Optimal Outcomes of Negotiations over Resources

Ulle Endriss, Nicolas Maudet, Fariba Sadri & Francesca Toni
Imperial College London
{ue,maudet,fs,ft}@doc.ic.ac.uk

Recent Work at Imperial

Recent work at Imperial approaches the problem domain of *resource allocation by negotiation* at three levels:

- [WP1] Concrete negotiation *strategies* for agents (with Bologna) expressed as integrity constraints in abductive logic programming [AISB-01, ATAL-01, JELIA-02] — see Paolo's talk
- [WP2] Communication *protocols* to specify "rules of encounter", also expressed as integrity constraints (not just for this scenario) [UKMAS-02, AAMAS-03] — see Nicolas' talk
- [WP1/2/5] Study of necessary/sufficient *patterns of resource exchange* and of the notion of *optimal outcomes* of negotiation [UKMAS-02, AAMAS-03] — this talk

Resource Allocation by Negotiation

We consider scenarios where agents negotiate deals to exchange resources in order to benefit either themselves or society as a whole.

Main Questions

- What does an *individual agent* want? How can we formally characterise its attitude towards negotiation?
- What do we consider a positive (or even optimal) outcome of a negotiation process from the viewpoint of *society*?
- To what extent are the interests of individuals and society compatible?
- What are minimal requirements on the *kinds of resource exchanges* a system needs to allow for to enable society to reach optimal outcomes (where this is at all possible)?

Talk Overview

- Basic scenario of resource allocation by negotiation
- Optimal outcomes for scenarios with money
  [based on work by Sandholm (1998) on task contracting]
- Optimal outcomes for scenarios without money
- Additional results for special cases
- Resource allocation in egalitarian agent societies
- Future work

Resource Allocation by Negotiation

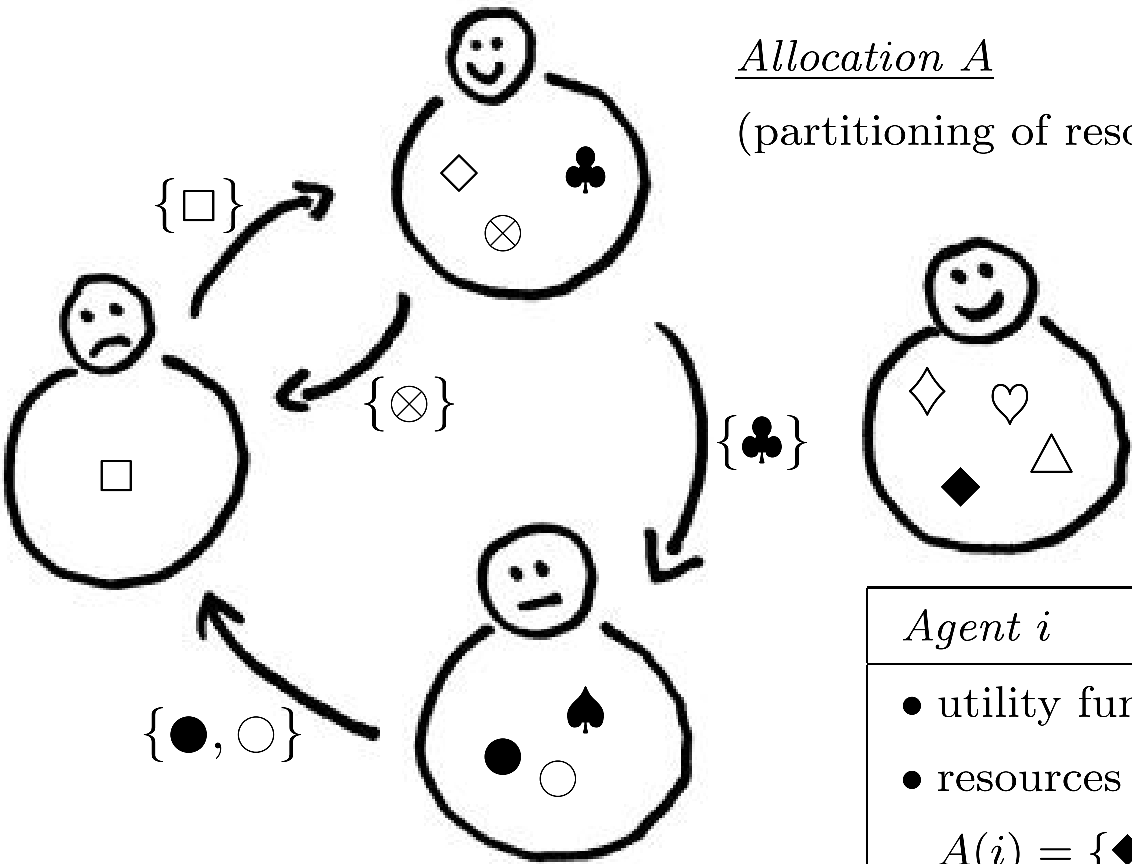

Allocation A

(partitioning of resources amongst agents)

Deal $\delta = (A, A')$

(pair of allocations: before/after)

Agent i
• utility function $u_i : 2^{\mathcal{R}es} \to \mathbb{R}$ • resources for allocation A: $A(i) = \{\blacklozenge, \lozenge, \heartsuit, \triangle\}$ • happiness for allocation A: $u_i(A) = u_i(A(i)) = 510.7$

Deal Types

Following Sandholm (1998), we can distinguish different deal *types:*

- *One-resource-at-a-time deals:* one resource changes hands
- *Cluster deals:* one agent gives a set of resources to another agent
- *Swap deals:* one resource is exchanged for another
- *Multiagent deals:* n agents, at most one resource per pair

Deals and Payments

A deal may be accompanied by a *payment* to compensate agents for otherwise disadvantageous deals. A payment function p is a function from agents to $\mathbb{R}$ (money) with $\sum_i p(i) = 0$ (i.e. the overall amount of money in the system stays constant).

What does an individual agent i want?

We assume agents are *individually rational* in the sense that they will never accept a disadvantageous deal, i.e.:

- agent i's gain in utility has to be more than the price it pays, or
- agent i's loss in utility has to be less than the amount of money it receives, respectively.

Formally: deal $\delta = (A, A')$ is acceptable to i iff we have:

$$u_i(A') - u_i(A) \quad > \quad p(i)$$

What does society "want"?

We may, for instance, assume that it is in the interest of society to maximise *social welfare* as defined in the utilitarian tradition, i.e. to maximise the sum of all individual utilities:

$$sw(A) \quad = \quad \sum_i u_i(A)$$

Society and the Individual

It is possible to show that a deal is individually rational *iff* it increases social welfare.

▶ Our notion of individual rationality seems appropriate for a society based on utilitarian principles (namely this particular definition of social welfare).

Sufficiency Result

By a variant of a result due to Sandholm (1998), the class of individually rational deals is *sufficient* to guarantee an optimal outcome of negotiation:

Theorem. Any sequence of deals (with money) that are individually rational will eventually culminate in an allocation of resources with maximal social welfare.

Consequences: Negation always pays off for both individual agents and society; we won't get stuck in local minima.

Example

Agent 1	Agent 2
$A_0(1) = \{r_1, r_2\}$	$A_0(2) = \{\,\}$
$u_1(\{\,\}) = 0$	$u_2(\{\,\}) = 0$
$u_1(\{r_1\}) = 2$	$u_2(\{r_1\}) = 3$
$u_1(\{r_2\}) = 3$	$u_2(\{r_2\}) = 3$
$u_1(\{r_1, r_2\}) = 7$	$u_2(\{r_1, r_2\}) = 8$

Social welfare for allocation A_0 is 7, but it could be 8. However, by moving a *single* resource from agent 1 to agent 2, the former would lose more than the latter would gain (not individually rational). The only possible deal would be to move the whole *set* $\{r_1, r_2\}$.

▶ Hence, *one-resource-at-a-time deals are not sufficient* to guarantee outcomes with maximal social welfare.

Necessity Result

Also due to Sandholm:

Theorem. Any given deal $\delta = (A, A')$ may be *necessary*, i.e. there are utility functions and an initial allocation such that any sequence of individually rational deals leading to an allocation with maximal social welfare would have to include δ.

Consequences: Simple swap deals etc. are not enough to guarantee optimal outcomes. $\Rightarrow$ We need richer negotiation protocols.

Additive Scenarios

Theorem. If all utility functions in the system are *additive* then *one-resource-at-a-time deals* (with money) are sufficient to guarantee maximal social welfare.

Scenarios without Money

A problem with the framework presented so far is that agents may require *unlimited amounts of money* to be able to agree to every beneficial deal.

► For a similar negotiation framework *without money* we get the following results:

Theorem. Any sequence of deals that are *cooperatively rational* will eventually culminate in a *Pareto optimal* allocation of resources.

Theorem. Again, any given deal may be necessary to guarantee Pareto optimal outcomes.

0-1 Scenarios

Theorem. If all utility functions are *"0-1"*, i.e. additive and utilities for single resources are either 1 (*need it*) or 0 (*don't need it*) then *one-resource-at-a-time deals* are sufficient to guarantee maximal social welfare (even without money).

▶ Some of the negotiation strategies put forward by Sadri, Toni & Torroni (2001) may be regarded as implementations of this result:

$$\begin{array}{lcl} \mathit{request}(R,T) \wedge \mathit{have}(R,T) \wedge \neg \mathit{need}(R) & \Rightarrow & \mathit{accept}(T+1) \\ \mathit{request}(R,T) \wedge \mathit{need}(R) & \Rightarrow & \mathit{refuse}(T+1) \\ \mathit{request}(R,T) \wedge \neg \mathit{have}(R,T) & \Rightarrow & \mathit{refuse}(T+1) \end{array}$$

If agents follow this strategy, then negotiation will always terminate and *if* there is a solution (an allocation where everyone gets what they need) then the final allocation will be such a solution.

Egalitarian Agent Societies

- The *utilitarian* social welfare function $sw(A) = \sum_i u_i(A)$ is usually taken for granted in the MAS literature ...
- In an *egalitarian* system, social welfare is tied to the welfare of the (currently) weakest agent: $sw_e(A) = min\{u_i(A) \,|\, i \in \mathcal{A}\}$. This may be more appropriate for *some* applications.
- In analogy to the notion of individual/cooperative rationality of the utilitarian framework, we have developed a local criterion for egalitarian agents to decide whether a given deal is acceptable. The corresponding class of deals is both sufficient and necessary to guarantee outcomes with maximal egalitarian social welfare.
- But note, for instance, that one-resource-at-a-time deals are *not* sufficient for 0-1 scenarios in the egalitarian framework.

Future Work

- What *types of deals* are required to guarantee optimal outcomes for what *classes of utility functions*, and vice versa?
  - Our results for additive and 0-1 scenarios are first steps in this direction.
- Investigate other notions of social welfare.
  - E.g., in *elitist agent societies* social welfare would depend on the welfare of the agent currently best off: agents cooperate to support their "champion" (so at least one agent may achieve its goal).
- Develop *protocols* for multi-item/multi-agent trading of resources.
  - see also Nicolas' talk

References

[1] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Communication Protocols for Logic-based Agents. In *Proceedings of the 5th UK Workshop on Multiagent Systems (UKMAS-2002)*, 2002. Abstract.

[2] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Conditions for Optimal Outcomes of Negotiations about Resources. In *Proceedings of the 5th UK Workshop on Multiagent Systems (UKMAS-2002)*, 2002.

[3] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Aspects of Protocol Conformance in InterAgent Dialogue. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, 2003. Poster. To appear.

[4] U. Endriss, N. Maudet, F. Sadri, and F. Toni. On Optimal Outcomes of Negotiations over Resources. In *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-2002)*, 2003. To appear.

[5] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Resource Allocation in Egalitarian Agent Societies. 2003. Submitted.

[6] F. Sadri, F. Toni, and P. Torroni. Dialogues for Negotiation: Agent Varieties and Dialogue Sequences. In *Intelligent Agents VIII: Agent Theories, Architectures, and Languages*, volume 2333 of *LNAI*. Springer-Verlag, 2001.

[7] F. Sadri, F. Toni, and P. Torroni. Logic Agents, Dialogues and Negotiation: An Abductive Approach. In M. Schroeder and K. Stathis, editors, *Proceedings of the AISB Symposium on Information Agents for E-commerce*, 2001.

[8] F. Sadri, F. Toni, and P. Torroni. An Abductive Logic Programming Architecture for Negotiating Agents. In *Proceedings of JELIA-2002*, 2002.

[9] T. W. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. In *AAAI Spring Symposium: Satisficing Models*, 1998.

Cooperative Rationality and Pareto Optimality

▶ For scenarios without money, we call a deal $\delta = (A, A')$ *cooperatively rational* iff $u_i(A) \leq u_i(A')$ for all agents i and $u_j(A) < u_j(A')$ for at least one agent j (the proposer).

▶ An allocation of resources is called *Pareto optimal* iff there is no other allocation that is better for some agents in the society without being worse for any of the others.