Computational Social Choice: Spring 2008

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

What is MARA?

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.*

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

Outline

- Concerning the *specification* of MARA problems:
  - Overview of different *types of resources*
  - Representation of the *preferences* of individual agents (done)
  - Notions of *social welfare* to specify the quality of an allocation (partly done already)
- Concerning methods for *solving* MARA problems:
  - Discussion of *allocation procedures* ($\leadsto$ future lectures)
  - Some *complexity results* concerning allocation procedures

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.

Types of Resources

- A central parameter in any resource allocation problem is the nature of the resources themselves.
- There is a whole range of different *types of resources*, and each of them may require different techniques . . .
- Distinguish properties of the *resources* themselves and characteristics of the chosen *allocation mechanism*. Examples:
  - Resource-inherent property: Is the resource perishable?
  - Characteristic of the allocation mechanism: Can the resource be shared amongst several agents?

Continuous vs. Discrete Resources

- Resources may be *continuous* (e.g. energy) or *discrete* (e.g. fruit).
- *Discrete* resources are *indivisible*; *continuous* resources may be treated either as (infinitely) *divisible* or as *indivisible* (e.g. only sell orange juice in units of 50 litres $\rightsquigarrow$ *discretisation*).
- *Representation* of a single bundle:
  - Several continuous resources: vector over non-negative reals
  - Several discrete resources: vector over non-negative integers
  - Several distinguishable discrete resources: vector over $\{0, 1\}$
- Classical literature in economics mostly concentrates on a single continuous resource; recent work in AI and Computer Science focusses on discrete resources.

Divisible or not

- Resources may be treated as being *divisible* or *indivisible*.
- Continuous/discrete: *physical property* of resources
  Divisible/indivisible: feature of the *allocation mechanism*

Sharable or not

- A *sharable* resource can be allocated to a number of different agents at the same time. Examples:
  - a photo taken by an earth observation satellite
  - path in a network (network routing)
- More often though, resources are assumed to be *non-sharable* and can only have a single owner at a time. Examples:
  - energy to power a specific device
  - fruit to be eaten by the agent obtaining it

Static or not

Resources that do not change their properties during a negotiation process are called *static* resources. There are at least two types of resources that are *not* static:

- *consumable* goods such as fuel
- *perishable* goods such as food

In general, resources cannot be assumed to be static. However, in many cases it is reasonable to assume that they are as far as the negotiation process at hand is concerned.

Single-unit vs. Multi-unit

- In *single-unit* settings there is exactly one copy of each type of good; all items are distinguishable (e.g. several houses).
- In *multi-unit* settings there may be several copies of the same type of good (e.g. 10 bottles of wine).
- Note that this distinction is only a matter or *representation:*
  - Every multi-unit problem can be translated into a single-unit problem by introducing new names for the items (inefficient, but possible).
  - Every single-unit problem is in fact also a (degenerate) multi-unit problem.
- Multi-unit problems allow for a more *compact* representation of allocations and preferences, but also require a richer *language* (variables ranging over integers, not just binary values).

Resources vs. Tasks

- *Tasks* may be considered resources with *negative utility*.
- Hence, *task allocation* may be regarded a MARA problem.
- However, tasks are often coupled with *constraints* regarding their coherent combination (timing).

Remark: From now on (for this and the next lecture), we are going to deal with the allocation of static indivisible resources that are available in single units and that cannot be shared . . .

Setting

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 $\mathcal{A}$.

Each agent $i \in \mathcal{A}$ has got a *utility function* $u_i : 2^{\mathcal{R}} \to \mathbb{R}$. We usually write $u_i(A)$ as a shorthand for $u_i(A(i))$.

We shall sometimes refer to the preference relation $\preceq_i$ induced by the utility function u_i: $R \preceq_i R'$ iff $u_i(R) \leq u_i(R')$.

Remark: MARA with indivisible resources is a prime example for a combinatorial domain. We have seen how to represent agent preferences in such domains in earlier lectures.

Social Welfare

An important parameter in the specification of a MARA problem concerns our goals: what kind of allocation do we want to achieve?

- Success may depend on a single factor (e.g. revenue of an auctioneer), but more often on an *aggregation of preferences* of the individual agents in the system.
- Concepts from social choice theory and welfare economics can be useful here ("multiagent systems as *societies of agents*").
- Here we use the term *social welfare* in a broad sense, to describe the quality of an allocation in view of a suitable aggregation of the individual agent preferences.

Pareto optimality is probably the most basic criterion for social optimality, but there are many others . . .

Social Welfare Orderings

A collective utility function is a function $W : \mathbb{R}^n \to \mathbb{R}$ mapping utility vectors to the reals. Here we define them over allocations A (inducing utility vectors):

- The *utilitarian* social welfare is defined as the sum of utilities:

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

- The *egalitarian* social welfare is given by the utility of the agent that is currently worst off:

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

- The *Nash product* is the product of the individual utilities:

$$sw_N(A) \quad = \quad \prod_{i \in \mathcal{A}gents} u_i(A)$$

Social Welfare Orderings (cont.)

- The *elitist* social welfare is given by the utility of the agent that is currently best off:

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

- Let $\vec{u}_A$ be the *ordered utility vector* induced by allocation A. Then the *k-rank dictator* CUF sw_k is defined as follows:

$$sw_k(A) \quad = \quad (\vec{u}_A)_k$$

  Recall that sw_k is the same as the egalitarian CUF for $k = 1$ and the same as the elitist CUF for $k = n$ (number of agents).

- The *leximin-ordering* $\preceq_\ell$ is a social welfare ordering that may be regarded as a refinement of the egalitarian CUF:

  $A \preceq_\ell A' \;\Leftrightarrow\; \vec{u}_A$ lexically precedes $\vec{u}_{A'}$ (not necessarily strictly)

Normalised Utility

It can be useful to *normalise* utility functions before aggregation:

- If A_0 is the initial allocation, then we may restrict attention to allocations A that Pareto-dominate A_0 and use *utility gains* $u_i(A) - u_i(A_0)$ rather than $u_i(A)$ as problem input.

- We could evaluate an agent's utility gains *relative* to the gains it could expect in the best possible case. Define an agent's *maximum utility* wrt. a set Adm of admissible allocations:

$$\widehat{u}_i \quad = \quad \max\{u_i(A) \mid A \in Adm\}$$

  Then define the *normalised* individual utility of agent i:

$$u'_i(A) \quad = \quad \frac{u_i(A)}{\widehat{u}_i}$$

  The optimum of the *leximin*-ordering wrt. normalised utilities is known as the *Kalai-Smorodinsky solution*.

Ordered Weighted Averaging

We can build families of parametrised CUFs that induce several SWOs. An example are the *ordered weighted averaging operators*.

Let $w = \langle w_1, w_2, \ldots, w_n \rangle$ be a vector of real numbers. Define:

$$sw_w(A) \quad = \quad \sum_{i \in \mathcal{A}gents} w_i \cdot \vec{u}(A)_i$$

This generalises several other SWOs:

- If w is such that $w_i = 0$ for all $i \neq k$ and $w_k = 1$, then we have exactly the k-rank dictator CUF.
- If $w_i = 1$ for all i, then we obtain the utilitarian CUF.
- If $w_i = \alpha^{i-1}$, with $\alpha > 0$, then the *leximin*-ordering is the limit of the SWO induced by sw_w as α goes to 0.

Envy-Freeness

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

$$A(i) \quad \succeq_i \quad A(j)$$

Recall that $A(i)$ is the bundle allocated to agent i in allocation A.

Note that envy-free allocations do not always *exist* (at least not if we require either complete or Pareto optimal allocations).

Example

Consider the following example with two agents and three goods: $\mathcal{A} = \{1, 2\}$ and $\mathcal{R} = \{a, b, c\}$. Suppose utility functions are additive:

$$u_1(\{a\}) = 18 \quad u_1(\{b\}) = 12 \quad u_1(\{c\}) = 8$$
$$u_2(\{a\}) = 15 \quad u_2(\{b\}) = 8 \quad u_2(\{c\}) = 12$$

Let A be the allocation giving a to agent 1 and b and c to agent 2.

- A has maximal *egalitarian* social welfare (18); *utilitarian* social welfare is not maximal (38 rather than 42); and neither is *elitist* social welfare (20 rather than 38).
- A is *Pareto optimal* and *leximin-optimal*, but not *envy-free*.
- There is no allocation that would be both Pareto optimal *and* envy-free. But if we change $u_1(\{a\}) = 20$ (from 18), then A becomes Pareto optimal and envy free.

Degrees of Envy

As we cannot always ensure envy-free allocations, another approach would be to try to *reduce* envy as much as possible.

But what does that actually mean?

A possible approach to systematically defining different ways of measuring the *degree of envy* of an allocation:

- Envy between two agents:
  $\max\{u_i(A(j)) - u_i(A(i)), 0\}$ [or even without max]
- Degree of envy of a single agent:
  0-1, max, sum
- Degree of envy of a society:
  max, sum [or indeed any SWO/CUF]

Allocation Procedures and Complexity

- We have now seen the various components that are needed to *specify* a MARA problem (type of resource, agent preferences, optimality criterion).
- In the next lecture we are going to see how agents can negotiate optimal allocations in a *distributed* manner and later on in the course we are going to see a *centralised* allocation procedure (combinatorial auctions).
- Now we are going to look into the *computational complexity* of the problem of finding an optimal allocation, independently from any specific allocation procedure.

Resource Allocation Problems

For the purpose of formally stating the resource allocation problems for which we want to analyse the complexity, let a resource allocation setting $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$ be given by:

- $\mathcal{A} = \{1, 2 \ldots, n\}$ is a set of n agents;
- $\mathcal{R} = \{r_1, r_2, \ldots, r_m\}$ is a collection of m resources; and
- $\mathcal{U} = \{u_1, u_2, \ldots, u_n\}$ describes the utility function $u_i : 2^{\mathcal{R}} \to \mathbb{Q}$ for the agent $i \in \mathcal{A}$.

The set of *allocations* A is the set of partitionings of $\mathcal{R}$ amongst $\mathcal{A}$ (or equivalently, the set of total functions from $\mathcal{R}$ to $\mathcal{A}$).

Welfare Optimisation

How hard is it to find an allocation with maximal social welfare? Rephrase this *optimisation problem* as a *decision problem:*

> WELFARE OPTIMISATION (WO)
> **Instance:** $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$; $K \in \mathbb{Q}$
> **Question:** Is there an allocation A such that $sw_u(A) > K$?

Unfortunately, the problem is intractable:

Theorem 1 WELFARE OPTIMISATION *is NP-complete.*

The proof (following slides) uses a reduction from a standard reference problem (SET PACKING) known to be NP-complete.

In the context of MARA, this kind of result seems to have first been stated by Rothkopf *et al.* (1998).

M.H. Rothkopf, A. Pekeč, and R.M. Harstad. Computationally Manageable Combinational Auctions. *Management Science*, 44(8):1131–1147, 1998.

Proof of NP-hardness

We are going to reduce our problem to SET PACKING, one of the standard problems known to be NP-complete:

> SET PACKING
> **Instance:** Collection $\mathcal{C}$ of finite sets and $K \in \mathbb{Q}$
> **Question:** Is there a collection of disjoint sets $\mathcal{C}' \subseteq \mathcal{C}$ s.t. $|\mathcal{C}'| > K$?

Given an instance $\mathcal{C}$ of SET PACKING, consider this MARA setting:

- Resources: each item in one of the sets in $\mathcal{C}$ is a resource
- Agents: one for each set in $\mathcal{C}$ + one other agent (called 0)
- Utilities: $u_C(R) = 1$ if $R = C$ and $u_C(R) = 0$ otherwise; $u_0(R) = 0$ for all bundles R

That is, every agent values "its" bundle at 1 and every other bundle at 0. Agent 0 values all bundles at 0.

Proof of NP-hardness (cont.)

Observe that not every allocation immediately corresponds to a valid solution of SET PACKING: the bundles owned by individual agents may not all be sets in $\mathcal{C}$.

But: for every given allocation there exists an(other) allocation with equal social welfare that does directly correspond to a valid solution for SET PACKING — just assign any goods owned by an agent with utility 0 to agent 0 (this reallocation does not affect social welfare). Note that social welfare is equal to $|\mathcal{C}'|$.

Hence, any algorithm for WO can also solve SET PACKING problems; so WO must be at least NP-hard. ✓

Proof of Membership in NP

This part is in fact very easy ...

Recall that a problem belongs to NP if it is possible to verify the correctness of a candidate solution in polynomial time.

This is clearly the case here: Given an allocation A, we can compute $sw_u(A)$ in polynomial time. And A constitutes a correct solution iff $sw_u(A) > K$. ✓

Remarks

- To be precise, we have proved NP-hardness wrt. *the number of pairs of agents and bundles with non-zero utility*. This corresponds to the number of sets involved in Set Packing.
- Observe that this number itself may already be very high (exponential in the number of goods).
- In other words, we have proved NP-completeness wrt. the *explicit form* of representing utility functions.

Representation Issues

- As for all complexity results, the *representation* of the input problem is crucial: if the input is represented inefficiently (e.g. using exponential space when this is not required), then complexity results (expressed with respect to the size of the input) may seem much more favourable than they really are.
- NP-completeness of Welfare Optimisation has been shown with respect to several *representations of utilities* (such as the k-additive form).
- In the sequel, the focus is on demonstrating *what questions* people have been asking rather than on exact results.

  Therefore, we do not give details regarding the representation (but most results apply to a variety of representation forms).

Welfare Improvement

The following problem is also NP-complete:

> Welfare Improvement (WI)
> **Instance:** $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$; allocation A
> **Question:** Is there an allocation A' such that $sw_u(A) < sw_u(A')$?

Given the close connection to Welfare Optimisation, this is not very surprising.

Pareto Optimality

A decision problem is said to be in coNP iff its complementary problem ("is it *not* the case that . . . ") is in NP.

Checking whether a given allocation is Pareto optimal is an example for a coNP-complete decision problem:

> PARETO OPTIMALITY (PO)
> **Instance:** $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$; allocation A
> **Question:** Is A Pareto optimal?

Envy-Freeness

Checking whether a given setting admits an envy-free allocation (if all goods need to be allocated) is again NP-complete:

> ENVY-FREENESS (EF)
> **Instance:** $\langle \mathcal{A}, \mathcal{R}, \mathcal{U} \rangle$
> **Question:** Is there a (complete) allocation A that is envy-free?

Checking whether there is an allocation that is both Pareto optimal and envy-free is even harder: Σ_2^p-complete (NP with NP oracle).

S. Bouveret and J. Lang. Efficiency and Envy-freeness in Fair Division of Indivisible Goods: Logical Representation and Complexity. Proc. IJCAI-2005.

Summary

We have given a first overview of the MARA research area . . .

- Specifying a MARA problem requires fixing at least the following parameters: *type of resource*, *agent preferences*, and *social welfare* or similar concept used to define global aims
- Decision problems arising in MARA are often intractable.
  - Successful *algorithm design* is still possible, but ad-hoc methods or brute-force algorithms won't work.
  - Sometimes negative complexity results can be circumvented by imposing *restrictions* (say, on utility functions).
- More details can be found in the MARA Survey.

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.

What next?

Next we are going to review different types of *allocation procedures* for finding a good allocation of resources to agents:

- *Distributed allocation procedures:* In distributed MARA the emphasis is on understanding under what circumstances we can expect socially optimal allocations to emerge when autonomous agents negotiate a sequence of local deals.
- *Cake-cutting procedures:* The classical literature on fair division has mostly addressed the case of a single divisible good (a.k.a. a "cake"). Here the challenge is to devise an interactive procedure that will guarantee, say, envy-free outcomes.
- *Combinatorial auctions:* Finding an allocation that maximises utilitarian social welfare is equivalent to determining the winners in a CA. We will look into optimisation algorithms for CAs and also discuss strategic (game-theoretical) questions.