Carnegie Mellon University

THE ROBOTICS INSTITUTE

THESIS DEFENSE

Intra-Robot Replanning and Learning for Multi-Robot Teams in Complex Dynamic Domains

Wednesday, April 24, 2019

6501 Gates Hillman Center

11:00 a.m.

Thesis Committee:

Manuela Veloso
Chair

Maxim Likhachev

Stephen Smith

Daniel Borrajo Millán
Universidad Carlos III de Madrid

Philip Cooksey

Abstract

In complex dynamic multi-robot domains, there is a set of individual robots that must coordinate together through a centralized planner that inevitably makes assumptions based on a model of the environment and the actions of the individual. Eventually, the individuals may encounter failures, because the centralized planner's models of the states and actions are incomplete and the assumptions it makes are incorrect. In this thesis, we address the problem of what an individual robot must do when faced with such failures and can no longer execute the plan generated by the centralized planner.

While previous work has exclusively explored centralized approaches or decentralized approaches for dynamic multi-robot problems, it lacks the combination of a centralized approach with intelligent planning individuals, whom are often found in decentralized approaches. In centralized approaches, the focus has been on removing the need for replanning through conditional planning and policy generations, on hierarchical decomposition to simplify the multi-robot problem, or on predicting the informational needs of teammates. In decentralized approaches, the focus has been on improving auctioning algorithms, task decomposition, task assignment, and policy generation. In this thesis, I contribute a novel intra-robot replanning algorithm for the individual robots that autonomously handle failures with a set of pre-defined plans. To make local replanning feasible, I introduce a rationale-driven plan that provides the reasoning behind the choices made by the centralized planner. The intra-robot replanning algorithm then has a choice of how to fix the plan, given the set pre-defined plans and provided rationales, or of invoking the centralized planner. We can improve this process by learning scores for the pre-defined plans that are used by the intra-robot replanning algorithm to improve the performance of the robot.

This thesis is motivated by my previous work with individual robot replanning and centralized planners and the inability of their individual robots to handle failures. With autonomous underwater vehicles (AUVs), their approach to failures is to rise to the surface of the ocean, message the centralized controller, and wait for a new plan. With the Small-Size League soccer robots, their approach to failures is to continue, blindly executing a failing team plan, because taking at least some course of action, even if currently failing, can be better than inaction in an adversarial domain. Of course, continuing to execute a failing plan is an inadvisable approach, but the individual soccer robots lacked the necessary individual intelligence to fix the problem. These domains share the common thread of relying on the centralized planner to provide them with a new solution to handle failures. And, in doing so, they act inefficiently, are slow to react, and oftentimes have a higher cost in regard to the team's performance.

In this thesis, we describe our rationale-driven plan which details the reasoning for the actions and parameters chosen within the plan. We then explain our intra-robot replanning algorithm which uses the rationale-driven plan to replan locally. We evaluate the work with multiple domains in different environments to provide evidence of the effectiveness and generality of our approach. We then describe our method for learning and improving the intra-robot replanning algorithm. Lastly, we discuss possible future work that can expand upon the work in this thesis.