Tracking Progress in Multi-Agent Path Finding

Bojie Shen, Zhe Chen, Muhammad Aamir Cheema, Daniel D. Harabor, Peter J. Stuckey

Faculty of Information Technology, Monash University, Melbourne, Australia

Problem

Multi-Agent Path Finding (MAPF) is a combinatorial problem that asks us to compute collision-free paths for teams of co-operative agents. Many works appear on this topic each year, and a large number of substantial advancements and improvements have been reported. However, measuring overall progress in MAPF is difficult due to the following reasons:

- There are many potential competitors.
- The computational burden for comprehensive experimentation is prohibitively large.
- The detailed data from past experimentation is usually unavailable.

Our objectives are to lower the barrier of entry for new researchers and to further promote the study of MAPF

Introduction

In this work, we introduce a new set of methodological and visualization tools to facilitate comparisons between a wide range of MAPF methods on standardised benchmarks available at [1]. In broad strokes:

- To track the progress of different types of algorithms, our system focuses on two important results reported by MAPF algorithms: (i) the best (i.e., largest) lower-bound value, and (ii) the best (i.e., smallest SIC) solution.
- Based on the collected results, we also propose clear indicators that can help establish state-of-the-art performance at each level of the benchmarks (see the right panels).
- Our system allows other researchers to participate by submitting their algorithms/results and collectively establish the state-of-the-art. We make all the results publicly available and provide the option for other researchers to download the results at each level.

To begin, we conducted a large set of experiments using several currently leading optimal and suboptimal solvers, aiming to map the current Pareto frontier. We implemented the system as an online platform, and our website is accessible at [2]. A demo video providing an overview of the system is also available at [3]. Furthermore, a full-length manuscript is available at [4].

Reference

[1] `https://movingai.com/benchmarks/mapf/index.html`.

[2] `http://tracker.pathfinding.ai`.

[3] `https://youtu.be/qtG6-h4FZxU`.

[4] B. Shen, Z. Chen, M. A. Cheema, D. D. Harabor, and P. J. Stuckey, "Tracking progress in multi-agent path finding." `http://tracker.pathfinding.ai/systemDemo`, 2023.

Domain and Map-level Tracking

- **Progress on Maps and Domains:** Our system generates plots to track the percentages of *closed*, *solved*, and *unknown* instances across all maps (see the example on the right). Additionally, it provides summaries of the related maps for each domain to create domain-level plots.
- **Comparison:** To facilitate researchers in evaluating their progress against other attempts, we provide tools for automatic comparison of algorithms across every level (see the examples at the bottom for domain-level comparison). Our principal evaluation criteria are: # of instances a given algorithm closed; # of instances that the algorithm solved; # of instances for which it achieved the best lower-bound; # of instances for which it reported the best solution.

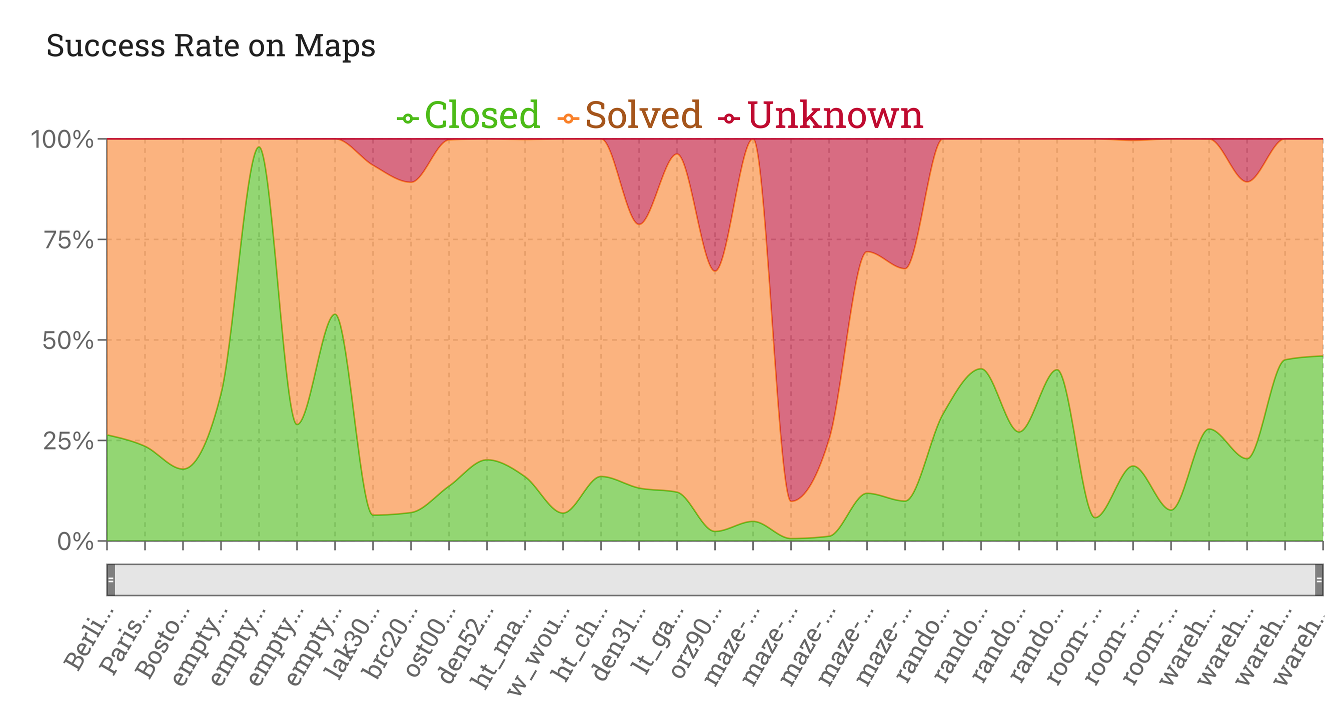

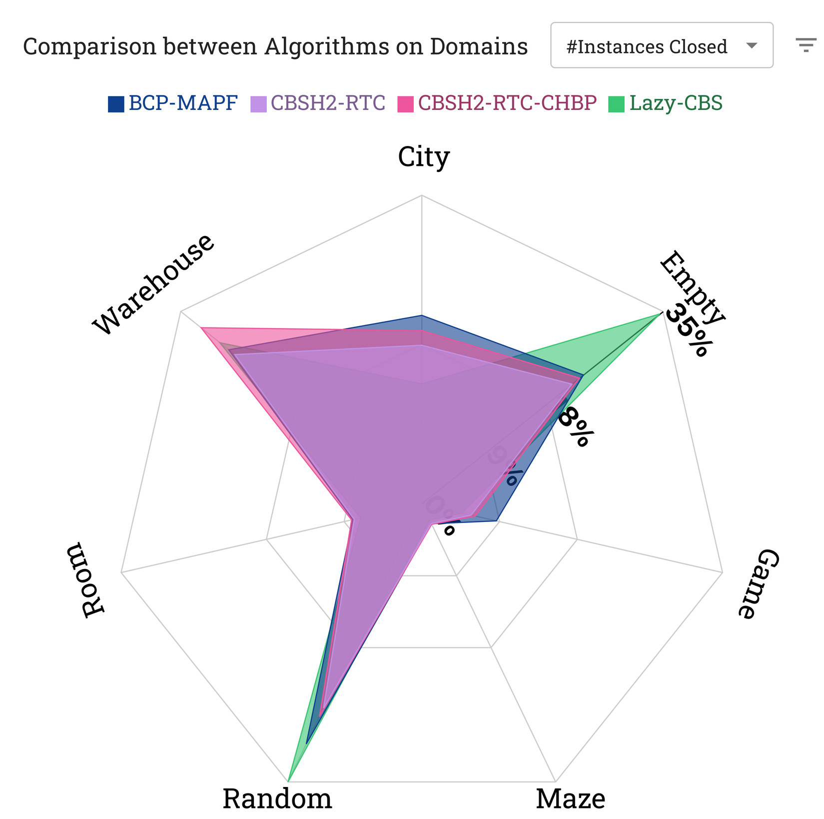

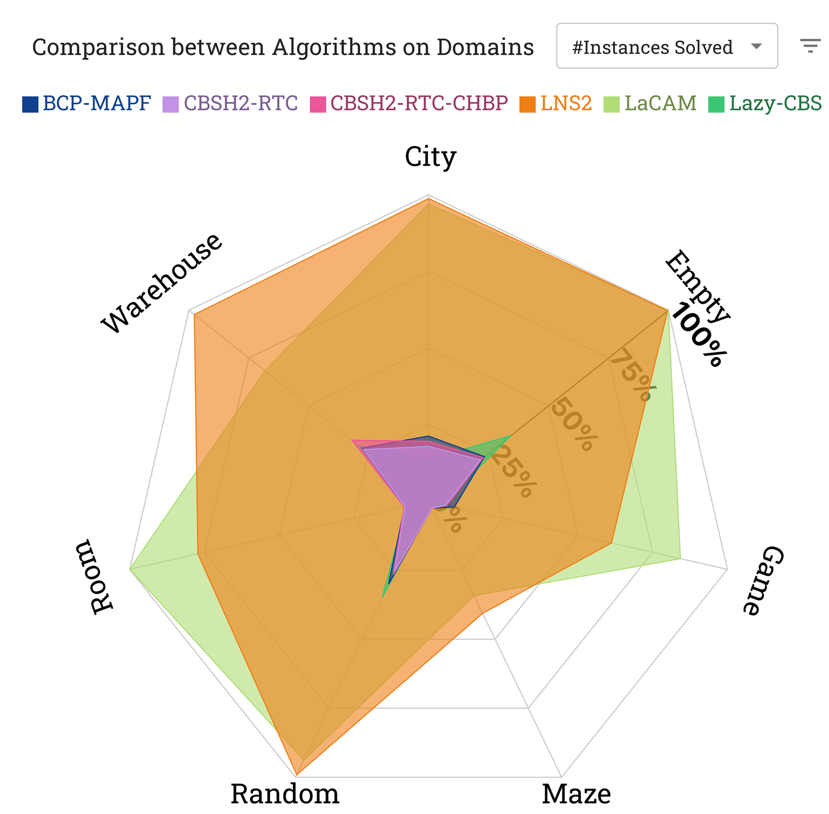

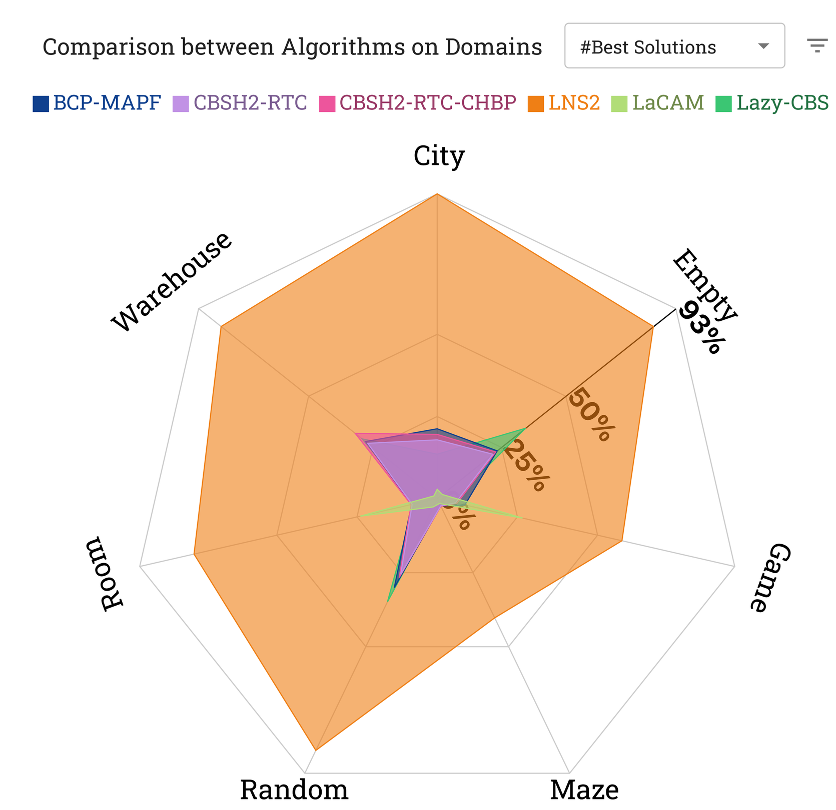

Scenario-level Tracking

- **Progress on Scenarios:** For a given map, our system automatically generates plots that show the percentage of closed, solved, and unknown instances for each scenario (e.g., see the example on the right).
- **Progress on # of Agents:** Each scenario contains instances with different numbers of agents. It is important to understand the scalability of MAPF algorithms across all scenarios, specifically at what number of agents progress stops. Therefore, our system includes the percentages of *closed*, *solved*, and *unknown* instances for different numbers of agents on the same map.

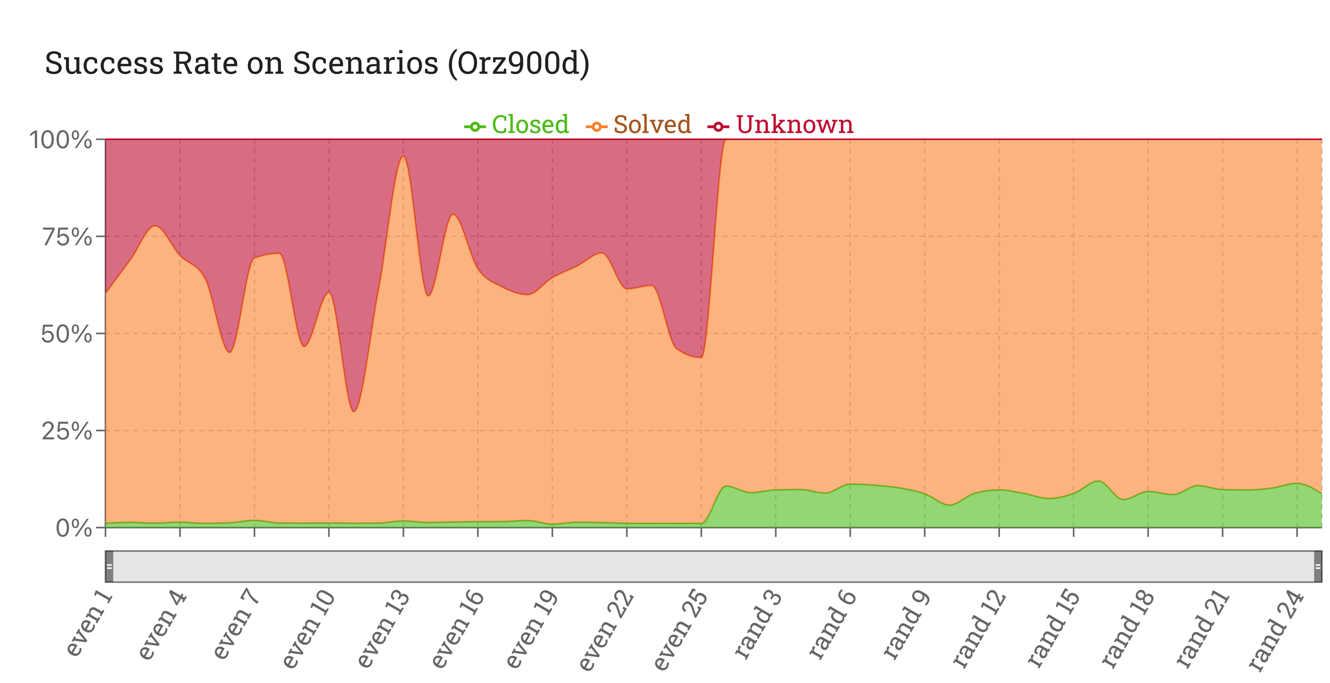

Instance-level Tracking

- **Concrete Plan:** Our system records a concrete plan for each best-known solution cost and provides a visualizer to better understand those solutions (e.g., see the example on the right).
- **Suboptimality Gap:** We automatically track and visualize the suboptimality ratio of each instance with different # of agents. The suboptimality ratio is defined as $(\mathcal{S} - \mathcal{L}) / \mathcal{L}$, where $\mathcal{L}$ and $\mathcal{S}$ represent the best known lower-bound and solution of the instance (e.g., see the example below).