Fairness and Efficiency in Multiagent Resource Allocation

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

Computational Social Choice

- Research at the interface of
  - mathematical economics: social choice, game theory, decision theory
  - computer science and AI, multiagent systems, logic
- Some examples:
  - voting: computational hardness as a barrier against manipulation
  - preference representation in combinatorial domains
  - logic-based modelling of social choice procedures
  - multiagent resource allocation and fair division

Y. Chevaleyre, U. Endriss, J. Lang and N. Maudet. A Short Introduction to Computational Social Choice. Proc. SOFSEM-2007.

Talk Overview

- Introduction to Multiagent Resource Allocation
- Preliminaries on Distributed Negotiation Framework
- A couple of examples for so-called *convergence* results. Issues:
  - Under what circumstances can we hope that a system where agents negotiate autonomously and locally will converge to a state considered optimal from a global point of view?
  - How should we actually define such a notion of global optimality? $\leadsto$ "fairness and efficiency"

Multiagent Resource Allocation (MARA)

A tentative definition would be the following:

> *MARA is the process of distributing a number of items amongst a number of interested parties.*

What items? This talk is about the allocation of *indivisible goods*.

Some questions to think about:

- *How* are these items being distributed (allocation procedure)?
- *Why* are they being distributed? What's a "good" allocation?

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.

Choice of Allocation Procedure

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).

Distributed Negotiation Framework

- Set of *agents* $\mathcal{A} = \{1..n\}$ 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
- Each agent $i \in \mathcal{A}$ has got a *valuation function* $v_i : 2^{\mathcal{R}} \to \mathbb{R}$.
  Example: $v_i(A) = v_i(A(i)) = 577.8$ — 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 valuation. 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* $\$5$, 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 (IR) *A deal $\delta = (A, A')$ is called individually rational iff there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ 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 if it results in a gain in value (or money) that strictly outweighs a possible loss in money (or value).

The Global/Social Perspective

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

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

Example

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

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

Furthermore, suppose the initial allocation of goods 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* good 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* $\{chair, table\}$.

Negotiating Socially Optimal Allocations

So, under what circumstances can we hope that agents may be able to negotiate a socially optimal allocation?

If we do not impose any structural restrictions on deals, then we can get a strong convergence result:

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

How can we explain this positive result? . . .

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

Linking the Local and the Global Perspective

IR deals turn out to be exactly those deals that increase social welfare:

Lemma 2 (Individual rationality and social welfare) *A deal $\delta = (A, A')$ is IR iff $sw(A) < sw(A')$.*

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

Convergence then follows from the fact that the overall space of allocations is finite.

Remark: Lemma 2 also suggests that our notion of rationality is somehow "appropriate" for agents living in a utilitarian society ...

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

Efficiency and Fairness

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

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

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

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*).

Research Agenda

Theorem 1 has been an example (the first) for a convergence result. We would like to better understand under what circumstances convergence is possible for different scenarios. Parameters:

- The class of *valuation functions* considered:
  arbitrary or, say, only additive valuations
- The type of *structural restrictions* imposed on deals:
  general multilateral deals, bilateral deals, 1-resource deals, . . .
- The *rationality criterion* used to define individual agent behaviour:
  individual rationality or something else
- The *social welfare criterion* used to define the global goal:
  utilitarian or egalitarian social welfare, envy-freeness, . . .

It's particularly hard to achieve envy-freeness, because a local deal in one part of society can make another agent somewhere else envious . . .

Envy-free States

Unfortunately, there are cases where envy-free allocations do not *exist*.

Example: 2 agents, 1 good desired by both

We can try to circumvent this problem by taking the balance of past side payments into account when defining envy-freeness:

- Associate each allocation A with a payment *balance* $\pi : \mathcal{A} \to \mathbb{R}$, mapping agents to the sum of payments they have made so far.
- A *state* (A, π) is a pair of an allocation and a payment balance.
- Each agent $i \in \mathcal{A}$ has got a (quasi-linear) *utility function* $u_i : 2^{\mathcal{R}} \times \mathbb{R} \to \mathbb{R}$, defined as follows: $u_i(R, x) = v_i(R) - x$.
- A state (A, π) is *envy-free* iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $i, j \in \mathcal{A}$. An *efficient envy-free* (EEF) state is an envy-free state maximising utilitarian social welfare.

Envy-freeness and Individual Rationality

By a known result from social choice theory, EEF states always exist (in the presence of money). But we want to find them by means of *rational* negotiation. Unfortunately, this is generally *impossible*.

Example: 2 agents, 1 item r with $v_1(\{r\}) = 4$ and $v_2(\{r\}) = 7$. Agent 1 owns r to begin with; giving it to agent 2 would be efficient.

- An *IR deal* would require a payment within interval $(4, 7)$.
- But to ensure *envy-freeness*, the payment should be in $[2, 3.5]$.

Compromise: We shall enforce an *initial equitability payment* $\pi_0(i) = v_i(A_0) - sw(A_0)/n$ before beginning negotiation.

Globally Uniform Payments

Because of the "non-local effects of local deals" in view of envy-freeness, to have any chance of getting a convergence result for EEF states, we will have to restrict the freedom of agents a little by fixing a specific payment function (still IR!):

- Let $\delta = (A, A')$ be an IR deal. The payments as given by the *globally uniform payment function* (GUPF) are defined as follows: $p(i) \;=\; [v_i(A') - v_i(A)] \;-\; [sw(A') - sw(A)]/n$.

That is, we distribute the (positive!) social surplus to all agents.

Convergence to EEF States

A final restriction concerns agent valuations. *Supermodular* valuations are valuations satisfying the following condition for all $R_1, R_2 \subseteq \mathcal{R}$:

$$v(R_1 \cup R_2) \ \geq \ v(R_1) + v(R_2) - v(R_1 \cap R_2)$$

We are now ready to state the result:

Theorem 3 (Convergence) *If all valuations are supermodular and if initial equitability payments have been made, then any sequence of IR deals using the GUPF will eventually result in an EEF state.*

Generalisation: if agents are nodes on a *graph* and can only envy and negotiate with agents they a linked to, then we get a similar result.

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

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

Conclusions

- Examples of recent work in Multiagent Resource Allocation
- Technical results about *convergence* to socially optimal states
- Two issues that are special about this line of work:
  - *distributed* resource allocation via *multilateral deals*
  - consideration of *fairness* criteria, not just *efficiency*
- Many open questions and topics for future work . . .
- Papers are available from my website:

  http://www.illc.uva.nl/∼ulle/