Multiagent Resource Allocation: What to optimise, how, and why?

Ulle Endriss
Imperial College London

Talk Overview

This talk examines the following question:

- *What are the main parameters that characterise a system for Multiagent Resource Allocation (MARA)?*

I shall consider three issues in more detail:

- Choice of *allocation procedure*
- Choice of language for representing *agent preferences*
- Choice of overall *performance criteria* (social welfare)

Parameters

- Nature of *resources:*
  - Can resources be shared by several agents?
  - Are resources continuous, discrete or mixed (e.g. discrete goods and one continuous resource to model "money")?
  - If discrete, are they available in single or in multiple units?
- Nature of *agent preferences* (more later):
  - What do they depend on and how should they be represented?
- Choice of *performance criteria* (more later):
  - How do we assess the quality of allocations?
- Choice of *allocation procedure:*
  - Centralised (auctions) or distributed (local negotiation steps)?
  - If centralised, is the "auctioneer" a seller (auction), a buyer (reverse auction), or a matchmaker (combinatorial exchanges)?

Choice of Allocation Procedure

To date, most work in MARA has concentrated on *centralised* allocation procedures (auctions). Advantages:

- simple communication protocols
- well-studied by economists
- pushed by recent advances in algorithm design

In the *distributed* approach, allocations evolve as a consequence of *local* negotiation steps. Advantages:

- potential to distribute computational burden
- trust in the "auctioneer"?
- seems more natural in cases with initial and/or evolving allocations
- strict interpretation of the MAS paradigm

Correspondences

Combinatorial auctions	**Distributed negotiation**
Bidders submitting (several) bids	agents with utility functions
Bidding language	representation of utilities
Revenue for the auctioneer	sum of individual utilities
Winner determination problem	finding an “optimal” allocation
One large computational effort	local negotiation
(Usually) free disposal	no free disposal (depends on agents)
No initial allocation	initial allocation

Choice of Preference Representation

- Agent preferences: *ordinal* relations or *cardinal* utility functions?
- Languages for representing preferences:
  - decision-theoretic or logic-based ($\leadsto$ see talk by Jérôme Lang)
  - utility functions or bidding languages (more later)
- *Expressiveness* versus *succinctness* of representing preferences
  - more later ($\leadsto$ see also talk by Jérôme Lang)
- Do we only model preferences over *bundles* or over entire resource *allocations*? Examples for such *externalities* include:
  - Envy ($\leadsto$ see talk by Sylvain Bouveret)
  - Also resource-dependent: in shared networks, the payoff depends on the number of agents accessing the same resource.
- Strategic considerations: do agents report their preferences truthfully and how does this affect the design of the system?

Expressiveness and Succinctness

- Generally, the more expressive a language the better.
- But: if other factors prevent us from fully exploiting such expressive power, then "more is better" may not be true ($\rightsquigarrow$ see talk by Yann Chevaleyre)
- Succinctness is particularly important in combinatorial domains such as multiagent resource allocation.

Alternative Representation of Utility Functions

- Problem: The *"bundle form"* of representing utility functions can be problematic if there are too many bundles with non-zero values.
- A utility function is called *k-additive* iff the utility assigned to a bundle R can be represented as the sum of basic utilities assigned to subsets of R with cardinality $\leq k$ (*limited synergies*).
- The *k-additive form* of representing utility functions:

$$u(R) \quad = \quad \sum_{T \subseteq R} \alpha^T \quad \text{with } \alpha^T = 0 \text{ whenever } |T| > k$$

  Example: $u = 3.r_1 + 7.r_2 - 2.r_2.r_3$ is a 2-additive function

- Note that *any* utility function is representable as a k-additive function for some $k \leq |\mathcal{R}|$.

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Multiagent resource allocation with k-additive utility functions*. DIMACS-LAMSADE Workshop 2004.

Separation Results

Proposition 1 (Efficiency of the k-additive form) *The bundle form cannot polynomially simulate the k-additive form.*

Proof. Consider the utility function $u(R) = |R|$. Representing u requires $|\mathcal{R}|$ non-zero coefficients in the k-additive form (*linear*), but $2^{|\mathcal{R}|}-1$ non-zero values in the bundle form (*exponential*). □

Proposition 2 (Efficiency of the bundle form) *The k-additive form cannot polynomially simulate the bundle form.*

Proof. Consider the utility function $u(R) = \begin{cases} 1 & \text{if } |R| = 1 \\ 0 & \text{otherwise} \end{cases}$

Requires $|\mathcal{R}|$ non-zero values in the bundle form (*linear*), but $2^{|\mathcal{R}|}-1$ non-zero coefficients in the k-additive form (*exponential*): namely $\alpha^T = 1$ for $|T| = 1$, $\alpha^T = -2$ for $|T| = 2$, $\alpha^T = 3$ for $|T| = 3$, ... □

Adding Negation

Hence, neither bundle nor k-additive form are strictly more succinct in general (although the k-additive form seems more useful in practice).

▶ The *k-additive form with negation* of representing utility functions:

$$u(R) \;=\; \sum_{P \subseteq R} \sum_{N \subseteq \mathcal{R} \setminus R} \alpha^{(P,N)} \quad \text{with } \alpha^{(P,N)} = 0 \text{ whenever } |P \cup N| > k$$

Clearly,

- the bundle form cannot polynomially simulate the k-additive form with negation either; and
- the k-additive form with negation form *can* polynomially simulate the k-additive form.

To see this, set $N = \{\,\}$ (in both cases).

More Separation Results

The following propositions show that adding negation makes the representation of utility functions *strictly* more succinct:

Proposition 3 (Efficiency of adding negation) *The k-additive form cannot polynomially simulate the k-additive form with negation.*

Proof. Consider the utility function u with $u(\{\,\}) = 1$ and $u(R) = 0$ for $R \neq \{\,\}$. Requires only a single non-zero coefficient if negation is available, namely $\alpha^{(\{\,\},\mathcal{R})} = 1$, but $2^{|\mathcal{R}|}$ non-zero coefficients in the k-additive form without negation, namely $\alpha^T = (-1)^{|T|}$. □

Proposition 4 (Simulation of the bundle form) *The k-additive form with negation can polynomially simulate the bundle form.*

Proof. Let u be any utility function given in bundle form. Now define $\alpha^{(T,\mathcal{R}\setminus T)} := u(T)$ for all bundles T with $u(T) \neq 0$ and set all other coefficients to 0. These coefficients define the same function u. □

Utility Functions and Bidding Languages

In combinatorial auctions, agents report their preferences (which may be distorted by strategic considerations) through bids. Different *bidding languages* correspond to different classes of utility functions:

- The *XOR-language* corresponds to the bundle form:
  - can specify prices for different (mutually exclusive) bundles
  - fully expressive (which is not the case for all bidding languages)
  - not very succinct (as we have seen)
- The *OR-language* is the "standard" bidding language:
  - to specify prices for (non-exclusive) bundles
  - not fully expressive
  - does not correspond to a natural class of utility functions
- Languages corresponding to the k-additive form (with negation):
  - yet to be explored by auction designers

System Performance

- How can we measure the *performance* of a MARA system? (performance as in quality of the final allocation, not about speed)
- Example: revenue for the auctioneer in combinatorial auctions
- In the case of distributed negotiation (without a central authority) the level of performance should depend on all agents.
- Multiagent systems are often described as *"societies of agents"*. This suggests to use tools from microeconomics and social choice theory to assess the performance of the overall system ("society").

Social Welfare

A *social welfare ordering* formalises the notion of a society's "preferences" given the preferences of its members (the agents).

- The *utilitarian* social welfare $sw_u(A)$ of an allocation of resources A is defined as follows:

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

  That is, anything that increases average (and thereby overall) utility is taken to be socially beneficial.

- In the *egalitarian* approach, on the other hand, social welfare is tied to the welfare of society's weakest member:

$$sw_e(A) \quad = \quad \min\{u_i(A) \mid i \in \mathcal{A}gents\}$$

Utilitarianism versus Egalitarianism

- In the MAS literature the utilitarian viewpoint (that is, social welfare = sum of individual utilities) is usually taken for granted.
- In philosophy/sociology/economics not.
- John Rawls' *"veil of ignorance"* (*A Theory of Justice*, 1971):

> *Without knowing what your position in society (class, race, sex, . . .) will be, what kind of society would you choose to live in?*

- Reformulating the *veil of ignorance for multiagent systems:*

> *If you were to send a software agent into an artificial society to negotiate on your behalf, what would you consider acceptable principles for that society to operate by?*

- Conclusion: worthwhile to investigate egalitarian (and other) social principles also in the context of multiagent systems.

Other Egalitarian Approaches

- Every allocation A gives rise to an *ordered utility vector $\vec{u}(A)$*: compute $u_i(A)$ for all agents i and rearrange in ascending order.

  Example: $\vec{u}(A) = \langle 0, 5, 20 \rangle$ means that the weakest agent enjoys utility 0, the strongest utility 20, and the middle one utility 5.

- The *leximin-ordering* $\prec$ over allocations is defined as follows:

$$A \prec A' \quad \text{iff} \quad \vec{u}(A) \text{ lexically precedes } \vec{u}(A')$$

  Example: $A \prec A'$ for $\vec{u}(A) = \langle 0, 6, 7, 29 \rangle$ and $\vec{u}(A') = \langle 0, 6, 9, 25 \rangle$

- *Kalai-Smorodinsky (KS) solution* (relative egalitarian solution):
  - Let $u_i^{max} = \max\{u_i(A) \mid A \in Allocations\}$ for each agent i.
  - The KS solution is defined as the maximum of the leximin-ordering with respect to $(\frac{u_i(A)}{u_i^{max}})$.

Further Notions of Social Welfare

- *Pareto optimality:* no other allocation is better for some agents without being worse for others
- *Lorenz optimality:* the sum of utilities of the k weakest agents cannot be maintained for all and increased for some $k \leq |\mathcal{A}gents|$
- *Nash product:* product of utilities $sw_N(A) = \prod_{i \in \mathcal{A}gents} u_i(A)$
- We have also proposed a notion of *elitist* social welfare $sw_{el}(A)$:

$$sw_{el}(A) = \max\{u_i(A) \mid i \in \mathcal{A}gents\}$$

- Remark: In some cases it may be more appropriate to use $u_i(A) - u_i(A_{init})$ instead of $u_i(A)$ for some of the notions of social welfare discussed (or we could normalise utility functions such that $u_i(A_{init}) = 0$).

Constraints on Allocations

For some applications, we may want to restrict the range of allocations that can be chosen. Examples:

- We could restrict the allocation space to those allocations that Pareto-dominate the initial allocation:
  - non-negative utility functions
  - easier to justify the enforcement of an egalitarian rule
- In the case of reverse combinatorial auctions, the auctioneer may have constraints such as not to buy all products of a certain type from the same supplier, even when that would be cheaper. ($\leadsto$ see talk by Juan Rodríguez)

Envy

- An allocation is called *envy-free* iff no agent would rather have one of the bundles allocated to any of the other agents:

$$u_i(A(i)) \quad \geq \quad u_i(A(j))$$

  Note that envy-free allocations do not always exist.

- As we cannot always ensure envy-free allocations, maybe we should aim at *reducing* envy as far as possible.

- What would be a reasonable definition of *minimal envy*?
  - minimise the number of envious agents
  - minimise the average degree of envy (distance to the most envied competitor) of all envious agents

Welfare Engineering

- Choice (and possibly design) of *social welfare orderings* that are appropriate for specific agent-based applications.
  - Example: The *elitist* collective utility function sw_{el} seems unethical for human society, but may be appropriate for a distributed application where each agent gets the same task.
  - Slogan: "welfare economics for *artificial* agent societies"
- Design of suitable *rationality criteria* for agents participating in negotiation in view of different notions of social welfare.
  - Example: To achieve *Lorenz optimal* allocations in *0-1 domains without money*, ask agents to negotiate *cooperatively rational* or *inequality-reducing* deals over *one resource* at a time.
  - Slogan: "*inverse* welfare economics" ($\leadsto$ mechanism design)

U. Endriss and N. Maudet. *Welfare engineering in multiagent systems*. ESAW-2003.

Criteria for Social Welfare Choice

We have tried to identify criteria that determine what social welfare ordering is appropriate for which application (work in progress):

- What does the income of the system provider depend on?
  - *Utility-dependent* ("tax on gain") $\leadsto$ utilitarian
  - *Membership-dependent* ("joining fee") $\leadsto$ "fair" approach
  - *Transaction-dependent* ("pay as you go") $\leadsto$ not clear (but note the connections to communication complexity)
- Can agents join or leave the society *during* negotiation?
  Yes: review definitions (e.g. utilitarian welfare as average utility)
- Can agents participate in *more than one* negotiation?
  Yes: strong point for fair approaches (egalitarian, envy-reducing)

Y. Chevaleyre, U. Endriss, S. Estivie and N. Maudet. *Welfare engineering in practice: On the variety of multiagent resource allocation problems*. ESAW-2004.

Conclusion

- I have discussed some of the design parameters in MARA, giving particular consideration to three important issues:
  - the choice of allocation procedure (centralised or distributed)
  - the representation of agent preferences (succinctness)
  - the choice of suitable social welfare measures to assess overall system performance
- I think an interesting question to consider would be:

  Is it possible to give a (reasonably) general definition of "MARA system" and to derive any concrete system by instantiating the relevant design parameters?