Efficiency and Fairness in Distributed Resource Allocation

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

[joint work with Nicolas Maudet, Yann Chevaleyre,
Sylvia Estivie, Jérôme Lang, Fariba Sadri and Francesca Toni]

Talk Overview

I will start with some general remarks:

- Multiagent Resource Allocation
- Efficiency and Fairness
- Distributed Negotiation

Then I will present two concrete results in detail:

- Finding efficient (maxisum) allocations of resources by means of distributed negotiation amongst self-interested agents.
- Finding efficient and fair (envy-free) allocations of resources by means of distributed negotiation amongst self-interested agents.

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.

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*).

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?
- 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*.

This talk is about a particular distributed negotiation framework . . .

Resource Allocation by Negotiation

- 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.

Individual Rationality

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 *iff* it results in a gain in value (or money) that strictly outweighs a possible loss in money (or value).

We also call this the *local perspective* ...

Social Welfare

As for the *global perspective*, we first concentrate on *efficiency:*

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

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

Observe that there's no need to include the agents' *monetary balances* into this definition, because they'll 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.

▶ How well *are* we doing?

Example

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

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

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 individually rational).

The only possible deal would be to move the whole *set* $\{chair, table\}$.

Linking the Local and the Global Perspective

Lemma 1 (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. $\Box$

We can now prove a first result on negotiation processes:

Lemma 2 (Termination) *There can be no infinite sequence of IR deals; that is, negotiation must always terminate.*

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

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

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 IR deals will eventually result in an efficient allocation (with max. 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.

Related Results

This (positive) convergence result heavily relies on the *multilateral* character of the negotiation framework. Some related issues:

- *Deal restrictions:* Efficient outcomes cannot be guaranteed unless the negotiation protocol allows for deals involving *any number of agents* and *resources*.
- *Valuation restrictions:* If all valuations are *modular*, then *1-deals* (deals over one resource at a time) suffice.
- *Maximality:* For any class of valuations that strictly includes the modular functions, convergence cannot be guaranteed by 1-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.

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

Fairness

So far, the results have been all about *efficiency*. In fact, we have seen that the local criterion of *individual rationality* perfectly fits our global efficiency criterion of *maximal social welfare* (recall Lemma 1).

If we take the individual agent behaviour as a given (we do, today), then we cannot possibly hope to always achieve *fair* outcomes.

The remainder of this talk is about exploring how far we can get nevertheless, for one particular interpretation of fairness . . .

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Reaching envy-free states in distributed negotiation settings*. IJCAI-2007.

Envy-freeness

An allocation is called envy-free if nobody wants to change bundle with any of the others:

Definition 3 (EF allocations) *An allocation A is envy-free iff $v_i(A(i)) \geq v_i(A(j))$ for all agents $i, j \in \mathcal{A}$.*

If we require all goods to be allocated, then envy-free allocations may not always exist. Example: 2 agents, 1 good, liked by both agents

There has been some work on the computational complexity of checking whether a given scenario admits an envy-free solution: it's NP-hard in the simplest cases. So, it's pretty difficult.

But the above definition does not take money into account ...

S. Bouveret and J. Lang. *Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity*. IJCAI-2005.

Envy-freeness in the Presence of Money

We refine our negotiation framework as follows ...

- Associate each allocation A with a *balance function* $\pi : \mathcal{A} \to \mathbb{R}$, mapping agents to the sum of payments they've 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: $u_i(R, x) = v_i(R) - x$.

... and adapt the definition of envy-freeness:

Definition 4 (EF states) *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 EF state with an efficient allocation. Does such a state exist under all circumstances?

Existence of EEF States

Unlike for the case without money, EEF states always exist (Alkan *et al.*, 1991). There's a simple proof for supermodular valuations:

Theorem 4 (Existence of EEF states) *If all valuations are supermodular, then an EEF state always exists.*

Note that *supermodular* valuations are valuations satisfying the following condition for all bundles $R_1, R_2 \in \mathcal{R}$:

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

For ease of presentation (not technically required), from now on we assume that all valuations are *normalised:* $v(\{\,\}) = 0$.

A. Alkan, G. Demange and D. Gale. *Fair allocation of indivisible goods and criteria of justice*. Econometrica, 59(4):1023–1039, 1991.

Proof of Theorem 4

Of course, there's always an *efficient* allocation; let's call it A^*.

We'll try to fix a payment balance π^* such that (A^*, π^*) is EEF:

$$\pi^*(i) \quad = \quad v_i(A^*) \; - \; sw(A^*)/n$$

Note: the $\pi^*(i)$ add up to 0, so it's a *valid* payment balance. ✓

Now let $i, j \in \mathcal{A}$ be any two agents. As A^* is efficient, giving both $A^*(i)$ and $A^*(j)$ to i won't increase social welfare any further:

$$v_i(A^*(i)) \; + \; v_j(A^*(j)) \; \geq \; v_i(A^*(i) \cup A^*(j))$$

Now apply the supermodularity condition ... and rewrite:

$$\begin{aligned} v_i(A^*(i)) + v_j(A^*(j)) \; &\geq \; v_i(A^*(i)) + v_i(A^*(j)) \\ v_i(A^*(i)) - [v_i(A^*) - sw(A^*)/n] \; &\geq \; v_i(A^*(j)) - [v_j(A^*) - sw(A^*)/n] \\ u_i(A^*(i), \pi^*(i)) \; &\geq \; u_i(A^*(j), \pi^*(j)) \end{aligned}$$

That is, i does not envy j. Hence, (A^*, π^*) is *envy-free* (and EEF). ✓

Envy-freeness and Individual Rationality

Now that we know that EEF states always exist, we want to find them by means of *rational* negotiation.

Unfortunately, this is *impossible*. Example: 2 agents, 1 resource

$$v_1(\{r\}) = 4 \qquad v_2(\{r\}) = 7$$

Suppose agent 1 owns r to begin with.

The efficient allocation would be where agent 2 owns r.

An individually rational deal would require a payment within $(4, 7)$.

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

Realise just how unlikely it seems that our goal of guaranteeing EEF outcomes for distributed negotiation amongst self-interested (IR) agents could succeed ("non-local effects of local deals") . . .

We will have to restrict the freedom of agents a little by fixing a specific payment function (still IR!):

Definition 5 (GUPF) *Let $\delta = (A, A')$ be an IR deal. The payments as given by the globally uniform payment function are defined as:*

$$p(i) \;=\; [v_i(A') - v_i(A)] \;-\; [sw(A') - sw(A)]/n$$

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

Convergence in Supermodular Domains

After having deciphered all the acronyms, this should be rather surprising:

Theorem 5 (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.*

Proof: First try to show that this invariant holds for all states (A, π):

$$\pi(i) \; = \; v_i(A) - sw(A)/n \quad (*)$$

True initially by definition (initial equitability payments). Now let $\delta = (A, A')$ be a deal, with payment balances π and π'. Compute:

$$\begin{aligned} \pi'(i) \quad &= \quad \pi(i) + [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n \\ &= \quad v_i(A') - sw(A')/n \quad \leadsto \; (*) \text{ holds by induction} \end{aligned}$$

Theorem 3 shows that the system must *converge* to an *efficient* allocation A^* (whatever the payment function). Then the proof of Theorem 4 demonstrates that $(*)$ implies that (A^*, π^*) must be an *EEF* state. □

Conclusions

- Examples of recent work in Multiagent Resource Allocation
- Two issues that are special about this line of work:
  - *distributed* resource allocation via *multilateral deals*
  - consideration of *fairness* criteria, not just *efficiency*
- Technical results about *convergence* to socially optimal states:
  - convergence, restricted deals, restricted valuations, maximality
  - with respect to different (structural) classes of deals
  - with respect to different notions of social optimality
- Many open questions and topics for future work . . .
- Papers are available at my website:

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