Multiagent Resource Allocation

Michal Jakob
Artificial Intelligence Center,
Dept. of Computer Science and Engineering,
FEE, Czech Technical University

AE4M36MAS Autumn 2019 - Lecture 13

Motivating Example

Taxi is a limited resources

The need for using the taxi varies among passengers and can be depend on the time of the ride

10:00 slot: Passenger X?
10:30 slot: Passenger X?
...

...

Preferences over times, prices (and taxis)

Who should get the taxi and when (and possibly at which price)?

...

...

Passenger 2

Passenger 3

Passenger 4

Multiagent Resource Allocation (MARA)

What is Multiagent Resource Allocation?

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?
- **Why** are they being distributed?

Classification of MARA

1. **Resources** (What)
2. Agent (i.e. individual) **preferences** (Why)
3. Social (i.e. collective) **welfare** (Why)
4. Allocation **mechanism** (How)

Link to **social choice**: allocations are alternatives over which agents express their preferences.

Link to **game theory**: allocation mechanisms are games (that needs to be designed and for which strategies can be studied).

Type of Resources*

Central parameter in any resource allocation problem.

Different **types** of resources may require different resource allocation **techniques**.

Inherent **properties** of the **resource** vs. **characteristics** of the chosen **mechanism**.

* also termed **goods**

Types of Resources

Continuous vs. **Discrete**
Continuous resource can be arbitrarily divided.

Divisible vs. **Indivisible**
Discrete resources indivisible; continuous can be treated either way.

Sharable vs. **Non-Sharable**
Sharable can be assigned multiple times: e.g. a path in a network.

Static vs. **Non-Static**
static = properties do not change; non-static = properties do change e.g. perishable goods.

Single-Unit vs. **Multi-Unit**
One copy vs. multiple copies (ten trucks of the same type).

Resources vs. Tasks

Tasks may be considered resources with **negative** utility (cost).

Task allocation may be regarded a multiagent resource allocation problem.

- However, tasks are often coupled with **constraints** regarding their **coherent combination** (timing and ordering).

Preference Representation

Preference Representation

Agents may have **preferences** over

- the bundle of resources they receive
- the bundles of resources received by others (**externalities**)

What are suitable **languages** for representing agent **preferences**?

Notation

Set of **agents** $\mathcal{A} = \{1, \dots, n\}$

Set of **resources** $\mathcal{R}$

Agents have **preferences over allocations** $X \in \mathcal{X}$

Allocation X is a *partial* mapping of $\mathcal{R}$ to $\mathcal{A}$ (not all resources need to be allocated)

Preference Representation Languages

Expressive power
Can the chosen language encode all the preference structures we are interested in?

Succinctness
Is the representation of (typical) preference structures succinct? Is one language more succinct than the other?

Complexity
What is the computational complexity of related decision problems, such as comparing two alternatives?

Cognitive relevance
How close is a given language to the way in which humans would express their preferences?

Elicitation
How dicult is it to elicit the preferences of an agent so as to represent them in the chosen language?

Cardinal vs. Ordinal Preferences

A **preference structure** represents an agent's preferences over allocations $X \in \mathcal{X}$.

Cardinal preferences	Ordinal preferences
Cardinal preference structure is a function $u\colon \mathcal{X} \mapsto Val$, where Val is usually a set of numerical values such as $\mathbb{N}$ or $\mathbb{R}$ (and typically non-negative)	Ordinal preference structure is a *binary relation* $\preccurlyeq$ over the set of alternatives, that is *reflexive* and *transitive* (and connected).

If the alternatives over which agents have to express preferences are *bundles of indivisible resources* from the set $\mathcal{R}$, then we have $\mathcal{X} = 2^{\mathcal{R}}$.

Example

Hanging a picture with a **frame** (f), a **hammer** (h) and a **nail** (n)

Cardinal

X	$u(X)$
$\{\}$	0
$\{f\}$	0
$\{h\}$	0
$\{n\}$	10
$\{f,h\}$	0
$\{f,n\}$	20
$\{h,n\}$	15
$\{f,h,n\}$	50

Ordinal

$\succcurlyeq$	$\{\}$	$\{f\}$	$\{h\}$	$\{n\}$	$\{f,h\}$	$\{f,n\}$	$\{h,n\}$	$\{f,h,n\}$
$\{\}$	1	0	0	0	0	0	0	0
$\{f\}$	1	1	1	0	1	0	0	0
$\{h\}$	1	1	1	0	1	0	0	0
$\{n\}$	1	1	1	1	1	0	0	0
$\{f,h\}$	1	1	1	0	1	0	0	0
$\{f,n\}$	1	1	1	1	1	1	1	0
$\{h,n\}$	1	1	1	1	1	0	1	0
$\{f,h,n\}$	1	1	1	1	1	1	1	1

Cardinal can always be translated to ordinal.
Ordinal cannot be always translated to cardinal.

Preferences Properties

	Cardinal	Ordinal
Intrapersonal comparison	yes	Yes
Interpersonal comparison ("Ann likes x more than Bob likes y")	yes	No
Preference intensity	yes	No
Cognitive relevance	lower	higher
Explicit representation	$\mathcal{O}(\lvert\mathcal{X}\rvert)$	$\mathcal{O}(\lvert\mathcal{X}\rvert^2)$

Representation can be an issue ➔ compact representations

Social Welfare

Social Welfare

A third parameter in the specification of a MARA problem concerns our goals: **What kind of allocation do we want to achieve?**

We use the term **social welfare** in a very broad sense to describe **metrics** for assessing the **quality** of an **allocation** of resources.

Efficiency and Fairness

Two key indicators of social welfare.

Aspects of **efficiency*** 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:

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

*not in the computational sense

Utilitarian Social Welfare

Utilitarian Social Welfare

The **utilitarian** social welfare function (also called collective utility function) sw_u is defined as the sum of individual utilities:

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

Maximizing utilitarian CUF improves **efficiency**.

The utilitarian CUF is **zero-independent**: adding a constant value to your utility function won't a affect social welfare judgements.

Egalitarian Social Welfare

Egalitarian Social Welfare

The **egalitarian** social welfare function sw_e is defined as the sum of individual utilities:

$$sw_e(X) = \min_{i \in \mathcal{A}} u_i(X)$$

Maximising this function amounts to improving the situation of the weakest members of society (→ **fairness**).

Nash Product Social Welfare

Nash Social Welfare

The **Nash** social welfare function sw_e is defined as the sum of individual utilities:

$$sw_e(X) = \prod_{i \in \mathcal{A}} u_i(X)$$

This is a useful measure of social welfare as long as all utility functions can be assumed to be **positive**.

Nash CUF favours increases in overall utility, but also inequality-reducing redistributions $(2 \cdot 6 < 4 \cdot 4)$ ➔ **proportional fairness**.

The Nash CUF is **scale independent**: whether a particular agent measures their own utility in euros or dollars does not affect social welfare judgements.

Efficiency vs. Fairness Example

Consider an allocation problem

- **Agents** $\mathcal{A} = \{Alice, Bob\}$
- **Items** $\mathcal{R} = \{phone, bike, shoes, watch\}$.
- **Alice's utility** for an allocation X: 20 for the phone, 10 for the bike, 10 for the shoes, 0 for the watch
- **Bob's utility** for an allocation X: $5 \times$ the number of items in the allocation

Efficient allocation:

Bob gets the watch; Alice gets the rest
➔ total utility: 45

Fair allocation:

Alice gets the phone; Bob gets the rest
➔ minimum utility (Bob's): 15

Efficiency vs. Fairness Trade-off

Efficient Allocation

We assume **cardinal preferences**.

Utilitarian welfare function is considered to **measure the efficiency** of an allocation.

An allocation is called **efficient** (also **utilitarian**) if it maximizes the sum of utilities of all agents.

We denote the social value of an efficient allocation as $\mathrm{EFFICIENT}(\mathcal{X})$, i.e.,

$$\mathrm{EFFICIENT}(\mathcal{X}) = \sup\{\textstyle\sum_{i\in\mathcal{A}} u_i(X) \,|X \in \mathcal{X}\}$$

α-Fair Allocation

Constant Elasticity Social Welfare Function

Constant Elasticity Social Welfare Function sw_α with **inequality aversion parameter** α is defined as

$$sw(X,\alpha) = \begin{cases} \displaystyle\sum_{i \in \mathcal{A}} \frac{u_i(X)^{1-\alpha}}{1-\alpha} & \text{for } \alpha \geq 0, \alpha \neq 1 \\ \displaystyle\sum_{i \in \mathcal{A}} \log u_i(X) & \text{for } \alpha = 1 \end{cases}$$

$\alpha = 0$: Utilitarin SWF

$\alpha = 1$: Proportional fairness (~Nash SWF)

$\alpha \rightarrow \infty$: Egalitarian (Max-min) SWF

α-Fair Allocation

α-fair allocation $X^*(\alpha)$ is an allocation that maximizes the constant elasticity social welfare function for the corresponding value of α, i.e,

$$X^*(\alpha) = \operatorname*{argmax}_{X \in \mathcal{X}} sw(X, \alpha)$$

We denote the social value of the α-fair allocation as $\mathrm{FAIR}(\mathcal{X}, \alpha)$, i.e.,

$$\mathrm{FAIR}(\mathcal{X}, \alpha) = sw(X^*, \alpha)$$

Price of Fairness

Quantifies the **loss of efficiency** due to the requirement for fairness.

Price of Fairness

$$\mathrm{POF}(\mathcal{X}, \alpha) = \frac{\mathrm{EFFICIENT}(\mathcal{X}) - \mathrm{FAIR}(\mathcal{X}, \alpha)}{\mathrm{EFFICIENT}(\mathcal{X})}$$

Price is a fairness is always between zero and one, and corresponds to the **percentage efficiency loss** compared to the maximum system efficiency.

Note: $\mathrm{POF}(\mathcal{X}, 0) = 0$

Price of Fairness

Theorem

Consider a resource allocation problem with $n \geq 2$ agents where all agents have non-negative utilities with the same maximum achievable utility and the set of all feasible utility allocation is convex.
Then for the α-fair allocations, $\alpha \geq 0$, the price of fairness is bounded by

$$\mathrm{POF}(\mathcal{X}, \alpha) \leq 1 - \Theta(n^{-\frac{\alpha}{1+\alpha}})$$

Generalization to heterogeneous utilities possible

- the price then increases with the ratio between the highest and lowest achievable utility

Price of Fairness

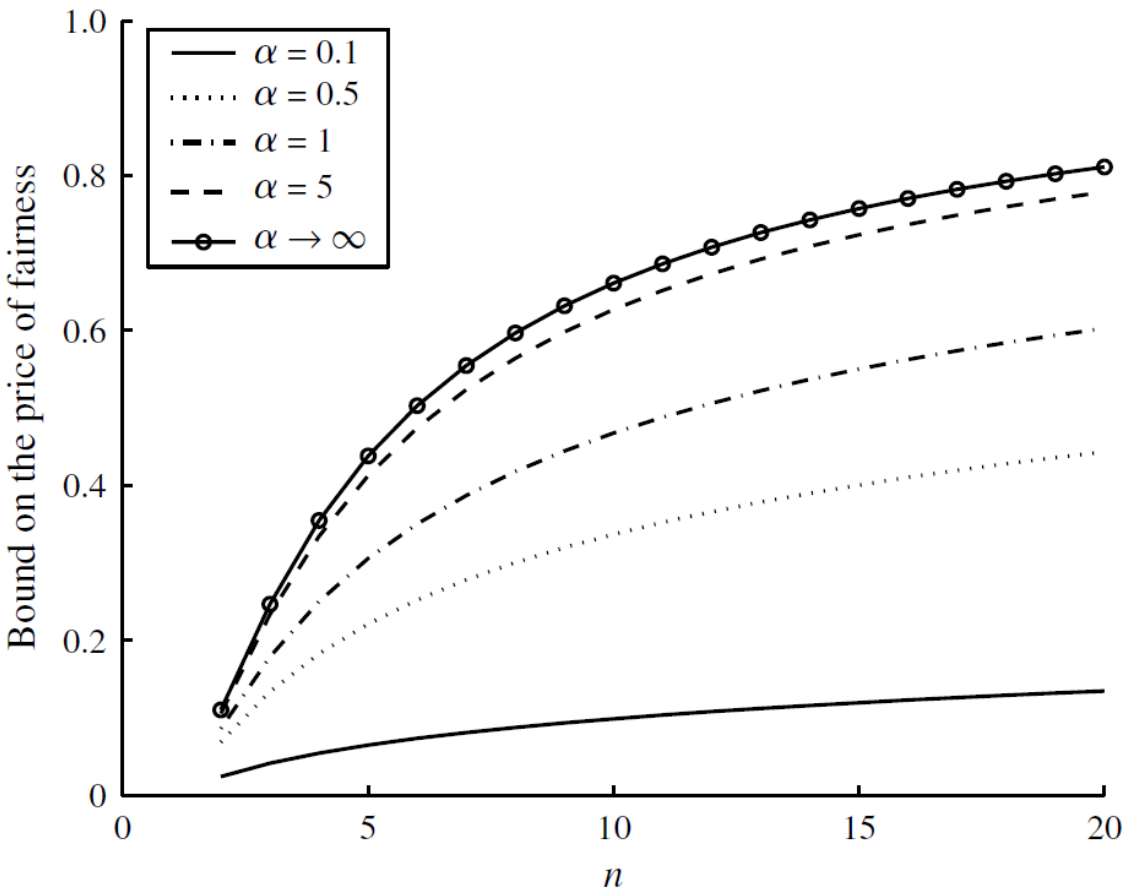

The worst-case price is increasing with the number of players and the value of α.

Bounds are very strong, near-tight.

Price of Efficiency

Quantifies the **loss of fairness** due to the requirement for efficiency.

We adopt the **minimum utility** egalitarian social welfare function as the fairness metric.

Price of Efficiency

$$\mathrm{POE}(\mathcal{X},\alpha) = \frac{\max_{X\in\mathcal{X}} \min_{i\in\mathcal{A}} u_i(X) - \min_{i\in\mathcal{A}} u_i(X_i^*(\alpha))}{\max_{X\in\mathcal{X}} \min_{i\in\mathcal{A}} u_i(X)}$$

(where $X^*(\alpha)$ is the α-fair allocation)

Price of Efficiency

Theorem

Consider a resource allocation with $n \geq 2$ agents where all agents have non-negative utilities and the same maximum achievable utility and the set of all feasible utility allocations is convex.
Then for the α-fair allocations, $\alpha \geq 0$, the price of efficiency is bounded by

$$\mathrm{POE}(\mathcal{X}, \alpha) \leq 1 - \Theta(n^{-\frac{1}{\alpha}})$$

Price of Efficiency

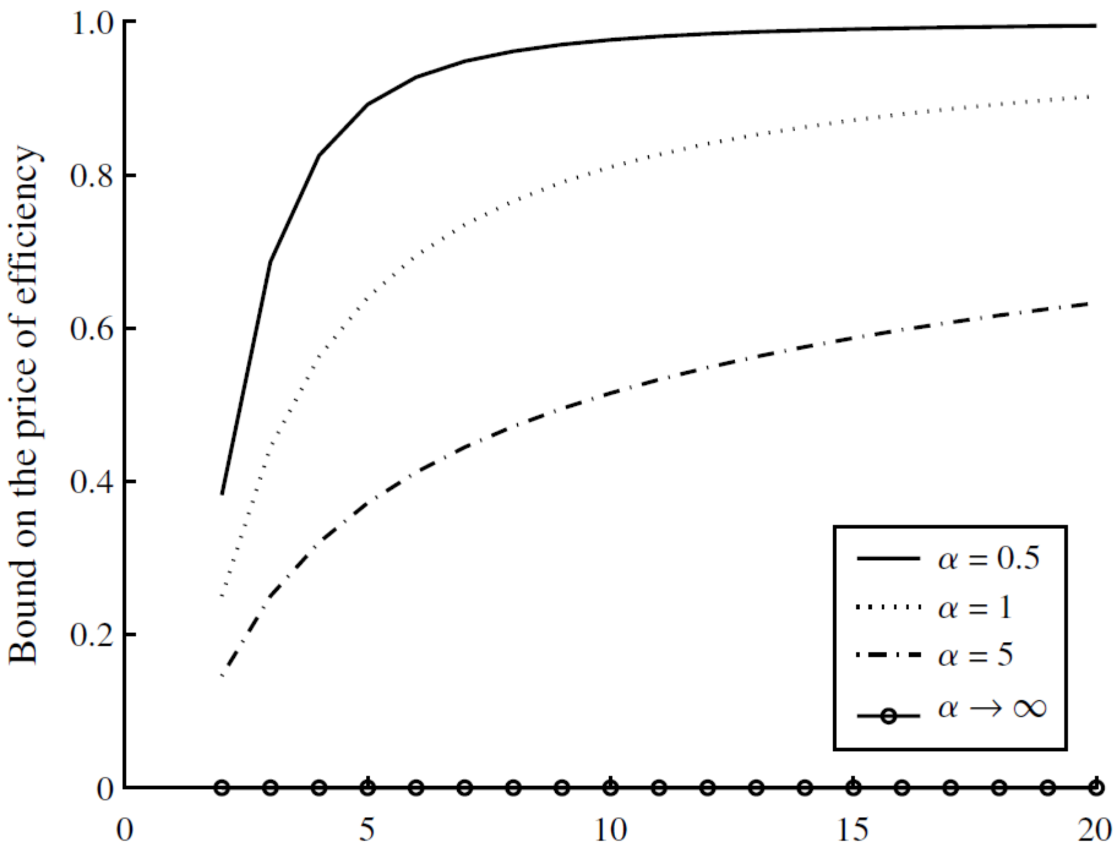

The worst-case price of efficiency is increasing with the number of players and the value of α.

Bounds are very strong, near-tight.

Example for four agents

Bounds Bounds on the Price of Fairness (Solid) and the Price of Efficiency (Dashed) of α-Fair Allocations for $n = 4$ Players

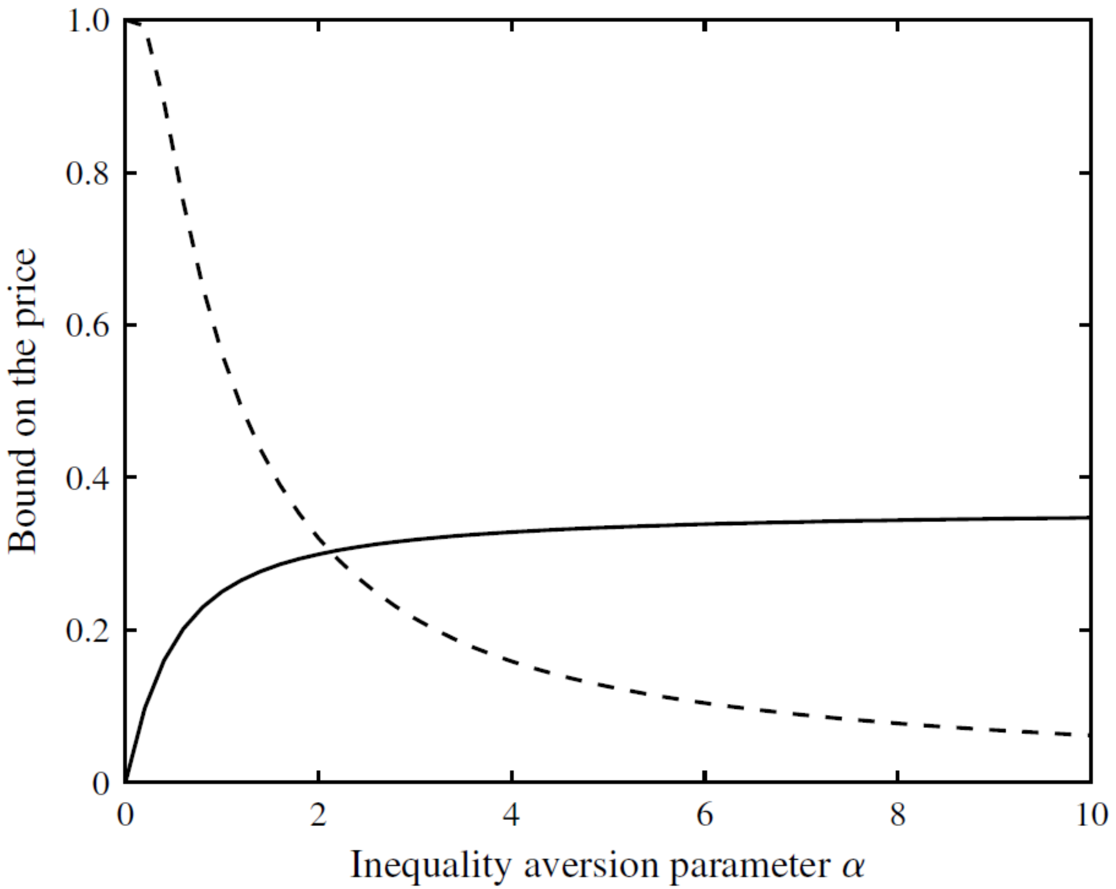

$\alpha = 1$: **Maximizes fairness** while guaranteeing max 20% loss of system efficiency.

$\alpha = 3$: **Maximizes efficiency** while guaranteeing maximum 20% drop in the utility for the worst-off agent.

$\alpha = 2$: **Balances efficiency and fairness**

Allocation Procedures

Allocation Procedures

Protocols: What messages do agents have to exchange and in which order?

Strategies: What strategies may an agent use for a given protocol? How can we give incentives to agents to behave in a certain way?

Algorithms: How do we solve the computational problems faced by agents when engaged in negotiation?

Centralised vs. Distributed Allocation

Centralised case	Distributed case
• A **single entity decides** on the final allocation, possibly after having elicited the preferences of the other agents. • Example: auctions	• **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?

Centralised vs. Distributed Comparison

Centralised	Distributed
• The **communication** protocols required are relatively **simple**. • Many **results** from **economics** and **game theory** , in particular on mechanism design, can be exploited. • **Powerful algorithms** for winner determination in combinatorial auctions. • Possible **trust** issues. • Difficult to deal with **unbounded problems**.	• Avoids **trust** issues. • Inherently **scalable**. • Can take an **initial allocation** into account. • More natural to model **step-wise improvements** over the status quo. • Can deal with **unbounded domains**. • More **complex** protocols significantly more **difficult** to analyse (convergence etc.)

➔ Auctions

Conclusions

Solving allocation problems requires defining 1) resources, 2) agents and their preferences, 3) system/social preferences and 4) mechanism.

There is an inherent trade-off between efficiency and fairness in allocation.

Auctions are a widely adopted centralized allocation mechanism which (typically) aims to optimize efficiency and is neutral toward fairness.

Reading:

- Chevaleyre, Y., Dunne, P.E., Endriss, U., Lang, J., Lemaitre, M., Maudet, N., Padget, J., Phelps, S., Rodríguez-Aguilar, J.A. and Sousa, P., 2006. Issues in multiagent resource allocation.
- Bertsimas, D., Farias, V.F. and Trichakis, N., 2012. On the efficiency-fairness trade-off. *Management Science*, *58*(12), pp.2234-2250.

Course Wrap-Up

Topics covered:

- single agent: agent architectures, BDI
- cooperative multi-agent: DCSP, DCOP
- competetive multi-agent:
  - agent perspective: non-cooperative game theory, coalition game theory
  - system designer perspective: social choice, auctions, resource allocation

AIC looking for talented Ph.D. students to pursue research in MAS-related topics:
http://aic.fel.cvut.cz/positions/phdpositions.pdf

Exams:

- 17., 24.1. and 7.2. early afternoon