Knowledge Engineering

Semester 2, 2004-05

Michael Rovatsos
mrovatso@inf.ed.ac.uk

Lecture 13 – Distributed Rational Decision-Making
25th February 2005

Where are we?

Last time . . .

- Agent interaction & communication
- Speech act theory
- Interaction Protocols
- But how should agents behave in interaction situations?

Today . . .

- Distributed Rational Decision-Making

Basic Considerations

- In entirely cooperative systems, we can impose constrains on agent behaviour to achieve global system objective
- In open systems, this is impossible!
  - We do not own all the agents in the system
  - We don't know anything about their internal design
  - Ultimately, they might be malicious
- But there is (some) hope . . .
  if we assume agents to be rational
- In this case, they can be considered "selfish", rather than "malevolent" or "randomly behaving"
- Question: How can we design interaction mechanisms that achieve some global objective *despite* agents being selfish?

Decision Theory

- A theory of (single-agent) rational decision making
- Based on a set of alternatives, what is the optimal decision an agent may make?
- Informally speaking, this depends on how desirable an alternative see and how likely we think it is
  - decision theory = utility theory + probability theory
- Let $O = \{o_1, \ldots o_n\}$ a set of possible outcomes (e.g. possible "runs" of the system until final states are reached)
- A **preference ordering** $\succ_i \subseteq O \times O$ for agent i is an antisymmetric, transitive relation on O, i.e.
  - $o \succ_i o' \Rightarrow o' \not\succ_i o$
  - $o \succ_i o' \wedge o' \succ o'' \Rightarrow o \succ_i o''$
- Such an ordering can be used to express strict preferences of an agent over O (write $\succeq_i$ if also reflexive, i.e. $o \succeq_i o$)

Decision Theory

- Preferences are often expressed through a **utility function** $u_i : O \Rightarrow \mathbb{R}$:
$$u_i(o) > u_i(o') \Leftrightarrow o \succ o', \quad u_i(o) \geq u_i(o') \Leftrightarrow o \succeq o'$$
- Principle of **expected utility maximisation**:
$$a^* = \arg\max_{a \in A} \sum_{o \in O} P(o|a)u(o)$$
where $a \in A$ are the actions/decisions an agent may take
- Generally accepted criterion, but also problems:
  - Incomplete information (wrt outcomes, probabilities, preferences)
  - Risk aversion attitude (value of additional utility depending on current "wealth", e.g. money)
  - Quantification problem (optimal=maximising average utility?, comparability of different utility values)

Game Theory

- Application of decision-theoretic principles to interaction among several agents
- Basic model: agents perform simultaneous actions (potentially over several stages), the actual outcome depends on the combination of action chosen by all agents
- **Normal-form games**: final result reached in single step (in contrast to **extensive-form games**)
  - Agents $\{1, \ldots, n\}$, S_i=set of (pure) **strategies** for agent i, $S = \times_{i=1}^{n} S_i$ space of **joint strategies**
  - Utility functions $u_i : S \to \mathbb{R}$ map joint strategies to utilities
  - A probability distribution $\sigma_i : S_i \to [0, 1]$ is called a **mixed strategy** of agent i (can be extended to joint strategies)
- Game theory is concerned with the study of this kind of games (in particular developing solution concepts for games)

Dominance and Best Response Strategies

- Two simple and very common criteria for rational decision making in games
- Strategy $s \in S_i$ is said to **dominate** $s' \in S_i$ iff
$$\forall s_{-i} \in S_{-i} \quad u_i(s, s_{-i}) \geq u_i(s', s_{-i})$$
$(s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$, same abbrev. used for $S)$
- Dominated strategies can be safely deleted from the set of strategies, a rational agent will never play them
- Some games are solvable in **dominant strategy equilibrium**, i.e. all agents have a single (pure/mixed) strategy that dominates all other strategies

Dominance and Best Response Strategies

- Strategy $s \in S_i$ is a **best response** to strategies $s_{-i} \in S_{-i}$ iff
$$\forall s' \in S_i, s' \neq s \quad u_i(s, s_{-i}) \geq u_i(s', s_{-i})$$
- Weaker notion, only considers optimal reaction to a *specific* behaviour of other agents
- Unlike dominant strategies, best-response strategies (trivially) always exist
- Strict versions of the above relations require that ">" holds' for at least one s'
- Replace s_i/s_{-i} above by σ_i/σ_{-i} and you can extend the definitions for dominant/best-response strategies to mixed strategies

Nash Equilibrium

- Nash (1951) defined the most famous equilibrium concept for normal-form games
- A joint strategy $s \in S$ is said to be in (pure-strategy) **Nash equilibrium** (NE), iff

$$\forall i \in \{1, \ldots n\} \forall s_i' \in S_i \quad u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i})$$

- Intuitively, this means that no agent has an incentive to deviate from this strategy combination
- Very appealing notion, because it can be shown that a (mixed-strategy) NE always exists
- But also some problems:
  - Not always unique, how to agree on one of them?
  - Proof of existence does not provide method to actually find it
  - Many games do not have pure-strategy NE

Example

Two men are collectively charged with a crime and held in separate cells, with no way of meeting or communicating. They are told that:

- if one confesses and the other does not, the confessor will be freed, and the other will be jailed for three years;
- if both confess, then each will be jailed for two years.

Both prisoners know that if neither confesses, then they will each be jailed for one year.

Example

The Prisoner's Dilemma: Nash equilibrium is not Pareto efficient (or: no one will dare to cooperate although mutual cooperation is preferred over mutual defection)

2 / 1	C	D
C	(3,3)	(0,5)
D	(5,0)	(1,1)

Problem: $DC \succ CC \succ DD \succ CD$ (from first player's point of view) and $u(CC) > \frac{u(DC)+u(CD)}{2}$

The Evolution of Cooperation?

- Typical **non-zero sum game**: there is a potential for cooperation but how should it emerge among self-interested agents?
- This situation occurs in many real life cases:
  - Nuclear arms race
  - Tragedy of the commons
  - "Free rider" problems
- In (infinitely) iterated case, cooperation is the rational choice in the PD (but "backward induction" problem)
- Axelrod's tournament (1984): Iterated Prisoner's Dilemma with lots of strategies (how to play against different opponents?)
- TIT FOR TAT strategy (don't be envious, be nice, retaliate appropriately, don't hold grudges) very successful

Example

The Coordination Game: No temptation to defect, buy two equilibria (hard to know which one will be chosen by other party)

2 / 1	A	B
A	(1,1)	(-1,-1)
B	(-1,-1)	(1,1)

Game Theory & Multiagent Systems

- Game theory = foundation for **mechanism design**
- Design of negotiation protocols for automated negotiation (i.e. coordination in the presence of a conflict of interest)
- Find protocols that satisfy certain properties
- *Individual Rationality:* for all agents, the negotiated solution should offer at least as much utility as not participating in the protocol
  - Necessary precondition for any viable protocol
- *Social Welfare:* the sum of all agents' utilities under some solution
  - Somewhat arbitrary, inter-agent utilities might not be comparable
- *Pareto Efficiency*
  - No agent could be better off than in current solution without at least one other agent being worse off

Criteria for Negotiation Protocols

- Stability: motivation for agents to behave in the desired manner
  - Dominant strategy equilibrium: very stable but does not always exist
  - Nash equilibrium
    - Pure Nash equilibria do not exist in all games
    - There might be more than one. How to pick the right one?
    - Sometimes not Pareto efficient
    - Not stable against deviation of a group of agents in coordinated manner
    - Doesn't necessarily hold in later stages of a sequential game
  - Computational efficiency
  - Distribution, communication efficiency

Revelation Principle

- An example of the kind thing that can be proven using game theory
- Let $\Theta = \{\theta_1, \ldots, \theta_n\}$ "types" of agents i that totally determine their preferences, $f : \Theta \rightarrow O$ a **social choice** function that calculates social outcome given agent types
- Problem: agents might not reveal their types truthfully
- A protocol **implements** f if the protocol has an equilibrium (dominant strategy/Nash) whose outcome is the same as that of f if agents revealed types truthfully
- **Revelation principle**:

  Suppose protocol p implements f in Nash/DS equilibrium. Then f is implementable in Nash/DS equilibrium via a single-step protocol where agents reveal their entire types truthfully.

Revelation Principle

- Proof idea:
  - add additional step to p in which agents' potentially insincere strategies are computed automatically
  - simulate original protocol after this step
  - ⇒ motivation for agents to reveal their true type in single step (protocol lies optimally on agents' behalf)
- Significance: enables us to restrict search for desirable protocol to ones where truthful revelation occurs in one step
- However, only existence result
  - What if there are other equilibria?
  - What if "lying" step is hard to compute?
  - What if agents don't play equilibrium strategies?

Electronic Auctions

- Auctions = preference-based method for allocating goods
- Most common types of auctions:
  - English (first-price open-cry)
  - Dutch (reverse)
  - First-price sealed bid
  - Vickrey auction (second-price sealed bid)
- Additional variations depending on following characteristics:
  - private-value vs. public-value (also: correlated value)
  - risk-neutral, risk-seeking, risk-averse bidders/auctioneer
- Some interesting issues/problems:
  - Lying bidders
  - Lying auctioneer
  - Bidder collusion
  - Incentive for speculation

The English Auction (EA)

- Each bidder raises freely his bid (in public), auction ends if no bidder is willing to raise his bid anymore
- Bidding process public ⇒ in correlated auctions, it can be worthwhile to counter-speculate
- In correlated auctions, often auctioneer increases price at constant/appropriate rate, also use of reservation prices
- Dominant strategy in private-value EA: bid a small amount above one's own valuation

The English Auction (EA)

- Advantages:
  - Truthful bidding is individually rational & stable
  - No lying auctioneer
- Disadvantages:
  - Can take long to terminate in correlated/common value auctions
  - Information is given away by bidding in public
  - Use of shills (in correlated-value EA) and "minimum price bids" possible
  - Bidder collusion self-enforcing (once agreement has been reached, it is safe to participate in a coalition) and identification of partners easily possible

Dutch/First-Price Sealed Bid Auctions

- Dutch (descending) auction: seller continuously lowers prices until one of the bidders accepts the price
- First-price sealed bid: bidders submit bids so that only auctioneer can see them, highest bid wins (only one round of bidding)
- DA/FPSB strategically equivalent (no information given away during auction, highest bid wins)
- Advantages:
  - Efficient in terms of real time (especially Dutch)
  - No information is given away during auction
  - Bidder collusion not self-enforcing, and bidders have to identify each other

Dutch/First-Price Sealed Bid Auctions

- Disadvantages:
  - No dominant strategy, individually optimal strategy depends on assumptions about others' valuations
  - Ideally bid less than own valuation but just enough to win
  - Incentive to counter-speculate
  - ⇒ no incentive to bid truthfully
  - This might incur loss of computational resources in the system
  - Lying auctioneer

The Vickrey Auction (VA)

- Second-price sealed bid: Highest bidder wins, but pays price of second-highest bid
- Advantages:
  - Truthful bidding is dominant strategy
  - No incentive for counter-speculation
  - Computational efficiency
- Disadvantages:
  - Bidder collusion self-enforcing
  - Lying auctioneer
- Unfortunately, VA is not very popular in real life
- But very successful in computational multiagent systems

Further Issues

- Pareto efficiency: all protocols alocate auction item to the bidder who values it most (in isolated private value/common value auctions)
  - But this result requires risk-neutrality if there is some uncertainty about own valuations
- Revenue equivalence in terms of expected revenue among all protocols if valuations independent, bidders risk-neutral and auction is private value
- **Winner's curse** in correlated/common value auctions
  - If I win, I always know I won't get to re-sell at the same price, because others value the goods less!

Further Issues

- Some properties of protocols change
  - if there is uncertainty about own valuations
  - if one can pay to obtain information about others' valuations
  - if we are looking at sequential (multiple) auctions
- Undesirable private information revelation
  - Example: truthful bidding in EA/VA may lead sub-contractors to re-negotiate rates after finding out that price was lower than they thought
- In terms of communication, auctions are not a very expressive method of negotiation!
  - Solely concerned with determining a selling price for some item

Other Methods

- Voting: determining an optimal "social choice" given individual preferences
- Bargaining: different set of possible agreements ("deals"), but conflict of interest regarding these
- Market Equilibrium Mechanisms: how to derive optimal production and consumption plans in a market
- Contract Nets: determining optimal task allocations among a set of agents
- Coalition Formation: how to find the best coalition structure in an agent society (if different coalitions can ensure different payoffs) and how to reward coalition participants

Critique

While game-theoretic/decision-theoretic approaches are currently very popular, there is also some criticism:

- How far can we get in terms of cooperation while assuming purely self-interested agents?
  - Good for economic interactions but how about other social processes?
  - In a sense, these approaces assume "worst case" of possible agent behaviour and disregard higher (more fragile) levels of cooperation
- Although mathematically rigorous,
  - ... the proofs only work under simplifying assumptions
  - ... often don't consider irrational behaviour
  - ... can only deal with a "utilitised" world
- Relationship to goal-directed, rational reasoning (e.g. BDI) and to deductive reasoning complex and not entirely clear

Summary

- Discussed rational decision-making mechanisms in societies of self-interested agents
- Idea of "mechanism design": design protocols that ensure global properties despite agents' self-interest under certain rationality assumptions
- Discussed foundations and fundamental problems of decision theory and game theory
- Looked at auctions as a particular method for automated negotiation
- Next time: **Semantic Web** (probably)