論文 / 著書情報
Article / Book Information

題目(和文)	
Title(English)	Planning, Execution, Representation, and their Integration for Multiple Moving Agents
著者(和文)	奥村圭祐
Author(English)	Keisuke Okumura
出典(和文)	学位:博士(工学), 学位授与機関:東京工業大学, 報告番号:甲第12460号, 授与年月日:2023年3月26日, 学位の種別:課程博士, 審査員:DEFAGO XAVIER,渡部 卓雄,石井 秀明,村田 剛志,下坂 正倫
Citation(English)	Degree:Doctor (Engineering), Conferring organization: Tokyo Institute of Technology, Report number:甲第12460号, Conferred date:2023/3/26, Degree Type:Course doctor, Examiner:,,,,
学位種別(和文)	博士論文
Category(English)	Doctoral Thesis
種別(和文)	審査の要旨
Type(English)	Exam Summary

（博士課程）

論文審査の要旨及び審査員

報告番号	甲第　号			学位申請者氏名	奥村　圭祐	
論文審査 審　査　員		氏　名	職　名		氏　名	職　名
	主査	DEFAGO, Xavier	教授	審査員	下坂　正倫	准教授
	審査員	石井　秀明	教授			
		村田　剛志	教授			
		渡部　卓雄	教授			

論文審査の要旨（2000 字程度）

The dissertation entitled *"Planning, Execution, Representation, and Their Integration for Multiple Moving Agents"* consists of 13 chapters and several appendices. The nine chapters presenting the original contributions (chap.4-12) are grouped into three parts: planning, execution, and representation.

Chapter 1 "Introduction" provides an overview of the research problem, its motivations, the state-of-the-art, and clearly states the main contributions made in the dissertation. The chapter presents the global picture and puts the contributions in the context of multi-agent and multi-robot systems.
Chapter 2 "Preliminaries" provides a synthetic reminder of important concepts used throughout the thesis and defines the relevant notation. It contributes to keeping the dissertation self-contained.
Chapter 3 "Background" details the general concept of multi-agent navigation and its multiple variants. The chapter presents the core problems of multi-agent path finding (MAPF), multi-robot motion planning (MRMP), planning with dynamic task assignment (unlabeled-MAPF), and planning execution robust to timing uncertainties or faults. The chapter provides a thorough review of the state-of-the-art organized into a solid taxonomy.

Part I "Planning" deals with the methods of computing optimal or high-quality plans for all agents within a reasonable runtime.
Chapter 4 "Short-Horizon Planning for MAPF" presents the PIBT algorithm which relies on priority inheritance and backtracking to solve the MAPF problem iteratively. The chapter identifies and formally proves important properties such as deterministic reachability. In addition, the algorithm is evaluated both by simulations and experimentally on real robots.
Chapter 5 "Short-Horizon Planning for Unlabeled-MAPF" studies a variant of the MAPF problem wherein goals are not assigned to any agent a priori. As a result, the problem of Unlabeled-MAPF combines the assignment of goals to each agent with the computation of their paths. The chapter present the TSWAP algorithm (and several variants) which is based on similar principles to PIBT. The paper proves important properties of the algorithm (and variants), such as termination. The quality of the solutions and the scalability of the algorithm is evaluated by extensive simulation.
Chapter 6 "Short-Horizon Planning Guides Long-Horizon Planning" introduces the LaCAM algorithm which combines long-horizon planning with the short-horizon planning of the previous algorithms. The LaCAM algorithm uses PIBT (or similar algorithms) as a primitive to guide the search for an optimal solution. As a result, LaCAM guarantees that a solution is always found if it exists and that such a solution is found very quickly most of the time. The completeness is proved formally, and the claim of quick and scalable runtime is assessed experimentally. The LaCAM algorithm greatly improves the state-of-the-art in terms of scalability and runtime.

Chapter 7 "Improving Solution Quality by Iterative Refinement" considers the problem of so-called anytime algorithms, where a preliminary solution is computed quickly and then refined iteratively, so that the quality of the solution improves as time passes. This is a very useful property in real-time and interactive systems. The proposed framework integrates the finding of initial solution (e.g., PIBT) with a refinement method. The design choices are evaluated by simulation on well-accepted benchmarks.

Part II "Execution" deals with the reality gap that exists between the plans that are optimal in an ideal world with their execution that can be subjected to uncertainties such as timing issues or faults.
Chapter 8 "Online Planning to Overcome Timing Uncertainties in Execution" considers the problem of dynamically reevaluating the planning according to environmental changes. The proposed idea extends proposed plans to deal with time-independent parts. The chapter focuses on unlabeled-MAPF and presents a variant of the TSWAP algorithm called online-TSWAP. The concept can however be equally applied to classical MAPF. The method is evaluated on a testbed consisting of real robots and via a form of fault injection.
Chapter 9 "Offline Planning to Overcome Timing Uncertainties in Execution" observes that some plans are inherently robust to timing uncertainties because no race conditions can possibly lead them to a deadlock situation. The chapter formalizes this notion and formally establish the complexity of finding such solutions. While the general problem is intractable, the chapter presents an efficient algorithm for a practical subset of the problem space. The method is evaluated by conducting stress-tests in simulation and on two different platforms of real robots.
Chapter 10 "Offline Planning to Overcome Crash Faults in Execution" extends the previous work to study the problem of tolerating the crash of some of the agents during execution. The chapter proves necessary conditions for the existence of fault-tolerant plans and establishes the intractability of the general case. The chapter proposes an algorithm that can solve a subset of problem instances in acceptable time and evaluates its effectiveness by stress test conducted in a simulated environment.

Part III "Representation" deals with the problem of representing the environment in a way that helps with the planning. In particular, the problem is to provide a tessellation of a continuous environment such that the discretization remains small and hence allows for a reduced search space for the planning algorithms.
Chapter 11 "Building Representation from Learning" proposes the notion of cooperative timed roadmaps (CTRM) combined with the use of machine learning to reduce the search space for path planning. The chapter evaluates the approach via simulation with very promising results and finds that representation is indeed very important to effective planning.
Chapter 12 "Building Representation while Planning" integrates the representation of the environment with the planning and presents the SSSP algorithm *(simultaneous sampling and planning)*. Here, instead of dividing the problem into two independent phases, the roadmap construction and the agent planning are done together. The chapter applies the approach to multiple path planning contexts, in both 2D and 3D and various degrees of freedom (multi-robots, robotic arm, vehicles with kinematic constraints). The concept is also illustrated on a platform of real robots.

Finally, **Chapter 13 "Conclusion and Discussion"** reviews the conclusions of each chapter and puts them in a larger context. The chapter outlines several promising research questions that stem from this research.

As presented above, the dissertation makes outstanding and very significant contributions to the field of multi-agent planning. In particular, the results of this research are key to improving the performance, scalability, and robustness of multi-agent and multi-robot planning when used in practical contexts. The results are extensive, and their impact are significant both on scientific and on engineering grounds. For those reasons, the dissertation and its research fulfil all requirements for a doctoral degree in engineering.