

# Lifelong Multi-Agent Path Finding for Online Pickup and Delivery Tasks\*

Hang Ma  
University of Southern California  
hangma@usc.edu

T. K. Satish Kumar  
University of Southern California  
tkskwork@gmail.com

Jiaoyang Li  
Tsinghua University  
lijiaoyang13@mails.tsinghua.edu.cn

Sven Koenig  
University of Southern California  
skoenig@usc.edu

## ABSTRACT

The multi-agent path-finding (MAPF) problem has recently received a lot of attention. However, it does not capture important characteristics of many real-world domains, such as automated warehouses, where agents are constantly engaged with new tasks. In this paper, we therefore study a lifelong version of the MAPF problem, called the multi-agent pickup and delivery (MAPD) problem. In the MAPD problem, agents have to attend to a stream of delivery tasks in an online setting. One agent has to be assigned to each delivery task. This agent has to first move to a given pickup location and then to a given delivery location while avoiding collisions with other agents. We present two decoupled MAPD algorithms, Token Passing (TP) and Token Passing with Task Swaps (TPTS). Theoretically, we show that they solve all well-formed MAPD instances, a realistic subclass of MAPD instances. Experimentally, we compare them against a centralized strawman MAPD algorithm without this guarantee in a simulated warehouse system. TP can easily be extended to a fully distributed MAPD algorithm and is the best choice when real-time computation is of primary concern since it remains efficient for MAPD instances with hundreds of agents and tasks. TPTS requires limited communication among agents and balances well between TP and the centralized MAPD algorithm.

## Keywords

agent coordination; multi-agent path finding; path planning; pickup and delivery tasks; task assignment

## 1. INTRODUCTION

Many real-world applications of multi-agent systems require agents to operate in known common environments.

\*Our research was supported by NSF under grant numbers 1409987 and 1319966. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies or the U.S. government.

The agents are constantly engaged with new tasks and have to navigate between locations where the tasks need to be executed. Examples include aircraft-towing vehicles [18], warehouse robots [31], office robots [28], and game characters in video games [21]. In the near future, for instance, aircraft-towing vehicles might navigate autonomously to aircraft and tow them from the runways to their gates so as to reduce pollution, energy consumption, congestion, and human workload. Today, warehouse robots already navigate autonomously to inventory pods and move them from their storage locations to packing stations.

Past research efforts have concentrated mostly on a “one-shot” version of this problem, called the multi-agent path-finding (MAPF) problem, which has been studied in artificial intelligence, robotics, and operations research. In the MAPF problem, each agent has to move from its current location to its destination while avoiding collisions with other agents in a known common environment. The number of agents is the same as the number of destinations, and the MAPF task ends once all agents reach their destinations. Therefore, the MAPF problem does not capture important characteristics of many real-world domains, such as automated warehouses, where agents are constantly engaged with new tasks.

In this paper, we therefore study a “lifelong” version of the MAPF problem, called the multi-agent pickup and delivery (MAPD) problem. In the MAPD problem, agents have to attend to a stream of delivery tasks in a known common environment that is modeled as an undirected graph. Tasks can enter the system at any time and are modeled as exogenous events that are characterized by a pickup location and a delivery location each. An agent that is currently not executing any task can be assigned to an unexecuted task. In order to execute the task, the agent has to first move from its current location to the pickup location and then to the delivery location of the task while avoiding collisions with other agents. We first formalize the MAPD problem and then present two decoupled MAPD algorithms, Token Passing (TP) and Token Passing with Task Swaps (TPTS), both of which are based on existing MAPF algorithms. Theoretically, we show that they solve all well-formed MAPD instances [2], a realistic subclass of MAPD instances. Experimentally, we compare them against a centralized strawman MAPD algorithm without this guarantee in a simulated warehouse system.

















- [8] W. Hönig, T. K. S. Kumar, L. Cohen, H. Ma, H. Xu, N. Ayanian, and S. Koenig. Multi-agent path finding with kinematic constraints. In *International Conference on Automated Planning and Scheduling*, pages 477–485, 2016.
- [9] W. Hönig, T. K. S. Kumar, H. Ma, N. Ayanian, and S. Koenig. Formation change for robot groups in occluded environments. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4836–4842, 2016.
- [10] M. Khorshid, R. Holte, and N. Sturtevant. A polynomial-time algorithm for non-optimal multi-agent pathfinding. In *Annual Symposium on Combinatorial Search*, 2011.
- [11] H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2:83–97, 1955.
- [12] R. Luna and K. E. Bekris. Push and Swap: Fast cooperative path-finding with completeness guarantees. In *International Joint Conference on Artificial Intelligence*, pages 294–300, 2011.
- [13] H. Ma and S. Koenig. Optimal target assignment and path finding for teams of agents. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1144–1152, 2016.
- [14] H. Ma, S. Koenig, N. Ayanian, L. Cohen, W. Hönig, T. K. S. Kumar, T. Uras, H. Xu, C. Tovey, and G. Sharon. Overview: Generalizations of multi-agent path finding to real-world scenarios. In *IJCAI-16 Workshop on Multi-Agent Path Finding*, 2016.
- [15] H. Ma, T. K. S. Kumar, and S. Koenig. Multi-agent path finding with delay probabilities. In *AAAI Conference on Artificial Intelligence*, 2017.
- [16] H. Ma, C. Tovey, G. Sharon, T. K. S. Kumar, and S. Koenig. Multi-agent path finding with payload transfers and the package-exchange robot-routing problem. In *AAAI Conference on Artificial Intelligence*, pages 3166–3173, 2016.
- [17] P. MacAlpine, E. Price, and P. Stone. SCRAM: Scalable collision-avoiding role assignment with minimal-makespan for formational positioning. In *AAAI Conference on Artificial Intelligence*, pages 2096–2102, 2015.
- [18] R. Morris, C. Pasareanu, K. Luckow, W. Malik, H. Ma, S. Kumar, and S. Koenig. Planning, scheduling and monitoring for airport surface operations. In *AAAI-16 Workshop on Planning for Hybrid Systems*, 2016.
- [19] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015.
- [20] G. Sharon, R. Stern, M. Goldenberg, and A. Felner. The increasing cost tree search for optimal multi-agent pathfinding. *Artificial Intelligence*, 195:470–495, 2013.
- [21] D. Silver. Cooperative pathfinding. In *Artificial Intelligence and Interactive Digital Entertainment*, pages 117–122, 2005.
- [22] T. S. Standley and R. E. Korf. Complete algorithms for cooperative pathfinding problems. In *International Joint Conference on Artificial Intelligence*, pages 668–673, 2011.
- [23] N. R. Sturtevant and M. Buro. Improving collaborative pathfinding using map abstraction. In *Artificial Intelligence and Interactive Digital Entertainment*, pages 80–85, 2006.
- [24] P. Surynek. A novel approach to path planning for multiple robots in bi-connected graphs. In *IEEE International Conference on Robotics and Automation*, pages 3613–3619, 2009.
- [25] P. Surynek. Reduced time-expansion graphs and goal decomposition for solving cooperative path finding sub-optimally. In *International Joint Conference on Artificial Intelligence*, pages 1916–1922, 2015.
- [26] C. Tovey, M. Lagoudakis, S. Jain, and S. Koenig. The generation of bidding rules for auction-based robot coordination. In F. S. L. Parker and A. Schultz, editors, *Multi-Robot Systems. From Swarms to Intelligent Automata*, volume 3, chapter 1, pages 3–14. Springer, 2005.
- [27] J. P. van den Berg, J. Snape, S. J. Guy, and D. Manocha. Reciprocal collision avoidance with acceleration-velocity obstacles. In *IEEE International Conference on Robotics and Automation*, pages 3475–3482, 2011.
- [28] M. Veloso, J. Biswas, B. Coltin, and S. Rosenthal. CoBots: Robust symbiotic autonomous mobile service robots. In *International Joint Conference on Artificial Intelligence*, pages 4423–4429, 2015.
- [29] G. Wagner. *Subdimensional Expansion: A Framework for Computationally Tractable Multirobot Path Planning*. PhD thesis, Carnegie Mellon University, 2015.
- [30] K. Wang and A. Botea. MAPP: a scalable multi-agent path planning algorithm with tractability and completeness guarantees. *Journal of Artificial Intelligence Research*, 42:55–90, 2011.
- [31] P. R. Wurman, R. D’Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.
- [32] J. Yu and S. M. LaValle. Planning optimal paths for multiple robots on graphs. In *IEEE International Conference on Robotics and Automation*, pages 3612–3617, 2013.
- [33] X. Zheng and S. Koenig. K-swaps: Cooperative negotiation for solving task-allocation problems. In *International Joint Conference on Artificial Intelligence*, pages 373–378, 2009.