



ARTICLE

Intelligent planning for large-scale multi-agent systems

Hang Ma

Simon Fraser University, Burnaby,
Canada**Correspondence**Hang Ma, Simon Fraser University,
Burnaby, Canada.
Email: hangma@sfu.ca**Funding information**Huawei Technologies Canada; Canada
Foundation for Innovation; Natural
Sciences and Engineering Research
Council of Canada, Grant/Award
Number: RGPIN2020-06540**Abstract**

This article summarizes the New Faculty Highlights talk with the same title at AAAI 2021. Intelligent agents such as different types of robots will soon become an integral part of our daily lives. In real-world multi-agent systems, the most fundamental challenges are assigning tasks to multiple agents (task-level coordination problems) and planning collision-free paths for the agents to task locations (motion-level coordination problems). This article surveys four directions of our research on using intelligent planning techniques for the above multi-agent coordination problems. Link to video abstract: https://youtu.be/HDAFcatq9_I

INTRODUCTION

Self-driving cars, autonomous drones, autonomous aircraft towing vehicles, automated warehouse robots, automated-guided port vehicles, home and office service robots, and other intelligent agents will become an integral part of our daily lives. For example, in the coming years, hundreds of autonomous aircraft towing vehicles will tow their assigned aircraft between runways and terminal gates (Morris et al. 2016). Today, hundreds of warehouse robots already navigate fully autonomously in automated fulfillment and sortation centers to deliver inventory shelves and express parcels to fulfill online orders (Wurman, D'Andrea, and Mountz 2008; Kou et al. 2020). For these real-world applications of large-scale multi-agent systems, the basic building blocks include assigning tasks to agents and planning paths for the agents to reach task locations. Agents must avoid collisions in a congested environment while reaching their task locations as promptly as possible and completing a large number of tasks as quickly as possible. The resulting task- and motion-level coordination problems, which model the task-assignment and path-planning operations of the agents, are fundamental for these multi-agent systems

but, at the same time, computationally challenging since there are typically many agents in such a system and the operating time of the system is long.

In the artificial intelligence (AI) community, much attention has been placed on a simplified one-short version of the path-planning problem in the above multi-agent systems, known as *multi-agent path finding (MAPF)* (Ma and Koenig 2017; Stern et al. 2019; Ma 2022). The problem of MAPF is to move multiple agents from their start vertices to their target vertices in discrete time steps on a given graph that models the environment and let the agents wait in their target vertices. During each time step, each agent can wait in its current vertex or move to an adjacent vertex. Agents are not allowed to collide. Two agents collide if and only if, during the same time step, they move to the same vertex or traverse the same edge in opposite directions. A MAPF solution consists of a set of collision-free paths, one for each agent that specifies the vertex occupies by the agent at each time step. The objective is to minimize the flowtime (also known as the sum-of-costs; Felner et al. 2017), that is, the sum of the numbers of time steps for the agents to reach their target vertices, or the makespan, that is, the earliest time step when all agents are at their target vertices. MAPF is NP-hard to solve optimally for makespan

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

© 2022 The Authors. *AI Magazine* published by Wiley Periodicals LLC on behalf of the Association for the Advancement of Artificial Intelligence.

(Surynek 2010) or flowtime (Yu and LaValle 2013c) minimization. MAPF algorithms include reductions to Boolean satisfiability (Surynek et al. 2016), integer linear programming (Yu and LaValle 2013b), and answer set programming (Erdem et al. 2013) and specialized combinatorial search algorithms (Standley and Korf 2011; Sharon et al. 2013; Wagner and Choset 2015; Sharon et al. 2015).

Nevertheless, the following concerns should be addressed when we generalize MAPF to the above real-world applications of large-scale multi-agent systems (Ma et al. 2016): 1. MAPF algorithms need to be made more efficient, namely they need to compute a solution faster for the same number of agents or scale to a larger number of agents in the same amount of time. 2. New variants and extensions of MAPF need to be studied to tackle problems arising in practice, such as executing MAPF solutions with unmodeled motion delays and robot constraints, combining target assignment and path finding, and long-term and online planning.

This article surveys our research on intelligent planning for large-scale multi-agent systems in four research directions that addresses the above concerns:

- **Improved MAPF algorithms:** We propose several improvements to state-of-the-art optimal MAPF algorithms based on heuristic search techniques and leveraging insights from solvers of other combinatorial problems. The resulting improved MAPF algorithms are several orders of magnitude faster. We develop bounded-suboptimal and suboptimal MAPF algorithms that produce close-to-optimal MAPF solutions for hundreds of agents in seconds of computation time. Furthermore, we demonstrate the benefits of our algorithms by applying them to the navigation of drones and video characters.
- **Safe execution of MAPF solutions:** We study the problem of handling motion delays when robots execute MAPF solutions. We propose to use execution policies to guarantee safe execution of MAPF solutions. These execution policies determine the correct timing for robots to follow each step of a given MAPF solution. We show how we have combined insights from this study and MAPF techniques to win a railway scheduling competition that was held at a top-tier machine learning conference. We also develop a hierarchical framework for generating and executing MAPF solutions for real-world multirobot systems. This framework makes use of an efficient procedure that accounts for kinematic constraints of robots to transform a MAPF solution to a continuous-time plan-execution schedule that is safe for robots to execute. We demonstrate our framework using different types of simulated and real robots.
- **Combined target assignment and path finding:** We formalize and study a variant of MAPF, called *combined target-assignment and path finding (TAPF)* (Ma

and Koenig 2016), that models the joint problem of (1) which locations the robots go to next and (2) how the robots go to the locations. The problem of TAPF is to assign target vertices to multiple agents and move the agents from their start vertices to their target vertices on a given graph without collisions. We develop an optimal TAPF algorithm that computes solutions for hundreds of agents in minutes of computation time. We apply our TAPF algorithm to solving the formation control problem with teams of real robots.

- **Long-term task and path planning:** We formalize and study an extension to MAPF and TAPF, called *multi-agent pickup and delivery (MAPD)* (Ma et al. 2017). MAPD models a long-term problem where a system needs to repeatedly assign incoming tasks to a set of agents and plan paths for the agents to the targets of their assigned tasks. We develop complete and deadlock-free MAPD algorithms for both online and offline settings and techniques to account for kinematic constraints of robots when solving MAPD. These algorithms made decisions for hundreds of agents and thousands of tasks in seconds of computation time. We apply our techniques for solving MAPD to automated parcel sorting with warehouse robots and demonstrate the benefits of our techniques on an industrial simulator with real-world data.

IMPROVED MAPF ALGORITHMS

Modern large-scale real-world multi-agent systems such as the automated fulfillment and sortation centers constructed by Amazon (Wurman, D'Andrea, and Mountz 2008) and Alibaba (Kou et al. 2020) require solving MAPF efficiently, namely, planning high-quality paths for hundreds of agents in a short computation time. However, it is NP-hard to solve MAPF optimally (Surynek 2010; Yu and LaValle 2013c). In our theoretical study (Ma et al. 2016), we further prove that MAPF is NP-hard to approximate within any constant factor less than $4/3$ for makespan minimization by a reduction from a specialized NP-complete version (Tovey 1984) of the Boolean satisfiability problem.

Nevertheless, a state-of-the-art MAPF algorithm, called *Conflict-Based Search (CBS)* (Sharon et al. 2015) can compute optimal MAPF solutions for dozens of agents in a few minutes of computation time. CBS is a two-level combinatorial search algorithm. It first finds individually optimal paths for all agents (ignoring collisions). On the high level, CBS performs a best-first search to resolve each collision of the computed paths by imposing constraints on both the agents involved in the collision. The constraints forbid the agents from occupying a vertex or traversing an edge at a given time step. On the low level, CBS uses an A* search in



both the space and the time dimensions to find a path for an agent that obeys its constraints. The high-level search of CBS branches on which collision to resolve.

We have explored three directions to make CBS more efficient. First, we have developed admissible heuristics for the high-level best-first search of CBS to significantly reduce the size of the search tree (Felner et al. 2018; Li et al. 2019). Experimental results show that the resulting algorithm CBSH (Li et al. 2019) is up to 50 times faster than CBS and can thus compute solutions for up to 3 times more agents than CBS within 1 min of computation time. Second, we have proposed a high-level branching rule for CBS called disjoint splitting (Li et al. 2019). Disjoint splitting guarantees that CBS solution spaces explored under the child nodes resulting from each node expansion are disjoint, thereby reducing duplicate search effort. Experimental results show that CBS with disjoint splitting is up to two orders of magnitude faster than CBS. Third, we have developed several techniques (Li et al. 2019, 2020, 2021) to tackle “pairwise path symmetry,” which occurs when two agents each has multiple paths of the same cost but any paths of the two agents are pairwise incompatible since they result in collisions. Our recent work (Li et al. 2021) experimentally demonstrates that the combination of several of the above directions result in an improved version of CBS that is up to four orders of magnitude faster than CBS and can thus compute solutions for up to 30 times more agents than CBS within 1 min of computation time.

We have also explored five directions to make MAPF algorithms more suitable for practical applications of large-scale multi-agent systems. First, we have developed an anytime bounded-suboptimal version of CBS (Cohen et al. 2018) that can be stopped at any time and return a MAPF solution with a guaranteed suboptimality bound. The suboptimality bound become smaller as the algorithm runs longer. Second, we have developed *prioritized-based search (PBS)* (Ma et al. 2019), a MAPF algorithm that searches in the space of all possible orderings of agents. PBS is a two-level algorithm similar to CBS but performs a depth-first search on the high level. Unlike CBS, PBS is suboptimal and complete only for a realistic family of MAPF instances. However, empirical study shows that PBS always returns close-to-optimal solutions (with a cost no more than 105% of the optimal cost) in our experimental setting and can compute solution for 600 agents in half a minute of computation time. Third, we have addressed the MAPF problem where a deadline is given for agents to reach their target vertices (Ma et al. 2018b, 2018a). We have developed an integer linear programming-based algorithm and a CBS-based algorithm that both maximize the number of agents reaching target vertices before the deadline. Fourth, we have addressed the MAPF problem for large-size agents with different geometric shapes and volumes of

agents (Li et al. 2019). We tackle this challenge by allowing each agent to occupy more than one vertex and generalizing the definition of collisions since the agent can intersect with multiple vertices and edges at each time step. We have developed a version of CBS to solve this problem optimally and demonstrate it for the navigation of a drone fleet. Fifth, we have shown how a combination of swarm-based approaches from the robotics community and an adapted version of CBS can be applied to the navigation of multiple game characters that needs to keep a desired formation while moving (Li et al. 2020). This adapted version of CBS balances between the makespan and the cost of deviating from the desired formation.

SAFE EXECUTION OF MAPF SOLUTIONS

MAPF algorithms can be used to compute collision-free paths for multiple agents in real-world multi-agent systems. However, real-world agents such as warehouse robots cannot perfectly execute the computed solutions, namely following their paths, since they can be delayed unexpectedly when they move, they have unmodeled kinematic constraints (for example, they do not move at the same speed), and so forth existing AI research has not studied how the agents can safely execute the computed MAPF solutions.

In our recent work (Ma, Kumar, and Koenig 2017), we study a variant of MAPF where each agent is delayed and stays in its current vertex with a given probability whenever it intends to move to another vertex during plan execution. We propose several decentralized execution policies to guarantee safe execution of MAPF solutions under such delay uncertainty. Decentralized execution policies use a GO or STOP command to control, at each time step, whether an agent should follow its planned path to move to the next vertex or not. For example, a naive policy is to move all agents in locked time steps according to a given MAPF solution, namely it stops all other agents if an agent is delayed, which requires all agents to communicate with each other at every time step. We proposed *minimal communication policy (MCP)* to make the execution more efficient. The key idea of MCP is to respect the precedence constraints that, if two different agents visit the same vertex, they have to visit it in the same order as specified by the given MAPF solution. To do so, MCP constructs a directed acyclic graph whose nodes represent events of agents visiting vertices and whose arcs represent precedence constraints. During execution, MCP thus stops an agent only when a precedence constraint is not satisfied until it receives a signal from another agent for the constraint. MCP also attempts to minimize the number of arcs between events of different agents, namely precedence constraints between agents, and thus

the communication cost. Furthermore, we have developed a version of CBS that computes MAPF solutions with small expected makespan for given delay probabilities, assuming that MCP is used for plan execution. In our most recent work (Li et al. 2021), we combine MAPF algorithms and MCP to develop a software that won the NeurIPS-20 Flatland Challenge¹, a railway scheduling competition, which was held in partnership with German, Swiss, and French railway companies. Our software uses automated planning and combinatorial search techniques only but outperformed all other entries, including all reinforcement learning entries, to win overall first place in both rounds of the competition in the top machine learning conference NeurIPS 2020. The problem in the competition is a MAPF variant similar to above where the agents can be stopped at a random time step for a random duration during plan execution.

We have also extended the above idea to a hierarchical framework (Ma et al. 2017) for plan generation and execution to account for kinematic constraints of agents during planning and other system dynamics during plan execution. This framework uses a post-processing procedure called *MAPF-POST* (Hönig et al. 2016) that transforms a MAPF solution into a plan-execution schedule in continuous time using a simple temporal network (STN). The STN is also directed acyclic graph similar to the one constructed by MCP. The STN takes into account important kinematic constraints such as various edge lengths, different agent sizes, and translational and rotational velocity limits of agents by representing the kinematic constraints as temporal constraints between agents, which are arcs annotated with time bounds in the directed acyclic graph, in the computation. The resulting plan-execution schedule guarantees a user-specified safety distance among agents and avoids replanning, namely resolving MAPF problems, in many cases during execution. We have verified our framework using different types of simulated and real robots. The pipeline of our framework is as follows. First, our framework runs an optimal MAPF algorithm such as CBS to compute a MAPF solution. Second, it then uses *MAPF-POST* to construct an STN that transforms the MAPF solution into a plan-execution schedule, which specifies the expected execution time when a robot should arrive at each location. Third, all robots move along their paths to meet the expected execution time at each location. If the execution deviates from the plan in the third step, the framework constructs a new STN based on the real execution times and computes a new plan-execution schedule for the rest of the plan. The framework needs to perform the first step to replan, namely to solve a new MAPF instance, only when *MAPF-POST* fails to compute a plan-execution schedule in the second step. In practice, our experimental results show that the robots can

almost always execute the given MAPF solution safely without replanning.

COMBINED TARGET ASSIGNMENT AND PATH FINDING

Many real-world applications of multi-agent systems require the coordination of not only path-planning operations but also target-assignment operations. For example, an automated warehouse system needs to decide which robots to deliver which parcels and on what routes the robots should move. In our recent work (Ma and Koenig 2016), we formalize and study a variant of MAPF called TAPF that couples the target-assignment and the path-finding problems for multiple teams of agents. In TAPF, agents are partitioned into teams and each team is given the same number of target vertices as there are agents in the team. The problem of TAPF is to assign the target vertices of each team to agents in the same team and plan collision-free paths for the agents to their target vertices so that each agent reaches exactly one target vertex and each target vertex is reached by an agent. Existing AI research has considered only two extremes of TAPF. On one hand, the special case of TAPF with one team of agents can be solved optimally in polynomial time (Yu and LaValle 2013a) by solving a max-flow problem. On the other hand, MAPF algorithms assume that each agent forms a single-agent team, namely the agent is assigned the only target vertex in its team. It remains unclear how and how well one can solve the general case of TAPF with multiple teams of agents.

Therefore, we have developed an optimal TAPF algorithm called Conflict-Based MinCost Flow (CBM) (Ma and Koenig 2016). CBM breaks TAPF down to the NP-hard subproblem of coordinating different teams of agents and the polynomial-time solvable subproblems of coordinating the agents in every team. It then tackles these subproblems by using a combination of CBS for the NP-hard subproblem and a min-cost max-flow algorithm (Goldberg and Tarjan 1987) for the polynomial-time solvable subproblems. Our experimental results demonstrate that CBM can compute optimal TAPF solutions for more than 400 agents in minutes of runtime, showcasing its potential for large-scale multi-agent systems. In our empirical study (Hönig et al. 2016), we verify CBM using both simulated and real robots. We demonstrate how CBM can be combined with *MAPF-POST* to generate plan-execution schedules that allow for the safe execution of TAPF on these real-world agents. We apply the resulting techniques to solve a formation control problem where drones in different colors move to desired locations to display an English word in 3D space with each letter in a unique color.



LONG-TERM TASK AND PATH PLANNING

Agents in many multi-agent systems need to constantly attend to new tasks after they finish their current tasks. For example, a warehouse robot in an Amazon fulfillment center (Wurman, D'Andrea, and Mountz 2008) needs to pick up and deliver another inventory shelf after it finishes delivering its current inventory shelf. Existing AI research on MAPF and TAPF has focused mostly on *one-shot* problems only where each agent has one target vertex and the agents stop moving after they all reach their target vertices. Therefore, in our recent work (Ma et al. 2017), we formalize and study MAPD that generalizes the one-shot problems MAPF and TAPF to a *long-term* problem. In MAPD, agents have to attend to a stream of incoming tasks. Each task enters the system at an unknown time and is characterized by two target vertices, namely a pickup vertex and a delivery vertex. A free agent, namely one that is currently not executing any task, can be assigned an unexecuted task. To execute the task, the agent has to first move from its current vertex to the pickup vertex of the task, become occupied and start to execute the task upon reaching the pickup vertex of the task, and then move from the pickup vertex to the delivery vertex of the task, while avoiding collisions with other agents.

There are three benefits of modeling “pickup-and-delivery” tasks that have an intermediate target vertex (the pickup vertex) and a final target vertex (the delivery vertex) each instead of “navigation” tasks that have only one target vertex each. 1. Modeling “pickup-and-delivery” tasks results in a mix of both single-agent teams (each consisting of an occupied agent) and a team of free agents, while modeling “navigation” tasks results in only one team of agents, which limits its generalizability. 2. The resulting MAPD algorithms still apply directly to “navigation” tasks because each such task is a special case of a “pickup-and-delivery” task with the pickup and delivery vertices being the same vertex. (c) Modeling “pickup-and-delivery” tasks also makes it easy to explain the resulting MAPD algorithms, even though these algorithms can easily be generalized to cases where the tasks have multiple (ordered) intermediate target vertices and a subsequent final target vertex each.

We have considered an *online* setting (Ma et al. 2017) of MAPD where tasks can enter the system at any time and are not known until they have been added to the system. We develop decentralized and centralized MAPD algorithms that are deadlock-free. The key idea of these online MAPD algorithms is to decouple the long-term problem of MAPD into a sequence of one-shot subproblems at each time and repeatedly apply task-assignment, MAPF, and TAPF algorithms to these subproblems. Exper-

imental results show that these MAPD algorithms can determine the tasks and paths for 500 agents in seconds of computation time. We also demonstrate how one of our MAPD algorithms can account for kinematic constraints of robots directly during planning (Ma et al. 2019) in an online setting. The resulting algorithm can compute an executable solution for 30 min of operation of 250 robots and 2000 tasks in less than 10 s of computation time, which is more efficient than using MAPF-POST in a post-processing phase.

We have also considered an *offline* setting (Liu et al. 2019) of MAPD where all tasks are known a priori, thus affording opportunity to optimize the order in which agents execute tasks. We develop an offline MAPD algorithm that models the task scheduling problem as a specialized asymmetric version of the Traveling Salesman problem to compute a chronologically ordered task sequence for each agent. Experimental results show that the offline MAPD algorithm can compute solutions where agents finish all tasks by up to 46% sooner than online MAPD algorithms.

Finally, we have applied MAPD algorithms to an application of parcel sorting with warehouse robots in an automated sortation center (Kou et al. 2020) where the robots need to move to sorting stations to obtain an express parcel and deliver it to a correct sorting bin associated with the postal code of the shipping address of the parcel. A machine in each sorting station scans the shipping address of a parcel, determines the sorting bin the parcel should be delivered to, and load the parcel onto a robot as long as there are robots waiting in the queue of the sorting station. Our goal is to optimize the idle time of the sorting stations, namely the duration when there are no robots queuing, since it is often the throughput bottleneck of such automated sortation centers. The problem is to assign sorting stations to robots that are not delivering parcels and plan paths for all the robots. We develop an algorithm that solves the TAPF subproblem for robots that are not delivering parcels and solves the MAPF subproblem for robots that are delivering parcels. We test our algorithm using an industrial simulator with real-world data of online orders. Experimental results show that our algorithm can make decisions for 350 agents in no more than 2 s of computation time and improve throughput (measured in the average number of parcels obtained by robots per second) of a sortation center by up to 12%.

SUMMARY

We described our research on using intelligent planning techniques to tackle multi-agent coordination problems. We outlined four directions that generalize

task-assignment and MAPF research to real-world applications of large-scale multi-agent systems. For our ongoing research, we are currently developing a deeper theoretical understanding of using MAPF algorithms for long-term autonomy of such systems (Ma 2021), a learning-based distributed MAPF algorithm (Ma, Luo, and Ma 2021), and algorithms that can jointly solve MAPF and complex task-planning problems (Zhong et al. 2022; Xu et al. 2022). We hope that researchers working in this area can benefit from the insights provided in this article.

ACKNOWLEDGMENT

This work was supported by NSERC under Grant number RGPIN2020-06540 as well as a Canada Foundation for Innovation John R. Evans Leaders Fund award and a gift from Huawei Technologies Canada.

CONFLICT OF INTEREST

The author declares that there is no conflict.

ORCID

Hang Ma  <https://orcid.org/0000-0002-0602-0999>

ENDNOTE

¹<https://discourse.aicrowd.com/t/neurips-2020-flatland-winners/4010>

REFERENCES

- Cohen, L., M. Greco, H. Ma, C. Hernandez, A. Felner, T. K. S. Kumar, and S. Koenig. 2018. "Anytime focal search with applications." In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1434–41.
- Erdem, E., D. G. Kisa, U. Oztok, and P. Schueller. 2013. "A general formal framework for pathfinding problems with multiple agents." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 290–6.
- Felner, A., J. Li, E. Boyarski, H. Ma, L. Cohen, T. K. S. Kumar, and S. Koenig. 2018. "Adding heuristics to conflict-based search for multi-agent pathfinding." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 83–7.
- Felner, A., R. Stern, S. E. Shimony, E. Boyarski, M. Goldenberg, G. Sharon, N. R. Sturtevant, G. Wagner, and P. Surynek. 2017. "Search-based optimal solvers for the multi-agent pathfinding problem: summary and challenges." In *Proceedings of the International Symposium on Combinatorial Search*, 29–37.
- Goldberg, A. V., and R. E. Tarjan. 1987. "Solving minimum-cost flow problems by successive approximation." In *Proceedings of the Annual ACM Symposium on Theory of Computing*, 7–18.
- Hönig, W., T. K. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig. 2016. "Multi-agent path finding with kinematic constraints." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 477–85.
- Hönig, W., T. K. S. Kumar, H. Ma, N. Ayanian, and S. Koenig. 2016. "Formation change for robot groups in occluded environments." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and System*, 4836–42.
- Kou, N. M., C. Peng, H. Ma, T. K. S. Kumar, and S. Koenig. 2020. "Idle time optimization for target assignment and path finding in sortation centers." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9925–32.
- Li, J., E. Boyarski, A. Felner, H. Ma, and S. Koenig. 2019. "Improved heuristics for multi-agent path finding with conflict-based search." In *Proceedings of the International Joint Conference on Artificial Intelligence*, 442–9.
- Li, J., Z. Chen, Y. Zheng, S.-H. Chan, D. Harabor, P. J. Stuckey, H. Ma, and S. Koenig. 2021. "Scalable rail planning and replanning: winning the 2020 flatland challenge." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 477–85.
- Li, J., G. Gange, D. Harabor, P. J. Stuckey, H. Ma, and S. Koenig. 2020. "New techniques for pairwise symmetry breaking in multi-agent path finding." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 193–201.
- Li, J., D. Harabor, P. J. Stuckey, A. Felner, H. Ma, and S. Koenig. 2019. "Disjoint splitting for conflict-based search for multi-agent path finding." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 279–83.
- Li, J., D. Harabor, P. J. Stuckey, H. Ma, G. Gange, and S. Koenig. 2021. "Pairwise symmetry reasoning for multi-agent path finding search." *Artificial Intelligence* 301: 103574.
- Li, J., D. Harabor, P. J. Stuckey, H. Ma, and S. Koenig. 2019. "Symmetry-breaking constraints for grid-based multi-agent path finding." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 6087–95.
- Li, J., K. Sun, H. Ma, A. Felner, T. K. S. Kumar, and S. Koenig. 2020. "Moving agents in formation in congested environments." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 726–34.
- Li, J., P. Surynek, A. Felner, H. Ma, T. K. S. Kumar, and S. Koenig. 2019. "Multi-agent path finding for large agents." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7627–34.
- Liu, M., H. Ma, J. Li, and S. Koenig. 2019. "Target assignment and path planning for multi-agent pickup and delivery." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 2253–5.
- Ma, H. 2021. "A competitive analysis of online multi-agent path finding." In *Proceedings of the International Conference on Automated Planning and Scheduling*, 234–42.
- Ma, H. 2022. "Graph-based multi-robot path finding and planning." *Current Robotics Reports* 3: 1–8.
- Ma, H., D. Harabor, P. J. Stuckey, J. Li, and S. Koenig. 2019. "Searching with consistent prioritization for multi-agent path finding." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7643–50.
- Ma, H., W. Hönig, L. Cohen, T. Uras, H. Xu, T. K. S. Kumar, N. Ayanian, and S. Koenig. 2017. "Overview: a hierarchical framework for plan generation and execution in multi-robot systems." *IEEE Intelligent Systems* 32(6): 6–12.
- Ma, H., W. Hönig, T. K. S. Kumar, N. Ayanian, and S. Koenig. 2019. "Lifelong path planning with kinematic constraints for multi-agent pickup and delivery." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7651–8.
- Ma, H., and S. Koenig. 2016. "Optimal target assignment and path finding for teams of agents." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 1144–52.



- Ma, H., and S. Koenig. 2017. "AI buzzwords explained: multi-agent path finding (MAPF)." *AI Matters* 3(3): 15–9.
- Ma, H., S. Koenig, N. Ayanian, L. Cohen, W. Hönl, T. K. S. Kumar, T. Uras, H. Xu, C. Tovey, and G. Sharon. 2016. "Overview: generalizations of multi-agent path finding to real-world scenarios." In *Proceedings of the IJCAI-16 Workshop on Multi-Agent Path Finding*.
- Ma, H., T. K. S. Kumar, and S. Koenig. 2017. "Multi-agent path finding with delay probabilities." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3605–12.
- Ma, H., J. Li, T. K. S. Kumar, and S. Koenig. 2017. "Lifelong multi-agent path finding for online pickup and delivery tasks." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 837–45.
- Ma, H., C. Tovey, G. Sharon, T. K. S. Kumar, and S. Koenig. 2016. "Multi-agent path finding with payload transfers and the package-exchange robot-routing problem." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3166–73.
- Ma, H., G. Wagner, A. Felner, J. Li, T. K. S. Kumar, and S. Koenig. 2018a. "Multi-agent path finding with deadlines." In *Proceedings of the International Joint Conference on Artificial Intelligence*, 417–23.
- Ma, H., G. Wagner, A. Felner, J. Li, T. K. S. Kumar, and S. Koenig. 2018b. "Multi-agent path finding with deadlines: preliminary results." In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, 2004–6.
- Ma, Z., Y. Luo, and H. Ma. 2021. "Distributed heuristic multi-agent path finding with communication." In *Proceedings of the IEEE International Conference on Robotics and Automation*, (in press), 8699–8705.
- Morris, R., C. Pasareanu, K. Luckow, W. Malik, H. Ma, T. K. S. Kumar, and S. Koenig. 2016. "Planning, scheduling and monitoring for airport surface operations." In *Proceedings of the AAAI-16 Workshop on Planning for Hybrid Systems*, 608–14.
- Sharon, G., R. Stern, A. Felner, and N. R. Sturtevant. 2015. "Conflict-based search for optimal multi-agent pathfinding." *Artificial Intelligence* 219: 40–66.
- Sharon, G., R. Stern, M. Goldenberg, and A. Felner. 2013. "The increasing cost tree search for optimal multi-agent pathfinding." *Artificial Intelligence* 195: 470–95.
- Standley, T. S., and R. E. Korf. 2011. "Complete algorithms for cooperative pathfinding problems." In *Proceedings of the International Joint Conference on Artificial Intelligence*, 668–73.
- Stern, R., N. R. Sturtevant, A. Felner, S. Koenig, H. Ma, T. T. Walker, J. Li, D. Atzmon, L. Cohen, T. K. S. Kumar, E. Boyarski, and R. Barták. 2019. "Multi-agent pathfinding: definitions, variants, and benchmarks." In *Proceedings of the International Symposium on Combinatorial Search*, 151–9.
- Surynek, P. 2010. "An optimization variant of multi-robot path planning is intractable." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1261–3.
- Surynek, P., A. Felner, R. Stern, and E. Boyarski. 2016. "Efficient SAT approach to multi-agent path finding under the sum of costs objective." In *Proceedings of the European Conference on Artificial Intelligence*, 810–8.
- Tovey, C. 1984. "A simplified NP-complete satisfiability problem." *Discrete Applied Mathematics* 8: 85–90.
- Wagner, G., and H. Choset. 2015. "Subdimensional expansion for multirobot path planning." *Artificial Intelligence* 219: 1–24.
- Wurman, P. R., R. D'Andrea, and M. Mountz. 2008. "Coordinating hundreds of cooperative, autonomous vehicles in warehouses." *AI Magazine* 29(1): 9–20.
- Xu, Q., J. Li, S. Koenig, and H. Ma. 2022. "Multi-goal multi-agent pickup and delivery." In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and System*, (in press).
- Yu, J., and S. M. LaValle. 2013a. "Multi-agent path planning and network flow." In *Algorithmic Foundations of Robotics X*, Springer Tracts in Advanced Robotics, 86, 157–73. Heidelberg: Springer.
- Yu, J., and S. M. LaValle. 2013b. "Planning optimal paths for multiple robots on graphs." In *Proceedings of the IEEE International Conference on Robotics and Automation*, 3612–7.
- Yu, J., and S. M. LaValle. 2013c. "Structure and intractability of optimal multi-robot path planning on graphs." In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1444–9.
- Zhong, X., J. Li, S. Koenig, and H. Ma. 2022. "Optimal and bounded-suboptimal multi-goal task assignment and path finding." In *Proceedings of the IEEE International Conference on Robotics and Automation*, 10731–7.

AUTHOR BIOGRAPHY

Hang Ma is an Assistant Professor in the School of Computing Science at Simon Fraser University. His research interests lie in the intersection of artificial intelligence, robotics, and machine learning. Specifically, he is interested in topics on automated planning, multi-agent/robot systems, spatio-temporal and constraint reasoning, and applications of probabilistic methods. His research work has won multiple awards, including the ICAPS 2016 Outstanding Paper Award in the Robotics Track, a runner-up for the 2020 Victor Lesser Distinguished Dissertation Award, the ICAPS 2021 Best Dissertation Award, and the ICAPS 2021 Best System Demonstration Award.

How to cite this article: Ma, H. 2022. "Intelligent planning for large-scale multi-agent systems." *AI Magazine* 43: 376–382.
<https://doi.org/10.1002/aaai.12069>