

Computational Social Choice: Spring 2007

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

Plan for Today

Last week have discussed the various parameters pertaining to the specification of a MARA problem. Today we will talk about approaches to solving such a MARA problem ...

- Distinction of *centralised* and *distributed* allocation procedures
 - Brief mentioning of (centralised) auction protocols (to be discussed in depth in a future lecture)
 - Brief mentioning of the practical aspects of implementing distributed resource allocation systems (Contract Net)
- Properties of distributed resource allocation procedures:
 - Guaranteeing *convergence* to a socially desirable allocation by means of a sequence of local negotiation steps
 - Different aspects of *complexity* of the above

Allocation Procedures

To solve a MARA problem, we firstly need to decide on an allocation procedure. This is a complex issue, involving at least:

- *Protocols*: What types of deals are possible? What messages do agents have to exchange to agree on one such deal?
- *Strategies*: What strategies can agents use for a given protocol? How can we incentivise agents to behave in a certain way?
- *Algorithms*: How do we solve the computational problems faced by agents when engaged in negotiation?

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 (for instance, the protocol may only allow for bilateral deals).

Which approach is appropriate under what circumstances?

Advantages of the Centralised Approach

Much recent work on negotiation and resource allocation (particularly in the MAS community) 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 to the centralised approach:

- Can we *trust* the centre (the auctioneer)?
- Does the centre have the *computational* resources required? (but beware: distributing it doesn't dissolve NP-hardness)
- Less natural to take an *initial allocation* into account (in an auction, usually the auctioneer owns everything to begin with).
- Less natural to model *step-wise improvements*.
- Arguably, only the distributed approach is a serious implementation of the *MAS paradigm* (that is, while admittedly being difficult, we would really like to understand how to make distributed decision making work ...).

Auction Protocols

Auctions are centralised mechanisms for the allocation of goods amongst several agents. Agents report their preferences (bidding) and the auctioneer decides on the final allocation (and on prices).

- Distinguish *direct* and *reverse* auctions (auctioneer buying).
- Bidding may be *open-cry* (English) or by *sealed bids*.
- Open-cry: *ascending* (English) or *descending* bids (Dutch).
- Pricing rule: *first-price* or *second-price* (Vickrey).
- *Combinatorial auctions*: several goods, sold/bought in bundles.

(Auctions will be the subject of a future lecture.)

R.P. McAfee and J. McMillan. Auctions and Bidding *Journal of Economic Literature*, 25:699–738, 1987.

P. Cramton, Y. Shoham, and R. Steinberg (eds.). *Combinatorial Auctions*, MIT Press, 2006.

The Contract Net Protocol

Originally developed for task decomposition and allocation, but also applicable to *distributed negotiation* over resources. Each agent may assume the roles 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 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. Comp.*, 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 items.
- Allow for deals *without explicit utility transfers* (monetary payments). The announcement phase remains the same, but bids are now about offering resources in exchange, not money.
- Allow agents to negotiate several deals *concurrently* and to *decommit* from deals within a certain period.
- In *levelled-commitment contracts*, agents are also allowed to decommit, but have to pay a pre-defined *penalty*.

Refer to the MARA Survey for references to these works.

Properties of Allocation Procedures

We may study different properties of allocation procedures:

- *Termination*: Is the procedure guaranteed to terminate eventually?
- *Convergence*: Will the final allocation be optimal according to our chosen social welfare measure?
- *Incentive-compatibility*: Do agents have an incentive to report their valuations truthfully? (\leadsto *mechanism design*)
- *Complexity results*: What is the computational complexity of finding a socially optimal allocation of resources?

Next, we are going to see an example for a convergence property ...

Negotiating Socially Optimal Allocations

We are now 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 AI 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:

- 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).

The Global/Social Perspective

As emphasised last week, there are many different (fairness or efficiency) criteria that we could use to define our goals.

For now, suppose that as system designers we are interested in maximising *utilitarian social welfare*:

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

Observe that there is no need to include the agents' monetary balances into this definition, because they'd always add up to 0.

While the local perspective is driving the negotiation process, we use the global perspective to assess how well we are doing.

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 IR). 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:

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

Discussion: The lemma confirms that individually rational behaviour is “appropriate” in utilitarian societies.

Convergence

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

Theorem 1 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Proof: Termination follows from our lemma and the fact that the number of allocations is finite. So let A be the terminal allocation. Assume A is *not* optimal, *i.e.* there exists an allocation A' with $sw_u(A) < sw_u(A')$. Then, by our lemma, $\delta = (A, A')$ is individually rational \Rightarrow contradiction. \checkmark

Discussion: Agents can act *locally* and need not be aware of the global picture (convergence is guaranteed by the theorem).

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

Multilateral Negotiation

On the downside, outcomes that maximise social welfare can only be guaranteed if the negotiation protocol allows for deals involving *any number of agents* and *resources*:

Theorem 2 (Necessity of complex deals) *Any 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 (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.

Negotiation in Restricted Domains

Most work on negotiation in MAS is concerned with bilateral negotiation or auctions. \rightsquigarrow Multilateral negotiation is difficult!

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:

- Theorem 2 continues to hold even when all agents have to use *monotonic* utility functions. $[R_1 \subseteq R_2 \Rightarrow u_i(R_1) \leq u_i(R_2)]$
- Theorem 2 continues to hold even when all agents have to use *dichotomous* utility functions. $[u_i(R) = 0 \vee u_i(R) = 1]$

Modular Domains

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 its elements.

► Negotiation in modular domains *is* feasible:


Theorem 3 (Modular domains) *If all utility functions are *modular*, then individually rational 1-deals (each involving just one resource) suffice to guarantee outcomes with maximal social welfare.*

We also know that the class of modular utility functions is *maximal*: no strictly larger class could still guarantee the same convergence property.

Y. Chevaleyre, U. Endriss, and N. Maudet. On Maximal Classes of Utility Functions for Efficient one-to-one Negotiation. IJCAI-2005.

Digression: Simulation and Experiments

While we know from Theorem 3 that 1-deals (blue) guarantee an optimal result, an experiment (20 agents, 200 resources, modular utilities) suggests that general bilateral deals (red) achieve the same goal faster:



The graph shows how utilitarian social welfare (y -axis) develops as agents attempt to make more and more deals (x -axis).

The simulation platform used to generate this graph has been written by Hylke Buisman, Gijs Kruitbosch and Nadya Peek (BSc AI students).

Communication Complexity

- Last week's NP-completeness results concern the *computational* complexity of an *abstract* problem: finding a socially optimal allocation *somehow* (not necessarily by negotiation).
- What we are really interested in is the complexity of actual negotiation processes.
- So we should also consider the *communication complexity* of negotiating socially optimal allocations: focus on the length of negotiation processes and the amount of information exchanged, rather than just on computational aspects.

U. Endriss and N. Maudet. On the Communication Complexity of Multilateral Trading. *Journal of Autonomous Agents and MAS*, 11(1):91–107, 2005.

Aspects of Complexity

- (1) How many *deals* are required to reach an optimal allocation?
 - communication complexity as number of individual deals
 - technical results to follow
- (2) How many *dialogue moves* are required to make one such deal?
 - affects communication complexity as number of moves
- (3) How expressive a *communication language* do we require?
 - Minimum requirements: *propose*, *accept*, *reject*
 - + content language to specify multilateral deals
 - affects communication complexity as number of bits exchanged
- (4) How complex is the *reasoning* task faced by an agent when deciding on its next dialogue move?
 - computational complexity (local rather than global view)

Number of Deals

There are two results on *upper bounds* pertaining to the first variant of our negotiation framework (with side payments, general utility functions, and aiming at maximising utilitarian social welfare):

Theorem 4 (Shortest path) *A single rational deal is sufficient to reach an allocation with maximal social welfare.*

Proof: Use Lemma 1 [$\delta = (A, A')$ is IR iff $sw_u(A) < sw_u(A')$]. ✓

Theorem 5 (Longest path) *A sequence of rational deals can consist of up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ deals, but not more.*

Proof: No allocation can be visited twice (same lemma) and there are $|\mathcal{A}|^{|\mathcal{R}|}$ distinct allocations \Rightarrow upper bound follows.

To show that the upper bound is *tight*, we need to show that it is possible that all allocations have distinct social welfare ... ✓

Path Length in Modular Domains

If all agents are using modular utility functions and only negotiate 1-deals, then we obtain the following bounds:

- Shortest path: $\leq |\mathcal{R}|$
- Longest path: $\leq |\mathcal{R}| \cdot (|\mathcal{A}| - 1)$

There are similar results for a framework without monetary side payments (where the goal is to reach a Pareto optimal allocation).

Dunne (2005) has also worked on the topic of communication complexity in distributed negotiation, but generally this is still very much an under-explored area . . .

P.E. Dunne. Extremal Behaviour in Multiagent Contract Negotiation. *Journal of Artificial Intelligence Research*, 23:41–78, 2005.

More on Convergence

Generally, it is interesting to see for what kind of combination of deals and optimality criteria we can get convergence results:

- *Deals*: structural constraints and rationality criteria
- *Optimality criteria*: various SWOs, degrees of envy, . . .

For example, the result (Theorem 1) we have discussed in detail shows that by using the rationality criterion given by our definition of individual rationality and by not imposing any structural constraints, we can guarantee convergence with respect to the optimality criterion given by the notion of utilitarian social welfare.

More on Convergence (cont.)

Known results include the following:

- Pareto optimal outcomes can be guaranteed by means of rational deals without money
- Outcomes maximising egalitarian social welfare by means of a specifically designed class of deals (need to give up selfishness)
- Envy-free outcomes by means of IR deals, under various side conditions (supermodular utilities, specific payments, . . .)
- The work on envy-freeness can also be extended to scenarios with negotiation (and envy) “modulo a graph” . . .

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

Y. Chevaleyre, U. Endriss, S. Estivie, and N. Maudet. Reaching Envy-free States in Distributed Negotiation Settings. IJCAI-2007.

Y. Chevaleyre, U. Endriss, and N. Maudet. Allocating Goods on a Graph to Eliminate Envy. AAAI-2007.

Summary

We have discussed the following issues (see also the MARA Survey):

- Pros and cons of using a centralised/distributed approach
 - Most of the lecture has then been on distributed negotiation
 - Centralised allocation procedures (combinatorial auctions) will be the subject of a future lecture.
- Abstract negotiation framework for indivisible resources:
 - Unspecified how exactly deals are being agreed upon (could be something like Contract Net, but it’s definitely non-trivial)
 - Only defined minimal conditions for deal acceptability (“myopic”), rather than to worry about game-theoretical issues
- Technical results on convergence and (communication) complexity

Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J.A. Rodríguez-Aguilar and P. Sousa. Issues in Multiagent Resource Allocation. *Informatica*, 30:3–31, 2006.