Reaching Envy-free States in Distributed Negotiation Settings

Yann Chevaleyre • Ulle Endriss • Sylvia Estivie • Nicolas Maudet

Resource Allocation by Negotiation

- Set of **agents** $\mathcal{A} = \{1..n\}$ and set of indivisible **resources** $\mathcal{R}$.
- An **allocation** A is a partitioning of $\mathcal{R}$ amongst the agents.
  Example: $A(i) = \{r_5, r_7\}$ — agent i owns resources r_5 and r_7
- Each agent $i \in \mathcal{A}$ has got a **valuation function** $v_i : 2^{\mathcal{R}} \to \mathbb{R}$.
  Example: $v_i(A) = v_i(A(i)) = 577.8$ — agent i is pretty happy
- Agents may engage in negotiation to exchange resources in order to benefit either themselves or society as a whole.
- A **deal** $\delta = (A, A')$ is a pair of allocations (before/after).
- A deal may come with side payments to compensate some of the agents for a loss in valuation. A **payment function** is a function $p : \mathcal{A} \to \mathbb{R}$ with $\sum_{i \in \mathcal{A}} p(i) = 0$.

Individual Rationality and Social Welfare

- A deal $\delta = (A, A')$ is **individually rational** (IR) iff there exists a payment function p such that $v_i(A') - v_i(A) > p(i)$ for all $i \in \mathcal{A}$, except possibly $p(i) = 0$ for agents i with $A(i) = A'(i)$.
- An allocation A is called **efficient** iff it maximises (utilitarian) **social welfare**: $sw(A) = \sum_{i \in \mathcal{A}} v_i(A)$.

Convergence Results

A known result states that any sequence of IR deals will eventually result in an efficient allocation of resources.

Our aim in this paper has been to explore to what extent similar convergence results are attainable when we are interested in allocations that are not only efficient but also fair.

Envy-free States

A common interpretation of **fairness** is **envy-freeness**: no agent should want to switch bundles with any of the others.

Unfortunately, envy-free allocations do not always exist. We can circumvent this problem by taking the balance of past side payments into account when defining envy-freeness:

- Associate each allocation A with a **balance** $\pi : \mathcal{A} \to \mathbb{R}$, mapping agents to the sum of payments they've made so far.
- A **state** (A, π) is a pair of an allocation and a payment balance.
- Each agent $i \in \mathcal{A}$ has got a (quasi-linear) **utility function** $u_i : 2^{\mathcal{R}} \times \mathbb{R} \to \mathbb{R}$, defined as follows: $u_i(R, x) = v_i(R) - x$.
- A state (A, π) is **envy-free** iff $u_i(A(i), \pi(i)) \geq u_i(A(j), \pi(j))$ for all agents $i, j \in \mathcal{A}$. An **efficient envy-free** (EEF) state is an envy-free state with an efficient allocation.

Envy-freeness and Individual Rationality

By a known result, EEF states always exist (in the presence of money). But we want to find them by means of **rational** negotiation. Unfortunately, this is generally **impossible**.

Example: 2 agents, 1 resource with $v_1(\{r\}) = 4$ and $v_2(\{r\}) = 7$. Agent 1 owns r to begin with; giving it to agent 2 would be efficient.

- An **IR deal** would require a payment within $(4, 7)$.
- But to ensure **envy-freeness**, the payment should be in $[2, 3.5]$.

Compromise: We shall enforce an **initial equitability payment** $\pi_0(i) = v_i(A_0) - sw(A_0)/n$ before beginning negotiation.

Globally Uniform Payments

Because of the "non-local effects of local deals" in view of envy-freeness, to have any chance of getting a convergence result for EEF states, we will have to restrict the freedom of agents a little by fixing a specific payment function (still IR!):

- Let $\delta = (A, A')$ be an IR deal. The payments as given by the **globally uniform payment function** (GUPF) are defined as follows: $p(i) = [v_i(A') - v_i(A)] - [sw(A') - sw(A)]/n$.

That is, we distribute the (positive!) social surplus to all agents.

Convergence to EEF States

Theorem: If all valuations are **supermodular** and if **initial equitability payments** have been made, then any sequence of **IR** deals using the **GUPF** will eventually result in an **EEF** state.

Degrees of Envy

To be able to analyse how "close" to an EEF state we can get in case not all of the preconditions of our theorem are satisfied, we require a notion of the **degree of envy** of a state.

The paper proposes a systematic approach to possible definitions:

- Envy between two agents: $\max\{u_i(A(j), \pi(j)) - u_i(A(i), \pi(i)), 0\}$ (or even without max)
- Degree of envy of a single agent: 0-1 (no/yes), max, sum
- Degree of envy of a society: max, sum (or indeed any SWO)

Experiments

We have carried out several experiments. The following example shows the evolution of the sum of envies for 20 agents with modular valuation functions, negotiating over 150 goods using the LUPF (like the GUPF, but the social surplus is divided only amongst the agents involved in the deal), without initial payments:

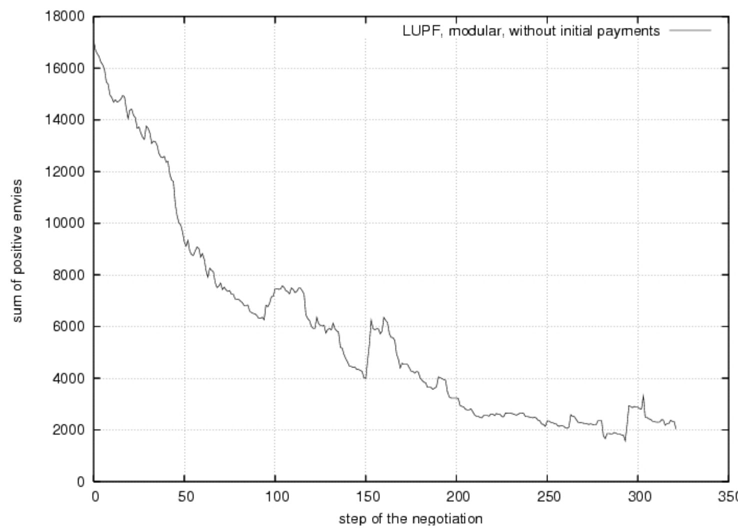

Contact: ulle@illc.uva.nl