Extended Abstract: A Competitive Analysis of Online Multi-Agent Path Finding

Hang Ma
Simon Fraser University
hangma@sfu.ca

Abstract

This is an extended abstract of a paper to be published at ICAPS 2021 (Ma 2021). We study online Multi-Agent Path Finding (MAPF), where new agents are constantly revealed over time and all agents must find collision-free paths to their given goal locations. We generalize existing complexity results of (offline) MAPF to online MAPF. We classify online MAPF algorithms into different categories. We present several complexity and competitiveness results for online MAPF and its algorithms, which provides theoretical insights into the effectiveness of using MAPF algorithms in an online setting for the first time.

Introduction

Online Multi-Agent Path Finding (MAPF) (Švancara et al. 2019) models the problem of finding collision-free paths for a stream of incoming agents in a given region. Its applications include autonomous intersection management (Dresner and Stone 2008), UAV traffic management (Ho et al. 2019), video games (Ma et al. 2017b), and automated warehouse systems (Wurman, D'Andrea, and Mountz 2008). Existing research has conducted empirical evaluations of several online MAPF algorithms (Švancara et al. 2019) based on recent techniques for (offline) MAPF (Stern et al. 2019). However, there is still a lack of theoretical understanding of online MAPF and its algorithms. In this paper, we thus perform a theoretical analysis from the points of view of competitive analysis and complexity theory.

Related Work

Offline MAPF: Online MAPF is an extension of the well-studied problem of (offline) MAPF (Ma and Koenig 2017; Stern et al. 2019), where all agents are known and start routing at the same time. MAPF is NP-hard to solve optimally for flowtime (the sum of the arrival times of all agents at their goal locations) minimization and to approximate within any constant factor less than 4/3 for makespan (the maximum of the arrival times of all agents at their goal locations) minimization (Surynek 2010; Yu and LaValle 2013b; Ma et al. 2016). It is NP-hard to solve optimally even on planar graphs (Yu 2015) and 2D 4-neighbor grids (Banfi, Basilico, and Amigoni 2017). MAPF algorithms include reductions to other combinatorial problems (Yu and LaValle 2013a; Erdem et al. 2013; Surynek et al. 2016) and specialized algorithms (Luna and Bekris 2011; Wang and Botea 2011; Sharon et al. 2013, 2015; Boyarski et al. 2015; Cohen et al. 2018; Ma et al. 2019a; Li et al. 2019a,b; Lam et al. 2019; Gange, Harabor, and Stuckey 2019; Li et al. 2020).

Online Problems: Ma et al. (2017a) and Ma et al. (2019b) have considered an online version of MAPF where a given set of agents must attend to a stream of tasks, consisting of (sub-)goal locations to be assigned to the agents, that appear at unknown times. This version considers the entire environment instead of a region of a system and thus does not consider the appearance and disappearance of agents. Švancara et al. (2019) and Ho et al. (2019) have considered another online version of MAPF, similar to the setting of this paper, where a stream of agents with preassigned goal location appear at unknown times. Algorithms for solving such online problems reduce each problem to a sequence of (offline) MAPF sub-problems that are solved by a MAPF algorithm. The effectiveness of these algorithms is characterized by objective functions that measure how soon the tasks are finished or the agents are routed to their goal locations. Existing study on online versions of MAPF has been empirical only. For example, both Ma et al. (2017a) and Švancara et al. (2019) have experimentally shown that algorithms that allow agents (that have paths already) to replan their paths and reroute tend to be more effective than those that do not. However, there is still a lack of theoretical understanding of solving MAPF in an online setting.

Assumptions and Contributions

We follow most of the notations of Švancara et al. (2019) and consider the setting where new agents can wait infinitely long before entering a given region and agents disappear upon exiting from the region because (1) existing online MAPF algorithms have been designed and tested only for this setting (Ho et al. 2019; Švancara et al. 2019), although other settings concerning what happens before agents enter the region and after agents leave the region have (only) been mentioned (briefly) by Stern et al. (2019); Švancara et al. (2019) and (2) queuing at entrances and exits of such an intersection region is handled by a task-level planner/scheduler with reserved queuing spaces (for example, queues in inventory stations and along the single-lane corridors in the storage region) in automated warehouses (Wurman, D'Andrea, and Mountz 2008; Kou et al. 2020) and many other real-world systems. We view the problem from the point of view of competitive analysis and thus assume that the algorithms have no knowledge of future arrivals of agents, as in the case of all existing online MAPF algorithms (Ho et al. 2019; Švancara et al. 2019), although such knowledge might be learned in practice.

As our first contribution, we formalize online MAPF as an extension of (offline) MAPF and demonstrate how to generalize existing NP-hardness and inapproximability results Ma and Pineau (2015); Ma et al. (2018) for MAPF to online MAPF. Specifically, we show that online MAPF is NP-hard to approximate within any constant factor less than 4/3 for makespan minimization, to solve optimally for flowtime minimization, and to approximate within any factor for latency minimization.

As our second contribution, we classify online MAPF algorithms based on (1) different controllability assumptions, namely at what time the system can plan paths for which sets of agents, into three categories: PLAN-NEW-SINGLE that plans only a path for one newly-revealed agent at a time, PLAN-NEW that plans paths only for newly-revealed agents, and PLAN-ALL that plans paths for all known agents and thus allows rerouting and (2) different rationality, namely how effective the planned paths are: optimally-rational algorithms that plan optimal paths for the given set of agents and rational algorithms (which are, in our opinion, the only algorithms worth considering, assuming no knowledge of future arrivals of agents) that plan paths at least asymptotically as effective as the naive baseline algorithm SEQUENCE that routes newly-revealed agents one at a time in sequence. These classifications cover all existing online MAPF algorithms in Švancara et al. (2019) and different settings, for example, where rerouting of robots is always allowed (Ma et al. 2017a) or disallowed (Ho et al. 2019), in real-world systems. The relationships between these algorithms are summarized in Figure 1.

As our third contribution, we study online MAPF algorithms under the competitive analysis framework. Specifically, we demonstrate how an arbitrary online MAPF algorithm can be rationalized and show that the competitive ratios of all rational online MAPF algorithms with respect to flowtime and makespan are both bounded from above by $\mathcal{O}(m)$ for an input sequence of m agents. We then show that (1) the bounds are tight for all rational algorithms in PLAN-NEW-SINGLE and PLAN-NEW, (2) the competitive ratio is at least 4/3 with respect to flowtime and 3/2 with respect to makespan for all rational algorithms in PLAN-ALL, and (3) the competitive ratio is infinite with respect to latency for all rational algorithms. The results hold even for optimally-rational algorithms and on 4-neighbor 2D grids. Therefore, for the first time, we provide theoretical insights into the effectiveness of using MAPF algorithms in an online setting (Salzman and Stern 2020) and address some of the long-standing open questions such as whether planning for multiple agents is more effective than planning for only one agent at a time in an online setting, whether algorithms that allow rerouting are more effective than those that disallow, and whether acting optimally rationally can improve the effectiveness. The results are summarized in Table 1.

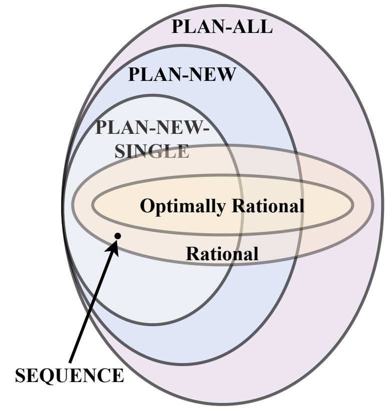

Figure 1: Relationships between online MAPF algorithms.

Controllability		PLAN-NEW-SINGLE			PLAN-NEW		PLAN-ALL	
Objective Function	Competitive Ratio Bounds	SEQUENCE	Rational	Optimally Rational	Rational	Optimally Rational	Rational	Optimally Rational
flowtime	upper	$\mathcal{O}(m)$	$\mathcal{O}(m)$					
	lower (even on 2D grids)	$\Omega(m)$					4/3	
makespan	upper	$\mathcal{O}(m)$	$\mathcal{O}(m)$					
	lower (even on 2D grids)	$\Omega(m)$					3/2	
latency	(even on 2D grids)	∞						

Table 1: Summary of main competitiveness results.

Conclusions

We conducted a theoretical study of online MAPF for the first time. Our results suggest that, if rerouting is disallowed, then planning for multiple agents is asymptotically (only) as effective as planning for one agent at a time and acting optimally rationally is asymptotically (only) as effective as acting rationally, which is also asymptotically (only) as effective as following the naive algorithm SEQUENCE. However, allowing rerouting can potentially result in high effectiveness, as indicated by the gap between the competitive ratio upper and lower bounds.

We refer the reader to the original paper (Ma 2021) for the detailed theoretical results.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) under grant number RGPIN-2020-06540.

References

Banfi, J.; Basilico, N.; and Amigoni, F. 2017. Intractability of time-optimal multirobot path planning on 2d grid graphs with holes. *IEEE Robotics and Automation Letters* 2(4): 1941–1947.

Boyarski, E.; Felner, A.; Stern, R.; Sharon, G.; Tolpin, D.; Betzalel, O.; and Shimony, S. E. 2015. ICBS: Improved Conflict-Based Search Algorithm for Multi-Agent Pathfinding. In *IJCAI*, 740–746.

Cohen, L.; Greco, M.; Ma, H.; Hernandez, C.; Felner, A.; Kumar, T. K. S.; and Koenig, S. 2018. Anytime Focal Search with Applications. In *IJCAI*, 1434–1441.

Dresner, K.; and Stone, P. 2008. A Multiagent Approach to Autonomous Intersection Management. *Journal of Artificial Intelligence Research* 31: 591–656.

Erdem, E.; Kisa, D. G.; Oztok, U.; and Schueller, P. 2013. A General Formal Framework for Pathfinding Problems with Multiple Agents. In *AAAI*, 290–296.

Gange, G.; Harabor, D.; and Stuckey, P. J. 2019. Lazy CBS: Implicit Conflict-Based Search Using Lazy Clause Generation. In *ICAPS*, 155–162.

Ho, F.; Salta, A.; Geraldes, R.; Goncalves, A.; Cavazza, M.; and Prendinger, H. 2019. Multi-agent Path Finding for UAV Traffic Management. In *AAMAS*, 131–139.

Kou, N. M.; Peng, C.; Ma, H.; Kumar, T. S.; and Koenig, S. 2020. Idle time optimization for target assignment and path finding in sortation centers. In *AAAI*, volume 34, 9925–9932.

Lam, E.; Le Bodic, P.; Harabor, D.; and Stuckey, P. J. 2019. Branch-and-Cut-and-Price for Multi-Agent Pathfinding. In *IJCAI*, 1289–1296.

Li, J.; Felner, A.; Boyarski, E.; Ma, H.; and Koenig, S. 2019a. Improved Heuristics for Multi-Agent Path Finding with Conflict-Based Search. In *IJCAI*, 442–449.

Li, J.; Gange, G.; Harabor, D.; Stuckey, P. J.; Ma, H.; and Koenig, S. 2020. New Techniques for Pairwise Symmetry Breaking in Multi-Agent Path Finding. In *ICAPS*, 193–201.

Li, J.; Harabor, D.; Stuckey, P. J.; Felner, A.; Ma, H.; and Koenig, S. 2019b. Disjoint Splitting for Conflict-Based Search for Multi-Agent Path Finding. In *ICAPS*, 279–283.

Luna, R.; and Bekris, K. E. 2011. Push and Swap: Fast Cooperative Path-finding with Completeness Guarantees. In *IJCAI*, 294–300.

Ma, H. 2021. A Competitive Analysis of Online Multi-Agent Path Finding. In *ICAPS*.

Ma, H.; Harabor, D.; Stuckey, P. J.; Li, J.; and Koenig, S. 2019a. Searching with Consistent Prioritization for Multi-Agent Path Finding. In *AAAI*, 7643–7650.

Ma, H.; Hönig, W.; Kumar, T. K. S.; Ayanian, N.; and Koenig, S. 2019b. Lifelong Path Planning with Kinematic Constraints for Multi-Agent Pickup and Delivery. In *AAAI*, 7651–7658.

Ma, H.; and Koenig, S. 2017. AI Buzzwords Explained: Multi-Agent Path Finding (MAPF). *AI Matters* 3(3): 15–19.

Ma, H.; Li, J.; Kumar, T. K. S.; and Koenig, S. 2017a. Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks. In *AAMAS*, 837–845.

Ma, H.; and Pineau, J. 2015. Information Gathering and Reward Exploitation of Subgoals for POMDPs. In *AAAI*, 3320–3326.

Ma, H.; Tovey, C.; Sharon, G.; Kumar, T. K. S.; and Koenig, S. 2016. Multi-Agent Path Finding with Payload Transfers and the Package-Exchange Robot-Routing Problem. In *AAAI*, 3166–3173.

Ma, H.; Wagner, G.; Felner, A.; Li, J.; Kumar, T. K. S.; and Koenig, S. 2018. Multi-Agent Path Finding with Deadlines: Preliminary Results. In *AAMAS*, 2004–2006.

Ma, H.; Yang, J.; Cohen, L.; Kumar, T. K. S.; and Koenig, S. 2017b. Feasibility Study: Moving Non-Homogeneous Teams in Congested Video Game Environments. In *AIIDE*, 270–272.

Salzman, O.; and Stern, R. 2020. Research Challenges and Opportunities in Multi-Agent Path Finding and Multi-Agent Pickup and Delivery Problems. In *AAMAS*, 1711–1715.

Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence* 219: 40–66.

Sharon, G.; Stern, R.; Goldenberg, M.; and Felner, A. 2013. The increasing cost tree search for optimal multi-agent pathfinding. *Artificial Intelligence* 195: 470–495.

Stern, R.; Sturtevant, N.; Felner, A.; Koenig, S.; Ma, H.; Walker, T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T. K. S.; Boyarski, E.; and Bartak, R. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In *SoCS*, 151–159.

Surynek, P. 2010. An optimization variant of multi-robot path planning is intractable. In *AAAI*, 1261–1263.

Surynek, P.; Felner, A.; Stern, R.; and Boyarski, E. 2016. Efficient SAT Approach to Multi-Agent Path Finding Under the Sum of Costs Objective. In *ECAI*, 810–818.

Švancara, J.; Vlk, M.; Stern, R.; Atzmon, D.; and Barták, R. 2019. Online Multi-Agent Pathfinding. In *AAAI*, 7732–7739.

Wang, K.; and Botea, A. 2011. MAPP: A Scalable Multi-Agent Path Planning Algorithm with Tractability and Completeness Guarantees. *Journal of Artificial Intelligence Research* 42: 55–90.

Wurman, P. R.; D'Andrea, R.; and Mountz, M. 2008. Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses. *AI Magazine* 29(1): 9–20.

Yu, J. 2015. Intractability of optimal multirobot path planning on planar graphs. *IEEE Robotics and Automation Letters* 1(1): 33–40.

Yu, J.; and LaValle, S. M. 2013a. Planning Optimal Paths for Multiple Robots on Graphs. In *ICRA*, 3612–3617.

Yu, J.; and LaValle, S. M. 2013b. Structure and Intractability of Optimal Multi-robot Path Planning on Graphs. In *AAAI*, 1444–1449.