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)	論文要旨
Type(English)	Summary

（博士課程）
Doctoral Program

論 文 要 旨

THESIS SUMMARY

系・コース： Department of, Graduate major in	情報工学 系 知能情報 コース	申請学位（専攻分野）： Academic Degree Requested	博士 （工学） Doctor of
学生氏名： Student's Name	奥村 圭祐	指導教員（主）： Academic Supervisor(main)	DEFAGO Xavier
		指導教員（副）： Academic Supervisor(sub)	

要旨（英文 800 語程度）
Thesis Summary（approx.800 English Words ）

Navigating a team of agents without colliding with each other plays a crucial role in the modern and coming automation era, including fleet operations in warehouses, as well as collaborative robotic construction, to name just a few. Aiming at providing foundations for such broad applications from computational aspects, this dissertation devotes to developing quick, scalable, near-optimal, robust, domain-independent, and end-to-end multi-agent navigation technologies. Multi-agent navigation is a very complicated art based on compositing many technical components from artificial intelligence and robotics. Therefore, the dissertation decomposes multi-agent navigation into three perspectives, that is, planning, execution, and representation, and respectively overcomes current limitations of cutting-edge technologies.

The first part studies planning, which asks how to determine a sequence of actions for agents. The corresponding problem of multi-agent navigation is often formulated as multi-agent pathfinding (MAPF) which asks for a list of collision-free paths on graphs for multiple agents. The primary challenge of MAPF is to maintain solvability and quality while suppressing computational effort. On one hand, state-of-the-art optimal MAPF algorithms have difficulty in solving instances from the grid MAPF benchmark, containing a few hundred agents, within realistic timeframes. On the other hand, sub-optimal algorithms can cope with massive instances in a short time (e.g., less than 30 seconds), meanwhile, such algorithms often lack completeness. Indeed, they often fail condensed or cluttered MAPF instances, even if less than ten agents are involved. Aiming at breaking this tradeoff, the dissertation first presents algorithms with short horizon planning called PIBT and TSWAP, respectively developed for solving MAPF iteratively, and, simultaneous target assignment and collision-free pathfinding. Then, the LaCAM algorithm is presented, which uses short-horizon planning like PIBT and TSWAP as a sub-procedure. LaCAM is complete for MAPF, furthermore, it eventually converges to optima, provided that solution cost is accumulative transition costs. As another direction, the dissertation also presents the framework of iterative refinement, enabling us to improve the quality of arbitrary MAPF solutions. Empirically, the dissertation demonstrates that these proposed methods have excellent performance in success rate, computation time, and solution quality, significantly outperforming existing MAPF technologies.

The second part studies execution, which asks how to achieve robust plan execution by agents under various uncertainties in the real world. The primary challenge here is, at runtime, how to ensure safety (i.e., no collision) and liveness (i.e., eventually reaching destinations) when something bad happens, unexpected from the planning phase. To this end, the dissertation studies a novel integration style of planning and execution, namely, deliberative offline planning assuming that agents reactively execute the plan at runtime. Two types of example studies are presented, called the OTIMAPP and MAPPCF problems, respectively for timing uncertainties and crash faults. For both proposed problems, theoretical foundations, such as computational complexities, as well as practical approaches to solving the problems are provided. As proofs-of-concept, demonstrations of decentralized execution with real robots are also included, while ensuring liveness, without any central intervention at runtime, and without any global interactions. Such things can be achieved neither by conventional centralized execution styles that rely on global monitoring systems nor by decentralized execution styles that lack centralized planners.

The third and last part studies representation, which asks how to model the world for agents from infinite design choices. Considering representation issues is necessary to realize end-to-end multi-agent navigation. The primary challenge here is how to construct small but effective search spaces for subsequent planning. It is necessary to construct a sparse representation of the workspace (i.e., roadmaps), otherwise, it becomes dramatically difficult to find a combination of plausible paths because of having to manage a higher number of inter-agent collisions. Nevertheless, roadmaps should be sufficiently dense to ensure high planning solvability and better solutions. To break this tradeoff, the dissertation provides two directions. The first approach is learning to construct sparse roadmaps from planning demonstrations, in short, data-driven roadmap construction. The second approach is combining roadmap construction and multi-agent search (i.e. collision-free path planning), making it possible to develop a small but effective search space such that the multi-agent search is willing to use. Both concepts are extensively tested in various scenarios, revealing their power, i.e., solving more planning instances much faster compared to existing methods.

Putting everything together, the dissertation presents a consistent story to realize multi-agent navigation technologies applicable to various domains.