Multi-Agent Path Finding for Large Agents

Jiaoyang Li, Pavel Surynek, Ariel Felner, Hang Ma, T. K. Satish Kumar, Sven Koenig
jiaoyanl@usc.edu, pavel.surynek@fit.cvut.cz, felner@bgu.ac.il,
hangma@usc.edu, tkskwork@gmail.com, skoenig@usc.edu
USC University of Southern California
Czech Technical University in Prague
Ben-Gurion University of the Negev

1 Abstract

Multi-Agent Path Finding (MAPF) has been widely studied in the AI community. For example, Conflict-Based Search (CBS) is a state-of-the-art MAPF algorithm based on a two-level tree-search. However, previous MAPF algorithms assume that an agent occupies only a single location at any given time, e.g., a single cell in a grid. This limits their applicability in many real-world domains that have geometric agents in lieu of point agents. Geometric agents are referred to as "large" agents because they can occupy multiple points at the same time. In this paper, we formalize and study LA-MAPF, i.e., MAPF for large agents. We first show how CBS can be adapted to solve LA-MAPF. We then present a generalized version of CBS, called Multi-Constraint CBS (MC-CBS), that adds multiple constraints (instead of one constraint) for an agent when it generates a high-level search node. Experimental results show that all MC-CBS variants outperform CBS by up to three orders of magnitude in terms of runtime.

2 Application

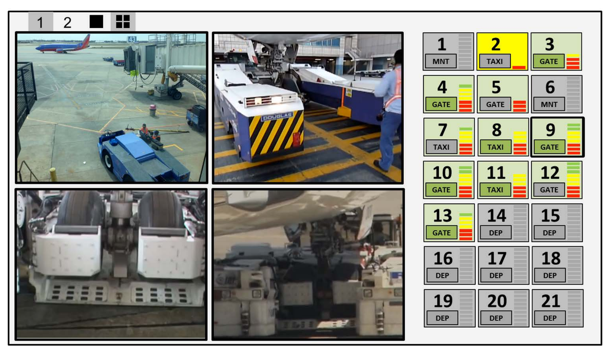

Figure 1: Autonomous towing vehicles for taxiing aircraft [1].

Figure 2: Autonomous warehouse robots [3].

Acknowledgements

The research at the University of Southern California was supported by the National Science Foundation (NSF) under grant numbers 1409987, 1724392, 1817189 and 1837779 as well as a gift from Amazon. The research was also supported by the United States-Israel Binational Science Foundation (BSF) under grant number 2017692 and the Czech Science Foundation (GACR) under grant number 19-17966S. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies or the U.S. government.

3 Multi-Agent Path Finding for Large Agent (LA-MAPF)

MAPF is the planning problem of finding a set of paths for a team of agents on a given graph. Each agent is required to move from a start location to a goal location, while avoiding collisions with other agents.

- MAPF: Agents are point agents. They collide when they are at the same vertex.
- LA-MAPF: Agents have extent. Different agents could have different shapes. They collide when their shapes overlap.

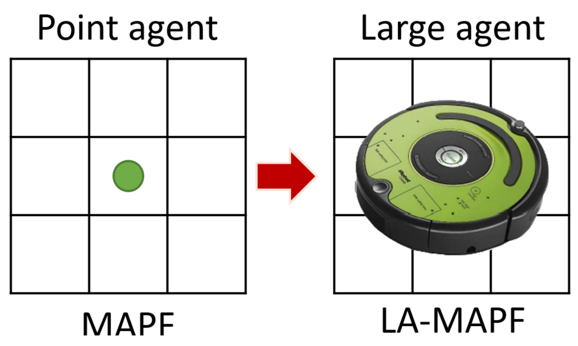

Figure 3: Agent model.

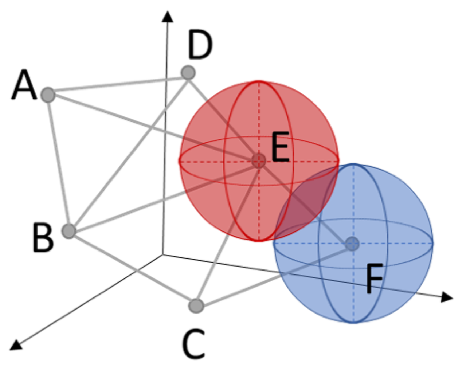

Figure 4: A collision example for LA-MAPF.

4 Conflict-Based Search (CBS) [2]

1. Find a path for every agent independently.
2. Check for collisions among paths.
3. If there is a collision <agent A, agent B, vertex v, time t>:
   - Option 1: prohibit A from being at v at t by adding constraint <A, v, t>.
   - Option 2: prohibit B from being at v at t by adding constraint <B, v, t>.
4. Repeat until finding collision-free paths.

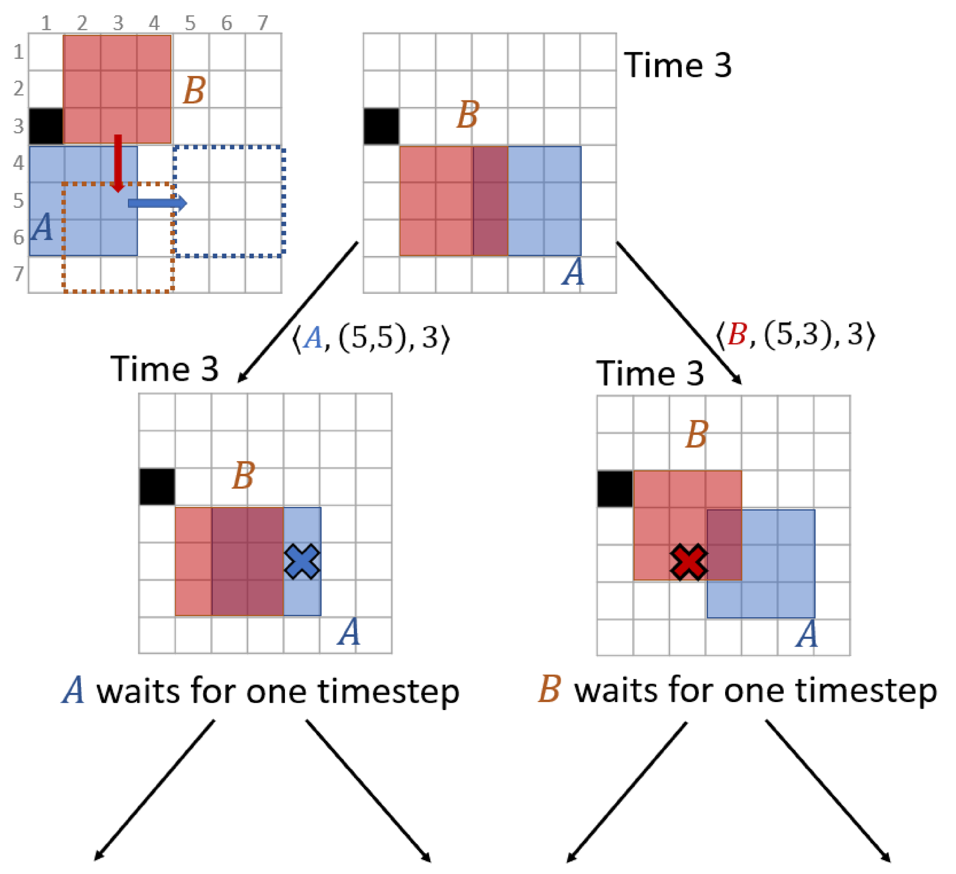

Figure 5: Applying CBS to solve an LA-MAPF instance.

* ✖/✖ means the center of A/B cannot be at cell ✖/✖.

Large agents usually have a large set of related collisions in close proximity. CBS has to resolve them one by one.

5 Multi-Constraint CBS (MC-CBS)

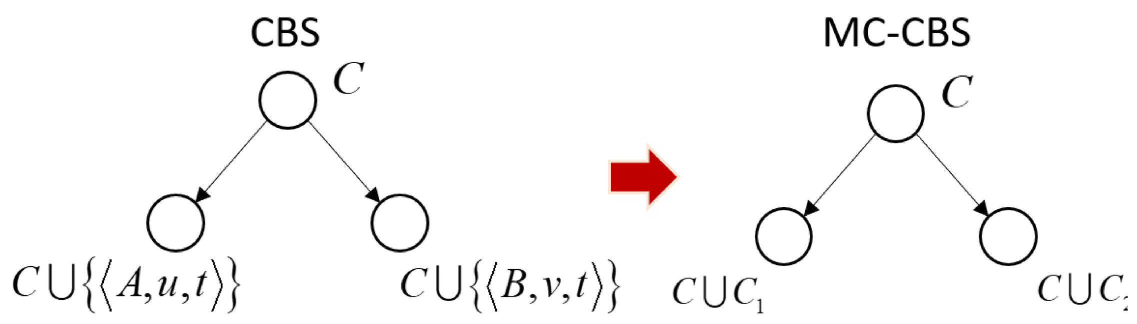

Figure 6: MC-CBS adds a constraint set, instead of a single constraint, to a child node.

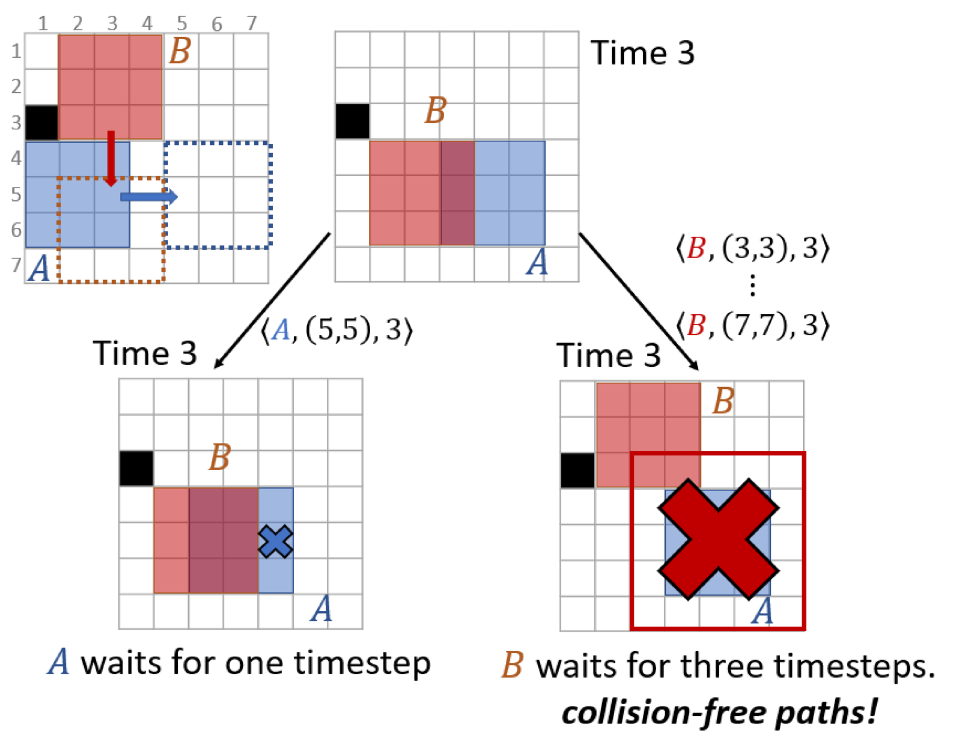

Figure 7: An MC-CBS example.

MC-CBS can resolve a set of related collisions in a single step.
To guarantee the optimality of MC-CBS, any pair of conflict-free paths should satisfy at least one of the constraint sets C_1 and C_2.

6 Mutually Disjunctive

Definition 1. Two constraints for A and B, respectively, are *mutually disjunctive* iff any pair of conflict-free paths for A and B satisfies at least one of the two constraints.

Definition 2. Two constraint sets are *mutually disjunctive* iff each constraint in one set is mutually disjunctive with each constraint in the other set.

Theorem 1. If two constraint sets C_1 and C_2 are mutually disjunctive, any pair of conflict-free paths satisfies at least one of the constraint sets.

Approaches for choosing mutually disjunctive constraint sets:
- ASYM: $|C_1| = 1$ and $|C_2| > 1$ (e.g., Fig. 5).
- SYM: $|C_1| \approx |C_2|$.
- MAX: C_1 and C_2 maximize the increased costs of the child nodes.

7 Experimental Results

Square agents on 2D grids.

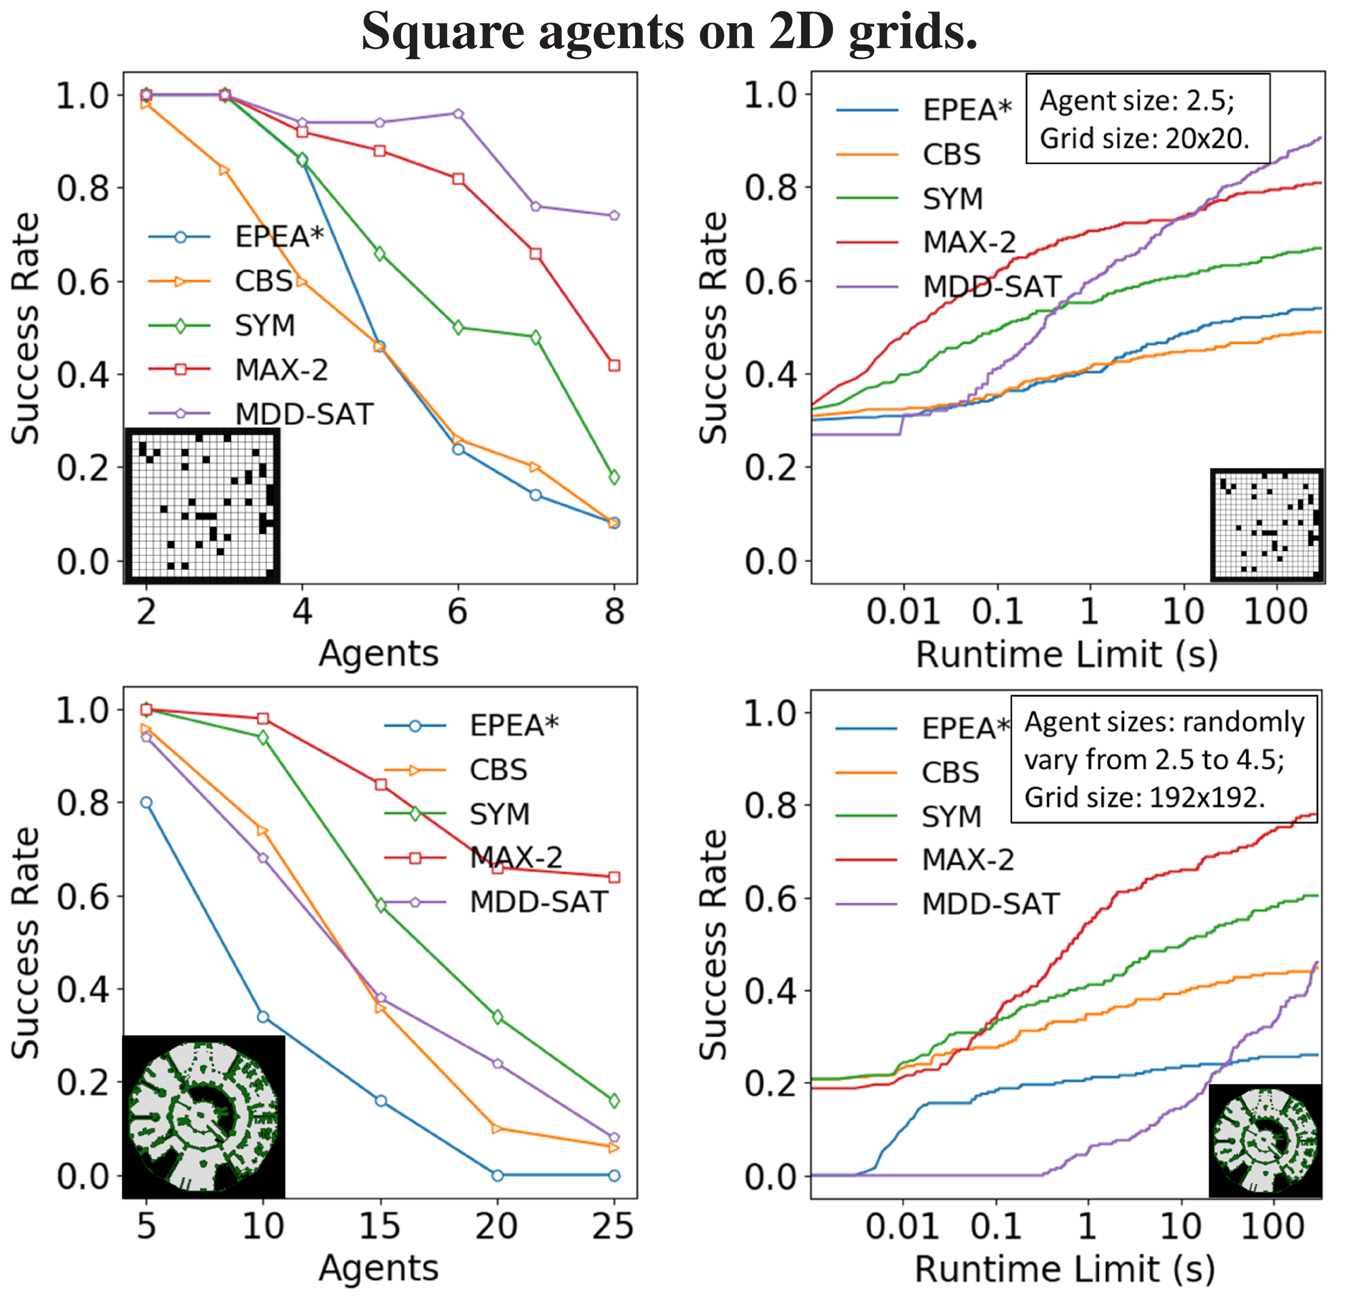

Figure 8: Success rates (=solved instances) on 2D grids. The runtime limit for left figures is 5 minutes. The right figures are results of all numbers of agents.

Ellipsoid agents on a 3D roadmap.

The environment is a 7.5 m × 6.5 m × 2.5 m space with (0.15 m, 0.15 m, 0.3 m) ellipsoid quadrotors. Quadrotors are required to fly through 3 windows in a wall in opposite directions. The roadmap has 869 vertices and 3,371 edges.

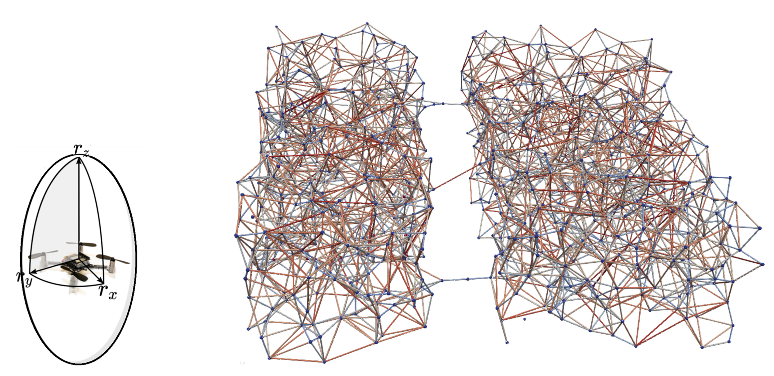

Figure 9: Agent shape and roadmap structure.

Figure 10: Runtime on the 3D roadmap.

Summary:
- All MC-CBS variants outperform CBS in all domains by up to 3 orders of magnitude in terms of runtime.
- MDD-SAT is strong for difficult problems in small domains.
- MAX is strong for easy problems or in large domains.

References

[1] R. Morris, M. L. Chang, R. Archer, E. Cross, S. Thompson, J. Franke, R. Garrett, W. Malik, K. McGuire, and G. Hemann. Self-driving aircraft towing vehicles: A preliminary report. In *Workshop on AI for Transportation*, pages 35–42, 2015.
[2] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015.
[3] P. R. Wurman, R. D'Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.