Towards Multi-Robot Coordination under Temporal Uncertainty

Doctoral Consortium

Charlie Street
Oxford Robotics Institute, University of Oxford
Oxford, UK
cstreet@robots.ox.ac.uk

ABSTRACT

When robots act in an environment, there will be *temporal uncertainty* over the execution of their actions, i.e. the duration of an action and the time it takes place will be *stochastic*. The presence of multiple robots in the environment contributes towards this uncertainty. Temporal uncertainty is often disregarded in multi-robot coordination, and so we aim to develop planning solutions that explicitly model this uncertainty to generate effective plans.

KEYWORDS

Multi-robot systems; Planning under uncertainty; Markov models; Temporal uncertainty; Probabilistic guarantees

ACM Reference Format:
Charlie Street. 2020. Towards Multi-Robot Coordination under Temporal Uncertainty. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 2 pages.

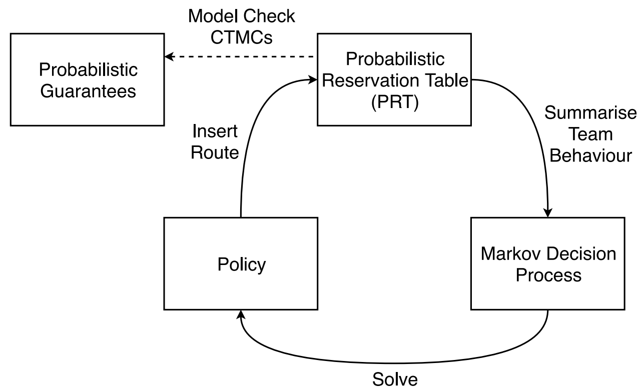

Figure 1: A high-level diagram of the congestion-aware planning framework.

1 INTRODUCTION

If there is *temporal uncertainty* over an action executed by a robot, the duration of that action and the time at which execution begins will vary. Almost any deployment of a multi-robot system will have a level of temporal uncertainty. For example, wheeled robots may suffer from tire slip when navigating, causing the robot to arrive late to its destination. The presence of multiple robots in the environment also contributes to this uncertainty. Existing methods for multi-robot planning commonly disregard temporal uncertainty [1, 3, 5]. Instead, these methods often make assumptions such as all actions having the same *fixed* duration, which simplifies planning at the cost of inefficient execution-time behaviour [8].

We wish to accurately model temporal uncertainty to enable effective multi-robot coordination, and will do so by modelling the duration of actions as continuous stochastic processes. With the introduction of rich continuous-time models of robot behaviour, we can then apply techniques from formal verification to obtain probabilistic guarantees over robot performance.

2 CONGESTION-AWARE PLANNING

We wish to solve multi-robot path planning problems, where robots have to reach their respective goals without colliding, in environments where the presence of multiple robots causes temporal uncertainty. Therefore, we present a novel planning framework that explicitly reasons over the effect the presence of multiple robots has on navigation performance [10]. We refer to this effect as *congestion*. An overview of the congestion-aware planning framework can be seen in Figure 1.

In this framework, a sequential planning assumption is made, with each robot considering those who planned before it. To summarise the effect of the other robots for the robot who is currently planning, we introduce a *probabilistic reservation table* (PRT). A PRT stores route information for robots who have planned, and so we can approximate the behaviour of robots without making assumptions about their behaviour, such as those seen in [5]. The PRT is inspired by the discrete deterministic reservation table presented in [9], although to consider continuous stochastic action durations the operation of the reservation table has to be changed significantly.

In this framework, we use continuous probability distributions to model continuous stochastic action durations. In particular we use phase-type distributions (PTD) to model action durations [4], as they are highly flexible in fitting distributions to empirical data, as well as being represented as continuous-time Markov chains (CTMC), which is useful for formal verification.

The planning framework in Figure 1 begins with an empty PRT. The first robot starts by building a Markov decision process (MDP) model of its environment, which can then be solved with heuristic search methods (e.g. [2]) to obtain a *policy*. This robot has no knowledge of the other robots as the PRT is empty. This policy informs the robot which edge to traverse given the robot's current location and time. Using the policy and PTDs, a CTMC is built and inserted into the PRT, which provides a rich continuous-time representation of a robot's policy, and is used to summarise their behaviour for future robots. For any subsequent robot, the first step is to obtain the route information of all robots who planned before

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

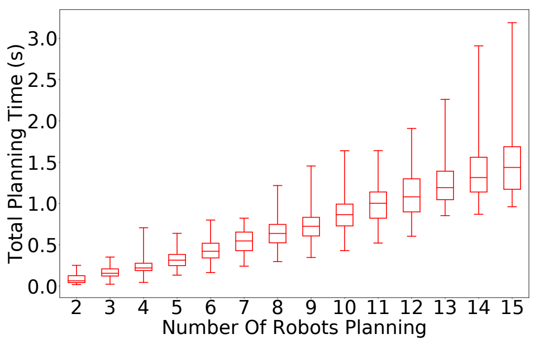

Figure 2: The scalability results for a warehouse-style map.

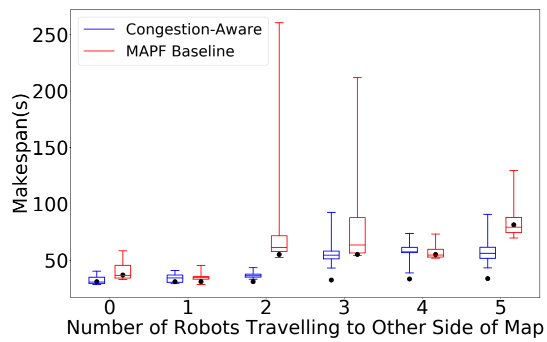

Figure 3: The execution-time performance results.

from the PRT. This information is then used in the construction of the robot's MDP, which will affect the generated policy.

After planning, we can use the CTMCS stored in the PRT to verify the continuous-time behaviour of the robot team. For example, we may wish to compute the probability that a robot arrives at its goal within 5 minutes. Due to the sequential planning assumption, an iterative verification procedure is required to update the CTMCs to take into account the effects of all other robots, which can then be verified using the PRISM model checker [7].

3 EXPERIMENTAL RESULTS

To demonstrate the effectiveness of the congestion-aware planning framework, we have analysed the scalability of planning, as well as the execution-time performance of the obtained policies. To test the scalability of the framework, we invoked the planner on a warehouse-style map using synthetic distributions to model the duration of navigation actions. For this experiment, 40 random configurations of robots were generated for 2-15 robot path planning problems. The results of this experiment can be seen in Figure 2, where the total planning time measures the total time for *all* robots to plan. This result shows there to be a sub-exponential increase in the total planning time, and so this framework mitigates the exponential state space increase seen when using joint planning models.

To test the performance of the obtained policies at execution time, we simulate a 5 robot setup in ROS using the Stage simulator on a warehouse-style map with two main sections linked by two tunnels, one of which is slightly longer. This allows robots to choose the longer tunnel if the shorter one is too congested. With duration distributions fit from empirical data, we generated 6 problem configurations, each one increasing the congestion by forcing an additional robot to travel through one of the tunnels. We compare our framework to a conservative baseline intended to emulate a multi-agent path finding (MAPF) solver [6]. We use the makespan to measure execution-time performance, i.e. the time taken for the last robot to arrive to its goal. The results of this experiment can be seen in Figure 3. These results show that the congestion-aware framework is able to route robots more effectively than the MAPF baseline, as it allows robots to be present in the same area simultaneously if the time cost incurred is lower than taking longer, less congested routes. In contrast, the MAPF baseline is conservative as it forces robots to stay away from each other.

4 CONCLUSIONS & FUTURE WORK

We have presented a framework for multi-robot planning under uncertainty that allows robots to reason over continuous stochastic action durations and congestion. Though we have focused on robot navigation, this framework is applicable to general actions and shared resources. We have also developed a procedure to compute probabilistic guarantees over the continuous-time behaviour of the robot team. Currently, we approximate the MDP models used for planning in order to highlight the benefits of the PRT. In future work, we will investigate more accurate planning approaches. We will also consider planning for collaborative tasks in this framework, while minimising unnecessary waiting time. Additionally, we plan to incorporate an online replanning mechanism into the framework, as well as demonstrate its effectiveness in real robot trials.

ACKNOWLEDGMENTS

This work is funded by the Honda Research Institute Europe GmbH.

REFERENCES

[1] Christopher Amato, George Konidaris, Ariel Anders, Gabriel Cruz, Jonathan P How, and Leslie P Kaelbling. 2016. Policy search for multi-robot coordination under uncertainty. *The International Journal of Robotics Research* 35, 14 (2016).
[2] Blai Bonet and Hector Geffner. 2003. Labeled RTDP: Improving the Convergence of Real-Time Dynamic Programming. In *ICAPS*.
[3] Craig Boutilier. 1996. Planning, Learning and Coordination in Multiagent Decision Processes. In *TARK '96*.
[4] Peter Buchholz, Jan Kriege, and Iryna Felko. 2014. *Input Modeling with Phase-Type Distributions and Markov Models: Theory and Applications*. Springer.
[5] Daniel Claes, Philipp Robbel, Frans A Oliehoek, Karl Tuyls, Daniel Hennes, and Wiebe Van der Hoek. 2015. Effective Approximations for Multi-Robot Coordination in Spatially Distributed Tasks. In *AAMAS*.
[6] Ariel Felner, Roni Stern, Solomon Eyal Shimony, Eli Boyarski, Meir Goldenberg, Guni Sharon, Nathan Sturtevant, Glenn Wagner, and Pavel Surynek. 2017. Search-Based Optimal Solvers for the Multi-Agent Pathfinding Problem: Summary and Challenges. In *SoCS*.
[7] M. Kwiatkowska, G. Norman, and D. Parker. 2011. PRISM 4.0: Verification of Probabilistic Real-time Systems. In *CAV '11*.
[8] João Vicente Messias, Matthijs Spaan, and Pedro Lima. 2013. GSMDPs for Multi-Robot Sequential Decision-Making. In *AAAI*.
[9] David Silver. 2005. Cooperative Pathfinding. In *AIIDE*.
[10] Charlie Street, Bruno Lacerda, Manuel Mühlig, and Nick Hawes. 2020. Multi-Robot Planning Under Uncertainty with Congestion-Aware Models. In *AAMAS*.