.
 - Fast in practice.
- eBay actually uses something slightly different (which they also call a Dutch auction) to sell multiple items.
- n identical items are sold, all winners pay the n th highest price.

24-5: Previously on “Let’s Make a Deal”

- Vickrey auction: one-shot, second-price
 - Truth-telling as the dominant strategy
- Generalizes to $n + 1$ st-price auction when n items are sold.

24-6: AdWords

- This is essentially how Google AdWords works
- You choose a keyword and an amount that you're willing to pay.
- Your bid is constructed from your willingness to pay and your ad's click-through rate.
- Google then uses a real-time Vickrey auction to determine which ads to place on a page.
- You pay based on a function of the second-highest bid and your click-through rate.
- Retains Vickrey's appealing properties.

24-7: Advantages and disadvantages

- Advantages of the Vickrey auction/Clarke tax:
 - Truth-telling as a dominant strategy
 - Easy for participants, no need for multiple rounds of bidding.
 - Most efficient solution is always discovered.
 - Disadvantages:
 - Leaves money 'on the table' (payments are more than cost of job)
 - Payments are a function of the quality of the second-best solution.
 - Not intuitive for humans.

24-8: Common and correlated-value auctions

- Everything we've said so far applies only to private value auctions.
- Common or correlated-value auctions are much less predictable.
- In particular, common-value auctions are subject to the *winner's curse*
 - As soon as you win a common-value auction, you know you've paid too much.

24-9: Winner's curse

- Example: Oil drilling
 - Suppose that four firms are bidding on drilling rights. Each has an estimate of how much oil is available in that plot.
 - A thinks \$5M, B thinks \$10M, C thinks \$12M, and D thinks \$20M.
 - Let's say it's really \$10 M, but the firms don't know this.
 - In an English auction, D will win for \$12M (plus 1 dollar)
 - They lose \$2M on this deal.
 - Problem: The winner is the firm who tended to overestimate by the most.
 - (Assumption: all firms have access to the same information.)

24-10: Winner's curse

- This also explains why sports free agents seem to underperform their contracts.
 - They're not underperforming, they're overpaid.
- How to avoid the winner's curse:
 - Better information gathering
 - Caution in bidding

24-11: Combinatorial auctions

- Often, goods that are being sold in an auction have complementarities.
 - Owning one good makes a second good more valuable.
- For example, let's say supercomputer access is sold in 1-hour increments.
- Lab 1 needs three hours before 5 pm - less time is worthless.
- Lab 2 needs two hours before noon.
- How to approach this:
 - 1. Separate auctions for each hour.
 - Complicated rules for backing out and reallocating needed.
 - 2. Auction combinations (or bundles) of goods.

24-12: Winner-determination problem

- Finding the winner for a single-item Vickrey auction is easy.
- Finding the winner for a combinatorial auction is (computationally) hard.
- Formulation:
 - Given: n bidders, m items
 - Let a bundle S be a subset of the m items.
 - A bid b is a pair (v,s) , where v is the amount an agent will pay for s .
 - An allocation $x_i(S)$ is described by a mapping from (i,s) into $\{0,1\}$. ($x_i(S) = 0$ if i does not get s , and 1 if he does.)

24-13: Winner-determination problem

- We can then write the winner-determination problem as an optimization problem:
 - Find the set of allocations that maximizes:
$$\sum_{i \in N} (v, s) x_i(s)$$
 - This problem can be solved in a number of ways; integer linear programming or backtracking search are the most common.

24-14: Winner-determination problem

- Problem: The size of the WDP is exponential in the number of items that can be sold.
 - Every possible bundle must be considered.
- Formulating the problem as ILP helps some
 - This problem has been studied since the 50s, so good heuristic techniques exist.

24-15: Winner-determination problem

- Other solutions:
 - Limiting the sorts of bundles allowed.
 - OR bids and XOR bids.
 - This transforms the problem into the knapsack problem.
 - Still NP-hard, but good heuristics exist
- Limiting size of bundles.
- Approximation algorithms

24-16: Current research issues

- Auctions are a particularly hot area of research.
- Topics include:
 - Information revelation - how can we preserve the truth-telling strategy of Vickrey without agents revealing their preferences to each other?
 - Winner determination.
 - Languages for expressing more complex constraints.
 - Preventing collusion and false-name bids.
 - 'online' auctions
 - Not "on the Internet" - meaning agents continuously arrive and leave.

24-17: Summary

- Vickrey auctions are particularly appealing from a computational standpoint.
 - Easy for participants to decide how to act.
 - Hard to manipulate.
- Resources always allocated to the agent that values them the most.
- Challenges:
 - Dealing with imperfect information
 - Combinatorial auctions run us up against NP-completeness (again).

24-18: Trust

- Another challenge in systems of self-interested agents is dealing with trust.
 - How can I ensure that an agent is who it says it is?
 - How can I ensure that the information it is giving is accurate?
- There are essentially two different approaches:
 - Cryptographic approaches
 - Reputation-based approaches

24-19: Cryptographic Trust

- How can I ensure that a message is from the person who claims to have sent it?
- The standard response is to use public-key encryption.
- I sign a message with my private key and send it to a recipient.
- They use my public key to decrypt the message.
- Since only someone with my private key could sign the message, it must be from me.
- Right?

24-20: Authorities

- Where did the recipient get my public key from?
- It could have been spoofed
 - An attacker distributes a fake public key for me, and then signs messages claiming to be from me with the corresponding fake private key.
- How can a recipient ensure that my public key is genuine?

24-21: Authorities

- I have someone else sign my public key with their private key
- If you already have their public key, then you can use that to verify that my key is genuine.
- If not, they will need to have a third organization sign the key with their private key.
- Two possible architectures:
 - Hierarchical: Some well-known organization acts as a key signer.
 - Verisign does this.
 - Web of trust: There is no central signer.
 - Instead, you find someone whose key you already know and try to construct a chain of signings that will allow you to verify new keys.

24-22: Reputation

- Encryption is useful for dealing with identity-related trust issues.
- It doesn't help at all with truth-related trust issues.
 - Just because you can verify that a message is from a particular person doesn't mean that you can believe its contents.
- This requires a different sort of trust - a belief that the the other party is not going to decieve or exploit you.
- In real life, we solve these problems via reputation.
 - If I've interacted with you before and it's gone well, I'll trust you in the future.
 - If someone I trust vouches for you, I'll trust you.

24-23: Reputation schemes in e-commerce sites

- Sites like eBay let users assign ratings to each other based on past experience.
 - Positive, negative, neutral
- This gives you some confidence that a seller is (or isn't) trustworthy.
- What could go wrong here?

24-24: Reputation schemes in e-commerce sites

- New identities
- Fake ratings
- Collusion among friends
- Unfair negative ratings
- Getting an accurate sample - how to make users provide feedback?
- Reputation can't be shared among sites

24-25: Current research problems

- How to prevent change of identity and false evaluations
- How to bootstrap users fairly
- Determining scale/granularity
- “transitive reputation”
- Sharing reputation across sites

24-26: Summary

- Mechanisms for ensuring trust on the Internet are still pretty primitive
- Identity can be dealt with, but more complex trust is difficult.
 - PKI works fine for verifying someone's identity, but requires the use of public keys.
 - Only verifies identity, not content
- Ensuring good behavior seems to require the development of communities and ongoing relationships (just like in real life)