

# SkyRover: A Modular Simulator for Cross-Domain Pathfinding

Wenhui Ma<sup>1</sup>, Wenhao Li<sup>2,\*</sup>, Bo Jin<sup>2</sup>, Changhong Lu<sup>3,4</sup>, Xiangfeng Wang<sup>1,4,\*</sup>

<sup>1</sup> School of Computer Science and Technology, East China Normal University

<sup>2</sup> School of Computer Science and Technology, Tongji University

<sup>3</sup> School of Mathematical Sciences, East China Normal University

<sup>4</sup> Key Laboratory of Mathematics and Engineering Applications

{wenhuima@stu, chlu@math, xfwang@cs}.ecnu.edu.cn, {whli, bjin}@tongji.edu.cn

## Abstract

Unmanned Aerial Vehicles (UAVs) and Automated Guided Vehicles (AGVs) increasingly collaborate in logistics, surveillance, inspection tasks, etc. However, existing simulators often focus on a single domain, limiting cross-domain study. This paper presents the *SkyRover*, a modular simulator for UAV-AGV multi-agent pathfinding (MAPF). *SkyRover* supports realistic agent dynamics, configurable 3D environments, and convenient APIs for external solvers and learning methods. By unifying ground and aerial operations, it facilitates cross-domain algorithm design, testing, and benchmarking. Experiments highlight *SkyRover*'s capacity for efficient pathfinding and high-fidelity simulations in UAV-AGV coordination. We believe the *SkyRover* fills a key gap in MAPF research. The project is available at <https://sites.google.com/view/mapf3d/home>.

## 1 Introduction

The rapid growth of automation and artificial intelligence has significantly broadened unmanned systems' application domains, especially in logistics and transportation. Companies like Amazon, JD, and Cainiao routinely deploy automated guided vehicles (AGVs) in large-scale warehouses [Qin *et al.*, 2022; Wurman *et al.*, 2008; Zhang *et al.*, 2020], while autonomous taxi services from Waymo and Baidu Apollo Go have begun trial operations globally [Sun *et al.*, 2020; Wang *et al.*, 2024b]. Meanwhile, the low-altitude economy has gained momentum with drone-based package delivery, low-altitude tourism, and urban air mobility. Commercial drone delivery services by Amazon Prime Air and Meituan already enhance last-mile logistics [Dorling *et al.*, 2016; Engesser *et al.*, 2023; Sun *et al.*, 2024].

Alongside these developments, collaborative UAV-AGV systems are increasingly vital in freight transportation [Gao *et al.*, 2020], search and rescue [Zhang *et al.*, 2024], precision agriculture [Tokekar *et al.*, 2016], and infrastructure inspection [Wu *et al.*, 2020], among other domains [Munasinghe *et*

*al.*, 2024]. AGVs follow terrestrial routes, while UAVs operate in three-dimensional or restricted airspaces. Their synergy can boost operational capacity: UAVs excel in rapid transport of lightweight cargo, whereas AGVs handle heavier loads or site-specific tasks. However, integrating these systems demands careful scheduling and planning.

Multi-Agent Path Finding (MAPF) [Stern *et al.*, 2019] is essential for collision-free, resource-efficient routing of large autonomous fleets. Although significant progress has been made for purely ground- or aerial-based MAPF [Alkazzi and Okumura, 2024; Ma, 2022; Salzman and Stern, 2020; Surynek, 2022], new complexities arise when UAVs and AGVs share or intersect operational zones. Their heterogeneity (in motion patterns, energy consumption, and environmental constraints) further complicates joint planning; for instance, UAVs are more weather-sensitive, whereas AGVs often contend with dynamically changing indoor settings.

Robust simulation tools are key to designing, training, and testing MAPF algorithms for UAV-AGV scenarios, especially for data-intensive learning-based approaches. Existing simulators handle single-domain movements effectively [Fernando, 2022; Nguyen *et al.*, 2019; Okumura *et al.*, 2021; Skrynnik *et al.*, 2024; Wang *et al.*, 2024a], but few offer native support for UAV-AGV collaboration.

To fill this gap, we introduce *SkyRover*, a modular simulator for cross-domain pathfinding with UAV and AGV coordination. To our knowledge, it is the first environment providing a unified toolkit for UAV-AGV MAPF research. *SkyRover* allows users to build complex scenarios (e.g., warehouses and park environments) with adjustable fidelity and integrates high-fidelity physics models alongside user-friendly APIs for algorithmic testing and rapid prototyping. Through customizable and extensible interfaces, *SkyRover* supports a wide spectrum of use cases—from basic scheduling experiments to sophisticated learning-based methods.

## 2 Related Work

Extensive progress has been made in solving MAPF for single-domain vehicles [Ma, 2022; Salzman and Stern, 2020; Surynek, 2022]. Previous simulators such as [Fernando, 2022; Nguyen *et al.*, 2019; Okumura *et al.*, 2021; Skrynnik *et al.*, 2024; Wang *et al.*, 2024a] have tackled various forms of agent-based simulation, yet few focus on integrated UAV-AGV scenarios. Moreover, many existing platforms concen-

\*Corresponding authors.

trate on either 2D grids or simplified 3D representations, limiting the study of aerial and ground interactions. By contrast, *SkyRover* explicitly targets these cross-domain concerns, offering realistic physics, 3D occupancy grids, and unified APIs. To our knowledge, it is the first environment to natively support collaborative UAV-AGV MAPF under a single, modular framework.

### 3 SkyRover Simulator

Following the motivations outlined in Section 1, we present *SkyRover*, a modular simulator designed to address cross-domain pathfinding for UAV-AGV coordination. It aims to unify the modeling of ground and aerial vehicles in 3D grids while maintaining ease of use, modularity, and compatibility with third-party robotic frameworks. Figure 1 depicts its overall architecture, which comprises five main modules: 1) the *Sim World Zoo*; 2) the *3D Grid Generator*; 3) the *Unified Algorithm Wrapper*; 4) the *Plan Executor*; 5) the *System Interface*. In this section, we detail each module and explain how they collectively enable rapid experimentation and deployment of cross-domain MAPF solutions.

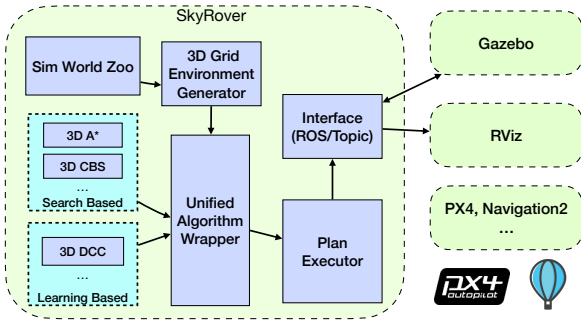


Figure 1: Main Architecture. *SkyRover* comprises multiple modules to support cross-domain MAPF research and development.

**Sim World Zoo:** The *Sim World Zoo* houses multiple Gazebo simulation worlds [Koenig and Howard, 2004], along with 3D models for UAVs and AGVs. Unlike simulators that rely solely on matrix-based grids, *SkyRover* offers more realistic environments by integrating detailed 3D worlds. Currently, *SkyRover* includes a warehouse and a park scenario, each featuring UAV and AGV models. These environments and virtual robots, sourced from <https://app.gazebosim.org/fuel/models>, can be further customized via the Gazebo UI. This design decision ensures that *SkyRover* users can scale from simple grid representations to complex settings that approximate real-world operational constraints.

**3D Grid Generator:** Although *SkyRover* supports high-fidelity worlds, it also maintains a 3D grid representation for MAPF algorithms. The *3D Grid Generator* automatically derives discrete map information from Gazebo environments. To achieve this, it employs a Gazebo plugin that creates a ROS2 service [Macenski *et al.*, 2022] for capturing a 2D .pgm map and a 3D .pcd (point cloud) file. These files

reflect the layout of objects and obstacles in the environment. Subsequently, the point cloud data is parsed to mark grid cells as free or occupied, generating a 3D occupancy grid suitable for both search- and learning-based algorithms.

**Unified Algorithm Wrapper:** MAPF algorithms generally fall into either search-based or learning-based categories [Hart *et al.*, 1968; Sharon *et al.*, 2015; Ma *et al.*, 2021]. They also vary in their internal map structures, with some operating on matrices and others on more general graphs. To standardize these approaches, *SkyRover* provides a *Unified Algorithm Wrapper* that abstracts the environment as a 3D grid with consistent interface calls, such as *init*: initializes the environment, obstacle data, agent start states, and goal locations; *step*: advances the simulation by one timestep, updating agent positions; and *reset*: resets the simulation for new tasks or training episodes. We provide 3D versions of popular MAPF algorithms, including 3D A\* [Hart *et al.*, 1968], 3D CBS [Sharon *et al.*, 2015], and a 3D extension of DCC [Ma *et al.*, 2021], ensuring that researchers and practitioners can rapidly evaluate both classical search and modern learning paradigms in a unified simulated environment.

**Plan Executor:** High-level path planners often produce routes in discrete grids or graphs, but real robots need continuous control inputs. The *Plan Executor* bridges this gap by translating each agent’s planned trajectory into commands interpretable by low-level controllers. The executor tracks each agent’s path, communicates with external control frameworks (e.g., PX4 [Meier *et al.*, 2015] or Navigation2 [Navigation2 Contributors, 2025]), and updates agent positions in the simulator. This arrangement allows *SkyRover* to support both abstract, high-level pathfinding and more realistic, hardware-oriented simulations.

**System Interface:** The *SkyRover* also offers multiple external interfaces to integrate with established robotics tools. Users can manipulate Gazebo models via Gazebo topics, visualize 3D occupancies with RViz [Kam *et al.*, 2015], and command PX4-based drones through ROS2 topics and MicroXRCE [eProsima, 2024]. This flexibility supports varied applications, ranging from hardware-in-the-loop testing to large-scale simulation. Ultimately, *SkyRover* allows researchers and practitioners to incorporate realistic dynamics, robust visualization, and real-time interactivity into customized UAV-AGV cooperative scenarios.

### 4 Experiment

In this section, we demonstrate how *SkyRover* supports diverse experimental setups. We showcase the simulator using two distinct Gazebo worlds: a warehouse (Figure 2a) and a park (Figure 2b). These worlds serve as representative testbeds, featuring varied layouts that allow us to highlight the simulator’s 3D mapping, pathfinding, and visualization capabilities. We define two typical AGV-UAV interaction tasks: a) **Inventory Scanning Task:** The AGV transports cargo to point A, where a UAV hovers above the AGV to scan and inventory the cargo. After the scanning process is completed, the AGV continues transporting the cargo to point B; b) **Aerial Cargo Transfer Task:** The AGV transports cargo to point A, where a UAV arrives and hovers above the AGV

to pick up the cargo. The UAV then lifts the cargo and transports it to an elevated point B. These tasks demonstrate the seamless coordination between AGVs and UAVs in logistics and transportation scenarios, showcasing the effectiveness of SkyRover in simulating complex robotic interactions.

**Environment Setup:** We begin by loading the chosen Gazebo world and generating a 3D occupancy grid using the gazebo-map-creator plugin [Mehmood, 2022]. This plugin extracts a point cloud of the environment, which is converted into a 0–1 grid where each cell spans 1 meter. Cells containing point cloud particles are treated as obstacles and labeled “1,” while free cells remain “0.” The resulting 3D grid is published to RViz for real-time visualization (Figure 2c).




Figure 2: (a) The warehouse Gazebo world, featuring multiple Holybro X500 drones and delivery AGVs; (b) The park scenario, offering more open space for UAV operations; (c) The 3D occupancy grid in RViz. Each dark cell has point cloud data and is thus considered an obstacle; (d) Example of integrating SkyRover with hardware-oriented controllers. Here, PX4 executes drone flight commands, while Navigation2 governs the TurtleBot.

**3D MAPF Examples:** After generating the occupancy grid, we deploy 6 Holybro X500 drones and 16 AGVs. We assign each agent a unique start and goal location in the warehouse world. Via the *init* interface of algorithm wrapper, we load the 3D-A\* and 3D-CBS solvers. During initialization, these solvers compute complete, conflict-free paths for all agents. At each timestep, the *step* function moves every agent to its next waypoint, ensuring collision-free navigation.

We also implement a 3D version of the learning-based method, DCC [Ma *et al.*, 2021]. Here, we adapt the original 2D convolution layers to 3D, and employ curriculum learning to train for 70,000 episodes in a random  $40 \times 40 \times 40$  grid. This trained model reaches 100% success under test conditions with sixteen agents. Unlike search-based algorithms, which plan entire routes beforehand, the learning-based approach invokes the *step* function to infer actions online after loading the model file. This approach suits dynamic environments where complete pre-planning is less feasible.

**Motion Control Integration:** For many studies, abstract position updates are sufficient for benchmarking algorithmic

performance. In these cases, we simply invoke Gazebo’s model-position service to teleport each agent, bypassing detailed dynamics. However, SkyRover also supports low-level motion control through PX4 [Meier *et al.*, 2015] and Navigation2 [Navigation2 Contributors, 2025]. Figure 2d shows an example involving a drone controlled by PX4 and a TurtleBot commanded by Navigation2. These finer-grained simulations accurately capture kinematic and dynamic constraints, which are essential for hardware-in-the-loop testing. Although more computationally intensive, these setups are valuable for research on real-time control, multi-robot coordination, and safety validation. Users can choose the approach—abstract or low-level—based on the complexity of their experiments and available computing resources.

**Preliminary Results and Comparison:** We conduct all experiments on Ubuntu 24.04 with ROS 2 Jazzy, Gazebo Harmonic, and the main branch of PX4, using a PC with an Intel i7 CPU and 32 GB RAM. Table 1 presents preliminary performance metrics comparing 3D variants of A\*, CBS and DCC on the warehouse world when 22 total agents (6 drones, 16 AGVs) must reach randomly assigned goals. **Computation time** measures the total time to compute or infer a path before the first move, while **Success rate** indicates the percentage of agents reaching their target without collision in a single run.

Algorithm	Comp. Time (s)	Success Rate (%)
3D-A*	54.7	100
3D-CBS	92.4	100
3D-DCC (Well-trained)	0.6	100

Table 1: Preliminary Results in the Warehouse (22 Agents).

These initial comparisons highlight SkyRover’s ability to benchmark MAPF algorithms in a unified and consistent framework. Future work could involve more in-depth evaluations, including scaling to larger agent teams and investigating runtime on different hardware setups.

## 5 Conclusion and Future Directions

In this paper, we introduce SkyRover, a modular simulator paving the way for integrated UAV-AGV MAPF in 3D environments. It combines realistic Gazebo worlds, a robust 3D grid generator, unified wrappers for classical and learning-based algorithms, and seamless integration with external robotics software. Experiments in the warehouse and park worlds confirmed its flexibility for discrete pathfinding and low-level control simulations.

**Limitations.** Current limitations include partial modeling of real-world effects (e.g., weather, sensor noise) and the computational load of large-scale simulations. Learning-based methods also require extensive training data, and the hyperparameter optimization might be time-intensive.

**Future Directions.** Going forward, we plan to integrate more realistic physics (wind perturbations, complex friction models), advanced sensor types (LiDAR, radar), and dynamic obstacle handling. Additionally, a systematic approach to large-scale distributed training could support reinforcement learning methods that tackle hundreds of agents in real time.

## Ethical Statement

SkyRover facilitates research on UAV-AGV coordination that could improve safety and efficiency in logistics, search and rescue, and infrastructure inspection. We acknowledge that such autonomous technologies must be developed responsibly. We designed this simulator as an open-source platform to democratize access to high-quality research tools while emphasizing that real-world deployments must adhere to local regulations and safety standards. Our commitment is to support the advancement of autonomous systems that benefit society while minimizing potential risks.

## Acknowledgments

This work was supported by the National Key R&D Program of China under grant 2021YFA1000302, National Natural Science Foundation of China 62406270, STCSM “Science and Technology Innovation Action Plan” 24YF2748800. We thank the reviewers for their valuable feedback. We also acknowledge the contributions of the open-source community behind ROS2, Gazebo, and PX4, which made SkyRover possible. Special thanks to our laboratory colleagues for their assistance with testing and valuable discussions.

## References

- [Alkazzi and Okumura, 2024] Jean-Marc Alkazzi and Keisuke Okumura. A comprehensive review on leveraging machine learning for multi-agent path finding. *IEEE Access*, 12:57390–57409, 2024.
- [Dorling *et al.*, 2016] Kevin Dorling, Jordan Heinrichs, Geoffrey G Messier, and Sebastian Magierowski. Vehicle routing problems for drone delivery. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47(1):70–85, 2016.
- [Engesser *et al.*, 2023] Valeska Engesser, Evy Rombaut, Lieslot Vanhaverbeke, and Philippe Lebeau. Autonomous delivery solutions for last-mile logistics operations: A literature review and research agenda. *Sustainability*, 15(3):2774, 2023.
- [eProsima, 2024] eProsima. eProsima Micro XRCE-DDS: A DDS-XRCE protocol implementation for resource constrained environments, 2024.
- [Fernando, 2022] Malintha Fernando. Mavswarm: A lightweight multi-aerial vehicle simulator. [https://github.com/malintha/multi\\_uav\\_simulator](https://github.com/malintha/multi_uav_simulator), 2022.
- [Gao *et al.*, 2020] Wei Gao, Junren Luo, Wanpeng Zhang, Weilin Yuan, and Zhiyong Liao. Commanding cooperative UGV-UAV with nested vehicle routing for emergency resource delivery. *IEEE Access*, 8:215691–215704, 2020.
- [Hart *et al.*, 1968] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [Kam *et al.*, 2015] Hyeong Ryeol Kam, Sung-Ho Lee, Taejung Park, and Chang-Hun Kim. Rviz: a toolkit for real domain data visualization. *Telecommunication Systems*, 60:337–345, 2015.
- [Koenig and Howard, 2004] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *IROS*, 2004.
- [Ma *et al.*, 2021] Ziyuan Ma, Yudong Luo, and Jia Pan. Learning selective communication for multi-agent path finding. *IEEE Robotics and Automation Letters*, 7(2):1455–1462, 2021.
- [Ma, 2022] Hang Ma. Graph-based multi-robot path finding and planning. *Current Robotics Reports*, 3(3):77–84, 2022.
- [Macenski *et al.*, 2022] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. Robot operating system 2: Design, architecture, and uses in the wild. *Science robotics*, 7(66):eabm6074, 2022.
- [Mehmood, 2022] Arshad Mehmood. ROS2 Gazebo world 2d/3d map generator. [https://github.com/arshadlab/gazebo\\_map\\_creator](https://github.com/arshadlab/gazebo_map_creator), 2022.
- [Meier *et al.*, 2015] Lorenz Meier, Dominik Honegger, and Marc Pollefeys. Px4: A node-based multithreaded open source robotics framework for deeply embedded platforms. In *ICRA*, 2015.
- [Munasinghe *et al.*, 2024] Isuru Munasinghe, Asanka Perera, and Ravinesh C Deo. A comprehensive review of uav-ugv collaboration: Advancements and challenges. *Journal of Sensor and Actuator Networks*, 13(6):81, 2024.
- [Navigation2 Contributors, 2025] Navigation2 Contributors. Navigation2 documentation. <https://github.com/ros-navigation/docs.nav2.org>, 2025. <https://docs.nav2.org>.
- [Nguyen *et al.*, 2019] Van Nguyen, Philipp Obermeier, Tran Son, Torsten Schaub, and William Yeoh. Generalized target assignment and path finding using answer set programming. In *SoCS*, 2019.
- [Okumura *et al.*, 2021] Keisuke Okumura, Yasumasa Tamura, and Xavier Défago. Iterative refinement for real-time multi-robot path planning. In *IROS*, 2021.
- [Qin *et al.*, 2022] Hengle Qin, Jun Xiao, Dongdong Ge, Linwei Xin, Jianjun Gao, Simai He, Haodong Hu, and John Gunnar Carlsson. Jd.com: Operations research algorithms drive intelligent warehouse robots to work. *INFORMS Journal on Applied Analytics*, 52(1):42–55, 2022.
- [Salzman and Stern, 2020] Oren Salzman and Roni Stern. Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems. In *AAMAS*, 2020.
- [Sharon *et al.*, 2015] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial intelligence*, 219:40–66, 2015.
- [Skrynnik *et al.*, 2024] Alexey Skrynnik, Anton Andreychuk, Anatolii Borzilov, Alexander Chernyavskiy, Konstantin Yakovlev, and Aleksandr Panov. POGEMA: A

- benchmark platform for cooperative multi-agent navigation. *arXiv:2407.14931*, 2024.
- [Stern *et al.*, 2019] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *SoCS*, 2019.
- [Sun *et al.*, 2020] Pei Sun, Henrik Kretschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020.
- [Sun *et al.*, 2024] Xuting Sun, Minghao Fang, Shu Guo, and Yue Hu. Uav-rider coordinated dispatching for the on-demand delivery service provider. *Transportation Research Part E: Logistics and Transportation Review*, 186:103571, 2024.
- [Surynek, 2022] Pavel Surynek. Problem compilation for multi-agent path finding: a survey. In *IJCAI*, 2022.
- [Tokekar *et al.*, 2016] Pratap Tokekar, Joshua Vander Hook, David Mulla, and Volkan Isler. Sensor planning for a symbiotic uav and ugv system for precision agriculture. *IEEE Transactions on Robotics*, 32(6):1498–1511, 2016.
- [Wang *et al.*, 2024a] Qian Wang, Rishi Veerapaneni, Yu Wu, Jiaoyang Li, and Maxim Likhachev. Mapf in 3d warehouses: Dataset and analysis. In *ICAPS*, 2024.
- [Wang *et al.*, 2024b] Shiqi Wang, Zhouye Zhao, Yuhang Xie, Mingchuan Ma, Zirui Chen, Zeyu Wang, Bohao Su, Wenrui Xu, and Tianyi Li. Recent surge in public interest in transportation: Sentiment analysis of baidu apollo go using weibo data. *arXiv:2408.10088*, 2024.
- [Wu *et al.*, 2020] Yu Wu, Shaobo Wu, and Xinting Hu. Co-operative path planning of uavs & ugvs for a persistent surveillance task in urban environments. *IEEE Internet of Things Journal*, 8(6):4906–4919, 2020.
- [Wurman *et al.*, 2008] Peter R Wurman, Raffaello D’Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–9, 2008.
- [Zhang *et al.*, 2020] Yi Zhang, Yu Qian, Yichen Yao, Haoyuan Hu, and Yinghui Xu. Learning to cooperate: Application of deep reinforcement learning for online agv path finding. In *AAMAS*, 2020.
- [Zhang *et al.*, 2024] Ying Zhang, Haibao Yan, Danni Zhu, Jiankun Wang, Cui-Hua Zhang, Weili Ding, Xi Luo, Changchun Hua, and Max Q-H Meng. Air-ground collaborative robots for fire and rescue missions: Towards mapping and navigation perspective. *arXiv:2412.20699*, 2024.