Summer School on Fair Division (FairDiv-2015): Tutorial on Protocols for Allocating Indivisible Goods

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

[http://www.illc.uva.nl/~ulle/teaching/fairdiv-2015/]

Bigger Picture: Fair Allocation of Goods (1)

Every allocation problem is defined by a number of characteristics:

- What goods?
  - today's focus is on indivisible, nonsharable, static goods (other options: see tutorial on cake cutting)
- What preferences?
  - cardinal/ordinal, representation lang. (see tutorial by Jérôme Lang)
  - domain restrictions: e.g., additivity/separability
  - impact of monetary side payments (if any)
- What social objective?
  - utilitarian/egalitarian/Nash social welfare, envy-freeness, . . . (see tutorials by Iannis Caragiannis and Christian Klamler)

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.

Bigger Picture: Fair Allocation of Goods (2)

Once the basic characteristics are clear, we also need to decide:

- What procedure?
  - base line: elicit all preference information and then *centrally* compute the socially optimal allocation
  - sometimes more attractive: interactive/*distributed* procedures

Topics not accounted for in this tutorial:

- *behavioural* considerations, beyond abstract notions of rationality (see tutorial by Dorothea Herreiner)
- *strategic* (game-theoretical) considerations (see tutorial by Gianluigi Greco)
- refinements of abstract models to account for specific *applications* (see application tutorials for examples)

Outline

This tutorial will continue the theme of Christian Klamler's tutorial: protocols for the fair allocation of indivisible goods ("objects").

- centralised approach: computational complexity of optimisation
- distributed approach: sequences of local exchanges

Our focus will be on settings with expressive (not just additive) preferences, where finding a good allocation is highly complex.

U. Endriss. *Lecture Notes on Fair Division*. ILLC, University of Amsterdam, 2009.

S. Bouveret, Y. Chevaleyre, and N. Maudet. Fair Allocation of Indivisible Goods. In F. Brandt et al. (eds.), *Handbook of COMSOC*. CUP, 2015. In press.

Allocation of Indivisible Goods

Notation and terminology:

- Set of *agents* $\mathcal{N} = \{1, \ldots, n\}$ and finite set of *objects* $\mathcal{O}$.
- An *allocation* A is a partitioning of $\mathcal{O}$ amongst the agents in $\mathcal{N}$.
  Example: $A(i) = \{a, b\}$ — agent i owns items a and b
- Each agent $i \in \mathcal{N}$ has got a *utility function* $u_i : 2^{\mathcal{O}} \to \mathbb{R}$.
  Example: $u_i(A) = u_i(A(i)) = 577.8$ — agent i is pretty happy

How can we find a socially optimal allocation of objects?

Social Objectives

There are many possible definition for social optimality:

- *Pareto optimality*: no (weak) improvements for all possible
- maximal *utilitarian social welfare*: $\sum_{i \in \mathcal{N}} u_i(A(i))$
- maximal *egalitarian social welfare*: $\min_{i \in \mathcal{N}} u_i(A(i))$
- maximal *Nash social welfare*: $\prod_{i \in \mathcal{N}} u_i(A(i))$
- *equitability*: $u_i(A(i)) = u_j(A(j))$ for all $i, j \in \mathcal{N}$
- minimal *inequality*, e.g., in terms of the *Gini index*
- *proportionality*: $u_i(A(i)) \geqslant \frac{1}{n} \cdot \max_{S \subseteq \mathcal{O}} u_i(S)$
- *envy-freeness*: $u_i(A(i)) \geqslant u_i(A(j))$ for all $i, j \in \mathcal{N}$
- ...

We will focus on maximising utilitarian social welfare, but all of this could also be attempted for other social objectives.

Base Line: Centralised Optimisation

Suppose all agents have sent us their preferences, expressed in a suitable language. How can we compute the social optimum?

Next:

- What the *computational complexity* of this problem?
- How much easier does it get for *restricted preferences*?

Remark: The results we will discuss concern simple cases where no compact preference representation language is required.

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:** $\langle \mathcal{N}, \mathcal{O}, \mathcal{U} \rangle$ and $K \in \mathbb{Q}$
> **Question:** Is there an allocation A such that $\mathrm{SW}_{\mathrm{util}}(A) > K$?

Unfortunately, the problem is intractable:

Theorem 1 WELFARE OPTIMISATION *is NP-complete, even when every agent assign nonzero utility to just a single bundle.*

Proof: NP-membership: we can check in polytime whether a given allocation A really has social welfare $> K$. NP-hardness: next slide. ✓

This 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

By reduction to Set Packing (known to be NP-complete):

> Set Packing
> **Instance:** Collection $\mathcal{C}$ of finite sets and $K \in \mathbb{N}$
> **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 allocation problem:

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

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

Then every set packing corresponds to an allocation (with SW $= |\mathcal{C}'|$). *Vice versa*, for every allocation there is one with the same SW corresponding to a set packing (give anything owned by agents with utility 0 to agent 0). ✓

Welfare Optimisation under Additive Preferences

Sometimes we can reduce complexity by restricting attention to problems with certain types of preferences.

A utility function $u : 2^{\mathcal{O}} \to \mathbb{R}$ is called *additive* if for all $S \subseteq \mathcal{O}$:

$$u(S) \quad = \quad \sum_{x \in S} u(\{x\})$$

For this restriction, we get a positive result:

Proposition 2 Welfare Optimisation *is in P in case all individual utility functions are additive.*

Exercise: Why is this true?

Remark: This does not (always) work for other social objectives (e.g., Iannis Caragiannis showed you that checking for proportional fairness is NP-complete even for two agents with additive utilities).

Protocols

Most of the protocols introduced in the first part of this tutorial (by Christian Klamler) make the explicit or implicit assumption that utilities are *additive*/separable, e.g.:

- *adjusted winner* (Brams & Taylor)
- *singles-doubles procedure* (Brams, Kilgour & Klamler)
- *picking sequences* (Bouveret & Lang)

Thus, even when the pure social welfare optimisation problem is easy, it still is nontrivial to design a good protocol (even for just 2 agents), e.g., because we may have other social objectives as well, want to minimise elicitation, are worried about strategic issues, etc.

The only protocol for *general preferences* discussed by Christian was the *descending demand procedure* (Herreiner & Puppe), which is computationally very demanding (sorting exponentially many bundles).

Distributed Approach

Instead of devising algorithms for computing a socially optimal allocation in a centralised manner, we now want agents to be able to do this in a distributed manner by contracting deals locally.

- We are given some *initial allocation* A_0.
- 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 utility. A *payment function* is a function $p : \mathcal{N} \to \mathbb{R}$ with $p(1) + \cdots + p(n) = 0$.

  Example: $p(i) = 5$ and $p(j) = -5$ means that agent i *pays* €5, while agent j *receives* €5.

Negotiating Socially Optimal Allocations

We won't talk about designing a concrete negotiation protocol, but rather study the framework from an abstract point of view.

The main question concerns the relationship between

- the *local view:* what deals will agents make in response to their individual preferences?; and
- the *global view:* how will the overall allocation of objects evolve in terms of social welfare?

We will go through this for one set of assumptions regarding the local view and one choice of desiderata regarding the global view.

U. Endriss, N. Maudet, F. Sadri and F. Toni. Negotiating Socially Optimal Allocations of Resources. *Journal of AI Research*, 25:315–348, 2006.

The Local/Individual Perspective

A rational agent (who does not plan ahead) will only accept deals that improve her individual welfare:

- A deal $\delta = (A, A')$ is called *individually rational* (IR) if there *exists* a payment function p such that $u_i(A') - u_i(A) > p(i)$ for all $i \in \mathcal{N}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.

That is, an agent will only accept a deal if it results in a gain in utility (or money) that strictly outweighs a possible loss in money (or utility).

The Global/Social Perspective

Suppose that, as system designers, we are interested in maximising *utilitarian social welfare:*

$$\mathrm{SW}_{\mathrm{util}}(A) \quad = \quad \sum_{i \in \mathcal{N}} u_i(A(i))$$

Observe that there is no need to include the agents' monetary balances into this definition, because they'd 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.

Exercise: How well/badly do you expect this to work?

Example

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

$$
\begin{array}{rclrcl}
u_{ann}(\emptyset) & = & 0 & u_{bob}(\emptyset) & = & 0 \\
u_{ann}(\{chair\}) & = & 2 & u_{bob}(\{chair\}) & = & 3 \\
u_{ann}(\{table\}) & = & 3 & u_{bob}(\{table\}) & = & 3 \\
u_{ann}(\{chair, table\}) & = & 7 & u_{bob}(\{chair, table\}) & = & 8
\end{array}
$$

Furthermore, suppose the initial allocation of objects is A_0 with $A_0(ann) = \{chair, table\}$ and $A_0(bob) = \emptyset$.

Social welfare for allocation A_0 is 7, but it could be 8. By moving only a *single* item 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* $\{chair, table\}$.

Convergence

The good news:

Theorem 3 (Sandholm, 1998) *Any sequence of IR deals will eventually result in an allocation with maximal social welfare.*

Discussion: Agents can act *locally* and need not be aware of the global picture (convergence is guaranteed by the theorem).

Discussion: Other results show that (a) arbitrarily complex deals might be needed and (b) paths may be exponentially long. Still NP-hard!

T. Sandholm. Contract Types for Satisficing Task Allocation: I Theoretical Results. Proc. AAAI Spring Symposium 1998.

So why does this work?

The key to the proof is the insight that IR deals are exactly those deals that increase social welfare:

- **Lemma 4** *A deal* $\delta = (A, A')$ *is individually rational if and only if* $\mathrm{SW_{util}}(A) < \mathrm{SW_{util}}(A')$.

  Proof: ($\Rightarrow$) Rationality means that overall utility gains outweigh overall payments (which are $= 0$).

  ($\Leftarrow$) The social surplus can be divided amongst all agents by using, say, the following payment function:

$$p(i) \;=\; u_i(A') \;-\; u_i(A) \;-\; \underbrace{\frac{\mathrm{SW_{util}}(A') - \mathrm{SW_{util}}(A)}{|\mathcal{N}|}}_{> 0} \qquad \checkmark$$

Thus, as SW increases with every deal, negotiation must *terminate*. Upon termination, the final allocation A must be *optimal*, because if there were a better allocation A', the deal $\delta = (A, A')$ would be IR.

Related Work

Many ways in which this can be (and has been) taken further:

- other social objectives? / other local criteria?
- what types of deals needed for what utility functions?
- path length to convergence?
- other types of goods: sharable, nonstatic, . . . ?
- negotiation on a social network?

For several combinations of the above there still are open problems.

Last Slide

We have seen that finding a fair/efficient allocation in case of indivisible goods gives rise to a combinatorial optimisation problem.

Two approaches:

- *Centralised:* Give a complete specification of the problem to an optimisation algorithm. Often *intractable*.
- *Distributed:* Try to get the agents to solve the problem.
  For certain fairness criteria and certain assumptions on agent behaviour, we can predict *convergence* to an optimal state.