Monotonic Concession Protocols for Multilateral Negotiation

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

Talk Overview

- The need for *multilateral* ("many-to-many") negotiation
- The *monotonic concession protocol* in the *two-agent case*
- Generalisation to the *multilateral case:* protocol structure
- Possible definitions for *multilateral concession criteria*
- Discussion of the *properties* of the resulting negotiation protocols
- Brief discussion of negotiation *strategies*

Multilateral Negotiation

- Most work on negotiation in MAS has considered either *bilateral* ("one-to-one") negotiation or *auctions* ("one-to-many").
- Modelling truly *multilateral* ("many-to-many") negotiation, where more than just two agents can come together and agree on a deal, is difficult but important. Example:

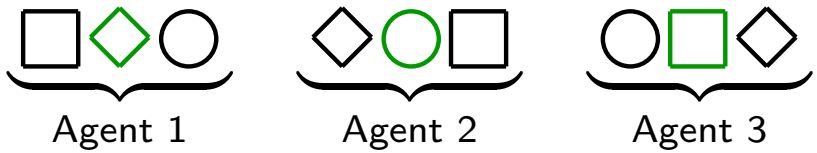

Each agent currently holds their second-favourite item; their lefthand neighbour holds their favourite item; and their righthand neighbour holds their least preferred item. This allocation is suboptimal, but no bilateral deal is feasible.

Monotonic Concession Protocol

A very natural form of negotiation is to first propose your preferred deal and then to make small concessions until agreement is reached.

For *two agents*, this has been formalised as the *monotonic concession protocol* (Zeuthen 1930; Harsanyi 1956; Rosenschein & Zlotkin 1994):

(1) In the first round, each agent makes an initial proposal.

(2) In each subsequent round, each agent can either make a *concession* or stick with their current proposal.

(3) Repeat this until *conflict* arises or an *agreement* is reached.

Here a *concession* is a proposal that is better for your opponent than your previous proposal. *Agreement* is reached if one agent makes a proposal that is better for their opponent than the opponent's own proposal. *Conflict* arises if there is a round where no agent concedes; this is considered the worst possible outcome.

Generalisation to the Multilateral Case

The definition of the overall *protocol* remains the same:

(1) In the first round, each agent makes an initial proposal.

(2) In each subsequent round, each agent can either make a *concession* or stick with their current proposal.

(3) Repeat this until *conflict* arises or an *agreement* is reached.

Here, *conflict* still means that nobody concedes during one round.

The notion of *agreement* is easily generalised: agreement is reached if one agent makes a proposal that everyone likes at least as much as their own proposal.

▶ What does it mean to make a *concession* to a group of opponents?

Possible Multilateral Concession Criteria

(1) *Strong concession:* Make a proposal that is strictly better for each of the other agents.

(2) *Weak concession:* Make a proposal that is strictly better for at least one of the other agents.

(3) *Pareto concession:* Make a proposal that is no worse for the other agents and strictly better for one of them.

(4) *Utilitarian concession:* Make a proposal such that the sum of utilities of the other agents increases.

(5) *Egalitarian concession:* Make a proposal such that the minimum utility amongst the other agents increases.

(6) *Nash concession:* Make a proposal such that the product of utilities of the other agents increases.

(7) *Egocentric concession:* Make a proposal that is worse for yourself.

Protocol Properties

All seven definitions are *faithful generalisations* of the two-agent case (the egocentric one only if we just consider non-dominated proposals).

The paper discusses several properties of the resulting protocols:

- *Termination:* would certain criteria permit an agent to make an infinite sequence of concessions?
- *Compositionality:* will the composition of two concessions each meeting a given criterion always meet that same criterion as well?
- *Deadlock-freedom:* is it possible that negotiation gets stuck, because no agent is *able* to make a valid concession?
- *Verifiability:* can the task of verifying conformance to the protocol be distributed amongst the agents?

The most interesting of these is deadlock-freedom ...

Deadlock-Freedom

A concession criterion is *deadlock-free* iff it guarantees that at least one agent can always make a concession satisfying the criterion, until an agreement has been reached.

Proposition 1 (Two-agent case) *In the two-agent case, all of our seven concession criteria are deadlock-free.*

Proposition 2 (General case) *The weak, the utilitarian, and the egocentric criteria are all deadlock-free. The Pareto, the strong, and the egalitarian criteria are not deadlock-free. The Nash criterion is deadlock-free iff utilities are required to be positive.*

In this context, we call a utility function positive iff all agreements but the conflict deal have strictly positive utility.

Negotiation Strategies

In the two-agent case, the *Zeuthen strategy* stipulates that the agent with the lower *willingness to risk conflict* should concede. This is defined as the ratio of the loss incurred by accepting your opponent's proposal and the loss incurred by causing conflict (utility 0):

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

Unclear how this could be generalised to the *multilateral* case. One option would be to evaluate willingness to risk conflict assuming the worst possible outcome in case of a concession:

$$Z_i \quad = \quad \frac{u_i(x_i) - \min\{u_i(x_k) \mid k \in \mathcal{A}gents\}}{u_i(x_i)}$$

The problem is that this strategy can lead to a *deadlock*. The agent with the lowest Z-value may simply not be *able* to make a concession that would tip the balance (see poster for an example).

Conclusions

- *Multilateral* negotiation is important: it is often not possible to decompose a complex deal into a sequence of bilateral deal.
- The two-agent *monotonic concession protocol* is a formalisation of a very natural form of negotiation.
- Generalising this idea to the multilateral case has given rise to *seven possible concession criteria* (more would be conceivable).
- The paper discusses the properties of the resulting protocols: *termination*, *compositionality*, *deadlock-freedom*, *verifiability*.
- Developing *negotiation strategies* is a difficult problem.
  A full game-theoretical analysis, in line with that of Harsanyi (1956) for the two-agent case, seems promising.