Scalable Mechanism Design for Multi-Agent Path Finding

Paul Friedrich, Yulun Zhang, Michael Curry, Ludwig Dierks,
Stephen McAleer, Jiaoyang Li, Tuomas Sandholm, Sven Seuken

Talk at IJCAI '24, 09 August 2024

University of Zurich[UZH]
ETH AI CENTER
HARVARD UNIVERSITY
Carnegie Mellon University
UIC UNIVERSITY OF ILLINOIS CHICAGO

Competing for paths: a new domain

Combines

- **large scale**
- **mechanism design:**
  self-interested agents
- **multi-agent path finding (MAPF):**
  to calculate allocation

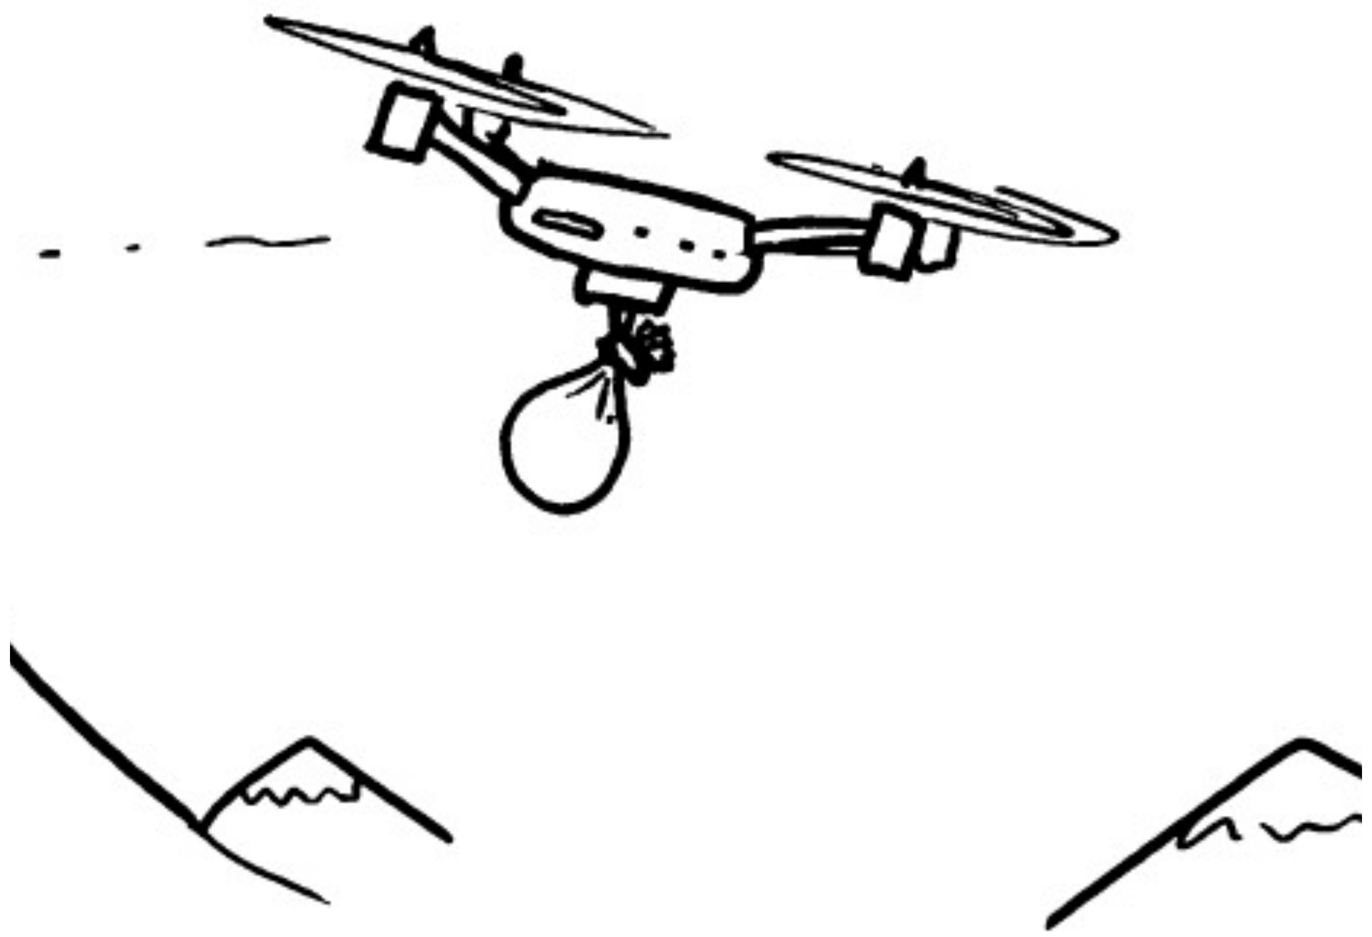

Traffic management: allocating

- road capacity to cars
- urban airspace to UAVs

Not your usual MAPF setting

The usual MAPF setting is **cooperative**:

- A central allocator system knows all true agent's characteristics

Our setting is **non-cooperative**:

- The system queries agents for their characteristics
- Agents will lie if they can secure a faster route

Without incentive awareness: **worse overall solutions!**

Problem summary

Design mechanisms for allocating collision-free paths that are:

- **strategyproof:** incentivize agents to be honest
- **efficient:** find allocations with high consumer welfare
- **scalable**

Up until now: “pick 2”…

… we can do all three!

Why don't we just use auctions?

Mechanism designers love auctions, especially **the VCG auction.**

Ask agents how much they value the resource → Find the allocation that's optimal for these reports → Charge clever payments that make lying unprofitable

Many possible paths / allocations / vertex combinations to ask

Optimal MAPF is hard!

Optimal MAPF is hard! ($\times N$)

The solution: maximal-in-range (MIR)

Ask agents for their start & goal, value for arrival, cost for delay → Suboptimal MAPF algo finds allocation that is **optimal within a range** → Charge VCG payments, with a twist that makes them "free"

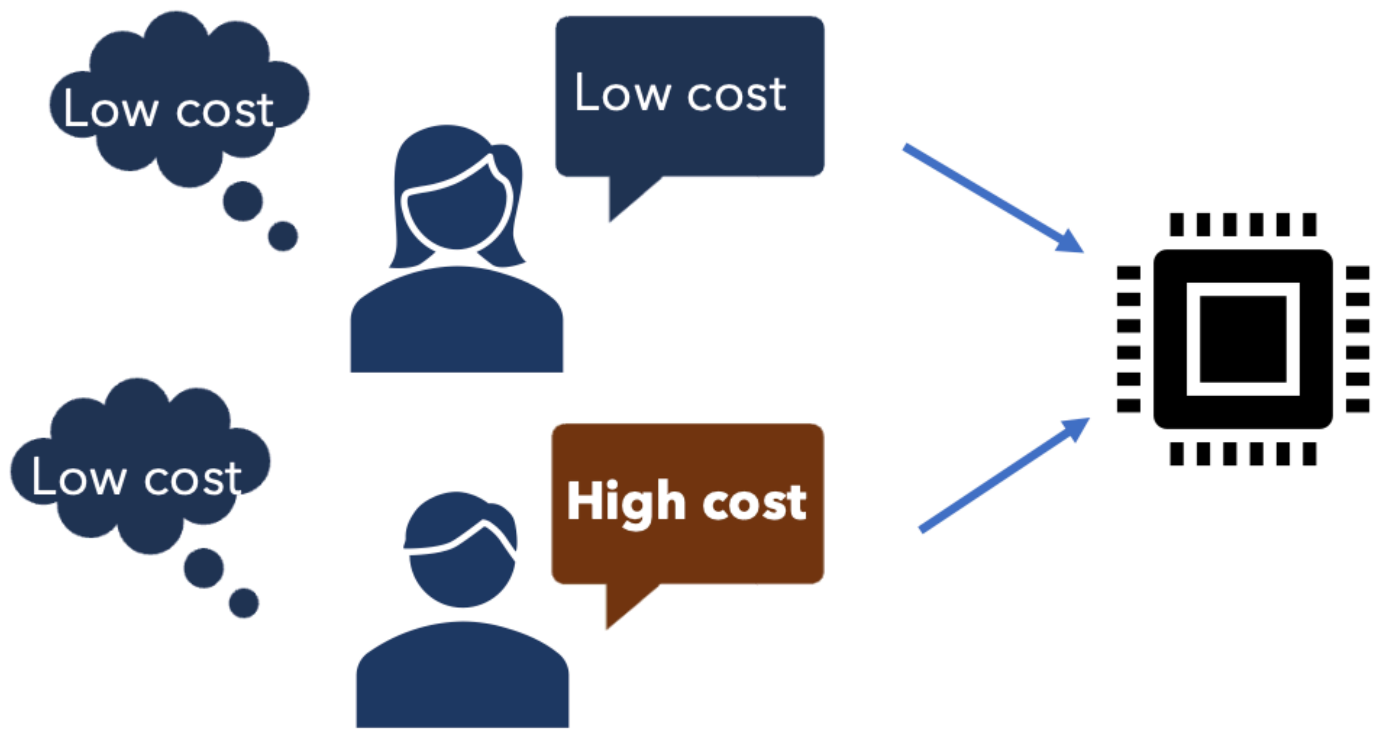

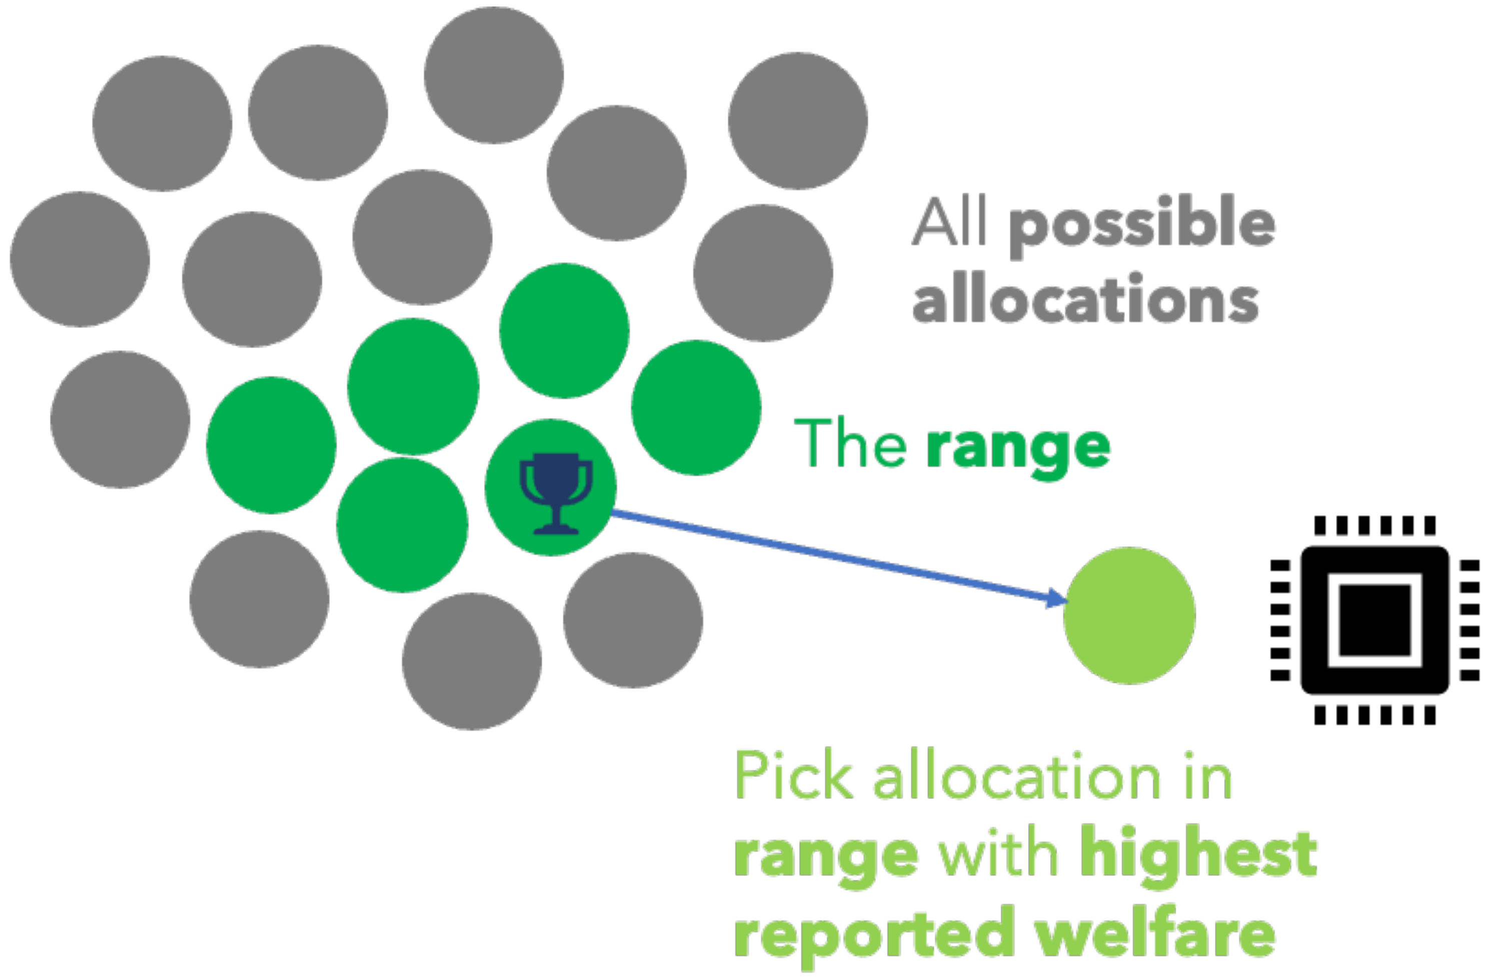

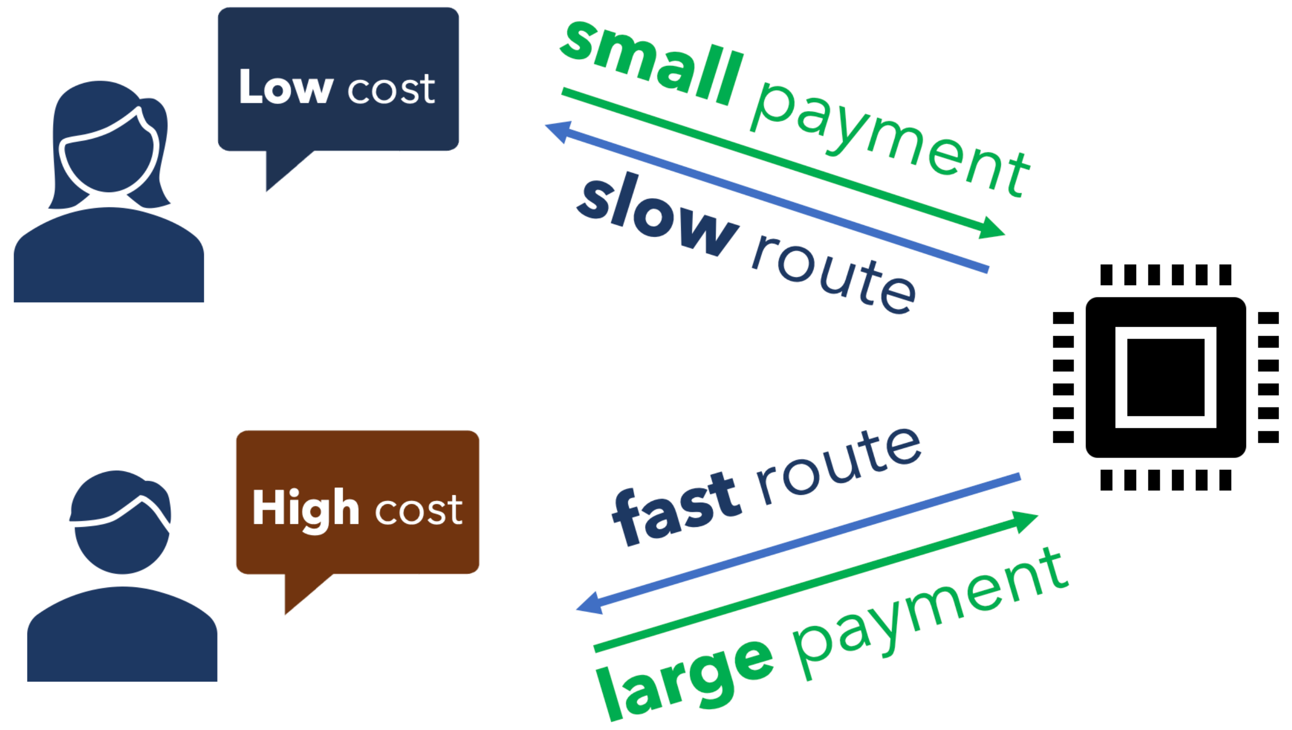

The mechanism is **strategy-proof** and **individually rational**, as long as agents **cannot influence the range by misreporting!**

Our three MAPF mechanisms

P-CBS: payment **conflict-based search**

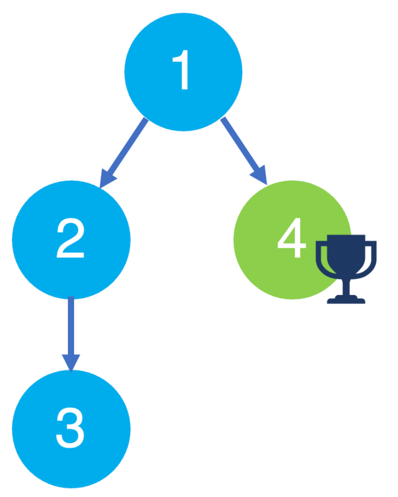

- Optimal MAPF
- VCG payments

E-PBS: exhaustive **priority-based search**

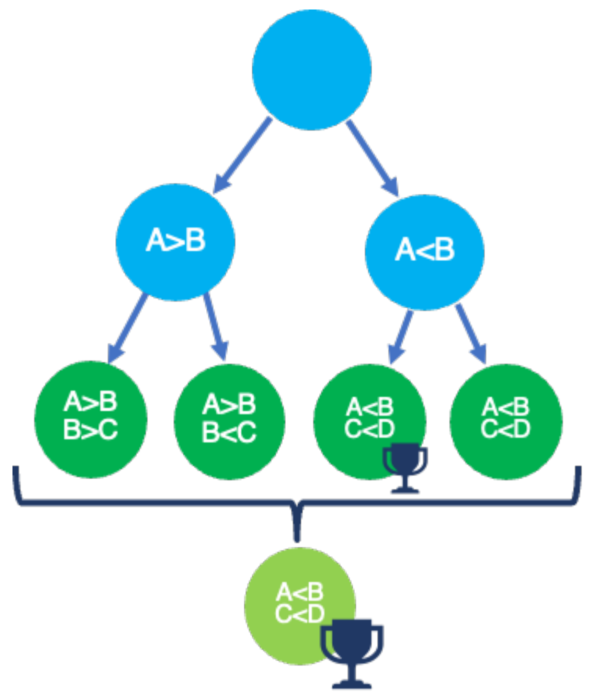

- Suboptimal, made strategyproof with MIR
- VCG-based payments "for free"

MC-PP: Monte-Carlo **prioritized planning**

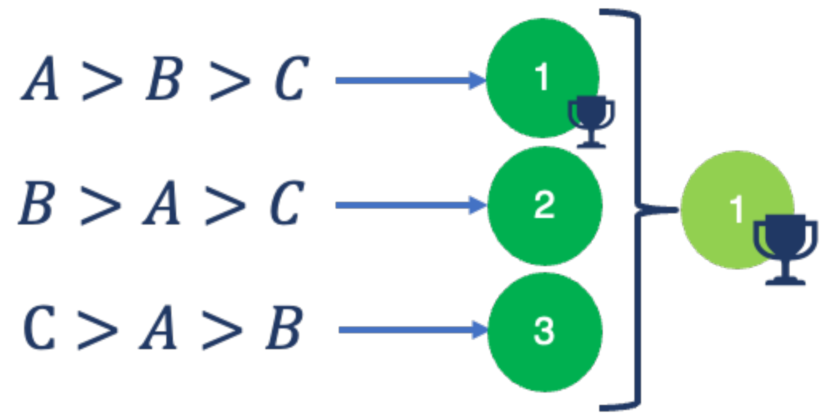

- Order agents by priority in m random ways
- Trade off scalability and optimality by choosing m
- VCG-based payments "for free"

Key takeaways

We highlight the **non-cooperative MAPF** domain

- **Standard MAPF** setting, large scale
- **Self-interested** agents

We make **fast MAPF algos strategyproof** with **maximal-in-range (MIR)**:

- Ensure that the allocation is optimal within a **fixed range of outcomes**
- Add **VCG-based payments** at **no computational cost**

➜ Experiments on 2D & 3D MAPF benchmark maps validate claims

Future work:

- Apply maximal-in-range (MIR) to other scalable MAPF algos
- Find better ranges to improve suboptimality, e.g. using learning

Scalable Mechanism Design for Multi-Agent Path Finding

Paul Friedrich, Yulun Zhang, Michael Curry, Ludwig Dierks,
Stephen McAleer, Jiaoyang Li, Tuomas Sandholm, Sven Seuken

Talk at IJCAI '24, 09 August 2024

University of Zurich UZH
ETH AI CENTER
HARVARD UNIVERSITY
Carnegie Mellon University
UIC UNIVERSITY OF ILLINOIS CHICAGO

Our mechanisms on MAPF benchmarks

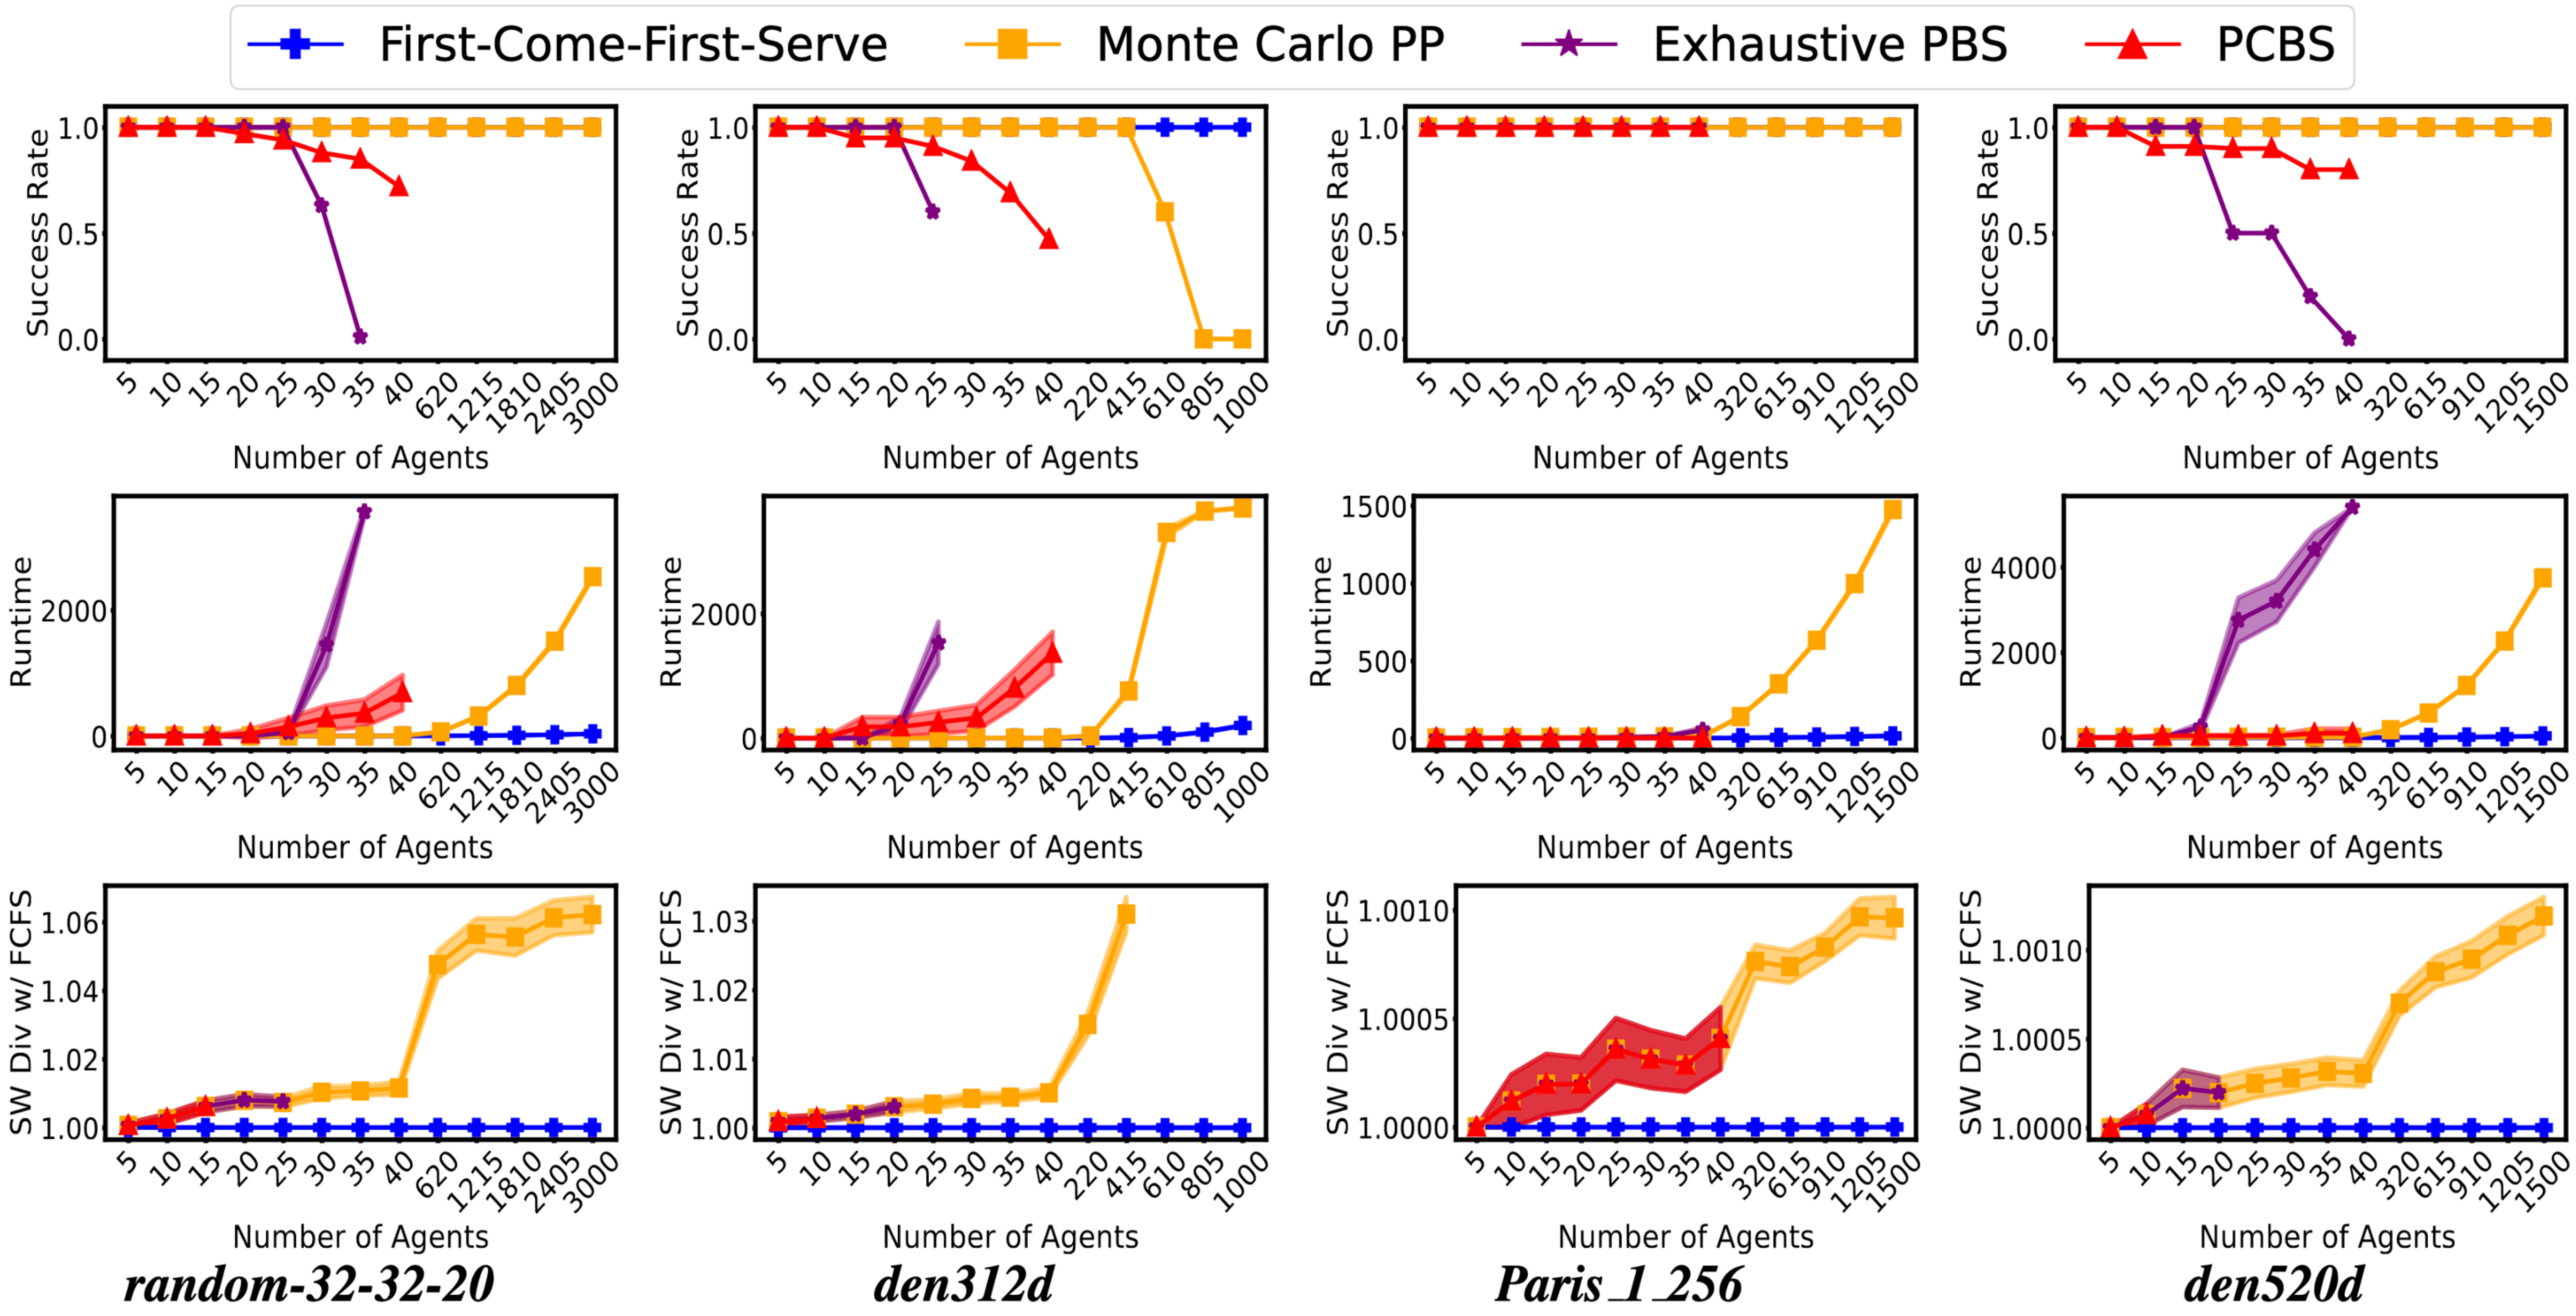

MCPP: trade off scalability and optimality

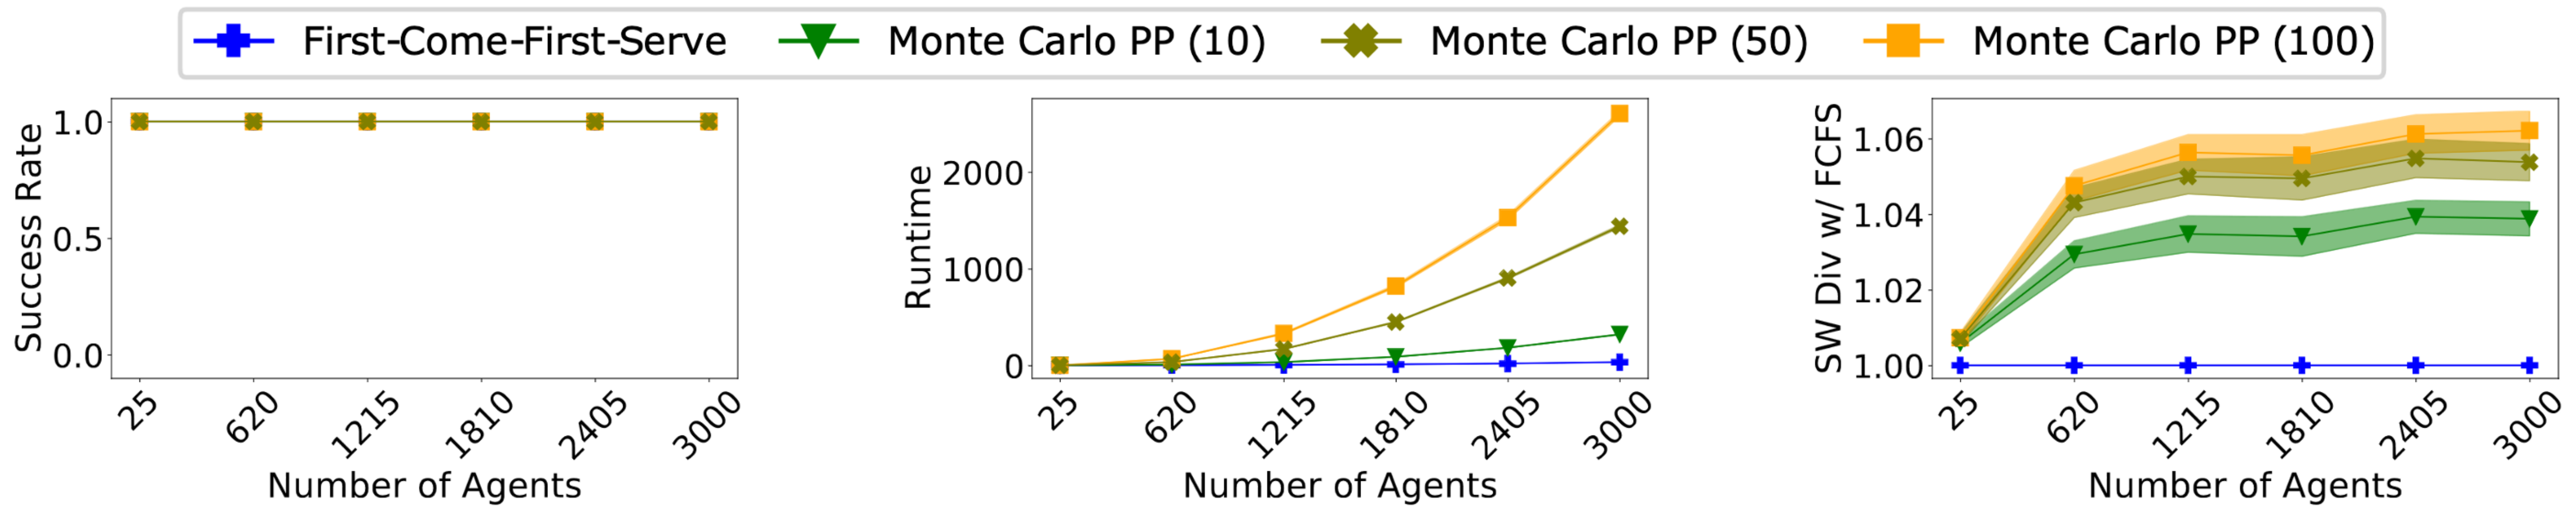

Our three MAPF mechanisms

Allocation | Payments | Purpose

PCBS: payment **conflict-based search**

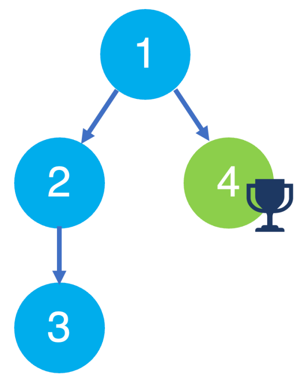

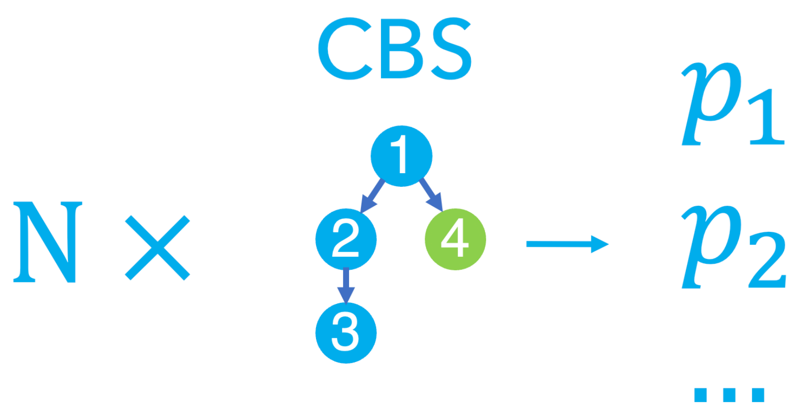

- Optimal benchmark
- Slow, but parallelizable

EPBS: exhaustive **priority-based** search

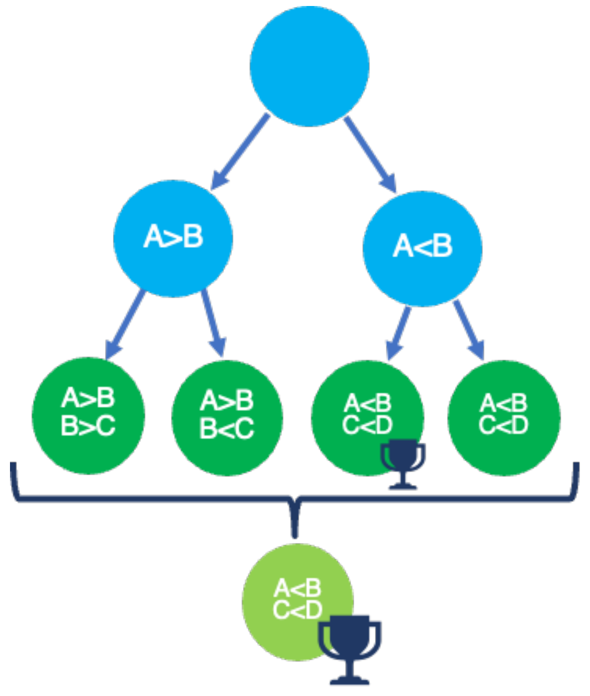

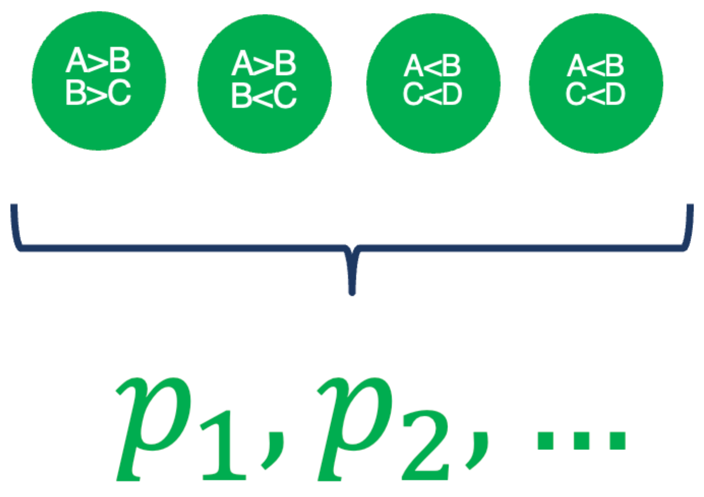

- How to make suboptimal MAPF strategy-proof with MIR
- Payments are “free"

MCPP: Monte-Carlo **prioritized planning**

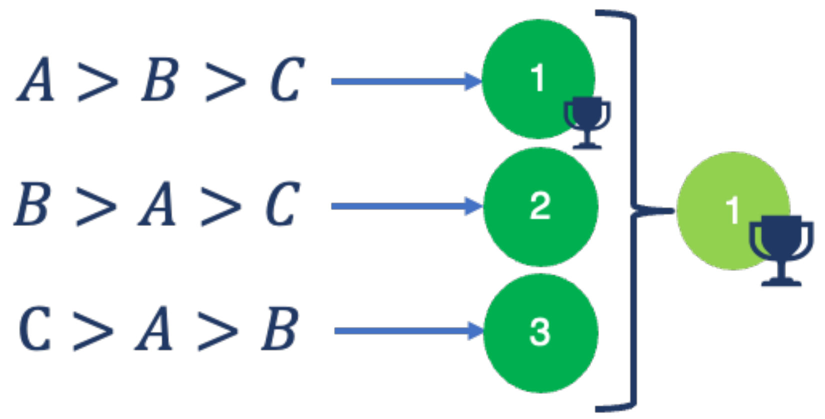

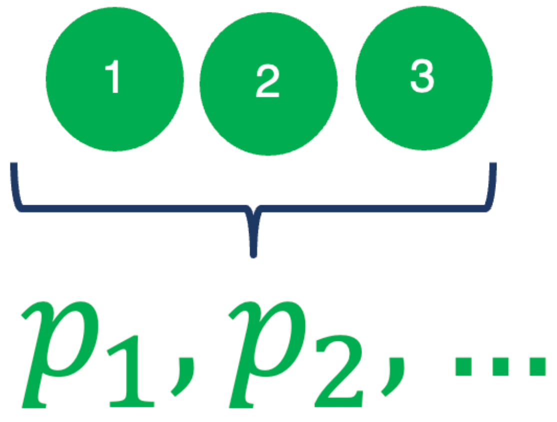

- Trade off scalability and optimality via choice of samples
- Parallelizable down to *n* times A*

Agent model & optimisation objective

Let $d = (\pi_1, \pi_2, \ldots, \pi_N) \in D$ an allocation, i.e. a set of conflict-free paths

MAPF:

- Minimise **sum-of-cost** (or "flowtime"): $\arg\min_{d \in D} \sum_{i=1}^{N} |\pi_i^d|$
- Minimize **makespan**: $\arg\min_{d \in D}(\max_{i \leq N} |\pi_i^d|)$

Our model: each agent reports a value for arrival v_i and cost per time travelled c_i

- Maximise **social welfare**: $\arg\max_{d \in D} \sum_{i=1}^{N} v_i - c_i |\pi_i^d|$

➜ Generalisation of sum-of-cost, identical for v=0, c=1