Multiagent Systems: Rational Decision Making and Negotiation

Ulle Endriss (ue@doc.ic.ac.uk)

Course website: http://www.doc.ic.ac.uk/~ue/mas-2005/

Introduction

- In multiagent systems (MAS), agents need to coordinate their actions, resolve conflicts, reach agreements ...
- Therefore, agents need to be able to *negotiate*.
- We will discuss different *protocols* for negotiation as well as *strategies* that agents may follow when using these protocols.
- Distinguish negotiation from communication: here we are not interested in the details of *how* agents manage to "talk" to each other, but rather *what* they talk about.

Aims and Objectives

Aims. To show how formal models for rational decision making and negotiation, developed mostly in the area of economics, have found important applications in multiagent systems.

Objectives. To give a brief introduction to both welfare economics and game theory, and to review several negotiation mechanisms.

- *Welfare Economics* (mathematical models of how the distribution of resources amongst agents affects social welfare)
- *Game Theory* (mathematical models of strategic behaviour in competitive interactions between rational agents)
- *Negotiation* (in particular one-to-one negotiation with the Monotonic Concession Protocol)
- *Auctions* (mechanisms for one-to-many negotiation)

Recommended Books

Most of the material presented in this part of the course is covered by chapters 6 and 7 of the following book:

- M. Wooldridge. *An Introduction to MultiAgent Systems*. John Wiley and Sons, 2002. See also http://www.csc.liv.ac.uk/~mjw/pubs/imas/.

Further reading:

- J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- T. Sandholm. *Distributed Rational Decision Making*. Chapter 5 in G. Weiß (editor), *Multiagent Systems*. MIT Press, 1999. Available at http://www.cs.cmu.edu/~sandholm/.

Welfare Economics

Rational Agents

- Before we can describe formal models of negotiation and interaction in multiagent systems, we require a suitable model that captures the relevant properties of an individual agent.
- We assume that agent are *rational:* their actions are directed towards *maximising their expected payoff.*
- In particular, we assume that agents are neither *altruistic* nor *malicious*.
- How can we model this concept of rationality? How can we formalise the notion of payoff?

Preferences over Alternative Agreements

- In general, agents negotiate in order to come to an *agreement* (an allocation of resources or tasks, a joint plan of action, a price or any other parameter of a commercial transaction, . . .)
- The *preference relation* of agent i over alternative agreements:

$$x \preceq_i y \quad \Leftrightarrow \quad \text{agreement } x \text{ is not better than } y \text{ (for agent } i\text{)}$$

- A preference relation $\preceq_i$ is usually required to be
  - *transitive:* if you prefer x over y and y over z, you should also prefer x over z; and
  - *connected:* for any two agreements x and y, you can decide which one you prefer (or whether you value them equally).
- Discussion: useful model, but not without problems (humans cannot always assign rational preferences . . .)

Utility Functions

- A utility function u_i (for agent i) is a mapping from the space of agreements to the reals.
- Example: $u_i(x) = 10$ means that agent i assigns a value of 10 to agreement x.
- A utility function u_i *representing* the preference relation $\preceq_i$:

$$x \preceq_i y \quad \Leftrightarrow \quad u_i(x) \leq u_i(y)$$

- Preferences are *qualitative*; utility functions are *quantitative*.
- Discussion: utility functions are very useful, but they suffer from the same problems as preference relations — even more so (humans typically do not reason with numerical utilities . . .)

Welfare Economics

- Welfare Economics is the branch of Economic Sciences that studies how the welfare distribution amongst the members of a society affects society as a whole.
- Multiagent systems are often described as *societies of agents*.
- The utility $u_i(x)$ assigned to agreement x by agent i may be interpreted as the level of "welfare" experienced by i.
- How does the welfare of individual agents affect the welfare of society as a whole?

Utilitarian Social Welfare

The *social welfare* associated with agreement x is defined as follows:

$$sw(x) \quad = \quad \sum_{i \in Agents} u_i(x)$$

This is the so-called *utilitarian* definition of social welfare, which is measuring the "sum of all pleasures" (Jeremy Bentham, ∼1820).

Observe that maximising this function amounts to maximising the *average utility* enjoyed by agents in the system.

Egalitarian Social Welfare

- The function sw is usually regarded as the most important *collective utility function* for MAS, but there are also others.
- The *egalitarian* collective utility function sw_e, for instance, measures social welfare as follows:

  $$sw_e(x) \quad = \quad min\{u_i(x) \mid i \in Agents\}$$

  Maximising this function amounts to improving the situation of the weakest members of society.
- The egalitarian variant of welfare economics has been developed, amongst others, by Amartya Sen since the 1970s (Nobel Prize in Economic Sciences in 1998).
- What interpretation of the term *social welfare* is appropriate depends on the application.

Nash Product

- The *Nash collective utility function* sw_n is defined as the product of individual utilities:

  $$sw_n(x) \quad = \quad \prod_{i \in Agents} u_i(x)$$

  This is a useful measure of social welfare as long as all utility functions can be assumed to be non-negative.
- Named after John F. Nash (Nobel Prize in Economic Sciences in 1994; Academy Award in 2001).
- Like the utilitarian collective utility function, the Nash product favours increases in overall utility, but also inequality-reducing redistributions of welfare ($2 \cdot 6 < 4 \cdot 4$).

Pareto Optimality

- An agreement x is called *Pareto optimal* iff there is no other agreement y that would be better for at least one agent without being worse for any of the others.
- Pareto optimal outcomes of negotiation are generally accepted to be desirable.
- Example: three agents and three possible agreements . . .

	Mallorca	New York	Cornwall
Peter	5	8	-3
Paul	0	2	25
Mary	3	7	2

Going to Mallorca is the only agreement that is *not* Pareto optimal. Cornwall gives maximal utilitarian social welfare. New York maximises egalitarian social welfare.

Game Theory

Game Theory

- Game Theory is the branch of Economic Sciences that studies the strategic behaviour of rational agents in the context of interactive decision-making problems.
- Given the *rules of the "game"* (the *negotiation mechanism*, the *protocol*), what *strategy* should a *rational* agent adopt?

Prisoner's Dilemma

Two partners in crime, A and B, are separated by police and each one of them is offered the following deal:

- only you confess $\Rightarrow$ free
- only the other one confesses $\Rightarrow$ 5 years in prison
- both confess $\Rightarrow$ 3 years in prison
- neither one confesses $\Rightarrow$ 1 year on remand

u_A/u_B	B confesses	B does not
A confesses	2/2	5/0
A does not	0/5	4/4

(utility $=$ 5 $-$ years in prison)

$\Rightarrow$ What would be a rational strategy?

Dominant Strategies

- A strategy is called *dominant* iff, independently of what any of the other agents do, following that strategy will result in a larger payoff than any other strategy.
- Prisoner's Dilemma: both agents have a dominant strategy, namely to confess:
  - from A's point of view:
    - if B confesses, then A is better off confessing as well
    - if B does not confess, then A is also better off confessing
  - similarly for B
- Terminology: for games of this kind, we say that each agent may either *cooperate* with its opponent (e.g. by not confessing) or *defect* (e.g. by confessing).

Nash Equilibria

- Introduced by John F. Nash in 1950.
- A *Nash equilibrium* is a set of strategies, one for each agent, such that no agent could improve its payoff by unilaterally deviating from their assigned strategy.
- In cases where there are no dominant strategies, a set of equilibrium strategies is the next best thing.
- Discussion: games with a Nash equilibrium are of great interest to MAS, because you do not need to keep your strategy secret and you do not need to waste resources on trying to find out about other agents' strategies.

Back to the Prisoner's Dilemma

- Unique Nash equilibrium: both agents confess:
  - if A changes strategy unilaterally, she will do worse
  - if B changes strategy unilaterally, she will also do worse
- Discussion: Our analysis shows that it would be *rational* to confess. However, this seems counter-intuitive, because both agents would be better off if both of them were to remain silent.
- Conflict: the *stable* solution given by the Nash equilibrium is not *efficient*, because the outcome is not Pareto optimal.
- Iterated Prisoner's Dilemma:
  - In each round, each agent can either cooperate or defect.
  - Because the other agent could retaliate in the next round, it is rational to cooperate.
  - But it does not work if the number of rounds is fixed ...

Game of Chicken

James and Marlon are driving their cars towards each other at top speed. Whoever swerves to the right first is a "chicken".

u_J/u_M	M drives on	M turns
J drives on	$0/0$	$8/1$
J turns	$1/8$	$5/5$

Analysing the Game of Chicken

- No dominant strategy (best move depends on the other agent)
- Two Nash equilibria:
  - James drives on and Marlon turns
    - if James deviates (and turns), he will be worse off
    - if Marlon deviates (and drives on), he will be worse off
  - Marlon drives on and James turns (similar argument)
- If you have reason to believe your opponent will turn, then you should drive on. If you have reason to believe your opponent will drive on, then you should turn.
- Mixed strategy: if you have no such information, then a rational strategy would be to play either one of the “pure” equilibrium strategies according to a suitable probability distribution ($\leadsto$ computing Nash equilibria).

Negotiation

Negotiation

- Negotiation is a central issue in multiagent systems: autonomous agents need to reach mutually beneficial agreements on just about everything . . .
- We can distinguish different types of negotiation:
  - *One-to-one* (or *bilateral*) negotiation
    Example: $\leadsto$ *Monotonic Concession Protocol*
  - *Many-to-one* negotiation
    Example: $\leadsto$ *Auctions*
  - *Many-to-many* (or *multilateral*) negotiation: difficult!

Mechanism Design

Some desirable properties of negotiation mechanisms:

- Rationality: it should be in the interest of individual agents to participate (no negative payoff)
- Efficiency: outcomes should be (at least) Pareto optimal
- Stability: agents should have no incentive to deviate from a particular desired strategy ($\leadsto$ Nash equilibrium)
- Fairness: no agent should have any *a priori* disadvantages
- Simplicity: the computational burden on each agent as well as the amount of communication required should be minimal
- Verifiability: it should be verifiable that agents follow the rules

General Setting for one-to-one Negotiation

- Two agents (agents 1 and 2) with utility functions u_1 and u_2
- Negotiation space: set of possible agreements x
- Protocol: the (public) "rules of encounter", specifying
  - what moves (e.g. proposals) are *legal* given a particular negotiation history;
  - when negotiation ends (with an agreement or in conflict);
  - and what the negotiated agreement is (if any).
- Strategy: private to each agent; specifies how an agent uses the protocol to get the best possible payoff (agreement) for herself

Monotonic Concession Protocol

- Example for a one-to-one negotiation protocol
- Assumption: the utility functions u_1 (of agent 1) and u_2 (of agent 2) are known to both agents
- The protocol proceeds in *rounds*; in each round both agents simultaneously make a proposal (by suggesting an agreement).
- First round: agents propose any agreements x_1 and x_2.
- Agreement is reached iff either $u_1(x_2) \geq u_1(x_1)$ or $u_2(x_1) \geq u_2(x_2)$; that is, if one agent proposes an agreement that is better for the other agent than its own proposal (in case both hold, flip a coin to decide outcome).
- In each round: either *concede* by making a proposal that is better for your opponent than your previous offer, or wait.
- *Conflict* arises when we get to a round where no one concedes.

Some Properties of the MCP

- Termination: guaranteed if the agreement space is finite
- Verifiability: easy to check that your opponent really concedes (only your own utility function matters)
- Criticism: you need to know your opponent's utility function to be able to concede (typical assumption in game theory; not always appropriate in MAS)

Strategies

- Question: What would be a good negotiation strategy when you use the MCP?
- The dangers of getting it wrong:
  - If you concede too often (or too much), then you risk not getting the best possible deal for yourself.
  - If you do not concede often enough, then you risk conflict (which is assumed to have utility 0).

Zeuthen Strategy

- Question: In each round, *who* should concede and *how much?*
- Idea: Evaluate agent i's *willingness to risk conflict*, given its own proposal x_i and its opponent's proposal x_j:

$$risk_i(x_i, x_j) \quad = \quad \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i) - u_i(conflict)} \quad = \quad \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i)}$$

  This is the ratio of the loss incurred by accepting x_j and the loss in case of conflict (both wrt. the utility of x_i).

- Strategy: start by proposing the best possible agreement; then
  - concede whenever your willingness to risk conflict is less or equal to your opponent's;
  - concede *just enough* to make your opponent's willingness to risk conflict less than yours.

Efficiency

If both agents use the Zeuthen Strategy, then the final agreement maximises the Nash product. This has first been observed by John C. Harsanyi in 1956 (Nobel Prize in Economic Sciences in 1994).

Proof sketch: Agent i concedes iff $risk_i(x_i, x_j) \leq risk_j(x_j, x_i)$, i.e. iff

$$\begin{aligned} \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i)} &\leq \frac{u_j(x_j) - u_j(x_i)}{u_j(x_j)} \\ \boxed{u_i(x_i) \cdot u_j(x_j)} - u_i(x_j) \cdot u_j(x_j) &\leq \boxed{u_j(x_j) \cdot u_i(x_i)} - u_j(x_i) \cdot u_i(x_i) \\ u_j(x_i) \cdot u_i(x_i) &\leq u_i(x_j) \cdot u_j(x_j) \end{aligned}$$

That is, agent i makes a (minimal) concession iff its current proposal does not yield the higher *product of utilities*.

Hence, the Zeuthen Strategy ensures a final agreement x that *maximises* this product. □

⇒ It follows that the final agreement will be Pareto optimal (why?).

Lack of Stability

Unfortunately, the mechanism where both agents use the Zeuthen Strategy is *not* stable. Agent 1 could exploit the following situation:

- Both current proposals maximise the product of utilities, i.e.:
  - we are one step away from an agreement; and
  - both agents have equal willingness to risk conflict.
- Then both agents *should* concede (in which case the protocol requires a coin to be flipped), although it is sufficient for one of them to concede to reach agreement.
- If agent 1 knows that agent 2 will play according to the Zeuthen Strategy, she could benefit from *defecting* (not conceding).

If both agents are prepared to exploit this weakness of the mechanism, they risk conflict ($\leadsto$ "Game of Chicken").

Extended Zeuthen Strategy

- Extended Zeuthen Strategy: play according to the Zeuthen Strategy and use the appropriate mixed equilibrium strategy in case the "last step situation" arises.
- Stability: the mechanism where both agents play according to the Extended Zeuthen Strategy is in Nash equilibrium (why?).
- Efficiency: in cases where no conflict arises, the extended strategy is still Pareto efficient.
- Discussion: lying about your own utility function may get the other agent to concede more often ...

A one-shot Negotiation Protocol

- Protocol: both agents suggest an agreement; the one giving a higher product of utilities wins (flip a coin in case of a tie)
- Obvious strategy: amongst the set of agreements with maximal product of utilities, propose the one that is best for you
- Properties: This mechanism is:
  - *efficient:* outcomes have maximal Nash product and are Pareto optimal (like MCP with Zeuthen Strategy)
  - *stable:* no agent has an incentive to deviate from the strategy (like MCP with extended Zeuthen Strategy)

  In addition, the one-shot protocol is also:
  - *simple:* only one round is required
- But why should anyone accept to use such a protocol?

Recap: How did we get to this point?

- Both agents making several small concessions until an agreement is reached is the most *intuitive* approach to one-to-one negotiation.
- The *Monotonic Concession Protocol* (MCP) is a straightforward formalisation of the above intuition.
- The extended *Zeuthen Strategy* is also motivated by intuition ("willingness to risk conflict") and constitutes a *stable* and (almost) *efficient* strategy for the MCP.
- The one-shot protocol (together with the obvious strategy) produces similar outcomes as MCP/Zeuthen, but it is a much *simpler* mechanism.

Auctions

Auctions

- Familiar from *Sotheby's* and *Bargain Hunt*.
- With the rise of the Internet, auctions have become popular in many *e*-commerce applications (e.g. *ebay*).
- In the context of MAS, auctions provide simple and implementable protocols for many-to-one negotiation.
- General setting for "simple" auctions:
  - one seller (the *auctioneer*)
  - many *buyers*
  - one single item to be sold, e.g.
    - ∗ a house to live in (*private value auction*)
    - ∗ a house that you may sell on (*correlated value auction*)
- There are many different *auction mechanisms* or *protocols*, even for simple auctions ...

English Auctions

- Protocol: auctioneer starts with the *reservation price*; in each round each agent can propose a higher bid; final bid wins
- Used to auction paintings, antiques, etc.
- Dominant strategy (for private value auctions): bid a little bit more in each round, until you win or reach your own valuation
- Counterspeculation (how do others value the good on auction?) is not necessary.
- Winner's curse (in correlated value auctions): if you win but have been uncertain about the true value of the good, should you be happy?

Dutch Auctions

- Protocol: the auctioneer starts at a very high price and lowers it a little bit in each round; the first bidder to accept wins
- Used at the flower wholesale markets in Amsterdam.
- 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)
- Also suffers from the winner's curse.

First-price Sealed-bid Auctions

- Protocol: one round; sealed bid; highest bid wins
(for simplicity, we assume no two agents make the same bid)
- Used for 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)
- Strategically equivalent to the Dutch auction protocol:
  - only the highest bid matters
  - no information gets revealed to other agents

Vickrey Auctions

- Proposed by William Vickrey in 1961 (Nobel Prize in Economic Sciences in 1996).
- Protocol: one round; sealed bid; highest bid wins, but the winner pays the price of the *second highest* bid
- Dominant strategy: bid your true valuation:
  - if you bid more, you risk to pay too much
  - if you bid less, you lower your chances of winning while still having to pay the same price in case you do win
- Problem: counterintuitive (problematic for humans)
- Antisocial behaviour: bid more than your true valuation to make opponents suffer (not "rational")
- For private value auctions, strategically equivalent to the English auction mechanism

Lying and Cheating

- Collusion (groups of bidders cooperate in order to cheat): none of the four auction protocols is collusion-proof.
- Lying auctioneer: problematic for Vickrey auctions, but not for any open-cry protocol or for first-price sealed-bid auctions.
- Shills: bidders placed by the auctioneer to artificially increase bids (English auction)

Pareto Efficiency

All four auction protocols guarantee a *Pareto optimal* outcome:

They result in an agreement x (the winner obtaining the good for the specified price from the auctioneer) such that there is no other agreement y that would be better for at least one of the agents without being worse for any of the others:

- paying a higher price would be worse for the winner
- paying a lower price would be worse for the auctioneer
- giving the good to a different buyer would be worse for the winner (who will pay a price less or at most equal to her private valuation, given agreement x)

Revenue for the Auctioneer

- Which protocol is best for the auctioneer?
- Revenue-equivalence Theorem (Vickrey, 1961):

  All four protocols give the same expected revenue for private value auctions where values are independently distributed.
- Intuition: revenue ≈ second highest valuation:
  - Vickrey: clear ✓
  - English: bidding stops just after second highest valuation ✓
  - Dutch/FPSB: because of the independent value distribution, top bid ≈ second highest valuation ✓
- If one bidder has a very high valuation, then Dutch and FPSB auctions are likely to be better for the auctioneer.
- Correlated value actions: English auctions are advantageous for the auctioneer.

Parameters of Simple Auction Protocols

To summarise, we can distinguish different types of auctions according to the following parameters:

- either *open-cry* or *sealed-bid*
- either *ascending* or *descending* or *one-shot*
- either *first-price* or *second-price*

We have seen the following examples:

- *English auctions:* first-price, open-cry, ascending
- *Dutch auctions:* first-price, open-cry, descending
- *First-price sealed-bid auctions:* first-price, sealed-bid, one-shot
- *Vickrey auctions:* second-price, sealed-bid, one-shot