Multilevel Fair Allocation under Additive Preferences

Maxime Lucet, Nawal Benabbou, Aurélie Beynier, and Nicolas Maudet
LIP6, Sorbonne Université — {firstname.lastname}@lip6.fr

Model and motivations

Motivations: Solving fair allocation problems in hierarchical structure (e.g. in university, across a territory, ...)

Model:

- a set of indivisible items $\mathcal{G} = \{g_1, \dots, g_m\}$
- a set of nodes $\mathcal{N} = \{1, \dots, n\}$ organized in a hierarchical structure represented by a tree $\mathcal{T}$.
- any node i has a weight $w_i \in \mathbb{R}_{\geq 0}$.

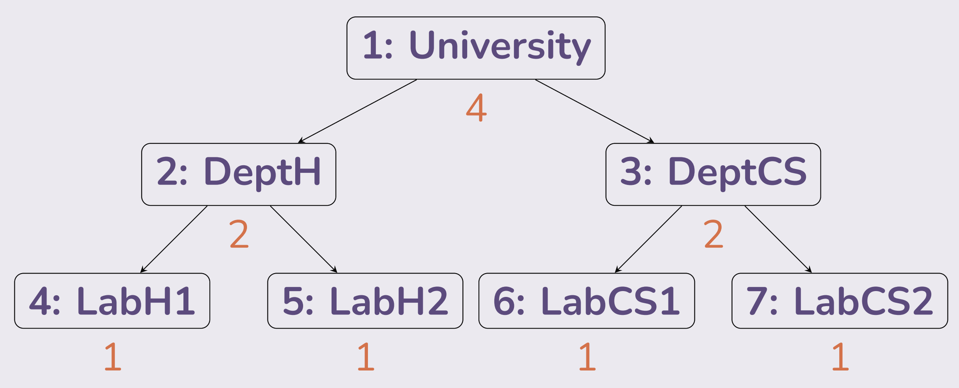

Figure 1: The hierarchical structure of a university. In orange, the weights.

A **complete multilevel allocation** π satisfies:

1. $\pi(1) = \mathcal{G}$,
2. $\pi(i) = \cup_{j \in \text{children}(i)} \pi(j)$ for an internal node i,
3. $\pi(i) \cap \pi(j) = \emptyset$ for two siblings i, j,

where $\pi(i)$ denotes the bundle of any node $i \in \mathcal{N}$

Utilities: there is a difference between internal nodes and leaves.

- Internal node i: utility v_i is the sum of the utility of its chidlren.
  $\Rightarrow$ it evaluates a multilevel allocation.
- Leaf x: utility u_x is a classical additive utility over the items.
  $\Rightarrow$ it evaluates a bundle of items.

Algorithm

We propose a **multilevel extension** of the Weighted Round Robin, denoted **MWRR** hereafter:

Input: Tree $\mathcal{T}$, items $\mathcal{G}$, nodes valuations
Output: Multilevel allocation π

Initialise an empty multilevel allocation π
Initialise picking scores to 0
Initialise a set of remainings items, i.e. $\mathcal{RI} \leftarrow \mathcal{G}$
while $\mathcal{RI} \neq \emptyset$ **do**
    Select leaf x recursively with the least picking score divided by its weight
    Let x choose her favorite item in $\mathcal{RI}$
    Update the bundles of x and its ancestors
    $\mathcal{RI} \leftarrow \mathcal{RI} \setminus \{g\}$
end while
return π

Fairness

Problem: Most fairness notions require agents to evaluate bundles, which our internal nodes cannot do!

We equip internal nodes with estimated utility functions to evaluate bundles. Given an internal node i, and a bundle $S \subseteq \mathcal{G}$, we propose 3 functions:

- Optimistic: $\hat{v}_i(S)$ – utility of S given the optimal allocation to the leaves of i,
- Agnostic: $\bar{v}_i(S)$ – weighted average utility of S,
- Pessimistic: $\check{v}_i(S)$ – utility of S given the worst allocation to the leaves of i.

Agnostic Fairness: A multilevel allocation is M[agno]-WEF1 if for any nodes j, k with the same parent, we have

$$\frac{v_j(\pi)}{w_j} \geq \frac{\bar{v}_j(\pi(k) \setminus \{g\})}{w_k}$$

for some item $g \in \pi(k)$.

Objective: computing a complete and M[pess/agno/opt]-WEF1 multilevel allocation.

Example: Let $\mathcal{G} = \{g_1, \dots, g_5\}$, and $\mathcal{N} = \{1, \dots, 7\}$ organized as in Fig. 1.

Leaves	g_1	g_2	g_3	g_4	g_5
$\{4, 6\}$	0.1	0.1	0.35	0.1	0.35
$\{5, 7\}$	0.2	0.2	0.2	0.2	0.2

Table 1: Preferences of the leaves.

- Optimistic: $\hat{v}_3(\{g_3, g_5\}) = 0.7$,
- Agnostic: $\bar{v}_3(\{g_3, g_5\}) = 0.55$,
- Pessimistic: $\check{v}_3(\{g_3, g_5\}) = 0.4$.

$\Rightarrow$: It will impact the fairness notions we can ensure.

Identical preferences

Identical Preferences: all leaves have the same preferences.

$\Rightarrow$ for any node i and any bundle S, $\hat{v}_i(S) = \bar{v}_i(S) = \check{v}_i(S)$.
$\Rightarrow$ Hence, the three fairness notions coincide.

Theorem: MWRR returns a complete M-WEF1 allocation.

Theorem: MWRR is also EF1 at the leaves if w_i is the number of leaves of the subtree rooted in i, for all nodes i.

General preferences

Theorem: MWRR returns a complete M[pess]-WEF1 allocation.

Theorem: MWRR may fail to return an M[agno/opt]-WEF1 allocation.

Theorem: MWRR cannot even guarantee to find an α-approximation of M[agno]-WEF1, for some $\alpha > 0$.

Experimentations

Protocol: We generated instances varying:

- the structure of the trees,
- the weights scheme of the agents,
- the preference generation methods.

$n = 15, m = 15$	$n = 127, m = 254$
0.0002 sec.	0.0191 sec.

Table 2: Average runtime for balanced trees.

Key Results

- Fast and scalable: see for example Table 2
- M[agno]-WEF1 almost always satisfied: < 50 unfair allocations out of $\simeq$200,000.
- M[opt]-WEF1 is challenging: up to 100% of unfair allocations.

Perspectives

The model could be extended further:

- Using a DAG to represent the hierarchical structure.
- Finding stronger fairness notions with provable guarantees than M[pess]-WEF1.
- Adapting other algorithms to the multilevel setting.

Full paper can be found here: