A Short Tutorial on Multiagent Resource Allocation

Nicolas Maudet

LAMSADE
Université Paris-Dauphine

05/06/08

Example (1)

- two agents want to allocate a set of four indivisible resources (of two different colours);
- one of them (A) wants as many as possible, the other one (B) really wants resources of the same colour (as many as possible);
- what is an optimal allocation?
  - give everything to the first agent?
  - give two of the same colour to B, the rest to A?
  - or maybe one of any colour to B and the rest to A?

Example (2 — Bachrach et. al.)

- a set of shareable resources (ex. machines);
- agents require access to exactly one resource;
- the more agents using a resource, the more productive it is (but marginal gain decreases);
- we want to maximize the overall production;
- agents are retributed wrt. marginal contribution;

Example (3 — Rosenschein and Zlotkin)

- a number of nodes are to be visited;
- a team of agents that can travel (at a cost) to visit the nodes;
- agents want to minimize the cost of their mission;
- minimize the max cost of the agents of the team.

Basic Resource Allocation Framework

We start with the following basic elements:

- allocations of $|\mathcal{R}|$ **resources** among $|\mathcal{A}|$ **agents**
- resources are **divisible** (or not) and **shareable** (or not);
- each agent has **preferences** over the bundles it may hold
  - utility functions $u_i(\{\heartsuit\}) = 12$
  - preference relations $\{\heartsuit\} \preceq_i \{\heartsuit, \diamondsuit\}$
- agents only care about their own bundle (no **externalities**)

Main Question

How to allocate the given set of resources amongst these agents, in a way that is socially optimal?

Social Outcomes

How to evaluate the well-being of the society?
Social welfare measures (welfare economics, social choice)

Definition (Pareto optimality)

No other allocation would make at least one of the agents better off without making any worse off

Definition (Utilitarian social welfare)

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

Social Outcomes

And fairness measures...

Definition (Egalitarian social welfare)

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

Definition (Envy-freeness)

No agent should prefer to take the bundle allocated to one of its peers rather than keeping their own

Example

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

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

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

- *A* has maximal *egalitarian* social welfare (18); *utilitarian* social welfare is not maximal (38 rather than 42);
- *A* is *Pareto 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.

Distributed Perspective

Unlike in centralized mechanisms, in particular (classical) **combinatorial auctions**...

- **no single** auctioneer computes the optimal allocation
- negotiation starts with an **initial allocation**
- agents asynchronously **negotiate** resources
- **deals** to move from one allocation to another, ie $\delta = (A, A')$
- deals may be enhanced with money (utility transfer);
- agents accept deals on the basis of a **rationality criterion** that we assume myopic

In What Sense are Decisions Local?

- the **individual rationality** criterion should refer to the agent's preferences only, e.g:

$$v_i(A') - v_i(A) > p(i)$$

- but sometimes would be too restrictive: we may consider those agents **involved in the deal**, e.g:
  for all i: $A' \succeq_i A$ and at least for one j: $A' \succ_j A$
- **deals** themselves maybe restricted by: (i) negotiation topology, (ii) number of agents involved, (iii) number of resources involved.

Properties of Allocation Procedures

We may study different properties of allocation procedures:

- **Termination**— Is the procedure guaranteed to terminate eventually?
- **Convergence**— Will the final allocation be optimal according to our chosen social welfare measure?
- **Incentive-compatibility**— Do agents have an incentive to report their valuations truthfully? ($\leadsto$ *mechanism design*)
- **Computational Complexity**— What is the computational complexity of finding a socially optimal allocation of resources?
- **Communication Complexity**— How long will the process be?

Outline of the rest of the talk

1. convergence results;
2. communication complexity;
3. other issues.

Linking the Local and the Global Perspectives

IR deals are exactly those deals that increase SW:

Lemma (Rationality and social welfare)

A deal $\delta = (A, A')$ with side payments is IR iff $sw_u(A) < sw_u(A')$.

Proof.

"$\Rightarrow$": Rationality means that overall utility gains outweigh overall payments (which are $= 0$).
"$\Leftarrow$": The social surplus can be divided amongst all deal participants by using the following payment function:

$$p(i) \ = \ u_i(A') \ - \ u_i(A) \ - \ \underbrace{\frac{sw_u(A') - sw_u(A)}{|\mathcal{A}|}}_{>0}$$

$\square$

Convergence

It is now easy to prove the following **convergence** result (originally stated by Sandholm in the context of distributed task allocation):

Theorem (Sandholm, 1998)

Any sequence of individually rational deals will eventually result in an allocation with maximal social welfare.

Proof.

Termination follows from our lemma and the fact that the number of allocations is finite So let A be the terminal allocation. Assume A is *not* optimal, *i.e.* there exists an allocation A' with $sw_u(A) < sw_u(A')$. Then, by our lemma, $\delta = (A, A')$ is individually rational $\Rightarrow$ contradiction. □

Linking the Local and the Global Perspectives

- In the framework w/o money, we may use instead the **cooperative rational** (CR) criterion.
- Observe then that we only have that CR implies SW increase;
- Instead CR deals characterize Pareto improvements;
- Convergence to Pareto-efficient states can be guaranteed similarly as before.

Example

$u_1(\{\}) = 0$	$u_2(\{\}) = 0$
$u_1(\{r_1\}) = 2$	$u_2(\{r_1\}) = 3$
$u_1(\{r_2\}) = 3$	$u_2(\{r_2\}) = 3$
$u_1(\{r_1, r_2\}) = 7$	$u_2(\{r_1, r_2\}) = 8$

Transaction Types [Sandholm98]

- **simple (1-deals)** —one resource moves from one agent to another;
- **cluster (*k*-deals)** —a bundle of resources moves from one agent to another;
- **swap** —an agent swap a resource with another agent;
- **multiagent** —any number of agents, each passing one resource at most;
- **combination** —any combination.

Restrictions on Preference Structures

- **dichotomic** —bundles are tagged good/bad;
- **additive utilities** —no synergies between the resources
  - superadditive: only positive synergies (complementary)
  - subadditive: only negative synergies (subsidiarity)
- **separable additive utilities** —synergies restricted to fixed subsets of resources
- **k-additive utilities** —synergies restricted to bundles of cardinality $\leq k$
- **monotone utilities** —an agent always prefer (or is indifferent) to hold a proper superset of the bundle he holds

Negative Result

In general, any deal may be (potentially) required.
Even worse:

Theorem (Necessity of Deals)

In monotonic or dichotomic domains, any deal may be required to guarantee convergence to a utilitarian sw opt. alllocation (or Pareto-efficient w/o money);

In other words, these restrictions do not buy us anything.

Negative Result

Proof.

Let $\delta = (A, A')$ be any deal.
We must show that an initial alloc. and a collection of utilities exist s.t. δ would be necessary to reach the optimal allocation.
That would be the case if $\forall B, sw(B) \leq sw(A) < sw(A')$, with A initial alloc.
Note: There is at least one agent s.t. j tq. $A'(j) \neq A(j)$.

$$u_i(R) = \begin{cases} |R| + \epsilon & \text{if } R = A'(i) \text{ or } (R = A(i) \text{ and } i \neq j) \\ |R| & \text{otherwise} \end{cases}$$

We have $sw(A') = |\mathcal{R}| + \epsilon.|\mathcal{A}|$, $sw(A) = sw(A') - \epsilon$, and $sw(B) \leq sw(A)$ □

Sufficiency Results in Restricted Domains

We focus on some restricted domains:

Theorem (Additive domains)

*In **additive** domains, any sequence of individually rational **one-resource-at-a-time** deals will result in an allocation with max utilitarian sw.*

Theorem (Additive separable domains)

*In **additive k-separable** domains, any sequence of individually rational **k-cluster** deals will result in an allocation with max utilitarian sw.*

The value of a bundle is simply obtained by adding the value of bundles belonging to the different topics:

$$u(R) \quad = \quad u(\{\,\}) + \sum_{j=1}^{q} [u(R \cap R_j) - u(\{\,\})]$$

Exemple: u({♣, ♠, ⊕, △}) = u({♣, ♠})+u({⊕})+u({△})

(Note: *k*-separable Domains)

1. the set of resources can be divided into preferentially independent subsets (**topics**)
2. each topic contains **at most** *k* **ressources**
3. agents share this partitioning.

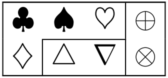

Necessity of these Conditions?

There cannot be. Example: **pseudo-constant** and a **modular**

$u_1(\{\,\}) =$	0	$u_2(\{\,\}) =$	0
$u_1(\{\spadesuit\}) =$	4	$u_2(\{\spadesuit\}) =$	$\boxed{1}$
$u_1(\{\clubsuit\}) =$	$\boxed{4}$	$u_2(\{\clubsuit\}) =$	3
$u_1(\{\spadesuit, \clubsuit\}) =$	4	$u_2(\{\spadesuit, \clubsuit\}) =$	4

sw is 5

Necessity of these Conditions?

Example: **pseudo-constant** and a **modular**

$u_1(\{\,\}) =$	0	$u_2(\{\,\}) =$	0
$u_1(\{\spadesuit\}) =$	$\boxed{4}$	$u_2(\{\spadesuit\}) =$	1
$u_1(\{\clubsuit\}) =$	4	$u_2(\{\clubsuit\}) =$	$\boxed{3}$
$u_1(\{\spadesuit,\clubsuit\}) =$	4	$u_2(\{\spadesuit,\clubsuit\}) =$	4

sw could be 7

Maximality

Prompted by negative results regarding necessity, we ask ourselves whether certain domains can be maximal (any domains including it would be lose the desired convergence properties);

Theorem (Maximality of modular domain)

Modular domain is maximal for 1-deals negotiations with payments.

In other words: as soon as the system contains an agent not having a modular function, convergence cannot be guaranteed any longer.

Convergence to Fair Allocations?

- typically these criteria are more difficult to optimize;
- equitable local deals can be designed

$$min\{u_i(A)|i \in \mathcal{A}^\delta < min\{u_i(A')|i \in \mathcal{A}^\delta$$

- but they may violate IR, and are local only to a certain extent;
- **idea:** use payments (surplus sharing) to compensate rational deals.

Convergence on Graphs?

Note of course that all of this crumble when we drop the assumption of fully connected systems (Yann to talk more about that...). Suppose the network is $a_2 - -a_1 - -a_3$.

$u_1(\{\}) = 0$	$u_2(\{\}) = 0$	$u_3(\{\}) = 0$
$u_1(\{r_1\}) = 2$	$u_2(\{r_1\}) = 3$	$u_3(\{r_1\}) = 0$
$u_1(\{r_2\}) = 3$	$u_2(\{r_2\}) = 3$	$u_3(\{r_2\}) = 0$
$u_1(\{r_1, r_2\}) = 7$	$u_2(\{r_1, r_2\}) = 8$	$u_3(\{r_1, r_2\}) = 100$

Outline of the rest of the talk

1. convergence results;
2. communication complexity;
3. other issues.

Communication Complexity in the Literature

Two agents hold an n-bit string and their goal is to communicate in order to compute the value of a (boolean) function over these two strings. What is the minimal number of bits that need to be exchanged to do so? [Yao,1979]

- Communication complexity of a **protocol**
  maximal number of bits exchanged when following the protocol in the worst case
- Communication complexity of a **function**
  communication complexity of the best protocol that computes that function

Aspects of Communication Complexity

(1) How many **deals** are required to reach an optimal allocation?
- communication complexity as number of individual deals

(2) How many **dialogue moves** are required to agree on one such deal?
- affects communication complexity as number of dialogue moves

(3) How expressive a **communication language** do we require?
- affects communication complexity as number of bits exchanged

Number of Deals (with money, utilitarian SW)

Upper bounds on the length of deal sequences

Theorem (Shortest path)

*A **single** rational deal is sufficient to reach an allocation with maximal social welfare.*

Theorem (Longest path)

A sequence of rational deals can consist of up to $|\mathcal{A}|^{|\mathcal{R}|} - 1$ *deals, but not more.*

Proof.

No allocation can be visited twice and there are $|\mathcal{A}|^{|\mathcal{R}|}$ distinct allocations $\Rightarrow$ upper bound follows □

Number of Deals (without money, Pareto optimality)

Upper bounds on the length of deal sequences

Theorem (Shortest path)

*A **single** cooperative rational deal is sufficient to reach a Pareto optimal allocation.*

Theorem (Longest path)

A sequence of rational deals can consist of up to $|\mathcal{A}| \cdot (2^{|\mathcal{R}|} - 1)$ *deals, but not more.*

Proof.

Each deal requires at least one agent having a strict improvement. No agent can hold a bundle he held previously and changed (strict improvement). □

Tightness of the bounds

Are these bounds tight?
(*i.e* can we really find a scenario where that many deals would be needed to reach the optimal allocation?)

- Framework With Money: **yes**
  reason: it is possible to construct utility functions such that distinct allocations have disctinct social welfare
- Framework Without Money: **no**
  reason: each deal involves at least two agents modifying their bundle

Further Results (Restriction on Deals / Preferences)

What happens if we concentrate on sequences of one-resource-at-a-time deals?
Then length of **shortest path** becomes (results hold when optimal outcome can be reached).

- $\leq |\mathcal{A}|^{|\mathcal{R}|} - |\mathcal{R}| \cdot (|\mathcal{A}| - 1)$ [Sandholm,98]
- if utility functions are **monotonic**: $\geq \frac{77}{128} 2^{\frac{|\mathcal{R}|}{2}} - 3$ [Dunne,04]
- if utility functions are **modular**: $|\mathcal{R}|$

Other Issues

- computational complexity problems that occur specifically in distributed settings;
- preference representation;
- concrete negotiation protocols (contract-net based, etc.);
- strategical aspects (manipulation, etc.);