Optimal Outcomes of Negotiations over Resources

Ulle Endriss[1], Nicolas Maudet[2], Fariba Sadri[1] and Francesca Toni[1]

[1] Department of Computing, Imperial College London
Email: {ue,fs,ft}@doc.ic.ac.uk

[2] School of Informatics, City University, London
Email: maudet@soi.city.ac.uk

Talk Overview

- Resource allocation by negotiation in multiagent systems
  definition of our negotiation framework (with money)

- Measuring social welfare
  what are optimal outcomes from the viewpoint of society?

- Results for scenarios with money
  what deals are sufficient to guarantee optimal outcomes?

- Negotiating over resources without money
  the problem of "unlimited money"; refinement of the framework

- Results for scenarios without money
  what deals are sufficient/necessary for optimal outcomes?

- Conclusion
  summary and future work

Resource Allocation by Negotiation

- Finite set of *agents* $\mathcal{A}$ and finite set of *resources* $\mathcal{R}$.
- An *allocation* A is a partitioning of $\mathcal{R}$ amongst the agents in $\mathcal{A}$.
  Example: $A(i) = \{r_3, r_7\}$ — agent i owns resources r_3 and r_7
- Every agent $i \in \mathcal{A}$ has got a *utility function* $u_i : 2^{\mathcal{R}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy
- Agents may engage in negotiation to exchange resources in order to benefit either themselves or society as a whole.
- A *deal* $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may be accompanied by a payment to compensate some of the agents for a loss in utility. A *payment function* is a function $p : \mathcal{A} \to \mathbb{R}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.
  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* AU\$5 while agent j *receives* AU\$5

The Local Perspective

A *rational* agent (who does not plan ahead) will only accept deals that improve its individual welfare:

Definition 1 *A deal $\delta = (A, A')$ is called individually rational iff there exists a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.*

The Global Perspective

A *social welfare function* is a mapping from the preferences of the members of a society to a preference profile for society itself.

Definition 2 *The (utilitarian) social welfare $sw(A)$ of an allocation of resources A is defined as follows:*

$$sw(A) \quad = \quad \sum_{i \in \mathcal{A}} u_i(A)$$

Linking the Local and the Global Perspective

Lemma 1 *A deal $\delta = (A, A')$ is individually rational iff it increases social welfare.*

Proof. '$\Rightarrow$': Use definitions.
'$\Leftarrow$': Every agent will get a positive payoff if the following payment function is used:

$$p(i) \;=\; u_i(A') \;-\; u_i(A) \;-\; \underbrace{\frac{sw(A') - sw(A)}{|\mathcal{A}|}}_{>0}$$

□

▶ This lemma confirms that individually rational behaviour is appropriate in utilitarian societies.

▶ In a related paper (MFI-2003), we investigate what deals are acceptable in *egalitarian agent societies*, where social welfare is tied to the well-being of the weakest agent.

Sufficient Deals (with Money)

The following result is due to Sandholm (1996):

Theorem 1 *Any sequence of individually rational deals will eventually result in an allocation with maximal social welfare.*

Discussion

- Agents can agree on deals *locally*; convergence towards a *global* optimum is guaranteed by the theorem. (+)
- Actually *finding* deals that are individually rational can be very complex. (–)
- Agents may require *unlimited amounts of money* to get through a negotiation. (–)

Scenarios without Money

If we do not allow for compensatory payments, we cannot always guarantee outcomes with maximal social welfare. Example:

Agent 1	Agent 2
$A_0(1) = \{r\}$	$A_0(2) = \{\,\}$
$u_1(\{\,\}) = 0$	$u_2(\{\,\}) = 0$
$u_1(\{r\}) = 4$	$u_2(\{r\}) = 7$

In the framework *with* money, agent 2 could pay AU\$5.5 to agent 1, but ...

► Trying to maximise social welfare is asking too much for scenarios without money. Let's try Pareto optimality instead ...

Pareto Optimality

Using the agents' utility functions and the notion of social welfare, we can define Pareto optimality as follows:

Definition 3 *An allocation A is called Pareto optimal iff there is no allocation A' such that $sw(A) < sw(A')$ and $u_i(A) \leq u_i(A')$ for all agents $i \in \mathcal{A}$.*

Still, if agents behave strictly individually rational, we cannot guarantee outcomes that are Pareto optimal either. Example:

Agent 1	Agent 2
$A_0(1) = \{r\}$	$A_0(2) = \{\,\}$
$u_1(\{\,\}) = 0$	$u_2(\{\,\}) = 0$
$u_1(\{r\}) = 0$	$u_2(\{r\}) = 7$

A_0 is not Pareto optimal, but it would not be individually rational for agent 1 to give the resource r to agent 2.

Cooperative Rationality

If agents are not only *rational* but also (a little bit) *cooperative*, then the following acceptability criterion for deals makes sense:

Definition 4 *A deal* $\delta = (A, A')$ *is called cooperatively rational iff* $u_i(A) \leq u_i(A')$ *for all agents* $i \in \mathcal{A}$ *and that inequality is strict for at least one agent (say, the one proposing the deal).*

Linking the local and the global view again:

Lemma 2 *Any cooperatively rational deal increases social welfare.*

Lemma 3 *For any allocation* A *that is not Pareto optimal there is an* A' *such that the deal* $\delta = (A, A')$ *is cooperatively rational.*

Sufficient Deals (without Money)

We get a similar sufficiency result as before:

Theorem 2 *Any sequence of cooperatively rational deals will eventually result in a Pareto optimal allocation of resources.*

Proof. (i) every deal increases social welfare + the number of distinct allocations is finite ⇒ termination ✓
(ii) assume A is a terminal allocation but not Pareto optimal ⇒ there still exists a cooperatively rational deal ⇒ contradiction ✓ □

Again, this means that cooperatively rational agents can negotiate *locally*; the (Pareto) optimal outcome for society is guaranteed.

▶ But complexity is still a problem . . .

Example

For simplicity, assume utility functions are *additive*, i.e. $u_i(R) = \sum_{r \in R} u_i(\{r\})$ for all agents i and resource bundles R.

Agent 1	Agent 2	Agent 3
$A_0(1) = \{r_2\}$	$A_0(2) = \{r_3\}$	$A_0(3) = \{r_1\}$
$u_1(\{r_1\}) = 7$	$u_2(\{r_1\}) = 4$	$u_3(\{r_1\}) = 6$
$u_1(\{r_2\}) = 6$	$u_2(\{r_2\}) = 7$	$u_3(\{r_2\}) = 4$
$u_1(\{r_3\}) = 4$	$u_2(\{r_3\}) = 6$	$u_3(\{r_3\}) = 7$

Any deal involving only two agents would require one of them to accept a loss in utility (not cooperatively rational!).

▶ Deals involving more than two agents can be *necessary* to guarantee optimal outcomes.

Necessary Deals (without Money)

Optimal outcomes can only be guaranteed if the negotiation protocol allows for deals involving any number of agents and resources:

Theorem 3 *Any given deal* $\delta = (A, A')$ *may be necessary, i.e. there are utility functions and an initial allocation such that any sequence of cooperatively rational deals leading to a Pareto optimal allocation would have to include* δ*.*

Proof. By systematically constructing of counterexamples. □

- There is a similar result for scenarios *with* money (see paper).

Conclusion: Future and Related Work

- We have shown that *cooperatively rational* deals are *sufficient* and *necessary* to guarantee *Pareto optimal* outcomes in negotiations over resources *without* money.
- How about scenarios with *limited* amounts of money?
- Can we reduce complexity by restricting utility functions? (some results for simple cases are in the paper)
- *Welfare engineering:* Given a suitable social welfare function, what kind of local behaviour will guarantee global optima? (see our paper on *egalitarian agent societies* for an example)
- Develop *protocols* for multi-agent/multi-item trading.