Introduction to Multi-Agent Programming

11. Auctions

English, Dutch, Vickrey, and Combinatorial Auctions

Alexander Kleiner, Bernhard Nebel

Contents

Introduction I

- With the rise of the Internet, auctions have become popular in many e-commerce applications (e.g. eBay)
- Auctions are an efficient tool for reaching agreements in a society of self-interested agents
  - For example, bandwidth allocation on a network, sponsor links
- Auctions can be used for efficient resource allocation within decentralized computational systems
  - Which do not necessarily consist of self-interested agents
  - They are frequently utilized for solving multi-agent and multi-robot coordination problems
  - For example, team-based exploration of unknown terrain

Introduction II

- An auction takes place between an agent known as the *auctioneer* and a collection of agents known as the *bidders*
  - The goal of the auction is for the auctioneer to allocate the *good* to one of the bidders
  - The auctioneer desires to maximize the price and bidders desire to minimize the price

- Dominant strategy: A strategy for bidding that leads in the long-term to a maximal payoff
- Payoff: *valuation - bid*
- Valuation: The money you are willing to spent

Mechanism Design

- *Mechanism design* is the design of protocols (e.g. auctions) for yielding multi-agent interactions with desirable properties, such as:
  - Guaranteed success: Agreement is certain
  - Maximizing social welfare: Agreement maximizes sum of utilities of all participating agents
  - Pareto efficiency: there is no other outcome that will make at least one agent better off without making at least one other agent worse off
  - Individual Rationality/Stability: Following the protocol is in best interest of all agents (no incentive to cheat, deviate from protocol etc.)
  - Simplicity: Protocol makes for the agent appropriate strategy „obvious“. (Agent can tractably determine optimal strategy)
  - Distribution: no single point of failure; minimize communication

Auction Parameters I

- Good/Item valuation
  - Private value: good has different value for each agent, e.g., grandpa's socks
  - Public (common) value: good has the same value for all bidders, e.g., one-dollar-Bill
  - Correlated value: value of good depends on own private value and private value for other agents, e.g., buy something with intention to sell it later
- Payment determination
  - First price: Winner pays his bid
  - Second price: Winner pays second-highest bid
- Secrecy of bids
  - Open cry: All agent's know all agent's bids
  - Sealed bid: No agent knows other agent's bids

Auction Parameters II

- Auction procedure
  - One shot: Only one bidding round
  - Ascending: Auctioneer begins at minimum price, bidders increase bids
  - Descending: Auctioneer begins at price over value of good and lowers the price at each round
  - Continuous: Internet
- Auctions may be
  - Standard Auction
    - One seller and multiple buyers
  - Reverse Auction
    - One buyer and multiple sellers
  - Double Auction
    - Multiple sellers and multiple buyers
- Combinatorial Auctions
  - Buyers and sellers may have combinatorial valuations for bundles of goods

English Auction

- English auctions are examples of *first-price open-cry ascending* auctions
- Protocol:
  - Auctioneer starts by offering the good at a low price
  - Auctioneer offers higher prices until no agent is willing to pay the proposed level
  - The good is allocated to the agent that made the highest offer

- Properties
  - Generates competition between bidders (generates revenue for the seller when bidders are uncertain of their valuation)
  - Dominant strategy: Bid slightly more than current bit, withdraw if bid reaches personal valuation of good
  - Winner's curse

Auction at Sotheby´s

The Winner's curse

- Termed in the 1950s:
  - Oil companies bid for drilling rights in the Gulf of Mexico
  - Problem was the bidding process given the uncertainties in estimating the potential value of an offshore oil field
  - "Competitive bidding in high risk situations," by Capen, Clapp and Campbell, *Journal of Petroleum Technology*, 1971
- For example
  - An oil field had an actual intrinsic value of $10 million
  - Oil companies might guess its value to be anywhere from $5 million to $20 million
  - The company who wrongly estimated at $20 million and placed a bid at that level would win the auction, and later find that it was not worth that much
- In many cases the winner is the person who has overestimated the most → "The Winner's curse"

Dutch Auction

- Dutch auctions are examples of first-price open-cry descending auctions
- Protocol:
  - Auctioneer starts by offering the good at artificially high value
  - Auctioneer lowers offer price until some agent makes a bid equal to the current offer price
  - The good is then allocated to the agent that made the offer

- Properties
  - Items are sold rapidly (can sell many lots within a single day)
  - Intuitive strategy: wait for a little bit after your true valuation has been called and hope no one else gets in there before you (no general dominant strategy)
  - Winner's curse also possible

Flower auction in Amsterdam

First-Price Sealed-Bid Auctions

- First-price sealed-bid auctions are *one-shot auctions*:
- Protocol:
  - Within a single round bidders submit a sealed bid for the good
  - The good is allocated to the agent that made highest bid
  - Winner pays the price of highest bid
- Often used in commercial auctions, e.g., public building contracts etc.
- Problem: the difference between the highest and second highest bid is “wasted money” (the winner could have offered less)
- Intuitive strategy: bid a little bit less than your true valuation (no general dominant strategy)
  - As more bidders as smaller the deviation should be!

Vickrey Auctions

- Proposed by William Vickrey in 1961 (Nobel Prize in Economic Sciences in 1996)
- Vickrey auctions are examples of second-price sealed-bid one-shot auctions
- Protocol:
  - within a single round bidders submit a sealed bid for the good
  - good is allocated to agent that made highest bid
  - winner pays price of second highest bid
- Dominant strategy: bid your true valuation
  - if you bid more, you risk to pay too much (winner's curse)
  - if you bid less, you lower your chances of winning while still having to pay the same price in case you win
- Antisocial behavior: bid more than your true valuation to make opponents suffer (not "rational")
- For private value auctions, strategically equivalent to the English auction mechanism

Expected Revenue

- Auctioneers want to maximize their revenue
  - Which auction protocol yields the highest possible price for them?
- *Risk-neutral bidders:*
  - The expected revenue to the auctioneer is provably identical in all four types of auctions (Sandholm 1999)
- *Risk-averse bidders* (i.e. bidders that would prefer to get the good even if they pay slightly more for it than their private valuation):
  - Dutch and first-price sealed-bid protocols lead to higher expected revenue for the auctioneer
  - Risk-averse agents can 'insure‘ themselves by bidding slightly more than risk-neutral bidders
- Risk-averse auctioneers do better with Vickrey or English auctions

Collusion and Lying

- Collusion (groups of bidders cooperate in order to cheat):
  - All four protocols are not collusion free
  - Bidders can agree beforehand to bid much lower than the public value
    - When the good is obtained, the bidders can then obtain its true value (higher than the artificially low price paid for it), and split the profits amongst themselves
    - Can be prevented by modifying the protocol so that bidders cannot identify each other
- Lying auctioneer:
  - Place bogus bidders (shills) that artificially increase the price
  - In Vickrey auction: Lying about second highest bid
  - Can be prevented by 'signing' of bids (e.g. digital signature), or trusted third party to handle bids
  - Not possible in English auctions!

Generalized first price auctions

Used by Yahoo for “sponsored links” auctions

- Introduced in 1997 for selling Internet advertising by Yahoo/Overture (before there were only “banner ads”)
- Advertisers submit a bid reporting the willingness to pay on a per-click basis for a particular keyword
  - Cost-Per-Click (CPC) bid
- Advertisers were billed for each “click” on sponsored links leading to their page
- The links were arranged in descending order of bids, making highest bids the most prominent
- Auctions take place during each search!
- However, auction mechanism turned out to be unstable!
  - Bidders revised their bids as often as possible

Generalized first price auctions II
Example

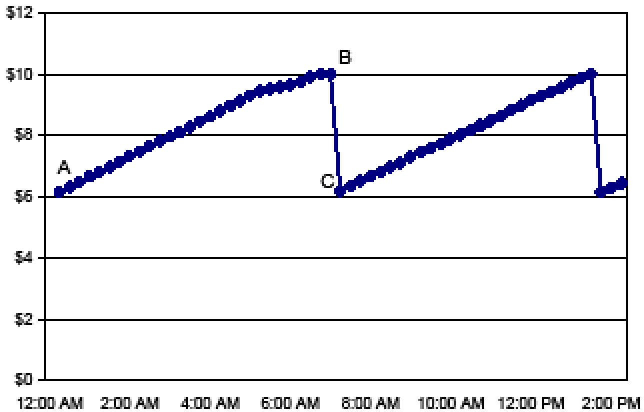

Top bids, in dollars, for a specific keyword (July 2002)

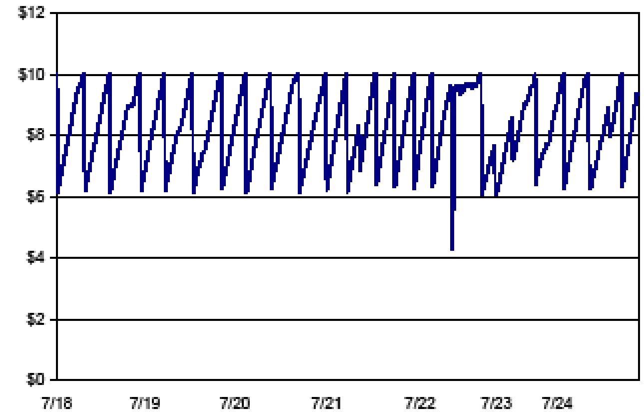

Continuation of this pattern for the same keyword for one week

1. Two advertiser agents (a1 & a2) compete for the top link position
2. Bidding starts with both of them below their maximum bids (A)
3. a1 recognizes an opportunity to win by raising the second bidder's bid by $0.01
4. a2 sees that it has been outbid, and raises its bid in turn
5. This process continues until the bids reach a1's maximum bid (B)
6. a1 can no longer increase, so it instead looks to avoid overspending by lowering its bid to $0.01 more than the third-place bidder (C)
7. a2 sees that it can still obtain the first place by bidding $0.01 more than a1's newly-lowered bid.
8. Bidding therefore begins to increase again …

Generalized second price auctions I

Used by Google for “sponsored link” auctions

- Introduced by Google for pricing sponsored links (AdWords Select)
- Observation: Buyers generally do not want to pay much more than the rank below them
  - Therefore: 2nd price auction
- Further modifications:
  - Advertisers bid for keywords and keyword combinations
  - Price consists of bid and quality score, e.g., rank = CPC_BID X quality score
- After seeing Google’s success, Yahoo also switched to second price auctions in 2002

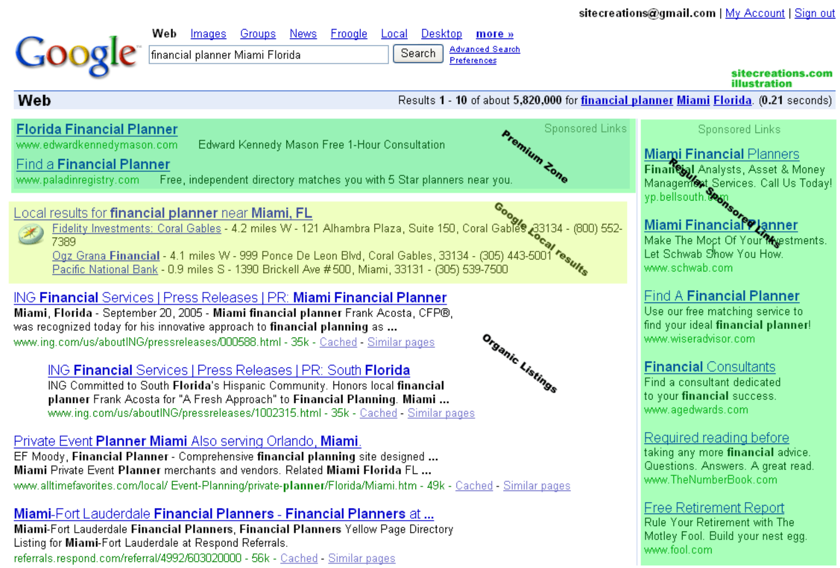

Advertiser	CPC Bid	Quality Score	Rank #	Position	CPC
A	$0.40	18	$0.40 x 18 = $\boxed{7.2}$	1	$0.37
B	$0.65	10	$0.65 x 10 = $\boxed{6.5}$	2	$0.39
C	$0.25	15	$0.25 x 15 = $\boxed{3.8}$	3	$0.10

Generalized second price auctions II

- Truthful bidding is not necessarily a dominant strategy if there is more than 1 slot!
- Payoff: The difference between the estimated value (valuation) of an object an the paid amount
- Example (without quality score):

	Valuation		Click-through rate
Bidder A	***7$***	***Slot 1***	***10***
Bidder B	***6$***	***Slot 2***	***4***
Bidder C	***1$***	***Slot 3***	***0***

Bidding of true valuation: A gets Slot 1 and payoff 7$*10 – 6$*10 = 10$

Lying, e.g. A bids '4': A gets Slot 2 and payoff 7$*4 – 1$*4 = 24$ > 10$

Better solution: Vickrey-Clarcke-Groves (VCG) auction → *see exercises*

Combinatorial Auctions
Introduction

- In a combinatorial auction, the auctioneer puts several goods on sale and the other agents submit bids for entire bundles of goods
- Given a set of bids, the winner determination problem is the problem of deciding which of the bids to accept
  - The solution must be feasible (no good may be allocated to more than one agent)
  - Ideally, it should also be optimal (in the sense of maximizing revenue for the auctioneer)
  - A challenging algorithmic problem

Complements and Substitutes

- The value an agent assigns to a bundle of goods may depend on the combination
  - Complements: The value assigned to a set is *greater* than the sum of the values assigns to its elements
    - *Example:* „a pair of shoes” (left shoe and a right shoe)
  - Substitutes: The value assigned to a set is *lower* than the sum of the values assigned to its elements
    - *Example:* a ticket to the theatre and another one to a football match for the same night
- In such cases an auction mechanism allocating one item at a time is problematic since the best bidding strategy in one auction may depend on the outcome of other auctions

Combinatorial Auctions
Protocol

- One auctioneer, several bidders, and many items to be sold
- Each bidder submits a number of package bids specifying the valuation (price) the bidder is prepared to pay for a particular bundle
- The auctioneer announces a number of winning bids
- The winning bids determine which bidder obtains which item, and how much each bidder has to pay
  - No item may be allocated to more than one bidder
- Examples of package bids:
  - Agent 1: ({a, b}, 5), ({b, c}, 7), ({c, d}, 6)
  - Agent 2: ({a, d}, 7), ({a, c, d}, 8)
  - Agent 3: ({b}, 5), ({a, b, c, d}, 12)
- Generally, there are $2^n - 1$ non-empty bundles for n items, how to compute the optimal solution?

Optimal Winner Determination Algorithms

- An auctioneer has a set of items $M = \{1,2,\ldots,m\}$ to sell
- Buyers submit a set of package bids $\mathbf{B} = \{B_1,B_2,\ldots,B_n\}$
- A package bid is a tuple $B_j = \langle S_j,p_j \rangle$, where $S_j \subseteq M$ is a set of items and $p_j > 0$ is a price
- $x_j \in \{0, 1\}$ is a decision variable for each bid B_j
- The *winner determination problem (WDP)* is to label the bids as winning or losing so as to maximize the sum of the accepted bid prices:

$$\max \sum_{j=1}^{n} p_j x_j \quad \text{s.t.} \quad \sum_{j|i \in S_j} x_j \leq 1, \quad \forall i \in \{1..m\}$$

$$x_j \in \{0,1\}$$

- This problem is computationally complex (*NP*-complete)
  - However, solvable for some problems with mixed integer program solvers, e.g. CPLEX and XPress-MP
  - … or by heuristic search

Solving WDPs by Heuristic Search I

- Two ways of representing the state space
  - Branch-on-items:
    - A state is a set of items for which an allocation decision has already been made
    - Branching is carried out by adding a further item
  - Branch-on-bids:
    - A state is a set of bids for which an acceptance decision has already been made
    - Branching is carried out by adding a further bid

Solving WDPs by Heuristic Search II
Branch-on-items

- Branching based on the question: “What bid should this item be assigned to?”
- Each path in the search tree consists of a sequence of disjoint bids
  - Bids that do not share items with each other
  - A path ends when no bid can be added to it
- Costs at each node are the sum of the prices of the bids accepted on the path

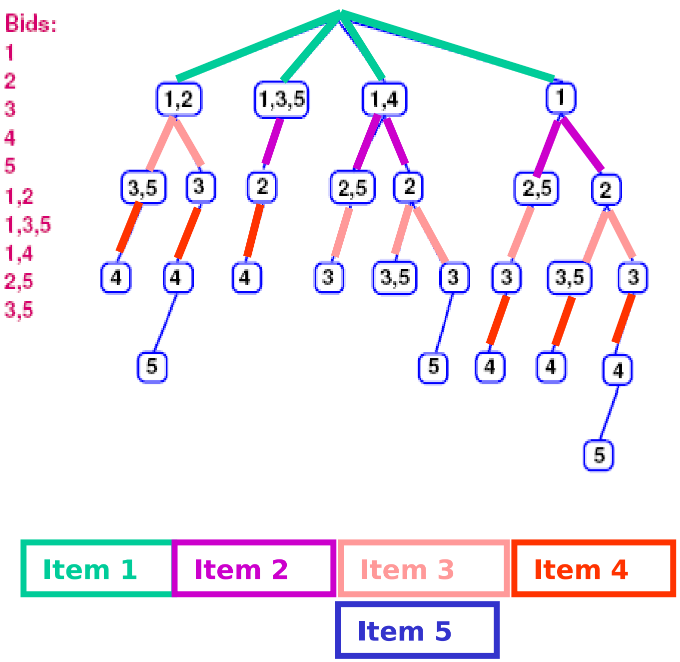

Solving WDPs by Heuristic Search III

Problem with branch-on-items

- What if the auctioneer's revenue can increase by keeping items?
- Example: Consider an auction of items 1 and 2
  - There is no bid for 1,
  - a $5 bid for 2,
  - and a $3 bid for {1;2}
  - → it is better to keep 1 and sell 2 than it would be to sell both
- The auctioneer's possibility of keeping items can be implemented by placing *dummy bids* of price zero on those items that received no 1-item bids (Sandholm 2002)
- For example, the following tree might be suboptimal for particular pricings:

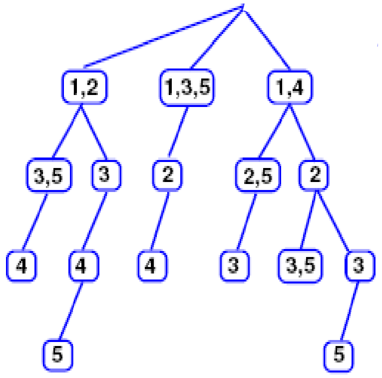

- Solution: Add dummy bid “1”

Solving WDPs by Heuristic Search IV
Branch-on-bids

- Branching is based on the question: "Should this bid be accepted or rejected?"
  → Binary tree
- When branching on a bid, the children in the search tree are the world where that bid is accepted (IN), and the world where that bid is rejected (OUT)
- No dummy bids are needed
- First a bid graph is constructed that represents all constraints between the bids
  - For example: Bids: {1,2};{2,3};{3};{1;3}

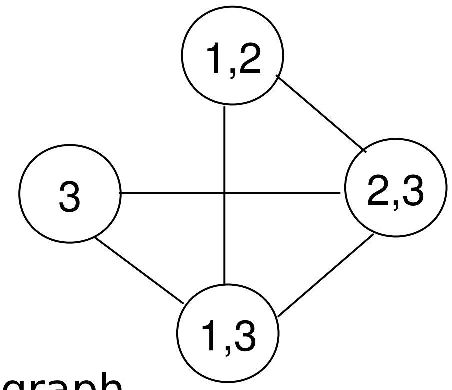

- Then, bids are accepted/rejected until all bids have been handled
  - On accept: remove all constrained bids from the graph
  - On reject: remove bid itself from the graph

Solving WDPs by Heuristic Search V

Branching on items vs. branching on bids

Bids in this example (only items of each bid are shown; prices are not shown):
{1,2}, {2,3}, {3}, {1,3}

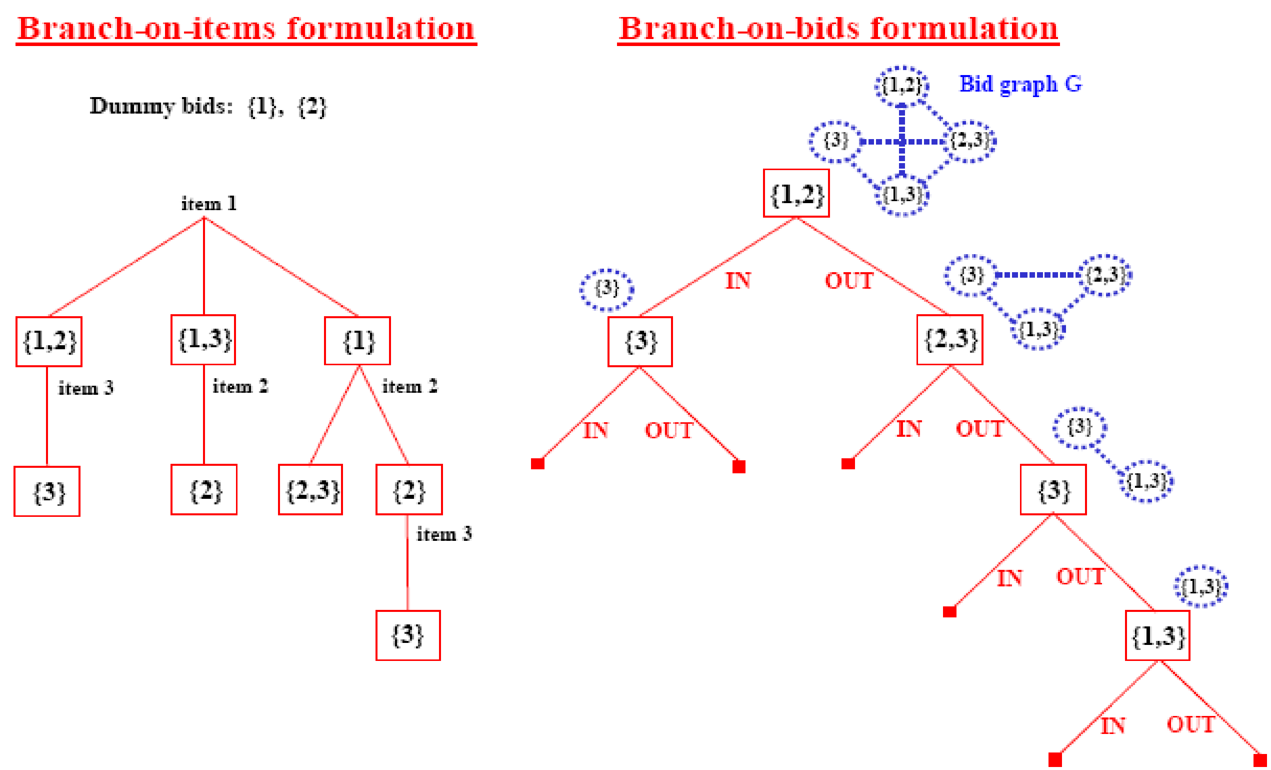

Source: Sandholm (2006)

Solving WDPs by Heuristic Search VI

Heuristic Function

- For any node N in the search tree, let g(N) be the revenue generated bids accepted according to N
- The heuristic function h(N) estimates for every node N how much additional revenue can be expected ongoing from N
- An upper bound on h(N) is given by the sum over the maximum contribution of the set of unallocated items *A*:

$$\sum_{i \in A} c(i), \qquad \text{where } c(i) = \max_{j \mid i \in S_j} \frac{p_j}{|S_j|}$$

- … and $B_j = \langle S_j, p_j \rangle \in \mathbf{B}$
- Tighter bounds can be obtained by solving the linear program relaxation of the remaining items (Sandholm 2006)

Auctions for multi-robot exploration I

Introduction

- Consider a team of mobile robots that has to visit a number of given targets (locations) in initially partially unknown terrain
- Examples of such tasks are cleaning missions, space-exploration, surveillance, and search and rescue
- Continuous re-allocation of targets to robots is necessary
  - For example, robots might discover that they are separated by a blockage from their target
- To allocate and re-allocate the targets among themselves, the robots can use auctions where they sell and buy targets
- Team objective is to minimize the sum of all path costs, hence, bidding prices are estimated travel costs
- The path cost of a robot is the sum of the edge costs along its path, from its current location to the last target that it visits

Auctions for multi-robot exploration II

Example

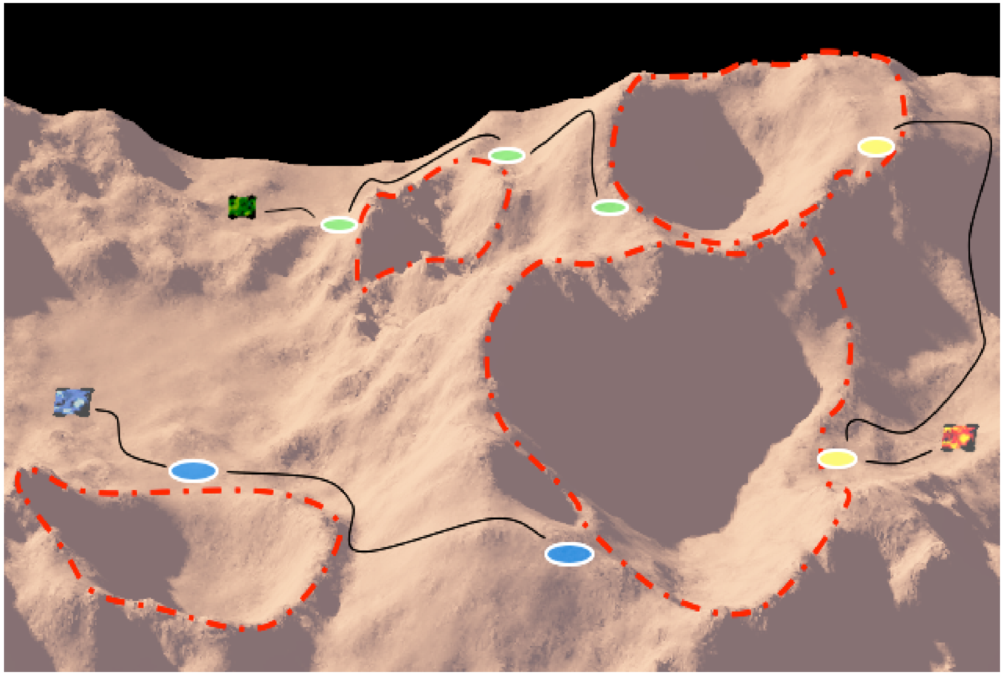

Three robots exploring Mars. The robots' task is to gather data around the four craters, e.g. to visit the highlighted target sites. Source: N. Kalra

Auctions for multi-robot exploration III

General Protocol

- Robot always follow a minimum cost path that visits all allocated targets
- Whenever a robot gains more information about the terrain, it shares this information with the other robots
- If the remaining path of at least one robot is blocked, then all robots put their unvisited targets up for auction
- The auction(s) close after a predetermined amount of time
  - Constraints: each robot wins at most one bundle and each target is contained in exactly one bundle
- After each auction, robots gained new targets or exchanged targets with other robots
- Then, the cycle repeats

Auctions for multi-robot exploration IV

Single-Round Combinatorial Auction

- Protocol:
  - Every robot bids on bundles of targets
  - The valuation is the estimated smallest path cost needed to visit all targets
  - A central auctioneer determines and informs the winning robots
- Optimal team performance:
  - Combinatorial auctions take all positive and negative synergies between targets into account
  - Minimization of the total path costs
- Three problems:
  - Robots cannot bid on all possible bundles of targets because the number of possible bundles is exponential in the number of targets
  - To calculate costs for each bundle requires to calculate the smallest path cost for visiting a set of targets (Traveling Salesman Problem)
  - Winner determination is NP-hard

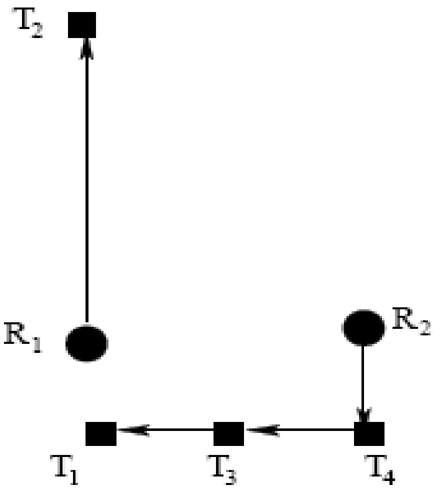

Source: M. Gini

Auctions for multi-robot exploration V

Parallel Single-Item Auctions

- Protocol:
  - Every robot bids on each target in parallel
  - The valuation is the smallest path cost needed to visit the target
  - The robot that that currently owns a target determines and informs the winning robot (the robot with the smallest bid) for the target
- Advantage:
  - Simple to implement and computation and communication efficient
- Disadvantage:
  - The team performance can be highly suboptimal since it does not take any synergies between targets into account

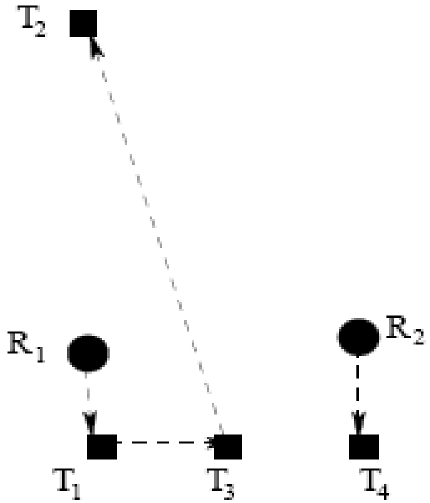

Source: M. Gini

Auctions for multi-robot exploration VI

Sequential Single-Item Auctions

- Protocol:
  - All targets are initially unallocated
  - Every robot bids on each unallocated target
  - The valuation is the increase in its smallest path cost that results from winning the auctioned target
  - The robot with the overall smallest bid is allocated the corresponding target
  - Each robot re-bids on unallocated targets, and the cycle repeats until all targets are owned by robots
  - Each robot then calculates the minimum-cost path for visiting all of its targets and moves along this path
- Advantages:
  - Hill climbing search: some synergies between targets are taken into account (but not all of them)
  - Simple to implement and computation and communication efficient
  - Since robots can determine the winners by listening to the bids (and identifying the smallest bid) the method can be executed decentralized

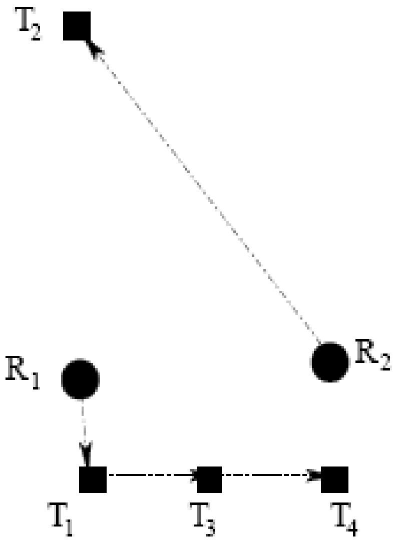

Source: M. Gini

Auctions for multi-robot exploration VII

Robot team exploration video

Two 2 E-Gators's given a mission with four named areas of interest in the Schenley Park
Source: R. Zlot

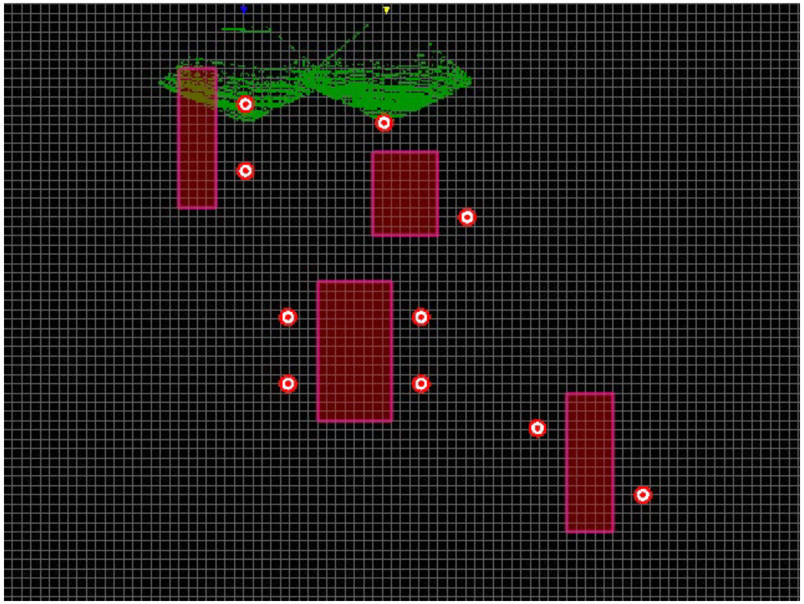

Maps built by the robots using their laser scanners (black areas are *unknown*, dark green areas are *free space*, and bright green areas are *obstacles*) Source: R. Zlot

Summary

- English, Dutch, First-Price Sealed-Bid, an Vickrey auctions are actively used for different types of situations
  - The expected revenue to the auctioneer is provably identical in all four types of auctions in case of r*isk-neutral bidders*

- Generalized second price auctions have shown good properties in practice, however, “truth telling” is not a dominant strategy

- Combinatorial auctions are a mechanism to allocate a number of goods to a number of agents
  - The WDP can be tackled using both integer programming and heuristic search
  - For real-time applications, such as robot exploration, single-item-auctions are the better choice

Literature

- General:
  - M. Woolridge: **An Introduction to Multi-Agent-Systems**, *Wiley*, 2001, 294 pages
- Sponsored Link Auctions:
  - B. Edelman, M. Ostrovsky, M. Schwarz Selling **Internet Advertising and the Generalized Second Price Auction: Billions of Dollars Worth of Keywords**, 2005
    - Link: http://rwj.berkeley.edu/schwarz/publications/gsp051003.pdf
- Winner Determination:
  - **Optimal Winner Determination Algorithms** (Tuomas Sandholm)
    - Link: http://www.cs.cmu.edu/~sandholm/windetalgs.pdf
  - T.W. Sandholm. **Distributed Rational Decision Making**. In G. Weiss (ed.), Multiagent Systems, MIT Press, 1999
  - P. Cramton, Y. Shoham, and R. Steinberg (eds.). **Combinatorial Auctions**. MIT Press, 2006.
- Multi-Robot exploration auctions:
  - Dias, M. B. and Stentz, A. 2001. **A Market Approach to Multirobot Coordination**. Technical Report, CMU-RI-TR-01-26, Robotics Institute, Carnegie Mellon University.
  - Zlot, R. et al. 2002. **Multi-Robot Exploration Controlled by a Market Economy**. IEE
  - S. Koenig, C. Tovey, M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, A. Kleywegt, A. Meyerson and S. Jain. **The Power of Sequential Single-Item Auctions for Agent Coordination** [Nectar Paper]. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI),* 1625-1629, 2006