Intelligent Planning for Large-Scale Multi-Robot Coordination

Jiaoyang Li

Carnegie Mellon University
jiaoyangli@cmu.edu

Robots will play a crucial role in the future and need to work as a team in increasingly more complex applications. Examples include fully automated manufacturing, where heterogeneous robots accomplish collaborative tasks without human intervention beyond the capabilities of single robots, and smart cities, where traffic control systems coordinate autonomous cars or drones on a large scale. Advances in robotics have laid the hardware foundations for building large-scale multi-robot systems. But how to coordinate robots intelligently is a difficult problem. We believe that graph-search-based planning can systematically exploit the combinatorial structure of multi-robot coordination problems and efficiently generate solutions with rigorous guarantees on correctness, completeness, and solution quality. Our long-term goal is to develop algorithms that enable large teams of autonomous agents to accomplish collaborative tasks intelligently in dynamic environments.

To achieve this goal, we started with one problem that is central to many multi-robot applications. Multi-Agent Path Finding (MAPF) is an NP-hard problem of planning collision-free paths for a team of agents while minimizing their travel times. We addressed the MAPF problem from both (1) a theoretical perspective by developing efficient algorithms to solve large MAPF instances with completeness and optimality guarantees via a variety of AI and optimization technologies, such as constraint reasoning, heuristic search, stochastic local search, and machine learning, and (2) an applicational perspective by developing algorithmic techniques for integrating MAPF with task planning and execution for various multi-robot systems, such as (a) mobile robot coordination for automated warehousing and, in general, non-holonomic robot teams, (b) traffic management for airports railway networks, and autonomous road intersections, (c) drone swarm control, (d) multi-arm assembly, and (e) character control in video games. We highlight some of our research contributions below.

1. Pushed the limits of MAPF solving by systematically reasoning about the collision-resolution space. We developed many AI and optimization techniques that sped up optimal algorithms by up to 4 orders of magnitude, improved the solution quality of suboptimal algorithms by up to 36 times, and boosted their scalability by orders of magnitude (see Figure 1a).
2. Advanced robot coordination on a large scale for automated warehousing by generalizing MAPF to lifelong scenarios where robots are constantly engaged with new tasks in real-time. Our algorithm RHCR, for example, can coordinate 1,000 robots with high throughput in a highly-congested simulated warehouse, significantly outperforming traditional methods (see Figure 1b).
3. Bridged the gap between MAPF and various applications by integrating MAPF with task planning and execution. For example, one of our MAPF-based systems has won a NeurIPS’20 competition on railway planning and can operate thousands of trains on dense railway networks under uncertainty in real-time (see Figure 1c).

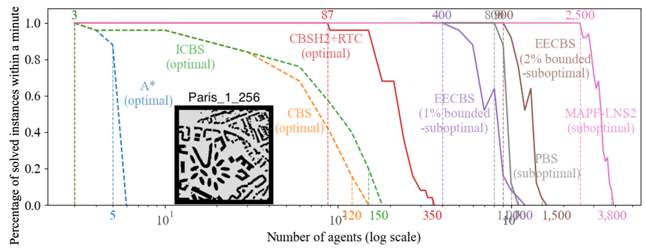

(a) MAPF benchmark instances.

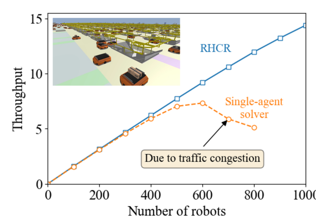

(b) Warehouse robots.

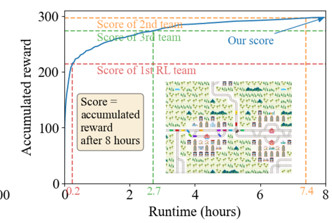

(c) Flatland Challenge for trains.

Figure 1: Comparing our MAPF-based solvers (solid lines) against other solvers (dashed lines) in various scenarios, i.e., (a) against previously leading MAPF solvers on MAPF benchmarks, (b) against traditional single-agent solvers in simulated warehouses, and (c) against other planning- or learning-based solvers on simulated railway networks.