Multi-Agent Systems

Albert-Ludwigs-Universität Freiburg

Bernhard Nebel, Felix Lindner, and Thorsten Engesser

Summer Term 2017

Course outline

Recap & Outlook

- Last time
  - Distributed Constraint Satisfaction as a means to implement distributed problem solving: Each agent tries to assign some value to its private variable under consideration of its known constraints and on the feedback it receives from the other agents.
- Today
  - Agents display selfishness, i.e., they have preferences (of probably different strengths) and compete with each other about rare goods. We will employ the metaphors of Auctions and Markets.

Auctions: Basic Terminology

- A seller is auctioning one item to a set of bidders.
- Each bidder has an intrinsic value (or true value) for the item.

Auctions: Basic Assumptions

- Bidders have independent, private values, i.e., agents are uncertain about the other agents' intrinsic values.
- Sellers do not have a good estimate of the bidders' true value of an item.

Opposing Goals

- Sellers want to maximize the price at which the good is allocated: "Which type of auction should I set up for my good?"
- Bidders want to minimize the price at which the good is allocated: "Which price should I bid?"

Auctions: Types

Examples

- Ascending-bid auctions (English auction)
- Descending-bid auction (Dutch auction)
- First-price sealed-bid auctions
- Second-price sealed-bid auctions (Vickrey auctions)
- and there are many more.

Dimensions

- Bidding rules: How are offers made?
- Clearing rules: When do which trades occur as a function of the bidding?
- Information rules: Who knows what when about the state of the auction?

English auction

- Bidding rule: Auctioneer starts off by an initial price for the good. Agents are invited to bid more.
- Clearing rule: No agent is willing to raise current bid, and the highest bidder wins.
- Information rule: Every agent can see what the others are bidding.

Dutch auction

- Bidding rule: Auctioneer starts off by an initial price for the good. Auctioneer continually lowers the offer price.
- Clearing rule: An agent makes a bid for the good which is equal to the current offer price.
- Information rule: Every agent can see what the others are bidding.

First-price sealed-bid auction

- Bidding rule: Bidders submit to the auctioneer a bid for the good.
- Clearing rule: Winner is the agent who made highest bid. Winner pays the price (s)he bid.
- Information rule: No agent can see what the others are bidding.

Second-price sealed-bid auction

- Bidding rule: Bidders submit to the auctioneer a bid for the good.
- Clearing rule: Winner is the agent who made highest bid. Winner pays the second-highest price bid.
- Information rule: No agent can see what the others are bidding.

Discussion

- Which auction protocol would you prefer as a bidder?
- Which auction protocol would you prefer as an auctioneer?

Outline of interesting results

1. Descending-bid auctions correspond to First-price sealed-bid auctions.
2. Ascending-bid auctions correspond to Second-price sealed-bid auctions.
3. In Second-price sealed-bid auctions bidding one's true value is a dominant strategy.
4. In First-price sealed-bid auctions $s(v_i) = (\frac{n-1}{n})v_i$ is an equilibrium strategy.
5. The money the seller can expect from a First-price sealed-bid auction is the same it can expect from a Second-price sealed-bid auction.

Correspondence I

Proposition

Descending-bid auctions correspond to First-price sealed-bid auctions.

Proof

- Consider a descending-bid auction:
- Seller is lowering the price from its high initial starting point.
- No bidder says anything until finally s.o. accepts the bid and pays the current price.
- For bidder i, there's a first price b_i (s)he accepts.
- The price b_i plays the role of bidder i's bid in a first-price sealed-bid auction.

Correspondence II

Proposition

Ascending-bid auctions correspond to Second-price sealed-bid auctions.

Proof

- Consider an ascending-bid auction.
- The seller steadily raises the price and bidders gradually drop out.
- The winner of the auction is the last bidder remaining, and (s)he pays the price at which the second-to-last bidder dropped out.

Outline of interesting results

1. Descending-bid auctions correspond to First-price sealed-bid auctions.
2. Ascending-bid auctions correspond to Second-price sealed-bid auctions.
3. In Second-price sealed-bid auctions bidding one's true value is a dominant strategy.
4. In First-price sealed-bid auctions $s(v_i) = (\frac{n-1}{n})v_i$ is an equilibrium strategy.
5. The money the seller can expect from a First-price sealed-bid auction is the same it can expect from a Second-price sealed-bid auction.

Second-price sealed-bid auctions as games: Basics

- Set of players $\mathcal{N} = \{1, 2, \dots, n\}$ (the bidders)
- The true value v_i of player i for the item
- Strategy $s(v_i) = b_i$ is the bid of player i as a function of v_i
- Payoff: If b_i is not the winning bid, then the payoff to i is 0. If b_i is the winning bid, and some other b_j is the second-place bid, then the payoff to i is $g(v_i) = v_i - b_j$.

Dominant strategy for second-price sealed-bid auctions

Proposition

In a second-price sealed-bid auction, it is a dominant strategy (i.e., best strategy regardless of what strategy everyone else is using) for each bidder i to choose a bid $b_i = v_i$.

Dominant strategy for second-price sealed-bid auctions: Proof

Proof

- Case i's valuation v_i is larger than the highest of the other bidders' bids.
  - If i bids v_i it wins and pays the next-highest bid amount.
  - If i bids more than v_i, i would still win and pay the same.
  - If i bids less than v_i, it either still wins and pays the same, or it looses and receives zero payoff.
- Case i's valuation v_i smaller than the highest of the other bidders' bids.
  - If i bids v_i, it looses and gains zero payoff.
  - If i bids more than v, i either looses and gains zero payoff or wins and gains zero (or less) payoff.
  - If i bid less than v, it still looses and gains zero payoff.

Outline of interesting results

1. Descending-bid auctions correspond to First-price sealed-bid auctions.
2. Ascending-bid auctions correspond to Second-price sealed-bid auctions.
3. In Second-price sealed-bid auctions bidding one's true value is a dominant strategy.
4. In First-price sealed-bid auctions $s(v_i) = (\frac{n-1}{n})v_i$ is an equilibrium strategy.
5. The money the seller can expect from a First-price sealed-bid auction is the same it can expect from a Second-price sealed-bid auction.

Equilibrium in First-Price sealed-bid Auctions: Two Bidders

UNI FREIBURG

Assumptions

- Common knowledge between two bidders: v_1, v_2 are independently and uniformly distributed between 0 and 1.
- Both bidders use the same strategy $s(\cdot)$
- $s(\cdot)$ is strictly increasing and differential (so, if two bidders have different values, then they will bid differently).
- $s(v) \leq v$ for all v, $s(0) = 0$

Sample Strategies

- $s(v) = v$: bid your true valuation
- $s(v) = cv, c < 1$: scale down your bid a bit
- ...

Equilibrium in First-Price sealed-bid Auctions: Two Bidders

UNI FREIBURG

Observations

- If bidder i has a true value of v_i (drawn from [0, 1]), the probability that i wins is v_i
- If bidder i wins, it receives the payoff $v_i - s(v_i)$
- Thus, the expected payoff for bidder i is

$$g(v_i) = v_i(v_i - s(v_i))$$

- Note that $s(v_i) = v_i$ is not a dominant strategy in first-price auctions: If i loses, it gets 0 payoff, and if i wins, the payoff is $(v_i - v_i) = 0$, too.

Equilibrium in First-Price sealed-bid Auctions: Two Bidders

UNI FREIBURG

Equilibrium strategy

- For $s(\cdot)$ to be an equilibrium strategy means that for each bidder i, there is no incentive to deviate from strategy $s(\cdot)$ if its competitor is also using strategy $s(\cdot)$.
- The incentive to deviate can be simulated by pretending other deviations v from v_i.
- Formally: The expected payoff $g(v) = v(v_i - s(v))$ should be maximal for $v = v_i$.

Equilibrium in First-Price sealed-bid Auctions: Two Bidders

UNI FREIBURG

Equilibrium strategy

- We set $g'(v) = v_i - s(v) - vs'(v) = 0$, and ask for $s(\cdot)$.
- $s'(v) = \frac{v_i - s(v)}{v}$
- $s'(v_i) = 1 - \frac{s(v_i)}{v_i}$
- $s(v_i) = v_i/2$ solves this differential equation.

Equilibrium in First-Price sealed-bid Auctions: n Bidders

Equilibrium strategy

- Expected payoff for bidder i

$$G(v_i) = v_i^{n-1}(v_i - s(v_i))$$

- Setting $G'(v_i) = 0$, we have

$$(n-1)v^{n-2}v_i - (n-1)v^{n-2}s(v_i) - v_i^{n-1}s'(v_i) = 0$$

- Then we obtain $s'(v_i) = (n-1)(1 - \frac{s(v_i)}{v_i})$
- This differential equation is solved by

$$s(v_i) = \left(\frac{n-1}{n}\right) v_i$$

Outline of interesting results

1. Descending-bid auctions correspond to First-price sealed-bid auctions.
2. Ascending-bid auctions correspond to Second-price sealed-bid auctions.
3. In Second-price sealed-bid auctions bidding one's true value is a dominant strategy.
4. In First-price sealed-bid auctions $s(v_i) = (\frac{n-1}{n})v_i$ is an equilibrium strategy.
5. The money the seller can expect from a First-price sealed-bid auction is the same it can expect from a Second-price sealed-bid auction.

Seller Revenue

The seller's perspective

Which auction type should the seller choose to get as much money as possible?

From what we know so far

- Two competing forces
  - Second-price sealed-bid auction: Seller charges second-highest bid
  - First-price sealed-bid auction: Bidders reduce their bids

Seller Revenue

Proposition

Suppose n numbers are drawn independently from the uniform distribution on the interval $[0, 1]$ and then sorted from smallest to largest. The expected value of the number in the kth position on this list is $\frac{k}{n+1}$.

Application to Second-price Auctions

The seller gets the second-highest value, i.e., at position $n-1$. Thus the expected value is $\frac{n-1}{n+1}$.

Application to First-price Auctions

The winning bidder bids $\frac{(n-1)}{n}v_i$ (Equilibrium strategy). Using the propositon, $v_i = \frac{n}{n+1}$. Hence $\frac{(n-1)}{n}\frac{n}{n+1} = \frac{n-1}{n+1}$

Summary of interesting results

1. Descending-bid auctions correspond to First-price sealed-bid auctions.
2. Ascending-bid auctions correspond to Second-price sealed-bid auctions.
3. In Second-price sealed-bid auctions bidding one's true value is a dominant strategy.
4. In First-price sealed-bid auctions $s(v_i) = (\frac{n-1}{n})v_i$ is an equilibrium strategy.
5. The money the seller can expect from a First-price sealed-bid auctions is the same it can expect from a Second-price sealed-bid auctions.

Discussion

- Application to ressource assignment: Is it wise to just sequentially run auctions to assign several items to several agents?

a	1	10, 9
b	2	5, 0

- Matching markets generalize this setting to many items and many bidders.

Matching Markets

- Origin: Economics, operation research
- Agents have different preferences for different kinds of goods
- Prices determine the allocation of goods . . .
- . . . and socially optimal allocations emerge.

Agenda

1. Bipartite graphs, paths, and matchings
2. Market Framework
3. Algorithm for computing market-clearing prices

Bipartite Graph

Definition

A graph $G = (V, E)$ is a bipartite graph iff there are two disjoint sets $A, B \subseteq V, A \cup B = V$ such that every edge $e \in E$ connects a vertex in A to one in B.
G is balanced iff $|A| = |B|$.

Definitions: Matching

Matching

Given a graph $G = (V, E)$, a matching $M \subseteq E$ in G is a set of edges such that no two edges share a common vertex.

Maximum matching

A matching is a maximum matching iff it contains the largest possible number of edges.

Perfect matching

A perfect matching in a balanced bipartite graph $G = (A \cup B, E)$ is a set of edges between nodes in A and nodes in B such that each node is endpoint of exactly one edge.

Definition: Free Vertex

Free Vertex

A vertex $v \in V$ is a free vertex w.r.t. to a matching M iff no edge from M is incident to v.

Definitions: Path

Path

A path $P \subseteq E$ in a graph $G = (V, E)$ is a sequence of edges which connect a sequence of distinct vertices.

Alternating Path

P is an alternating path w.r.t. M, iff P is a path in G and for every two subsequent edges on P it holds that one is in M and the other is not.

Augmenting Path

An alternating path P is an augmenting path iff its start and end vertices are free.

Maximum matching on bipartite graph

function MAXIMUM-MATCHING(G)
    $M = \emptyset$
    repeat
        $P =$ AUGMENTING-PATH(M, G)
        $M = M \otimes P$ [Symmetric difference, $A \otimes B = \{x | x \in A \oplus x \in B\}$]
    until $P = \emptyset$
end function

function AUGMENTING-PATH(M, G)
    Direct unmatched edges $A \to B$, matched $B \to A$
    Add vertices s, t, connect s to free vertices in A and the free vertices of B to t
    Run BFS to find a shortest path P from s to t
    return $P \setminus \{s, t\}$
end function

Example: Compute maximum matching

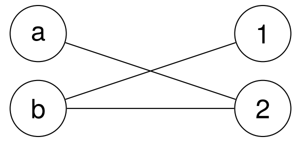

Example: Compute maximum matching II

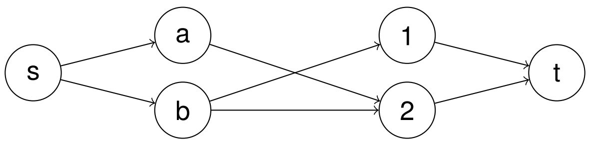

Example: Compute maximum matching III

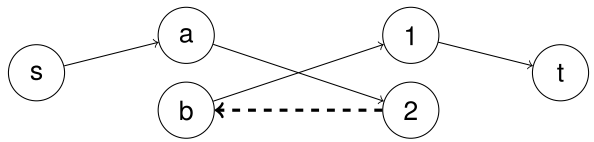

Example: Compute maximum matching IV

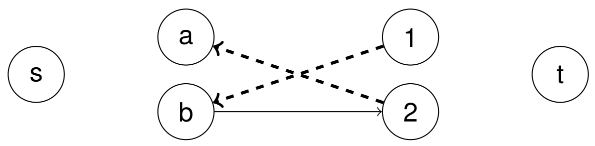

- If the graph has a perfect matching, a component can solve the allocation problem using the maximum-matching algorithm in a centralized way.

Agenda

The market framework

- Set $\mathcal{S}$ of Sellers (each with an object to sell)
- Set $\mathcal{B}$ of Buyers (each with valuation for each object)
- v_{ij}: Valuation of buyer j for the object held by seller i
- p_i: Price fixed by seller i
- $v_{ij} - p_i$: Buyer j's payoff if he/she/it buys from seller i
- Note: If $v_{ij} - p_i$ is negative for some j and all i then j does no transaction and receives payoff 0.

Preferred sellers and Market-clearing prices

Definition: Preferred sellers

The set of sellers that maximize j's payoff is called preferred sellers of buyer j (provided the payoffs are non-negative).

Definition: Preferred seller graph

The preferred seller graph (PSG) is the bipartite graph $G = (\mathcal{S} \cup \mathcal{B}, E)$ s.th. $(i,j) \in E$ iff i is a preferred seller of j.

Definition: Market-clearing prices

A set of prices is called market clearing iff the resulting preferred-seller graph has a perfect matching.

Neighbor Sets and Constricted Sets

Definition: Neighbor Set

Given a balanced bipartite graph $G = (\mathcal{S} \cup \mathcal{B}, E)$. For any sets of nodes $X \subseteq \mathcal{B}$ (resp. $\mathcal{S}$) a node $v \in \mathcal{S}$ (resp. $\mathcal{B}$) is a neighbor of X iff it has an edge to some node in $\mathcal{B}$.
The neighbor set of $\mathcal{B}$, $N(\mathcal{B})$, is the set of all neighbors of $\mathcal{B}$.

Definition: Constricted Set

A set $\mathcal{B}$ is a constricted set iff $|\mathcal{B}| > |N(\mathcal{B})|$.

The effect of prices on PSG: Illustrations I

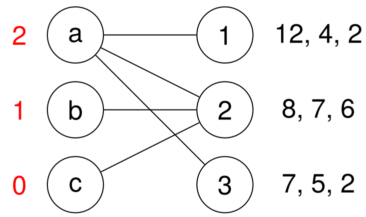

- Does the graph have a constricted set?
- Is there a perfect matching?
- Are the prices market clearing?

The effect of prices on PSG: Illustrations II

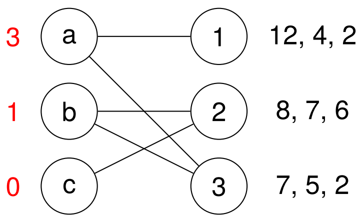

- PSG has perfect matching $\Rightarrow$prices clear the market.
- Some cooperation needed among sellers (or use the previously introduced centralized algorithm)
- Do the agents as a whole get what they want most?

The effect of prices on PSG: Illustrations III

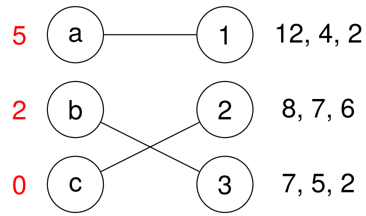

- Higher prices do not improve the allocation.
- Payoff: the sellers' gain = the buyers' loss.

Matching Theorem

Theorem (see [2])

A preferred seller graph has a perfect matching iff it contains no constricted set.

Market-Clearing Prices: Optimality

Theorem (see [2])

Given a set of market-clearing prices. A perfect matching in the resulting preferred-seller graph has the maximum total valuation.

Proof

- Consider a set of market-clearing prices *MCP*, and let *M* be a perfect matching in the PSG
- Total payoff of buyers = Total Valuation – Sum of all prices
- Since each buyer individually maximizes her payoff (PSG!), the total payoff of buyers is maximal, i.e., no other matching yields more payoff than *M* under the given *MCP*.
- The sum of prices is the same for all matchings $\Rightarrow$Matchings that maximize total payoff also maximizes total valuation.

Market-Clearing Prices: Optimality II

Theorem (see [2])

Given a set of market-clearing prices. A perfect matching in the resulting preferred-seller graph produces the maximum possible sum of payoffs to all sellers and buyers.

- Very similar argument.

Agenda

Market-Clearing Prices: Existence

Theorem

For any set of buyer valuations, there exists a set of market-clearing prices.

- Let's construct a algorithm for searching for a set of market-clearing prices.

Market-Clearing Prices: Algorithm

Algorithm (see [2])

1. All sellers initially set prices to 0
2. Construct preferred-seller graph *PSG*
3. If *PSG* has a perfect matching, it's done
4. Else there must be a constricted set of buyers $\mathcal{B}$
5. The neighbors $N(\mathcal{B})$ raise their prices by one unit
6. Scale prices s.th. the minimum price is 0
7. Goto 2

- This algorithm terminates $\Leftrightarrow$ it finds a set of prices that yield a PSG with a perfect matching $\Leftrightarrow$ (by definition) it finds a set of market-clearing prices.
- Will it terminate?

Example execution I

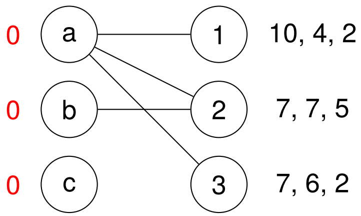

- Note: In this case there are two constricted sets $\{1,3\}$ and $\{1,2,3\}$. Either of them can be chosen.

Example execution II

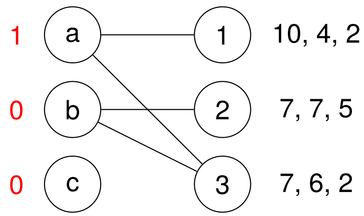

Example execution III

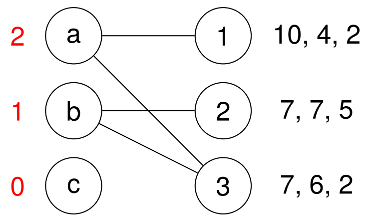

Example execution IV

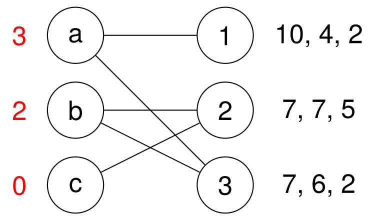

- Perfect matching: $(1, a), (2, c), (3, b)$

Market-Clearing Prices: Termination

Definition: Potentials of buyers, sellers, and of the auction

The potential of a buyer is the maximum payoff she can currently get from any seller. The potential of a seller is the current price she is charging. The potential energy of the auction is the sum of the potential of all buyers and sellers.

- Note: The minimum payoff of a buyer (its potential) is 0. The minimum price a seller can charge (its potential) is 0. Thus, the potential energy of the auction must be greater or equal to 0.

Market-Clearing Prices: Termination

Theorem (see [2])

The procedure always terminates.

Proof

- Initially, all sellers have potential 0 and buyers have potential according to their highest valuation. Thus, at the start $P_0 \geq 0$.
- In every round, when sellers in $N(\mathcal{B})$ ($\mathcal{B}$ constricted set) raise their prices, each of these sellers' potential increases and the potential of buyers in $\mathcal{B}$ decreases.
- Since $|\mathcal{B}| > |N(\mathcal{B})|$, the overall potential decreases.
- Because the potential cannot drop below 0, the procedure will terminate.

Single-item auction as special case

a	1	3, 0, 0
b	2	2, 0, 0
c	3	1, 0, 0

- Item *a* is the single item, the other items are dummies (to have a balanced bipartite graph)
- Item *a* increases its price until only one is left.
- Finally, 1 gets assigned to *a* and he/she/it must pay 2.
- Proposition: Market-Clearing Prices auction implements English auction (second-price sealed-bid auction).

Course outline

Literature I

M. Wooldridge, An Introduction to MultiAgent Systems, John Wiley & Sons, 2002.

D. Easley, J. Kleinberg, Networks, Crowds, and Markets: Reasoning about a Highly Connected World, Cambridge University Press, 2010.