ALMA: A Scalable, Privacy-preserving Algorithm for Multi-agent Allocations

Panayiotis Danassis, Aris Filos-Ratsikas, Aleksei Triastcyn, Boi Faltings
Artificial intelligence Laboratory, EPFL

- Multi-agent coordination in **massive scale**
  - Autonomous vehicles
  - Robotic agents
  - Intelligent infrastructure
  - IoT devices
  - Meeting scheduling
  - Etc.

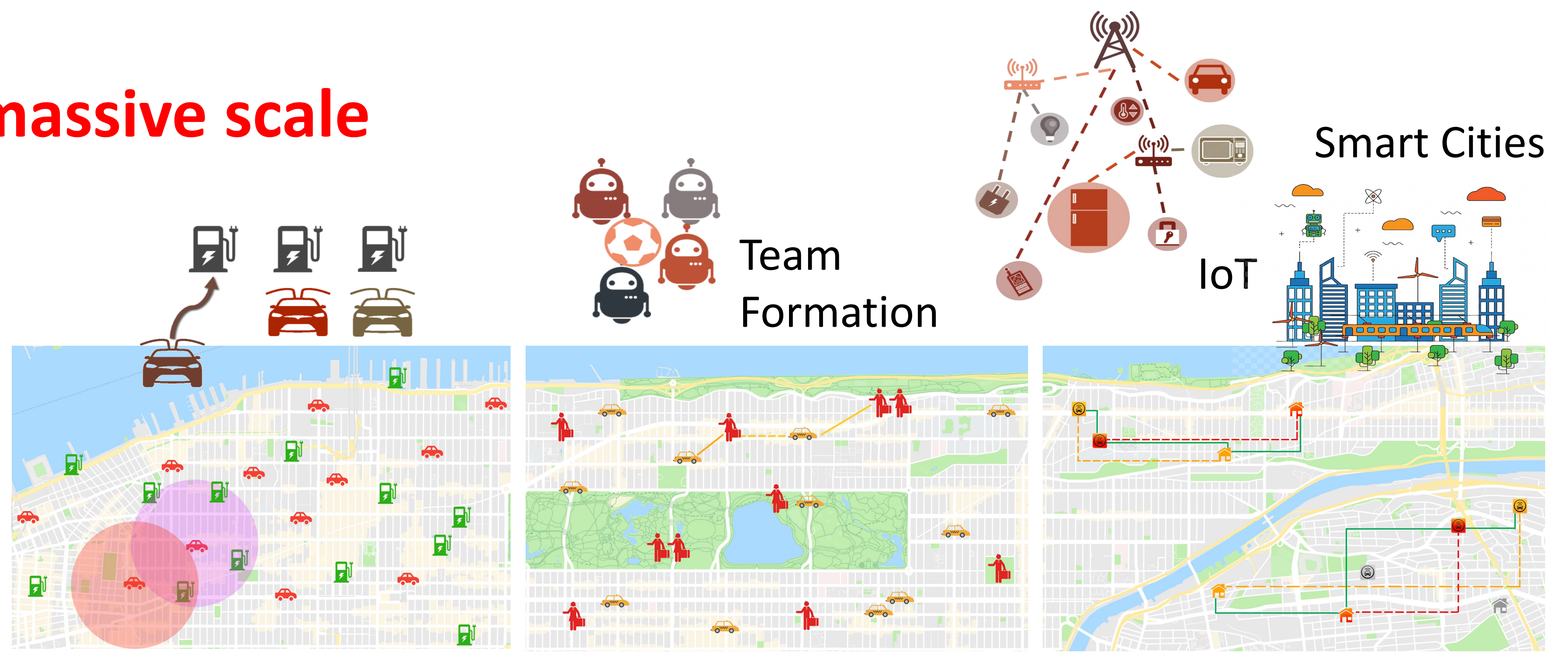

- Traditional approaches **break** in massively-large multi-agent environments:
  1. **Complexity** / scalability
  2. Infeasible **communication** and observability requirements
  3. No protection of **private** data

- **ALMA**: ALtruistic MAtching
  - Agents try their best options
  - If collision: better alternatives ⇒ more likely to back off
  - Back-off probability P_B^n that depends on the utility loss:
  - $P_B^n(loss) = 1 - loss,$
  - $loss = u(r_i) - u(r_{i+1})$

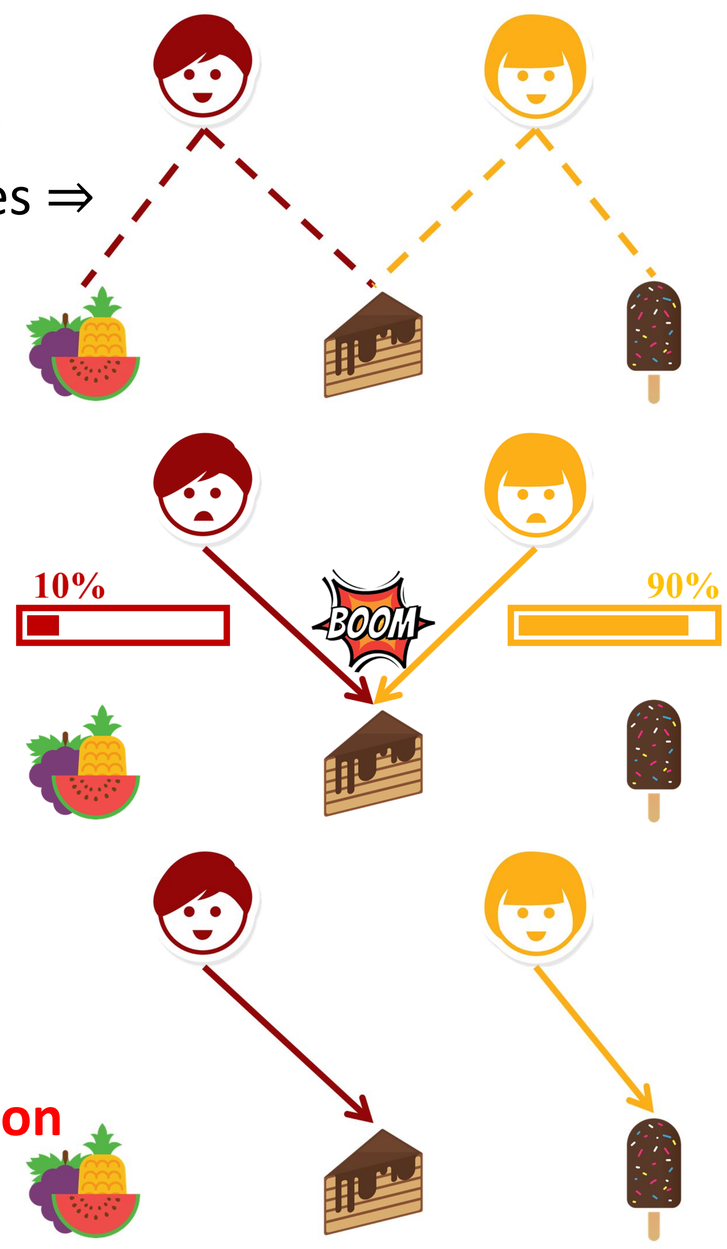

- Theorem
  - **Constant convergence** time
  - Independent of the total problem size

- Advantages
  - **No** inter-agent **communication**
  - Piece-wise **Local**
  - **Differentially Private**

- Convergence Time

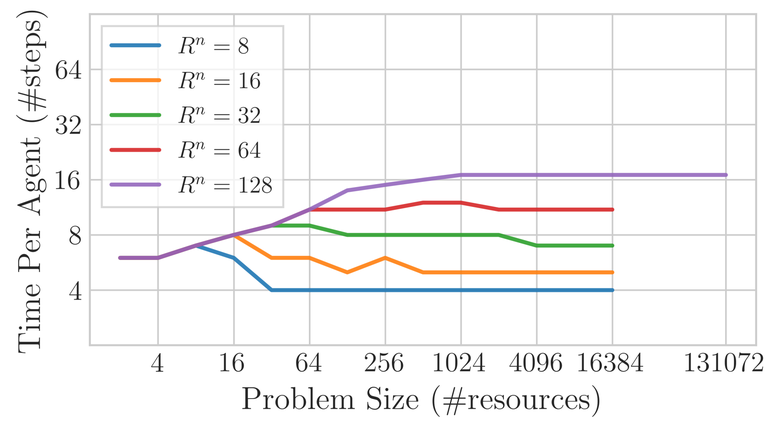

- Differential Privacy Cost

Smaller ε ⇒ stronger privacy
For ε = 0.5 (median) ⇒ An attacker can identify an individual with probability 0.62 (i.e., not much better than random)

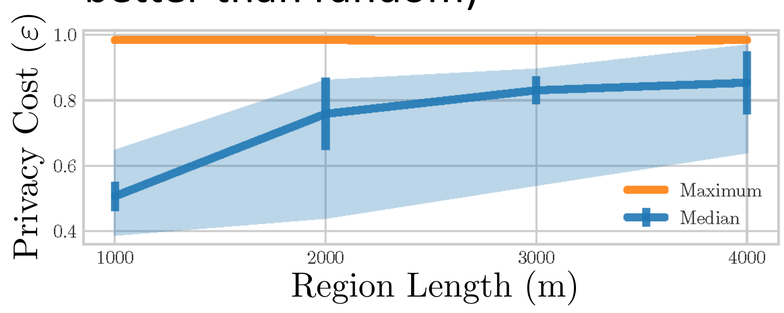

- Sum of Agent Utilities

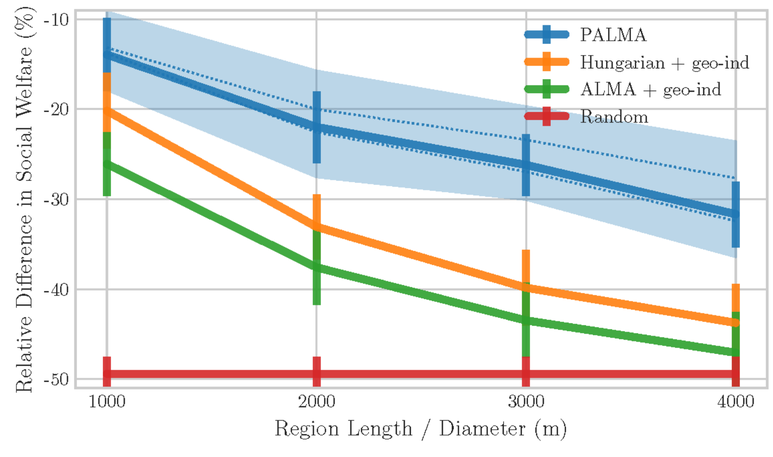

- **Stochastic conflict resolution** without communication
- **No backtracking** on allocations **avoids combinatorics** and enables constant time convergence
- **Decentralized** computation gives agents control of their privacy and allows for efficient parallel computation

Read More: ALMA | Privacy | Learning | MoD | Thesis

EPFL