A Survey of Multiagent Allocation of Indivisible Goods

Yann Chevaleyre
In collaboration with S. Bouveret, N. Maudet, U. Endriss

27 octobre 2014

Outline

- the Framework
- Representing Preferences
- Efficiency, Fairness
- Computing Fair Outcomes
- Protocols for Fair Allocations
- Convergence to efficient and fair states

Framework

- $N = \{1, \ldots, n\}$ will be a set of n *agents*
- $\mathcal{O} = \{o_1, \ldots, o_p\}$ a set of p (indivisible, non-shareable) *objects*.
- Each subset $\mathcal{S}$ of $\mathcal{O}$ is called a *bundle*.
- An *allocation* is a function $\pi : N \to 2^{\mathcal{O}}$ mapping each agent to the bundle she receives, such that $\pi(i) \cap \pi(j) = \varnothing$ when $i \neq j$ since the items cannot be shared.
- $\pi(i)$ will be called agent i's *bundle* (or *share*). When $\bigcup_{i \in N} \pi(i) = \mathcal{O}$
- The allocation is said to be *complete*. Otherwise, it is *partial*. The set of all allocations is denoted Π.

Representing Ordinal Preferences

- simplest ordinal preferences : a *linear order* $\rhd_i$, which basically means that each agent is able to rank each item from the best to the worst
- how do we compare sets of preferences ?
  We need assumptions to *lift* preferences on items $\unrhd_i$ to preferences on sets $\succsim_i$?

Example of Preferences on items

$a \rhd b \rhd c \rhd d$

- *separab*ility.
  A preference relation $\succsim$ on $2^{\mathcal{O}}$ is separable if for every pair of bundles $(\mathcal{S}, \mathcal{S}')$, and every bundle $\mathcal{S}''$ such that $(\mathcal{S} \cup \mathcal{S}') \cap \mathcal{S}'' = \varnothing$, we have : $\mathcal{S} \succsim \mathcal{S}' \Rightarrow \mathcal{S} \cup \mathcal{S}'' \succsim \mathcal{S}' \cup \mathcal{S}''$
  - e.g. $\{a, c\} \succsim \{b, c\}$
- *Monotonicity*.
  $\mathcal{S} \succsim \mathcal{S}' \Rightarrow \mathcal{S} \cup \mathcal{S}'' \succsim \mathcal{S}'$
  - e.g. $\{a, c\} \succsim \{b\}$
- *Dominance*.
  If there exists a mapping f from each object in $\mathcal{S}$ to each on in $\mathcal{S}'$ such that $o \succsim f(o)$ for these objects then $\mathcal{S} \succsim \mathcal{S}'$
  - e.g. $\{a, c\} \succsim \{b, d\}$

Cardinal Preferences

- Simplest cardinal preferences on items :
  A utility represented by a *weight function* $w_i : \mathcal{O} \to \mathbb{F}$, mapping each object to a score taken from a numerical set (that we will assume to be $\mathbb{N}$, $\mathbb{Q}$ or $\mathbb{R}$ for the sake of simplicity).
- Lift this to sets of items :
  - $u_i(\mathcal{S}) = \sum_{o \in \mathcal{S}} w_i(o)$
- *Modularity.*
  A utility function $u : 2^{\mathcal{O}} \to \mathbb{F}$ is *modular* if and only if for each pair of bundles $(\mathcal{S}, \mathcal{S}')$, we have $u(\mathcal{S} \cup \mathcal{S}') = u(\mathcal{S}) + u(\mathcal{S}') - u(\mathcal{S} \cap \mathcal{S}')$.
- Implies Separability, Dominance, etc..

Beyond Separability

- **Complementarities and Substitutabilities**
- Cardinal case : GAI Networks
  $u(B) = u_{\{o_1 o_2\}}(B \cap \{o_1, o_2\}) + u_{\{o_2 o_3\}}(B \cap \{o_2, o_3\})$
  where $u_{\{o_1 o_2\}}$ and $u_{\{o_2 o_3\}}$ are represented in tabular form.
- Ordinal case : CI-networks
  statements of the form $\mathcal{S}^+, \mathcal{S}^- : \mathcal{S}_1 \rhd \mathcal{S}_2$ This informally means : “if I have all the items in $\mathcal{S}^+$ and none of those in $\mathcal{S}^-$, I prefer obtaining all items in $\mathcal{S}_1$ to obtaining all those in $\mathcal{S}_2$ (*ceteris paribus*).”

CI-net example

Consider two CI-statements : $S1 = (o_1, \varnothing : o_4 \rhd o_2 o_3)$; $S2 = (\varnothing, o_1, : o_2 o_3 \rhd o_4)$.

- We can deduce $o_1 o_4 \succ o_1 o_2 o_3$ $(S1)$ and $o_2 o_3 \succ o_4$ $(S2)$.
- Note : $\succ$ is not separable, as having o_1 or not in the bundle reverses the preference between $o_2 o_3$ and o_4.

Fairness and Efficiency

- *Fairness :*
  - *MaxMin Allocations*
    An allocation is maxmin when the utility of the poorest agent is as high as posible, *i.e.*

$$\max_{\pi \in \Pi} \left\{ \min_{i \in N} u_i(\pi(i)) \right\}$$

  - *Envy-freeness.*
    An allocation is envy-free when $\pi(i) \succsim_i \pi(j)$ for all agents $i, j \in N$.
  - Proportionality
    Each agent should get from the allocation at least the n^{th} of the total utility she would have received if she were alone
- *Efficiency*
  - *Utilitarian optimality*
    allocation maximizing social welfare in cardinal setting
  - *Pareto efficiency*
    (in particular in ordinal setting)

Price of Fairness

- Efficiency and fairness are often incompatible...
- To measure this :

$$Price\,of\,Fairness = \frac{\sum_{i \in N} u_i(\pi^*)}{\sum_{i \in N} u_i(\pi^f)}$$

  where π^* is an efficient allocation and π^f is a fair allocation
- *Maximin*
  - *The price of fairness for maxmin allocations is unbounded (Caragiannis et al., 2012)*
- *Envy freeness*
  - EF optimality does not imply pareto optimality
  - The price of fairness for envy-freeness is $\Theta(n)$ *(Caragiannis et al., 2012)*

Computing Fair Allocations : Cardinal case

- Maximin :
  - hard to compute, even hard to approximate (Golovin 2005)
  - Approximable for GAI nets of arity 2 (Bezakoa and Dani, 2005)
- Envy-free
  - trivial setting : throw all objects away (when completeness not required)
  - Existence is hard (Lipton 2004)
  - Existence of Envy-free and Efficient alloction :
    Above NP for most compact preference languages (Bouveret, Lang 2008)
- What about allocation with bounded envy ?

$$\begin{aligned} e_{ij}(\pi) &= \max\{0, u_i(\pi(j)) - u_i(\pi(i))\} \\ e(\pi) &= \max\{e_{ij}(\pi) \mid i, j \in N\} \end{aligned}$$

- *Thm* : It is always possible to find an allocation whose envy is bounded by α, the maximal marginal utility of the problem (Lipton 2004)

Computing Fair Allocations : Ordinal case

- Envy freeness in ordinal setting :
  Many incomparabilities => need a relaxed notion of EF
- Agent i possibly (resp. necessarily) envies agent j if $\pi(i) \not\succ_i \pi(j)$ (resp. $\pi(j) \succ_i \pi(i)$).
- With separable and monotonic prefs
  - determining whether a possible envy-free efficient allocation exists is in P (Bouveret et al., 2010)
  - NP for necessary envy-freeness (Bouveret et al., 2010 ; Aziz et al., 2014)

Protocols for Fair Allocation

- ***Adjusted Winner Procedure***
  (works for 2 agents with additive utilities)
  1. Each good is assigned to the agent who values it the most
  2. While $u_1 \neq u_2$ or richest has become poorest do
     1. Let o be a good own by the richest, such that $\frac{u_{rich}(o)}{u_{poor}(o)}$ is minimal
     2. transfer o to poor.
  3. If richest has become poorest, then *split* last good between agents.
- If items can be split then
  The adjusted winner procedure returns an equitable, envy-free, and Pareto-optimal allocation
- ***Under cut procedure***
  Under some technical conditions
  an envy-free allocation exists and the undercut protocol returns it

Protocols for general prefernces and more than two agents

- For non modular preferences :
  - Thm :
    Any deterministic algorithm would require an exponential number of queries to compute any finite approximation for the minimal envy problem (Lipton et al., 2004), or maxmin allocation (Golovin, 2005).
- The descending demand procedure (Herreiner and Puppe, 2002)
  - linear ordering over all subsets of resources (satisfying monotonicity)
  - one by one, agents name their preferred bundle, then their next preferred bundle, and so on..
  - if a feasible complete allocation can be obtained, by combining only bundles mentioned so far in the procedure, stop
  - Thm :
    it returns a Pareto-effiicient and rank-maxmin-optimal allocation

Decentralized protocols

- Fair division based on local deals
  - agents will start from an initial allocation, and myopically contract local, deals, independently from the rest of the society.
  - e.g. restrict to deals diminishing inequality among involved agents
    - For separable criteria (e.g. leximin ordering), can reach optimum
    - Any kind of restriction on deal types ruins the guarantee of convergence in the general domain (Endriss et al., 2006)
    - The upper bound on the length of the sequence of deals can be exponential in the worst case (Sandholm, 1998 ; Endriss and Maudet, 2005)

Conclusion

- Even with separable/additive preferences, computing optimal allocation is very costly
- For protocols, in general communication is the bottleneck
- For decentralized protocols, require complex deals
- Q : How such protocols and algorithms will be adopted in practice, for instance whether agents may strategize, and whether suggested solutions can be easily understood and ac- cepted ?