Multiagent Resource Allocation

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

Multiagent Resource Allocation

A tentative definition would be the following:

> *Multiagent Resource Allocation (MARA) is the process of distributing a number of items amongst a number of agents.*

Now: *What* kind of items (resources) are being distributed? *How* are they being distributed? And finally, *why* are they being distributed?

Why am I interested in this sort of stuff?

- Many applications: electronic commerce, industrial procurement, satellite exploitation, grid computing, manufacturing, ...
- Highly interdisciplinary: game theory, decision theory, social choice, welfare economics, logic, complexity theory, algorithm design, specification and verification, operations research, ...

Issues

- Preference representation: How do we represent the preferences of individual agents? How should agents communicate their preferences?
- Preference aggregation (social welfare): How do we aggregate individual preferences to decide what is a "good" allocation?
- Allocation procedures: How do we find a good allocation? — options include (centralised) auctions and distributed negotiation.
- Algorithm design: How can we design fast algorithms for this?
- Complexity questions: What is the computational complexity of finding an optimal allocation? What about communication complexity?
- Mechanism design: How do we provide incentives to individual agents to play according to the rules (reveal their true preferences)?
- Simulation and experimentation: When theoretical tools fail us, how can we use simulations to gain further insights?

Talk Overview

- Some results on work in a distributed negotiation setting:
  - Convergence to an optimal allocation
  - Necessity of having to implement very complex deals
  - Positive and negative results in restricted settings
- Brief glance over some other issues of interest:
  - Complexity of negotiation
  - Mechanism design ("inverse game theory")
  - Preference representation in combinatorial domains
  - Fair division

Negotiating Socially Optimal Allocations

- Multiagent systems may be thought of as *"societies of agents"*. This suggests to use tools from microeconomics and social choice theory to assess the performance of the overall system ("society").
- Agents *negotiate* deals to exchange resources to benefit either themselves or society as a whole (*distributed* approach).
- Agents may use very simple rationality criteria to decide what deals to accept, but interaction patterns may be complex (*multilateral* deals).

U. Endriss, N. Maudet, F. Sadri and F. Toni. *Negotiating socially optimal allocations of resources*. Journal of Artificial Intelligence Research, 25:315–348, 2006.

Resource Allocation by Negotiation

- 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 (Rationality) *A deal $\delta = (A, A')$ is called 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

Definition 2 (Social welfare) *The (utlitarian) social welfare of an allocation of resources A is defined as follows:*

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

Example

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

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

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 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 rational iff $sw(A) < sw(A')$.*

Proof. "$\Rightarrow$": Rationality means that overall utility gains outweigh overall payments (which are $= 0$). "$\Leftarrow$": Using side payments, the social surplus can be divided amongst all deal participants. $\square$

▶ We can now prove a first result on negotiation processes:

Lemma 2 (Termination) *There can be no infinite sequence of 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 rational deals will eventually result in an allocation with maximal social welfare.*

▶ 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, i.e. there are utility functions and an initial allocation such that any sequence of 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 multiagent systems 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 4 continues to hold even when all agents are required to use *monotonic* utility functions. $[R_1 \subseteq R_2 \Rightarrow u_i(R_1) \leq u_i(R_2)]$
- Theorem 4 continues to hold even when all agents are required to use *dichotomous* utility functions. $[u_i(R) = 0 \vee u_i(R) = 1]$

Modular Domains

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

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

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 rational 1-deals (involving just one resource) suffice to guarantee outcomes with maximal social welfare.*

Sufficiency, Necessity, Maximality

- Theorem 5 says that the class of modular utility functions is *sufficient* for successful 1-deal negotiation.
- Is it also *necessary*?
  Answer: No. It's easy to construct examples.
- Is there *any* class of functions that is *sufficient* and *necessary*?
  Answer: No. Seems surprising at first, but for a proof it suffices to find two sufficient classes the union of which is not sufficient.
- As there can be no unique class of utility functions characterising all situations where 1-deal negotiation works, we have looked for *maximal* classes . . .

Maximality of Modular Utilities

We say that a class of utility functions $\mathcal{F}$ *permits 1-deal negotiation* iff any sequence of rational 1-deals will converge to a socially optimal allocation whenever all utility functions belong to $\mathcal{F}$.

Another surprising result:

Theorem 6 (Maximality) *Let $\mathcal{M}$ be the class of modular utility functions. Then for any class of utility functions $\mathcal{F}$ such that $\mathcal{M} \subset \mathcal{F}$, $\mathcal{F}$ does not permit 1-deal negotiation.*

▶ Are there other (interesting) classes of functions that are maximal?

Y. Chevaleyre, U. Endriss and N. Maudet. *On maximal classes of utility functions for efficient one-to-one negotiation*. IJCAI-2005.

Complexity Issues

Finding an optimal allocation is generally intractable. First observed for combinatorial auctions (easy reduction from SET PACKING):

Theorem 7 (Complexity) *Given $K \in \mathbb{Z}$, checking whether there exists an allocation A with $sw(A) > K$ is NP-complete.*

In the context of distributed negotiation schemes we may also ask other kinds of questions. Example:

- How hard is it to decide whether a given negotiation scenario permits 1-deal negotiation? (NP-hard, likely PSPACE-complete)

But neither of these are about the *process* of negotiation ...

M.H. Rothkopf, A. Pekeč, and R.M. Harstad. *Computationally manageable combinational auctions*. Management Science, 44(8):1131–1147, 1998.

P.E. Dunne, M. Wooldridge, and M. Laurence. *The complexity of contract negotiation*. Artificial Intelligence, 164(1–2):23–46, 2005.

Aspects of Complexity

(1) How many *deals* are required to reach an optimal allocation?

- communication complexity as number of individual deals

(2) How many *dialogue moves* are required to agree on one such deal?

- affects communication complexity as number of dialogue moves

(3) How expressive a *communication language* do we require?

- Minimum requirements: performatives *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)

U. Endriss and N. Maudet. *On the communication complexity of multilateral trading*. Journal of Auton. Agents and Multiagent Systems, 11(1):91–107, 2005.

Mechanism Design

Mechanism design is concerned with the *design of mechanisms* for collective decision making that favour particular outcomes despite agents pursuing their individual interests.

Mechanism design is sometimes referred to as *reverse game theory*. While game theory analyses the strategic behaviour of rational agents in a given game, mechanism design uses these insights to design games inducing certain strategies (and hence outcomes).

> Find out more tomorrow at Krzysztof's talk in the Computational Social Choice Seminar!

Preference Representation

- In *combinatorial domains*, preference representation is not trivial:
  - for instance, negotiation over n goods requires expressing preferences over 2^n bundles
  - *also:* multi-criteria decision making; voting for assemblies; ...
- Logic can help to design languages that allow for a *succinct* representation of preferences. For instance:
  - Model specific interests of agents as propositional formulas.
  - Associate each such *goal* with a *weight* (importance). Utility may then be defined as the sum of the weights of the satisfied goals (other option: *prioritised goals*).

Efficiency and Fairness

When assessing the quality of an allocation (or any other agreement) we can distinguish (at least) two types of indicators of social welfare.

Aspects of *efficiency* (*not* in the computational sense) include:

- The chosen agreement should be such that there is no alternative agreement that would be better for some and not worse for any of the other agents (*Pareto optimality*).
- If preferences are quantitative, the sum of all payoffs should be as high as possible (*utilitarianism*).

Aspects of *fairness* include:

- The agent that is going to be worst off should be as well off as possible (*egalitarianism*).
- No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own (*envy-freeness*).

Notions of Social Welfare

- *Utilitarian:* sum of utilities $sw_u(A) = \sum_{i \in \mathcal{A}} u_i(A)$
- *Nash product:* product of utilities $sw_N(A) = \prod_{i \in \mathcal{A}} u_i(A)$
- *Egalitarian:* utility of the weakest $sw_e(A) = min\{u_i(A) \,|\, i \in \mathcal{A}\}$
- *Elitist:* utility of the strongest $sw_{el}(A) = max\{u_i(A) \,|\, i \in \mathcal{A}\}$
- *Pareto optimality:* no other allocation is better for some agents without being worse for any of the others
- *Lorenz optimality:* the sum of utilities of the k weakest agents cannot be maintained for all and increased for some $k \leq |\mathcal{A}|$
- *Envy-freeness:* no agent would rather have the bundle allocated to one of the other agents $u_i(A(i)) \geq u_i(A(j))$
  - envy-free allocations are not always possible
  - could search for *envy-reducing* deals (for instance, with respect to the number of envious agents or the average degree of envy)

Conclusions

- Negotiation is a timely, exciting and fruitful area of research:
  - Knowledge transfer from economics to computer science and AI
  - Application of computational tools to problems in economics
  - Game theory is all over the place
  - Logic can and should play a bigger role
- Many open problems and scope for new directions of research
- For more information on the field in general, have a look at the MARA (Multiagent Resource Allocation) website:

  http://www.illc.uva.nl/~ulle/MARA/

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.