

Multiagent Systems: Spring 2006

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Negotiation

- Negotiation is a central issue in MAS: autonomous agents need to reach mutually beneficial agreements on just about anything . . .
- We can distinguish different types of negotiation:
 - *Bilateral* (*one-to-one*) negotiation:
Two agents negotiate with each other (\rightsquigarrow today's lecture).
 - *Auctions* (*one-to-many* negotiation):
One agent (the auctioneer) negotiates with several other agents (the bidders).
 - *Distributed* and *multilateral* (*many-to-many*) negotiation:
Many agents are involved, and different groups of agents can (concurrently) come to (a sequence of) agreements.

Plan for Today

We shall mostly concentrate on a particular negotiation mechanism:

- the *Monotonic Concession Protocol* in combination with
- the *Zeuthen Strategy*

We shall be interested in the formal properties of this negotiation mechanism, in particular:

- *efficiency* and *stability*

Rosenschein and Zlotkin (1994) have coined the terms “Monotonic Concession Protocol” and “Zeuthen Strategy”, but the basic ideas of what we are going to discuss have been around since the 1950s.

J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.

Desiderata

Some desirable properties of negotiation mechanisms:

- *Rationality*: it should be in the interest of individual agents to participate (no negative payoff)
- *Stability*: agents should have no incentive to deviate from a particular desired strategy (\rightsquigarrow Nash equilibrium)
- *Efficiency*: outcomes should be (at least) Pareto optimal
- *Fairness*: outcomes should satisfy appropriate fairness conditions (equity, egalitarianism, envy-freeness, . . .)
- *Symmetry*: 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 Bilateral Negotiation

- Two agents (agents 1 and 2) with utility functions u_1 and u_2
- Negotiation space: set of possible *agreements*
- *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 themselves

A Natural Negotiation Protocol

An example for a bilateral negotiation protocol:

*Both agents start by proposing a deal of their choosing.
If no agreement is reached, each agent may either make a small concession or decide to stick to their proposal.
This continues until either an agreement is reached that is acceptable to both agents, or until both agents refuse to make a concession and negotiation breaks down*

This very natural form of negotiation has been formalised in the shape of the so-called *Monotonic Concession Protocol* . . .

Notation and Assumptions

- Set of two *agents*: $\mathcal{A} = \{1, 2\}$
- Finite set \mathcal{X} of *potential agreements* (*proposals, deals, . . .*)
- Each agent $i \in \mathcal{A}$ is equipped with a *utility function*: $u_i : \mathcal{X} \rightarrow \mathbb{R}_0^+$
 - Note: By restricting attention to agreements with non-negative utilities we ensure individual rationality *a priori*: no agent will have a negative payoff.
- The set \mathcal{X} includes a specific agreement, called the *conflict deal*, that yields utility 0 for both agents.
 - Note: The conflict deal will be chosen in case negotiation breaks down. This is the worst possible outcome.

Monotonic Concession Protocol (MCP)

- The protocol proceeds in *rounds*; in each round both agents make simultaneous *proposals* (by suggesting an agreement from \mathcal{X}).
- In the *first round* each agent is free to make any proposal.
- In *subsequent rounds*, each agent $i \in \mathcal{A}$ has got two options (let $x_i \in \mathcal{X}$ be the most recent proposal of i):
 - Make a *concession* and propose a new deal x'_i that is preferable to the other agent j : $u_j(x_i) < u_j(x'_i)$
 - Refuse to make a concession and stick to proposal x_i .
- *Agreement* is reached iff if one agent proposes an agreement that is at least as good for the other agent as their own proposal:

$$u_1(x_2) \geq u_1(x_1) \quad \text{or} \quad u_2(x_1) \geq u_2(x_2)$$

In case both conditions hold, flip a coin to decide the outcome.

- *Conflict* arises when we get to a round where nobody concedes. In this case the *conflict deal* will be the outcome of the negotiation.

Some Properties of the MCP

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

Strategies

- **Question:** What would be a good negotiation strategy to adopt when you are participating in a negotiation regulated by 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 :

$$Z_i = \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i) - u_i(\text{conflict})} = \frac{u_i(x_i) - u_i(x_j)}{u_i(x_i)}$$

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). [$Z_i = 1$ if $u_i(x_i) = 0$]

- **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.

F. Zeuthen. *Problems of Monopoly and Economic Warfare*. Routledge, 1930.

Example

[...]

Why the Zeuthen Strategy?

The Zeuthen Strategy does have some *intuitive* appeal ... but why this strategy and not some other intuitively appealing approach?

John C. Harsanyi (Nobel Prize in Economic Sciences in 1994) has demonstrated how the Zeuthen Strategy can be derived from a small number of fundamental *axioms* ...

J.C. Harsanyi. *Approaches to the Bargaining Problem before and after the Theory of Games*. *Econometrica*, 24(2):144–157, 1956.

Harsanyi's Axioms

- (1) *Symmetry*: The two agents follow identical strategies.
- (2) *Perfect information*: Each agent can correctly estimate the probability that the other will definitely reject a certain proposal.
- (3) *Monotonicity*: The probability of agent i refusing to concede is a monotonic non-decreasing function in $u_i(x_i) - u_i(x_j)$.
- (4) *Expected-utility maximisation*: Each agent will make a concession iff this will give them higher expected utility than not conceding.

Deriving the Zeuthen Strategy

Suppose agent 1's latest offer is x_1 and agent 2's latest offer is x_2 .

Let p_1 be the probability that agent 1 will eventually reject x_2 .

Let p_2 be the probability that agent 2 will eventually reject x_1 .

Compute the expected payoff for agent 1:

- The expected payoff for agent 1 of *rejecting* x_2 is $(1 - p_2) \cdot u_1(x_1)$.
- The *certain* payoff associated with *accepting* x_2 is $u_1(x_2)$.

Hence (by *expected-utility maximisation*), agent 1 should accept iff

$$u_1(x_2) > (1 - p_2) \cdot u_1(x_1)$$

This is equivalent to: agent 1 should accept ($p_1 = 0$) iff

$$Z_1 = \frac{u_1(x_1) - u_1(x_2)}{u_1(x_1)} < p_2$$

The same kind of analysis applies to agent 2 ...

Deriving the Zeuthen Strategy (cont.)

So far we know (*):

$$Z_1 < p_2 \text{ entails } p_1 = 0 \quad Z_2 < p_1 \text{ entails } p_2 = 0$$

$$Z_1 > p_2 \text{ entails } p_1 = 1 \quad Z_2 > p_1 \text{ entails } p_2 = 1$$

Hence, p_1 must be a function of p_2 and Z_1 ; and p_2 must be a function of p_1 and Z_2 . By *symmetry*, these must be the *same* function:

$$p_1 = F(p_2, Z_1) \quad \text{and} \quad p_2 = F(p_1, Z_2)$$

Hence, there is another function G such that:

$$p_1 = G(Z_1, Z_2) \quad \text{and} \quad p_2 = G(Z_2, Z_1)$$

Also, because of (*), one of the following three cases must apply (**):

$$p_1 = 0 \ \& \ p_2 = 1 \quad \text{or} \quad p_1 = 1 \ \& \ p_2 = 0 \quad \text{or} \quad p_1 = Z_2 \ \& \ p_2 = Z_1$$

Deriving the Zeuthen Strategy (cont.)

But the function G is (almost) uniquely determined by the axiom of *monotonicity* together with (**). We obtain:

- $p_1 = 0$ and $p_2 = 1$ (that is, 1 concedes) if $Z_1 < Z_2$
- $p_1 = 1$ and $p_2 = 0$ (that is, 2 concedes) if $Z_1 > Z_2$
- $p_1 = 0$ and $p_2 = 0$ (that is, both concede) if $Z_1 = Z_2$

Strictly speaking, the final case only follows together with a variant of the expected-utility maximisation axiom covering the case of simultaneous concessions. See Harsanyi (1956) for details.

Efficiency

Theorem 1 (Harsanyi, 1956) *If both agents use the Zeuthen Strategy, then the final agreement maximises the Nash product.*

Proof: According to the strategy, agent i concedes iff $Z_i \leq Z_j$, i.e. iff

$$\frac{u_i(x_i) - u_i(x_j)}{u_i(x_i)} \leq \frac{u_j(x_j) - u_j(x_i)}{u_j(x_j)}$$

$$\boxed{u_i(x_i) \cdot u_j(x_j)} - u_i(x_j) \cdot u_j(x_j) \leq \boxed{u_j(x_j) \cdot u_i(x_i)} - u_j(x_i) \cdot u_i(x_i)$$

$$u_j(x_i) \cdot u_i(x_i) \leq u_i(x_j) \cdot u_j(x_j)$$

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. \square

► 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, it could benefit from defecting (not conceding).

If both agents are prepared to exploit this weakness of the mechanism, they risk conflict (\rightsquigarrow “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.

Note: The mixed strategy can be computed using the method introduced last week; it is *not* always $(\frac{1}{2}, \frac{1}{2})$.

- *Stability*: the profile where both agents play according to the Extended Zeuthen Strategy is a mixed Nash equilibrium (why?).
- *Efficiency*: in cases where no conflict arises, the extended strategy still maximises the Nash product (and still is Pareto efficient).

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 bilateral 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”), further backed up by an axiomatic derivation (Harsanyi), 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 much *simpler* a mechanism.

Manipulating the Protocol

So it makes sense to assume that agents are committed to negotiating a *Nash-optimal solution* (an agreement that maximises the Nash CUF).

So far, we have assumed that agents have *perfect knowledge*: not only regarding each other’s utility function, but also regarding the range of potential agreements \mathcal{X} .


What happens when we drop this latter assumption? If agents negotiate over the reallocation of some *tasks*, for instance, lying about their own initial tasks will affect the set \mathcal{X} .

This kind of *manipulation* has been studied in detail by Rosenschein and Zlotkin (1994). Here we shall only go through some examples ...

J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.

Remember the Postmen Domain


Our agents are (two) postal workers. They meet in the post office in the morning and discuss the fact that Ann has letters for Carol, Dick and Ed, while Bob has letters for *Frank, Gary and Hugh* ...




Let the *utility* of an agreement to an agent be the distance saved with respect to the initial allocation of tasks.

Hidden Tasks

The figure on the left shows the *true state of the world*: both agents have to deliver letters to both addresses. If Bob *hides* the fact that he has to deliver to the righthand node, we get the situation on the right:



There are two Nash-optimal solutions: one goes left (payoff 2), the other goes right (payoff 4). So the expected payoff is 3 for each agent.



In the unique Nash-optimal solution, Ann goes left (payoff 2) and Bob goes right (*true* payoff 4). Note that Bob can still deliver his hidden letter.


Summary

We have analysed negotiation between two self-interested agents:


- The *Monotonic Concession Protocol* (MCP) is a formalisation of natural step-wise negotiation behaviour.
- The *Zeuthen Strategy* for the MCP can be motivated in two ways:
 - *intuitively*, using the idea of “willingness to risk conflict”
 - *axiomatically*, by deriving it from more fundamental postulates
- We have seen that if willingness to risk conflict is identical for both agents in the final step, then either *efficiency* or *stability* need to be sacrificed (depending on the chosen strategy).
- We have also seen that a much simpler *one-shot protocol* can directly select Nash-optimal solutions.
- In task-oriented domains, such protocols can be manipulated by either *hiding* tasks or by producing *phantom tasks*.

Phantom Tasks

An alternative way for Bob to manipulate the protocol would be to declare a *phantom task* on top of his actual tasks:



There are two Nash-optimal solutions: one goes left (payoff 2), the other goes right (payoff 4). So the expected payoff is 3 for each agent.



There is a unique Nash-optimal solution: Ann goes left (payoff 2) and Bob goes right — but only to the first node (*true* payoff 4).

References

- J.S. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- J.C. Harsanyi. Approaches to the Bargaining Problem before and after the Theory of Games. *Econometrica*, 24(2):144–157, 1956.

What next?

Today we have only dealt with problems where *two* agents need to come to an agreement. Negotiation between n agents, in particular if every agent can talk to every other agent, is a lot more complicated.

If it is possible to put some restrictions on the “negotiation topology” the problem may become more manageable. A case of special interest are *auctions*. In an auction, one agent (the auctioneer) negotiates with many other agents (the bidders).

Over the next couple of weeks or so we’ll be talking about auctions:

- *Basic Auction Theory* (for a single good)
- *Combinatorial Auctions* (for bundles of goods)