

Multiagent Systems: Spring 2006

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

Distributed Negotiation

Negotiation is a central topic in MAS: whether agents are competing or collaborating, they do need to come to agreements over allocations of resources, joint plans of actions, and so on. So far we have discussed:

- Bilateral negotiation between two agents
- Auctions as a means of one-to-many negotiation

The most general scenario, however, would be many agents all negotiating with each other in a *distributed* manner being able to forge *multilateral* deals between more than just two agents each.

Plan for Today

- Centralised vs. distributed negotiation
- Contract Net protocol
- Distributed negotiation of socially optimal allocations

Centralised vs. Distributed Negotiation

An allocation procedure to determine a suitable allocation of resources may be either centralised or distributed:

- In the *centralised* case, a single entity decides on the final allocation, possibly after having elicited the preferences of the other agents. Example: combinatorial auctions
- In the *distributed* case, allocations emerge as the result of a sequence of local negotiation steps. Such local steps may or may not be subject to structural restrictions (say, bilateral deals).

Which approach is appropriate under what circumstances?

Advantages of the Centralised Approach

Much recent work in the MAS community on negotiation and resource allocation has concentrated on centralised approaches, in particular on combinatorial auctions.

There are several reasons for this:

- The *communication protocols* required are relatively simple.
- Many results from *economics* and *game theory*, in particular on mechanism design, can be exploited.
- There has been a recent push in the design of *powerful algorithms* for winner determination in combinatorial auctions.

Disadvantages of the Centralised Approach

But there are also some disadvantages of the centralised approach:

- Can we *trust* the centre (the auctioneer)?
- Does the centre have the *computational* resources required? (but beware: distributing the work does not dissolve NP-hardness)
- Less natural to take an *initial allocation* into account (in an auction, typically the auctioneer owns everything to begin with).
- Less natural to model *step-wise improvements* over the *status quo*.
- Arguably, only the distributed approach is a serious implementation of the *MAS paradigm*.

Challenges of Distributed Negotiation

Research on distributed negotiation has not yet reached the level of maturity we find in centralised approaches such as combinatorial auctions. There are many challenging questions how to best set up a framework for distributed (and multilateral) negotiation. Examples:

- What are appropriate *communication protocols*?
- There are *exponentially many groups* of agents that may want to forge a deal (not the case for bilateral negotiation or auctions). How can we master this *complexity*?
- What we have learned about *mechanism design* all relies on a centre computing allocations and prices. Can this be distributed?
- Recall the *monotonic concession protocol* for bilateral negotiation:
 - Is it possible to extend this idea to multilateral negotiation?
 - What does it mean to concede to a group of opponents?

The Contract Net Protocol

Originally developed for task decomposition and allocation, but also applicable to negotiation over resources.

Each agent may assume to role of *manager* and *bidder*. The Contract Net protocol is a one-to-many protocol matching an offer by a manager to one of potentially many bidders. There are four *phases*:

- *Announcement phase*: The manager advertises a deal to a number of partner agents (the bidders).
- *Bidding phase*: The bidders send their proposals to the manager.
- *Assignment phase*: The manager elects the best bid and assigns the resource(s) accordingly.
- *Confirmation phase*: The elected bidder sends a confirmation.

R.G. Smith. *The Contract Net Protocol: High-level Communication and Control in a Distributed Problem Solver*. IEEE Trans. on Computers, 29:1104–1113, 1980.

Extensions

The immediate adaptation of the original Contract Net protocol only allows managers to advertise a *single resource* at a time, and a bidder can only offer *money in return* for that resource (not other items). Possible extensions:

- Allow for negotiation over the exchanges of *bundles* of resources.
- Allow for deals *without explicit utility transfers* (monetary payments). The announcement phase remains the same, but bids are now about offering resources in exchange, rather than money.
- Allow agents to negotiate several deals *concurrently* and to *decommit* from deals for a certain period, to help them to negotiate better deals.
- In *levelled-commitment contracts*, agents are also allowed to decommit, but have to pay a pre-defined *penalty* in case they choose to do so.

Refer to the MARA Survey for references to these works.

Y. Chevaleyre *et al.* *Issues in Multiagent Resource Allocation*. Informatica, 30:3–31, 2006. Section on Allocation Procedures.

Negotiating Socially Optimal Allocations

For the remainder of this lecture, we are going to analyse a specific model of distributed negotiation (defined on the next slide).

We are not going to talk about designing a concrete negotiation protocol, but rather study the framework from an abstract point of view. The main question concerns the relationship between

- the *local view*: what deals will agents make in response to their individual preferences?; and
- the *global view*: how will the overall allocation of resources evolve in terms of social welfare?

U. Endriss, N. Maudet, F. Sadri and F. Toni. *Negotiating Socially Optimal Allocations of Resources*. Journal of Artif. Intelligence Research, 25:315–348, 2006.

An Abstract Negotiation Framework

- Finite set of *agents* \mathcal{A} and finite set of indivisible *resources* \mathcal{R} .
- An *allocation* A is a partitioning of \mathcal{R} amongst the agents in \mathcal{A} .
Example: $A(i) = \{r_5, r_7\}$ — agent i owns resources r_5 and r_7
- Every agent $i \in \mathcal{A}$ has got a *utility function* $u_i : 2^{\mathcal{R}} \rightarrow \mathbb{R}$.
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 come with a number of side payments to compensate some of the agents for a loss in utility. A *payment function* is a function $p : \mathcal{A} \rightarrow \mathbb{R}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.
Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* €5, while agent j *receives* €5.

The Local/Individual 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 $A(i) = A'(i)$.

That is, an agent will only accept a deal iff it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility).

Observe that this is weaker than the standard notion of rationality familiar from game theory.

The Global/Social Perspective

As a system designers we are interested in the quality of allocations at the social level. One interesting metric is utilitarian social welfare:

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

$$sw_u(A) = \sum_{i \in Agents} u_i(A)$$

This can serve as an indicator for the overall profit generated.

► Recall that we have seen that there are several alternative definitions of social welfare. Which is appropriate depends on the application.

Example

Let $\mathcal{A} = \{ann, bob\}$ and $\mathcal{R} = \{chair, table\}$ and suppose our agents use the following utility functions:

$$\begin{array}{ll} u_{ann}(\{\}) = 0 & u_{bob}(\{\}) = 0 \\ u_{ann}(\{chair\}) = 2 & u_{bob}(\{chair\}) = 3 \\ u_{ann}(\{table\}) = 3 & u_{bob}(\{table\}) = 3 \\ u_{ann}(\{chair, table\}) = 7 & u_{bob}(\{chair, table\}) = 8 \end{array}$$

Furthermore, suppose the initial allocation of resources is A_0 with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \{\}$.

► Social welfare for allocation A_0 is 7, but it could be 8. By moving only a *single* resource from agent *ann* to agent *bob*, the former would lose more than the latter would gain (not individually rational). The only possible deal would be to move the whole *set* $\{chair, table\}$.

Linking the Local and the Global Perspectives

It turns out that individually rational deals are exactly those deals that increase social welfare:

Lemma 1 (Rationality and social welfare) A deal $\delta = (A, A')$ with side payments is *individually rational* iff $sw_u(A) < sw_u(A')$.

Proof: “ \Rightarrow ”: Rationality means that overall utility gains outweigh overall payments (which are $= 0$).

“ \Leftarrow ”: The social surplus can be divided amongst all deal participants by using the following payment function:

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

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

Termination

We can now prove a first result on negotiation processes:

Lemma 2 (Termination) There can be no infinite sequence of individually rational deals, i.e. negotiation must always *terminate*.

Proof: Follows from the first lemma and the observation that the space of distinct allocations is finite. \square

Convergence

It is now easy to prove the following *convergence* result (originally stated by Sandholm in the context of distributed task allocation):

Theorem 3 (Sandholm, 1998) *Any sequence of individually rational deals will eventually result in an allocation with maximal social welfare.*

Proof: Termination is guaranteed by Lemma 2. So let A be the terminal allocation. Assume A is *not* optimal, i.e. there exists an A' with $sw_u(A) < sw_u(A')$. Then, by Lemma 1, $\delta = (A, A')$ is individually rational \Rightarrow contradiction. \square

► Agents can act *locally* and need not be aware of the global picture (convergence towards a global optimum is guaranteed by the theorem).

T. Sandholm. *Contract Types for Satisficing Task Allocation: I Theoretical Results.* AAAI Spring Symposium 1998.

Multilateral Negotiation

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

Theorem 4 (Necessity of complex deals) *Any deal $\delta = (A, A')$ may be *necessary*: 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 (unless δ is “independently decomposable”).*

The proof involves the systematic definition of utility functions such that A' is optimal and A is the second best allocation. Independently decomposable deals (to which the result does not apply) are deals that can be split into two subdeals concerning distinct sets of agents.

Independently Decomposable Deals

The set of agents *involved in a deal* $\delta = (A, A')$ is given by $\mathcal{A}^\delta = \{i \in \mathcal{A} \mid A(i) \neq A'(i)\}$.

The *composition* of two deals is defined as follows:

If $\delta_1 = (A, A')$ and $\delta_2 = (A', A'')$, then $\delta_1 \circ \delta_2 = (A, A'')$.

If a deal δ is the composition of two deals concerning disjoint sets of agents, then δ is said to be *independently decomposable*. Formally:

Definition 3 *A deal δ is called *independently decomposable* iff there exist deals δ_1 and δ_2 such that $\delta = \delta_1 \circ \delta_2$ and $\mathcal{A}^{\delta_1} \cap \mathcal{A}^{\delta_2} = \{\}$.*

Theorem 4 does *not* apply to independently decomposable deals δ :

- Let δ be the deal of moving r_1 from agent 1 to agent 2, and r_2 from agent 3 to agent 4.
- If δ is individually rational, so will be one of the two “subdeals”: moving r_1 from agent 1 to agent 2; or r_2 from agent 3 to agent 4.

Proof of Theorem 4

Let $\delta = (A, A')$ be any deal that is not independently decomposable.

Need to *construct utility functions* such that A' has maximal social welfare and A is second best. Then δ will be necessary if we make A the initial allocation (by Lemma 1).

There must be an agent $j \in \mathcal{A}$ such that $A(j) \neq A'(j)$. Define:

$$u_i(R) = \begin{cases} 1 & \text{if } R = A'(i) \text{ or } (R = A(i) \text{ and } i \neq j) \\ 0 & \text{otherwise} \end{cases}$$

We get $sw_u(A') = |\mathcal{A}|$ and $sw_u(A) = sw_u(A') - 1$.

Because $\delta = (A, A')$ is not independently decomposable, there exists no allocation B different from both A and A' such that $B(i) = A(i)$ or $B(i) = A'(i)$ for all agents $i \in \mathcal{A}$.

Hence, $sw_u(B) \leq sw_u(A)$ for any other allocation B . \square

Negotiation in Restricted Domains

Multilateral negotiation is difficult to implement . . .

Maybe we can guarantee convergence to a socially optimal allocation for structurally simpler types of deals if we restrict the range of utility functions that agents can use? First, two negative results:

- Our proof already shows that Theorem 4 continues to hold even when all agents are required to use *dichotomous* utility functions. $[u_i(R) = 0 \vee u_i(R) = 1]$
- The same is true when all agents are required to use *monotonic* utility functions. $[R_1 \subseteq R_2 \Rightarrow u_i(R_1) \leq u_i(R_2)]$

Modular Domains

Recall that a utility function u_i is called *modular* iff it satisfies the following condition for all bundles $R_1, R_2 \subseteq \mathcal{R}$:

$$u_i(R_1 \cup R_2) = u_i(R_1) + u_i(R_2) - u_i(R_1 \cap R_2)$$

That is, in a modular domain there are no synergies between items; you can get the utility of a bundle by adding up the utilities of the items in that bundle.

► Negotiation in modular domains *is* feasible:

Theorem 5 (Modular domains) *If all utility functions are modular, then individually rational 1-deals (involving just one resource) suffice to guarantee outcomes with maximal social welfare.*

Proof: Easy. □

Scenarios without Money

Agents may require unlimited amounts of money to get through a negotiation. So what happens if we do not allow for side payments?

Without money, we cannot always guarantee outcomes with maximal utilitarian 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 €5.5 to agent 1 to achieve a mutually beneficial deal, 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 4 *An allocation A is called Pareto optimal iff there is no A' such that $sw_u(A) < sw_u(A')$ and $u_i(A) \leq u_i(A')$ for all $i \in \mathcal{A}$.*

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 *without side payments* makes sense:

Definition 5 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 (easy proofs):

Lemma 6 Any cooperatively rational deal increases social welfare.

Lemma 7 For any allocation A that is not Pareto optimal there is an A' such that the deal $\delta = (A, A')$ is cooperatively rational.

Convergence (without Money)

We get a similar convergence result as before:

Theorem 8 (Convergence) *Any sequence of cooperatively rational deals will eventually result in a Pareto optimal allocation of resources.*

Proof: (1) every deal increases social welfare + the number of distinct allocations is finite \Rightarrow termination \checkmark

(2) assume A is a terminal allocation but not Pareto optimal \Rightarrow there still exists a cooperatively rational deal \Rightarrow contradiction $\checkmark \square$

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

► But the structural complexity of deals 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 9 (Necessity of complex deals) *Any deal $\delta = (A, A')$ that is not independently decomposable may be *necessary*: 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: Very similar to the case with money. \square

Many Open Questions

For *deal types* between 1-deals and general deals: for what classes of utility functions are they sufficient to *guarantee convergence*?

- There are some more results around, but so far nobody has managed to come up with a very catchy or general result.
- For instance, being able to characterise under what circumstances *bilateral deals* would be sufficient would be extremely useful.

What about negotiation outcomes that are optimal with respect to *other social welfare orderings*?

- There are some results (e.g. for *egalitarian* social welfare), but the required characterisation of the negotiation behaviour for individual agents is less attractive than the simple rationality concepts used here.
- Also interesting: try to understand how allocations evolve wrt. fairness criteria when rational agents negotiate (theoretically or experimentally).

Summary

- Distinguish *centralised* and *distributed* approaches to negotiation and resource allocation. Research on distributed negotiation is not yet as far as for centralised approaches: it's a very important topic, but there are also many challenging problems.
- The *Contract Net protocol* can be used to take care of the communication requirements in distributed negotiation and provides means to help agents to identify possible deals, at least for structurally simple deals (e.g. bilateral deals).
- Finally, we have analysed distributed negotiation from an abstract point of view: there are some nice correspondences between the *local* level (rational deals) and the *global* level (social improvements). *Convergence* to a social optimum can usually be guaranteed in theory, but requires very expressive negotiation protocols in practice.

References

For a general overview of allocation procedures for multiagent resource allocation, refer to the following paper:

- Y. Chevaleyre *et al.* Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006. Section on Allocation Procedures.

The convergence results and related issues in distributed resource allocation may be found here:

- U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of Artificial Intelligence Research*, 25:315–348, 2006.

What next?

The main theme in this course is *Multiagent Resource Allocation* and so far we have mostly discussed specific aspects of the general problem (e.g. representational, algorithmic, or game-theoretical aspects).

Next week's lecture will give an *overview* of the MARA research area as a whole and fill some of the gaps left open so far.