

# MAPF-HD: Multi-Agent Path Finding in High-Density Environments

Hiroya Makino and Seigo Ito

**Abstract**—Multi-agent path finding (MAPF) involves planning efficient paths for multiple agents to move simultaneously while avoiding collisions. In typical warehouse environments, agents are often sparsely distributed along aisles; however, increasing the agent density can improve space efficiency. When the agent density is high, it becomes necessary to optimize the paths not only for goal-assigned agents but also for those obstructing them. This study proposes a novel MAPF framework for high-density environments (MAPF-HD). Several studies have explored MAPF in similar settings using integer linear programming (ILP). However, ILP-based methods require substantial computation time to optimize all agent paths simultaneously. Even in small grid-based environments with fewer than 100 cells, these computations can take tens to hundreds of seconds. Such high computational costs render these methods impractical for large-scale applications such as automated warehouses and valet parking. To address these limitations, we introduce the phased null-agent swapping (PHANS) method. PHANS employs a heuristic approach to incrementally swap positions between agents and empty vertices. This method solves the MAPF-HD problem within a few seconds, even in large environments containing more than 700 cells. The proposed method has the potential to improve efficiency in various real-world applications such as warehouse logistics, traffic management, and crowd control.

## I. INTRODUCTION

INDUSTRY is increasingly adopting automated guided vehicles (AGVs) and conveyors to enhance efficiency and reduce reliance on human labor. In environments with agents, such as AGVs or pallets on conveyors, efficient collision-free path planning is essential. These path-planning problems for multiple agents are formalized as multi-agent path finding (MAPF) problems [1]. MAPF has been applied to various settings, including automated warehouses [2], airport taxiway management [3], automated parking systems [4], and video games [5].

In typical MAPF settings, agents are often sparsely distributed along aisles; however, increasing the agent density can improve space efficiency. For instance, high-density automated warehouses store significantly more goods than lower-density ones, reducing required floor space and operational costs (Fig. 1). Similarly, high-density automated valet parking systems allow vehicles to be parked and retrieved in a smaller area by minimizing the gaps between vehicles.

Despite these space-saving benefits, planning efficient paths using conventional techniques remains challenging. Typically, in high-density environments, only a subset of agents has active goals at any given time (e.g., moving some items for shipment). The remaining agents, although not assigned to any particular destination, still occupy space and obstruct the movement of the target agents. For the target agents to reach their objectives, the obstructing agents must be relocated




Fig. 1. MAPF in a high-density warehouse environment. Only a subset of target agents (those with matching-colored flags) are assigned goals, whereas the rest act as movable obstacles. In addition to planning paths for target agents to reach their goals, efficiently relocating obstructing agents is crucial for successful navigation under high-density conditions.

TABLE I  
METHODS FOR MAPF IN HIGH-DENSITY ENVIRONMENTS.

Methods	Target agents	Obstructing agents	Computational efficiency
ILP	✓(optimal)	✓(optimal)	✗
Heuristics	✓(suboptimal)	✗(not optimized)	✓
PHANS	✓(suboptimal)	✓(suboptimal)	✓

efficiently. One approach is to use integer linear programming (ILP) to plan all agent paths jointly [6], [7]. However, this approach generally incurs significant computation time because all paths are optimized simultaneously by ILP. Another approach is to set temporary goals for obstructing agents and then apply existing MAPF methods. However, this often extends the paths of the target agents because the relocation of the obstructing agents is not optimized. Table I highlights these limitations and the need for a tailored solution to accommodate obstructing agents efficiently.

This study refines MAPF for high-density environments (MAPF-HD) and develops an efficient algorithm tailored to this context. This algorithm employs phased heuristic techniques to relocate agents by swapping their positions with empty vertices (null agents). This approach, termed phased null-agent swapping (PHANS), solves MAPF-HD within a few seconds even in a densely packed environment, as shown in Fig. 2. The implementation is available at <https://github.com/ToyotaCRDL/MAPF-in-High-Density-Envs>.

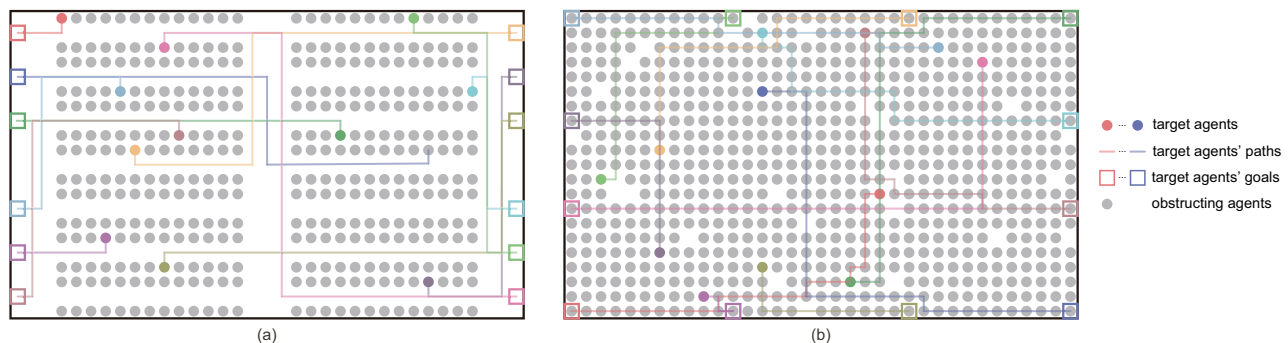


Fig. 2. Simulation examples of MAPF in (a) low-density and (b) high-density environments. Colored lines represent the paths of the target agents. Unlike low-density environments, obstructing agents are present along these paths in high-density environments. The target agents are moved while the obstructing agents are appropriately evacuated.

## II. RELATED WORK

### A. Optimal Solvers

MAPF involves determining optimal collision-free paths for multiple agents [1], [8]. Conflict-based search (CBS) [9], a widely used optimal solver, plans agents' paths to their goals (low-level search) and resolves collisions by branching (high-level search). However, CBS faces scalability challenges because MAPF is NP-hard [10]. Researchers have introduced improvements, such as Improved CBS, to prune the search more effectively [11]. Other optimal frameworks recast MAPF as a combinatorial optimization problem. For example, Lam et al. [12] applied the branch-and-cut-and-price ILP method. Similarly, Gómez et al. [13] used answer-set programming.

### B. Suboptimal Solvers

Several bounded-suboptimal and heuristic algorithms have been developed to improve the scalability of MAPF [14]. Enhanced CBS (ECBS) [15] is a bounded-suboptimal solver that employs focal search. To enhance scalability, explicit estimation CBS (EECBS) [16] replaces the focal search in ECBS with explicit estimation search. Priority inheritance with backtracking (PIBT) [17], a decentralized MAPF algorithm, assigns dynamic movement priorities to agents and allows backtracking to resolve conflicts. Unlike static priority planning, PIBT enables agents to inherit and adjust priorities dynamically, preventing deadlocks and improving coordination in dense environments. Backtracking allows agents to reconsider past actions when conflicts arise, thereby ensuring flexible and adaptive navigation.

PIBT is both scalable and applicable to high-density environments. Chen et al. [18] scaled PIBT for 2,000–12,000 agents occupying 90% of the available cells. Jiang et al. [19] showed that PIBT can solve a  $32 \times 32$  map with 97.7% agent density.

More recently, several reinforcement learning-based MAPF methods have been proposed [20]. These approaches require time for training; however, once a policy is learned, agents can select actions quickly even in large-scale environments.

### C. MAPF in High-Density Environments with Obstructing Agents

Even in dense environments, relocation problems can be solved when all agents have assigned destinations. However, to move target agents efficiently, the relocation of obstructing agents must be optimized.

Okoso et al. [6] introduced the cooperative automated valet parking problem (CoAVP). In CoAVP, agents are classified into three types: entering, departing, and relocating vehicles. Departing vehicles are assigned destinations. Entering and relocating vehicles can be anywhere in the parking lot, and their destinations are also optimized. Their ILP-based method extends the planar graph along the time axis and optimizes the paths within the time-expanded network. However, this approach presents two major challenges. First, ILP has high computational complexity and requires hundreds of seconds even in small environments (e.g.,  $15 \times 15$  grids). Consequently, this method is impractical in real-world applications. Second, collision types specific to high-density environments have not been sufficiently investigated. To save parking space, agents must be densely positioned. In such environments, two adjacent agents must move in almost perfect synchrony, which may not be realistic. Following conflicts [1] should be considered for high-density environments.

Makino et al. [7] defined a multi-agent and multi-rack path finding problem, where racks were classified as either target or nontarget. Agents convey target racks to their destinations while relocating nontarget racks. Similar to [6], their approach uses ILP, resulting in high time complexity.

This study redefines MAPF in high-density environments and develops a heuristic approach to mitigate the high time complexity associated with ILP.

## III. MULTI-AGENT PATH FINDING PROBLEM IN HIGH-DENSITY ENVIRONMENTS

This section defines MAPF-HD. Unlike standard MAPF, which focuses on the paths of the target (goal-assigned) agents toward designated goals, MAPF-HD operates in environments where obstructing agents are densely packed. The objective of MAPF-HD is to minimize the number of time steps required for all target agents to reach their goals without collisions. Fig. 3 shows an example of MAPF-HD.




Fig. 3. Example of the MAPF-HD problem: (a) initial positions, (b) evacuation of obstructing agents, and (c) movement of the target agent.

### A. Environment

An instance of MAPF-HD includes a four-connected grid graph,  $G = (V, E)$ , where  $V$  is the set of vertices, and  $E$  is the set of edges.  $V$  and  $E$  are defined based on the grid dimensions  $size_x$  and  $size_y$  as follows:

$$V = \{v_{x,y} \mid x \in \{0, \dots, size_x - 1\}, y \in \{0, \dots, size_y - 1\}\} \quad (1)$$

$$E = \{(v_{x,y}, v_{x',y'}) \mid v_{x,y}, v_{x',y'} \in V, |x - x'| + |y - y'| = 1\}, \quad (2)$$

where  $x$  and  $x'$  are row indices, and  $y$  and  $y'$  are column indices. These indices correspond to the grid coordinates. The Manhattan distance  $dst(v_{x,y}, v_{x',y'})$  between two points is calculated as follows:

$$dst(v_{x,y}, v_{x',y'}) = |x - x'| + |y - y'|. \quad (3)$$

### B. Agent

MAPF-HD involves two types of agents:

- Target agents:  $A^{tgt} = \{a_1^{tgt}, \dots, a_m^{tgt}\}$ , with  $m > 0$ . Each target agent  $a_i^{tgt}$  is specified by a tuple  $\langle s_i^{tgt}, g_i^{tgt} \rangle$ , where  $s_i^{tgt} \in V$  is the start vertex, and  $g_i^{tgt} \in V$  is the goal vertex.
- Obstructing agents:  $A^{obs} = \{a_1^{obs}, \dots, a_n^{obs}\}$ , with  $n \geq 0$ . Each obstructing agent  $a_i^{obs}$  is characterized solely by its start vertex, denoted as  $\langle s_i^{obs} \rangle$ , as it does not have a goal.

The initial positions of the target and obstructing agents do not overlap. Furthermore, the total number of agents is less than the number of vertices in the graph.

### C. Agent Action

Time is discretized into time steps, and at each time step  $t$ , every agent  $a_i \in A = A^{tgt} \cup A^{obs}$  occupies a vertex  $\pi_i[t] \in V$ . At each time step, an agent executes one of the following actions:

- Remain at the current vertex:  $\pi_i[t] = \pi_i[t + 1]$ .
- Move to an adjacent vertex:  $(\pi_i[t], \pi_i[t + 1]) \in E$ .

Agents avoid collisions at each time step by managing the two conflicts:

- Vertex conflict: Multiple agents simultaneously occupy the same vertex.
- Following conflict: An agent moves to a vertex occupied by another agent in the immediately preceding time step.

Some studies [9], [16], [17] are based on edge conflicts [1] rather than following conflicts. If following conflicts are not treated as conflicts, the agents must be spaced apart, and adjacent agents would need to move almost synchronously. However, in agent-dense environments, spacing the agents sufficiently apart may not be feasible. Moreover, perfect synchronization is not always realistic in real environments. Therefore, we treat following conflicts as prohibited. Note that following conflicts contain edge conflicts [1].

To prevent these conflicts, all agents  $a_i, a_j (i \neq j)$  must satisfy:

$$\begin{cases} \pi_i[t] \neq \pi_j[t] \\ \pi_i[t] \neq \pi_j[t + 1]. \end{cases} \quad (4)$$

### D. Solution and Objective

A solution to MAPF-HD is a set of collision-free paths  $\{\pi_1, \dots, \pi_{m+n}\}$  ensuring that all target agents reach their goals at a certain time step  $T$ . None of the obstructing agents have goals, and they are freely positioned. Formally, the problem assigns a path  $\pi_i = [\pi_i[0], \pi_i[1], \dots, \pi_i[T]]$  to each agent, which must satisfy the following:

$$\begin{cases} \pi_i[0] = s_i^{tgt} & \text{if } a_i^{tgt} \in A^{tgt} \\ \pi_i[0] = s_i^{obs} & \text{if } a_i^{obs} \in A^{obs} \\ \pi_i[T] = g_i^{tgt} & \text{if } a_i^{tgt} \in A^{tgt}. \end{cases} \quad (5)$$

The objective is to minimize makespan  $T$ , which is defined as the total number of time steps required for all target agents to reach their goals. In addition to makespan, classical MAPF often uses sum-of-costs [1], which is the sum of time steps required by each agent to reach its goal. In real-world applications such as warehouses and valet parking, the emphasis is on the speed of the target agents in reaching their goals, rather than on the total movement effort of all agents. Makespan captures this operational bottleneck, whereas sum-of-costs would be distorted by the large number of obstructing agents without goals.

## IV. PROPOSED METHOD

Current methods [6], [7] use ILP to optimize the paths of all agents simultaneously, resulting in high computational complexity. To reduce the computational costs, we propose a two-stage heuristic approach termed PHANS. For the target agents, the first stage plans paths from their starting positions to their goals using the A\* algorithm. The second stage sequentially evacuates the obstructing agents from these paths to clear the way for target agents. Section IV-A describes the methods used in both stages through an example with a single target agent, followed by a generalized algorithm for multiple target agents in Section IV-B.

### A. Single Target Agent

1) *First Stage*: The A\* algorithm is used to plan a target path from the start to the goal vertices for the target agent. Collisions with obstructing agents are allowed during this

stage. However, to minimize the costs of evacuating obstructing agents, the target agent’s movement should satisfy the following conditions:

- Minimize the movement time of the target agent to ensure that the target path is as short as possible.
- Minimize the waiting time for evacuating obstructing agents from the target path.

To satisfy the second condition, the A\* algorithm is modified by introducing an additional heuristic value,  $h_{add}$ . This modification incorporates the expected waiting cost for obstructing agents to clear the path. The cost function  $f$  is computed as the sum of  $g$ ,  $h$ , and  $h_{add}$  as follows:

$$f = g + h + h_{add}, \quad (6)$$

where  $g$  is the actual cost from the start to the current node, and  $h$  is the estimated cost from the current node to the goal.  $h_{add}$  is calculated as follows:

$$h_{add} = \max(0, 1 + h_{evac} - g). \quad (7)$$

Here,  $h_{evac}$  estimates the cost of evacuating an obstructing agent by calculating its distance to the nearest empty vertex using (3). If the estimate is accurate and the obstructing agents are evacuated by the time the target agent arrives,  $h_{add}$  is 0. Otherwise, the waiting time is incorporated into  $h_{add}$ . These computations are based on the initial positions; however, the set of empty vertices may change as obstructing agents are evacuated.

The proposed heuristic  $h + h_{add}$  is neither admissible nor consistent in general. When  $h_{add} = 0$ , it reduces to the standard admissible distance estimate. In other cases, it may overestimate the true evacuation cost, and its value can change nonmonotonically between neighboring states. Thus, the heuristic serves as a practical cost estimator that reflects the additional effort of relocating obstructing agents in high-density environments, rather than a guarantee of optimality.

2) *Second Stage*: The obstructing agents identified on the target path in the first stage are evacuated sequentially. This comprises five steps, as outlined in Fig. 4:

- 1) Select the next obstructing agent that blocks the movement of the target agent.
- 2) Choose the nearest (but not between the target and obstructing agents) empty vertex, termed the null agent.
- 3) Plan a path from the null vertex to the obstructing agent (null-agent path), ignoring potential collisions with the other obstructing agents.
- 4) Follow the null-agent path and swap the null agent with the neighboring agent (*null-agent swapping*).
- 5) Move the target agent once the obstructing agent is evacuated; if the goal is not reached, repeat the process from Step 1.

By iterating these steps, the obstructing agents are sequentially evacuated from the target path, enabling the target agent to progress toward its goal.

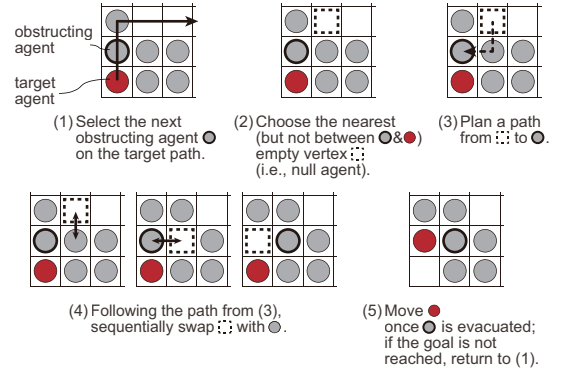


Fig. 4. Second stage of PHANS: Null-agent swapping mechanism.

### B. Multiple Target Agents

For multiple target agents, the first stage can be executed simultaneously for all targets. Additionally, the second stage allows the simultaneous evacuation of multiple obstructing agents to enhance efficiency. The pseudocode for this generalized approach is provided in Algorithm 1. We perform sequential A\* path planning using prioritized planning (PP), giving higher priority to targets with longer start-goal distances computed via (3).

Lines 1–4 in Algorithm 1 compute the distance from the start to the goal for each target agent in  $A^{tgt}$  using (3) and then prioritize the planning of target paths with the longest distances. As resolving conflicts during simultaneous path planning increases computational costs, paths are planned sequentially using PP.

Lines 6–10 identify obstructing agents on the target paths and add them to set  $B$  for subsequent processing. Line 11 orders the obstructing agents in  $B$  according to their remaining distance along the blocked target’s path to its goal. Agents that block a longer remaining segment are handled first to reduce the makespan.

Lines 12–16 assign a null agent to each obstructing agent and plan the path between them. The null-agent assignment proceeds in order from the head of  $B$ , stopping if no null agents remain. In general, the nearest null agent is assigned to each obstructing agent. However, because target agents are prioritized over obstructing agents, any null agents lying on the target agent’s path between the target agent and the obstructing agent are reserved for the target agent. Therefore, these reserved null agents are not assigned to obstructing agents. Lines 18–21 move the target agent after the obstructing agent is evacuated. If there are no collisions, lines 22–25 execute the sequential relocation of the obstructing agents along the paths planned with their assigned null agents.

### C. How Fast Is PHANS?

The A\* algorithm’s worst-case time complexity is  $\mathcal{O}(|E| + |V| \log |V|)$  [21]. The proposed method uses A\* to compute the target path and repeatedly replans the paths for obstructing agents along that path. With  $m$  target agents and the longest target path of length  $d$ , the worst-case time complexity of the proposed method is

---

**Algorithm 1** PHANS

---

**Class:** Agent

**pos:** Tuple  
**goal:** Tuple  
**dst:** Integer  
**path:** Sequence of Tuple

**Input:**

$A^{\text{tgt}} = \{(\text{pos: start position, goal: goal position, dst: None, path: None})$   
 $| \text{target agent } \in \text{ all target agents}\}$   
 $A^{\text{obs}} = \{(\text{pos: start position, goal: None, dst: None, path: None})$   
 $| \text{obstructing agent } \in \text{ all obstructing agents}\}$

**Output:** Agents' movements

```

1: Sort  $A^{\text{tgt}}$  in descending order of distance between start and goal
2: for all tgt in  $A^{\text{tgt}}$  do
3:   tgt.path = path from tgt.pos to tgt.goal
4:   tgt.dst = |tgt.path|
5: while not all target agents have reached their goals do
6:    $N =$  all null agents
7:    $B = \{\}$ 
8:   for all obs in  $A^{\text{obs}}$  do
9:     if obs is on the target paths then
10:       $B.add(\{\text{pos: obs.pos, goal: None, dst: len of path from}$ 
obs.pos to its goal,  $\text{path: None}\})$ 
11:   Sort  $B$  in descending order of dst
12:   for all bagt in  $B$  do
13:     if one or more null agents in  $N$  then
14:       bagt.goal = nearest (but not between the target and obstruct-
ing agent) null agent from bagt.pos
15:       bagt.path = null-agent path from bagt.goal to bagt.pos
16:       remove the selected null agent from  $N$ 
17:   while no target agents have moved do
18:     for all tgt in  $A^{\text{tgt}}$  do
19:       if |tgt.path| > 1 and tgt.path[1] is null then
20:         move the tgt from tgt.path[0] to tgt.path[1]
21:         tgt.path = tgt.path[1:]
22:     for all bagt in  $B$  do
23:       if bagt.path[1] is null then
24:         move the bagt from bagt.path[0] to bagt.path[1]
25:         bagt.path = bagt.path[1:]

```

---

$\mathcal{O}(m^2 d^2 (|E| + |V| \log |V|))$ , which accounts for repeated path recalculations.

Conversely, the methods proposed by Okoso et al. [6] and Makino et al. [7] extend the graph  $G = (V, E)$  along the time axis and solve the minimum-cost flow problem through edges for path planning. With a time-axis size of  $T$ , their worst-case time complexity is  $\mathcal{O}(2^{T|E|})$ .

Therefore, the proposed method operates in polynomial time, whereas the traditional approaches exhibit exponential time complexity. Consequently, the computational cost of the proposed method is expected to be significantly lower than that of the conventional methods.

*D. Why Does PHANS Terminate Within a Finite Number of Steps?*

Each target agent  $a_i^{\text{tgt}}$  has a target path of finite length, and the set of all obstructing agents on these paths is denoted as  $B$ .

For any vertex  $v$ , function  $d_{\text{null}}^t(v)$  returns the path length to the nearest empty vertex at time step  $t$ , where an empty vertex indicates no agent occupancy. If  $d_{\text{null}}^t(v) = 0$ , then vertex  $v$  is empty.

The proposed method prioritizes path planning for obstructing agents on the target paths by assigning a higher priority to those farthest from the goal. By definition of the environment, at least one empty vertex exists; therefore, this vertex can be assigned to the obstructing agent with the longest path to its goal. Let  $v^*$  denote the vertex of the obstructing agent.

$$v^* = \operatorname{argmax}_{v \in B} \text{dst}(v, \text{goal vertex of } v^{\text{'s}} \text{ obstructing target}). \quad (8)$$

A finite-length path exists between  $v^*$  and the nearest empty vertex (null agent), and null-agent swapping is performed at each time step along this path. Consequently,  $d_{\text{null}}^t(v^*)$  decreases monotonically at each time step, i.e.,  $d_{\text{null}}^t(v^*) > d_{\text{null}}^{t+1}(v^*)$ . This process continues until  $d_{\text{null}}^t(v^*)$  reaches zero after a finite number of steps. At this point,  $v^*$  becomes empty, allowing the farthest target agent to advance to a newly vacated vertex.

Given the finite lengths of the target paths, all target agents reach their goals after a finite sequence of null-agent swaps and movements toward empty vertices. Therefore, PHANS terminates within a finite number of steps.

V. NUMERICAL EXPERIMENTS

This section describes the four experiments conducted to evaluate the performance of the proposed method. Experiments 1 and 2 compared PHANS with PIBT, EECBS, and ILP in small-scale simulation environments where these methods could be applied. Experiment 3 assessed the scalability of PHANS in large environments, demonstrating its effectiveness. Additionally, Experiment 4 assessed the applicability of PHANS in complex environments with obstacles such as pillars. All experiments were conducted on an Intel Core i9-12900K processor with 128 GiB of RAM and Ubuntu 22.04.

A. Settings for Comparison Methods

We compared PHANS against ILP, PIBT, and EECBS. To the best of our knowledge, ILP is the only solver that optimizes the relocation of obstructing agents [6], [7]. PIBT is a state-of-the-art solver among the fast suboptimal algorithms. EECBS is an extension of the classical CBS. Li and Ma [22] used EECBS to relocate shelves in dense warehouse scenarios. Setting the suboptimality factor  $\omega = 1.0$  ensures that EECBS operates in its optimal mode, thereby enabling a direct comparison of optimal solutions while simplifying the experimental setup. Note that PIBT is far more scalable than EECBS [17].

Unlike PHANS and ILP, PIBT and EECBS require pre-defined goals for agents whose paths are planned. In the experiments, we assigned the obstructing agents goals that return them to their starting positions. However, the obstructing agents stopped moving once all target agents reached their goals. In other words, the makespan was computed only for the target agents.

ILP defines target and obstructing agents separately, and paths for all agents are planned simultaneously. As in [6], [7], we used Gurobi as the ILP solver.

PIBT can assign priorities to agents. Accordingly, we set each target's priority equal to its start-goal distance and




Fig. 5. Simulation environment for Experiment 1. The grid size is  $14 \times 7$ . Obstructing agents occupy 90% of the cells.




Fig. 6. Comparison of (a) success rates, (b) computation time, and (c) makespans between PHANS and conventional methods in Experiment 1. The density is the proportion of cells occupied by obstructing agents.

assigned a priority of 0 to obstructing agents. Because vanilla PIBT does not consider following conflicts, we extended it using the method described in [23].

### B. Experiment 1

We compared PHANS with conventional methods (PIBT, EECBS, and ILP) based on the success rate, computation time, and makespan. The computation time was the time each method took to complete the path-planning process, with a maximum of 180 s in each trial.

The simulation environment is shown in Fig. 5. The simulations involved two target agents, and the number of obstructing agents varied from 0% to 90% of the total number of cells. The objective was to minimize the makespan and move the target agents to their destinations without collisions.

Fig. 6 shows the results of fifty experimental runs for each density, with different initial positions for the agents in each run. The horizontal axis represents the density of obstructing agents. Fig. 6(a) shows the success rates, and Fig. 6(b) presents the computation time. Fig. 6(c) presents makespan comparisons for scenarios where feasible solutions were identified for at least half of the trials.

1) *Success Rate and Computation Time:* Figs. 6(a) and (b) show the success rate and computation time, respectively.

PHANS completed all path-planning tasks with the lowest computation times across various densities. Even at a density of 90%, the average computation time remained low at 0.04 s.

By contrast, ILP required an average computation time of approximately 50 s, even in low-density environments. This can be attributed to the time needed for preprocessing and constructing the time-expanded network and the inherently high computational cost of ILP. Despite its high computational cost, ILP achieved high success rates, with values of 1.0 up to a density of 60% and 0.86 at 80%, surpassing those of EECBS.

In low-density environments, the computation time of EECBS was less than that of ILP. However, starting at 20% density, the computation time of EECBS increased drastically, accompanied by a decline in the success rate. EECBS was efficient in low-density environments. By contrast, high-density environments increased the collision resolution cost, leading to lower success rates.

Compared with ILP and EECBS, PIBT required less computation time; at a density of 80%, it took only 0.11 s. The computation time increased considerably at a density of 90%, due to the computational complexity of PIBT. PIBT does not know where empty cells exist prior to the evacuation of obstructing agents. Agents in PIBT only know their goals, and the evacuating position (selecting empty cells) is not optimized. To avoid following conflicts, PIBT performs a random search to determine the evacuating position. If the probability that a cell is occupied is  $p$ , the probability of finding an empty cell in a given number of search attempts is  $p^{\text{times}-1} \cdot (1 - p)$ , where  $p$  is assumed to be fixed and independent of previously visited paths. At a density of 90%, the obstructing agents struggle to find the 10% unoccupied cells. The computational complexity at 90% density is much higher than at lower densities, which consequently increases the computation time.

2) *Makespan:* Fig. 6(c) compares the makespans of PHANS and conventional methods for scenarios where feasible solutions were identified for at least half of the trials.

PIBT has a much larger makespan than PHANS. PIBT can resolve local conflicts (around the target agents) via backtracking. However, backtracking cannot be applied to future conflicts, and moving the obstructing agents in advance is challenging. Therefore, the number of local-conflict resolutions increases, which increases the makespan.

Fig. 6(c) shows the makespan of ILP to be consistently smaller than that of PHANS. This difference is particularly notable in high-density environments, where ILP’s makespan was approximately 5–10% shorter. ILP is an optimal solver, whereas PHANS is suboptimal. The target paths, selected null agents, and paths of the null agents may not be optimal in PHANS; consequently, PHANS has a larger makespan.

### C. Experiment 2

Experiment 1 was conducted in an environment without obstacles. Experiment 2 was conducted in an environment with pillar-like obstacles (Fig. 7). We set the number of target agents to 2 and the cell occupation density varied from




Fig. 7. Simulation environment for Experiment 2. The grid size is  $14 \times 7$ . Obstructing agents and static obstacles occupy 90% of the cells.




Fig. 8. Comparison of (a) success rates, (b) computation time, and (c) makespans between PHANS and conventional methods in Experiment 2. The density is the proportion of cells occupied by obstructing agents and static obstacles.

20% to 90% of the cells. The densities included obstructing agents (varied) and static obstacles (fixed at 90 cells). As in Experiment 1, the maximum computation time for each trial was set to 180 s.

1) *Success Rate and Computation Time:* The success rates and computation times (Figs. 8(a) and (b)) of PHANS, EECBS, and ILP were similar to those in Experiment 1. PIBT achieved shorter computation times than in Experiment 1, particularly at 90% density. PIBT moved obstructing agents to empty cells via backtracking. As static obstacle cells were excluded from the search, the search space was reduced compared with environments without obstacles, ultimately resulting in a smaller computation time and improved success rate.

2) *Makespan:* The makespans (Fig. 8(c)) of EECBS and ILP were similar to those in Experiment 1, whereas PHANS and PIBT achieved shorter makespans under certain conditions. This is presumably because the presence of static obstacles reduces the search space, making it easier to evacuate obstructing agents to favorable locations.

#### D. Experiment 3

We evaluated the scalability of PHANS by testing its performance in an environment larger than that considered in



Fig. 9. Simulation environment for Experiment 3. The grid size is  $35 \times 21$ . Obstructing agents occupy 95% of the cells.

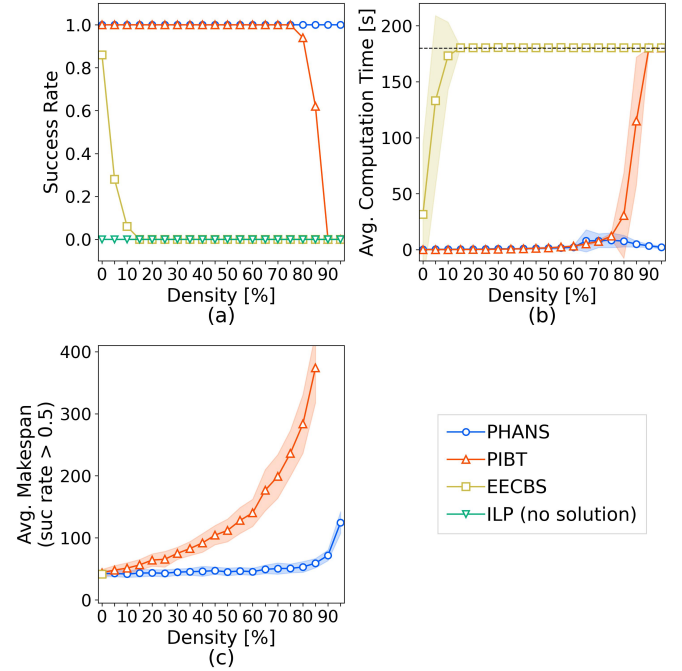


Fig. 10. Comparison of (a) success rates, (b) computation time, and (c) makespans between PHANS and conventional methods in Experiment 3. The density is the proportion of cells occupied by obstructing agents.

Experiments 1 and 2. Specifically, we assessed the makespan and computation time across varying densities of obstructing agents.

Fig. 9 shows the simulation environment, which is a free-location automated warehouse in which the picking locations were designed as the goals for the target agents along the perimeter. We set the number of target agents to 12, and the number of obstructing agents varied from occupying 0% to 95% of the cells. Fifty trials were performed at each density. As in Experiments 1 and 2, the maximum computation time for each trial was set to 180 s for all methods. Additionally, for the ILP baseline only, we ran a separate feasibility check at 95% density with an extended limit of 3600 s; ILP failed to find a solution within this time. All results reported in this section use the 180 s cap.

1) *Success Rate and Computation Time:* Figs. 10(a) and (b) show the success rate and computation time, respec-




Fig. 11. Simulation environment for Experiment 4. The grid size is  $35 \times 21$ . Obstructing agents and static obstacles occupy 95% of the cells.




Fig. 12. Comparison of (a) success rates, (b) computation time, and (c) makespans between PHANS and PIBT in Experiment 4. The density is the proportion of cells occupied by obstructing agents and static obstacles.

tively. PHANS completed all path-planning tasks with minimal computation time across various densities. Similar to Experiment 1, even at the highest density (95%), the average computation time remained low at 2.07 s.

Fig. 10(b) indicates that higher obstructing agent densities generally led to longer computation times. Increasing the number of obstructing agents on the target paths requires more evacuations. This increases the number of calculations required for null-agent paths, which increases the computation time.

Fig. 10(b) shows that the computation time of PHANS decreased when the density of obstructing agents exceeded 80%. Even at a density of 95%, the average computation time was only 2.07 s. PHANS assigns priorities to the obstructing agents, and null agents are systematically allocated until no further assignments are possible (Algorithm 1, lines 6–16). At extremely high densities, the limited availability of null agents simplifies the assignment process and reduces the computation

time. By contrast, ILP faces tighter global coupling on the time-expanded network; EECBS resolves more conflicts as density rises; and PIBT incurs heavier backtracking when feasible local moves vanish. Thus, only PHANS exploits the lack of empty cells to shrink the assignment space, which lowers the computation time at high occupation density.

The success rates and computation times of PIBT, EECBS, and ILP were similar to those in Experiment 1. The computational cost of PIBT increases because, in extremely high-density environments, the probability of having an empty neighboring cell decreases, which increases the backtracking computational cost. In ILP, the costs of preprocessing and constructing the time-expanded network increase exponentially with the scale of the environment. Consequently, even when the computation time limit was relaxed from 180 s to 3600 s, ILP failed to find a feasible solution.

2) *Makespan*: Fig. 10(c) compares the makespans of PHANS and conventional methods.

As in Experiment 1, the makespan of PHANS was significantly smaller than that of PIBT. In large-scale environments, target paths become longer, and the number of obstructing agents that need to be moved increases. At this point, moving the obstructing agents in advance has a significant effect.

The PHANS makespan remained below 60 for densities less than 90%. However, it increased drastically at higher densities. This increase is attributed to the additional time required to evacuate more obstructing agents, which leads to longer waiting times for the target agents and thus longer makespans.

### E. Experiment 4

We evaluated the applicability of PHANS in complex environments with obstacles such as pillars, which are common in large warehouse environments. For Experiment 4 (Fig. 11), we set static obstacles as in Experiment 2 and assessed the makespan and computation time across varying densities. We set the number of target agents to 12, whereas the occupation density varied from 15% to 95% of the cells. The densities included obstructing agents (varied) and static obstacles (fixed at 90 cells). Fifty trials were performed for each density. As in Experiments 1–3, the maximum computation time for each trial was 180 s.

1) *Success Rate and Computation Time*: As in Experiment 3, PHANS demonstrated high computational efficiency in the obstructed environment considered in Experiment 4. Compared with PIBT, PHANS successfully planned paths in most trials with significantly shorter computation time (Figs. 12(a) and (b)). However, at a density of 95%, PHANS failed to plan a path in one trial. PHANS prioritizes target agents based on their distance from the goal and statically determines target paths. Owing to high-priority target agents and static obstacles, a deadlock occurred in one trial during the evacuation of obstructing agents, causing path-planning failure.

2) *Makespan*: In the obstructed environment of Experiment 4, the PHANS makespans were similar to those in Experiment 3; however, the computation time increased at densities of 70–80% (Fig. 12(b) and (c)). The presence of static

obstacles increases the path-finding complexity for evacuating obstructing agents, increases the computational costs, and extends the paths for evacuating obstructing agents. However, by moving obstructing agents in advance based on the movement of target agents, the makespan was not significantly affected.

Interestingly, the makespan and computation time for PIBT were smaller in Experiment 4. As in Experiment 2, this is likely because the presence of static obstacles reduces the search space explored by PIBT during backtracking.

## VI. CONCLUSION

This study defines MAPF-HD and introduces a new heuristic method, PHANS, to solve MAPF-HD efficiently. PHANS operates in two phases: first, it plans paths for the target agents, and then it sequentially evacuates obstructing agents to clear the path. This approach involves swapping positions between agents and empty vertices, which we refer to as null agents.

Numerical experiments demonstrated that existing methods (e.g., ILP, EECBS, and PIBT) often incur substantial computational overhead or yield suboptimal makespans as the density of obstructing agents increases. By contrast, PHANS consistently delivers fast computation times and maintains competitive makespans even under extreme density conditions and in environments with static obstacles.

However, PHANS has limitations in certain complex scenarios. For instance, its performance can degrade in environments with a large number of static obstacles or a high ratio of target agents; nonetheless, such conditions are less typical in automated warehouses and valet parking applications. PHANS does not explicitly handle dynamic obstacles; however, its rapid planning speed allows for frequent re-planning when such obstacles appear.

Overall, our results highlight PHANS as a scalable and practical approach for real-world applications such as high-density warehouse logistics and valet parking systems. Future research will focus on adaptive target-path planning and null-agent assignment strategies. Additionally, we will extend PHANS to handle more complex scenarios and arbitrary graphs.

## REFERENCES

- [1] R. Stern, N. Sturtevant, A. Felner, S. Koenig, H. Ma, T. Walker, J. Li, D. Atzmon, L. Cohen, T. K. Kumar, R. Barták, and E. Boyarski, "Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 10, 2019, pp. 151–158.
- [2] P. R. Wurman, R. D'Andrea, and M. Mountz, "Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses," *AI Magazine*, vol. 29, no. 1, pp. 9–20, 2008.
- [3] J. Li, H. Zhang, M. Gong, Z. Liang, W. Liu, Z. Tong, L. Yi, R. Morris, C. Pasareanu, and S. Koenig, "Scheduling and Airport Taxiway Path Planning Under Uncertainty," in *Proceedings of the 2019 Aviation and Aeronautics Forum and Exposition*, 2019, pp. 1–8.
- [4] A. Okoso, K. Otaki, and T. Nishi, "Multi-Agent Path Finding with Priority for Cooperative Automated Valet Parking," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 2019, pp. 2135–2140.
- [5] D. Silver, "Cooperative Pathfinding," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 1, 2005, pp. 117–122.
- [6] A. Okoso, K. Otaki, S. Koide, and T. Nishi, "High Density Automated Valet Parking via Multi-agent Path Finding," in *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 2146–2153.
- [7] H. Makino, Y. Ohama, and S. Ito, "MARPF: Multi-Agent and Multi-Rack Path Finding," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [8] A. Felner, R. Stern, S. Shimony, E. Boyarski, M. Goldenberg, G. Sharon, N. Sturtevant, G. Wagner, and P. Surynek, "Search-Based Optimal Solvers for the Multi-Agent Pathfinding Problem: Summary and Challenges," *Proceedings of the International Symposium on Combinatorial Search*, vol. 8, no. 1, pp. 29–37, 2017.
- [9] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [10] J. Yu and S. LaValle, "Structure and Intractability of Optimal Multi-Robot Path Planning on Graphs," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, no. 1, pp. 1443–1449, 2013.
- [11] E. Boyarski, A. Felner, R. Stern, G. Sharon, O. Betzalel, D. Tolpin, and E. Shimony, "ICBS: The Improved Conflict-Based Search Algorithm for Multi-Agent Pathfinding," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 6, 2015, pp. 223–225.
- [12] E. Lam, P. L. Bodic, D. D. Harabor, and P. J. Stuckey, "Branch-and-Cut-and-Price for Multi-Agent Pathfinding," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, 2019, pp. 1289–1296.
- [13] R. N. Gómez, C. Hernández, and J. A. Baier, "Solving Sum-of-Costs Multi-Agent Pathfinding with Answer-Set Programming," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 06, pp. 9867–9874, 2020.
- [14] J. Gao, Y. Li, X. Li, K. Yan, K. Lin, and X. Wu, "A review of graph-based multi-agent pathfinding solvers: From classical to beyond classical," *Knowledge-Based Systems*, vol. 283, p. 111121, Jan. 2024.
- [15] M. Barer, G. Sharon, R. Stern, and A. Felner, "Suboptimal Variants of the Conflict-Based Search Algorithm for the Multi-Agent Pathfinding Problem," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 5, 2014, pp. 19–27.
- [16] J. Li, W. Ruml, and S. Koenig, "EECBS: A Bounded-Suboptimal Search for Multi-Agent Path Finding," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, 2021, pp. 12353–12362.
- [17] K. Okumura, M. Machida, and X. Défago, "Priority inheritance with backtracking for iterative multi-agent path finding," *Artificial Intelligence*, vol. 310, 2022.
- [18] Z. Chen, D. Harabor, J. Li, and P. J. Stuckey, "Traffic Flow Optimisation for Lifelong Multi-Agent Path Finding," *Proceedings of the International Symposium on Combinatorial Search*, vol. 17, pp. 265–266, 2024.
- [19] H. Jiang, Y. Zhang, R. Veerapaneni, and J. Li, "Scaling Lifelong Multi-Agent Path Finding to More Realistic Settings: Research Challenges and Opportunities," in *Proceedings of the International Symposium on Combinatorial Search*, 2024, pp. 234–242.
- [20] J.-M. Alkazzi and K. Okumura, "A Comprehensive Review on Leveraging Machine Learning for Multi-Agent Path Finding," *IEEE Access*, vol. 12, pp. 57390–57409, 2024.
- [21] M. L. Fredman and R. E. Tarjan, "Fibonacci heaps and their uses in improved network optimization algorithms," *Journal of the ACM*, vol. 34, no. 3, pp. 596–615, 1987.
- [22] B. Li and H. Ma, "Double-Deck Multi-Agent Pickup and Delivery: Multi-Robot Rearrangement in Large-Scale Warehouses," *IEEE Robotics and Automation Letters*, pp. 1–8, 2023.
- [23] K. Okumura, Y. Tamura, and X. Défago, "Time-Independent Planning for Multiple Moving Agents," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, pp. 11299–11307, 2021.

**Hiroya Makino** Toyota Central R&D Labs., Inc., 41-1, Yokomichi, Nagakute, Aichi, Japan. Email: hiro@mosk.tytlabs.co.jp

**Seigo Ito** Toyota Central R&D Labs., Inc., 41-1, Yokomichi, Nagakute, Aichi, Japan. Email: seigo@mosk.tytlabs.co.jp