

LNS2+RL: Combining Multi-Agent Reinforcement Learning with Large Neighborhood Search in Multi-Agent Path Finding

Yutong Wang¹, Tanishq Duhan¹, Jiaoyang Li², Guillaume Sartoretti¹

¹Department of Mechanical Engineering, National University of Singapore, Singapore

²Robotics Institute, Carnegie Mellon University, USA

e0576114@u.nus.edu, e1280621@u.nus.edu, jiaoyangli@cmu.edu, guillaume.sartoretti@nus.edu.sg

Abstract

Multi-Agent Path Finding (MAPF) is a critical component of logistics and warehouse management, which focuses on planning collision-free paths for a team of robots in a known environment. Recent work introduced a novel MAPF approach, LNS2, which proposed to repair a quickly obtained set of infeasible paths via iterative replanning, by relying on a fast, yet lower-quality, prioritized planning (PP) algorithm. At the same time, there has been a recent push for Multi-Agent Reinforcement Learning (MARL) based MAPF algorithms, which exhibit improved cooperation over such PP algorithms, although inevitably remaining slower. In this paper, we introduce a new MAPF algorithm, LNS2+RL, which combines the distinct yet complementary characteristics of LNS2 and MARL to effectively balance their individual limitations and get the best from both worlds. During early iterations, LNS2+RL relies on MARL for low-level replanning, which we show eliminates collisions much more than a PP algorithm. There, our MARL-based planner allows agents to reason about past and future information to gradually learn cooperative decision-making through a finely designed curriculum learning. At later stages of planning, LNS2+RL adaptively switches to PP algorithm to quickly resolve the remaining collisions, naturally trading off solution quality (number of collisions in the solution) and computational efficiency. Our comprehensive experiments on high-agent-density tasks across various team sizes, world sizes, and map structures consistently demonstrate the superior performance of LNS2+RL compared to many MAPF algorithms, including LNS2, LaCAM, EECBS, and SCRIMP. In maps with complex structures, the advantages of LNS2+RL are particularly pronounced, with LNS2+RL achieving a success rate of over 50% in nearly half of the tested tasks, while that of LaCAM, EECBS and SCRIMP falls to 0%.

Code — <https://github.com/marmotlab/LNS2-RL>

Introduction

Multi-Agent Path Finding (MAPF) focuses on the problem of generating a set of shortest collision-free paths for a team of agents from their initial positions to designated goals in a shared environment. It is a core problem in many applications, such as warehouse automation, where

balancing solution quality and speed is crucial. A recent paper (Li et al. 2022) introduced LNS2, a large neighborhood search-based MAPF algorithm that iteratively replans currently infeasible paths, differing from mainstream MAPF algorithms that rely on more expensive search/backtracking. LNS2 exhibits greater scalability than systematic search algorithms (Sharon et al. 2015) and prioritized planning algorithms (Erdmann and Lozano-Perez 1987). It first quickly generates an initial solution that may contain collisions, and then iteratively replans for carefully selected subsets of agents, using a fast, low-level Prioritized Planning algorithm based on Safe Interval Path Planning with Soft Constraints (termed PP+SIPPS) to finally yield a collision-free solution. However, in challenging tasks, the paths generated by PP+SIPPS still often contain a large number of collisions, inevitably increasing the cumulative time needed to address the overall MAPF problem.

In parallel to these developments, recent years have seen the development of many (slower) MAPF planners based on Multi-Agent Reinforcement Learning (MARL), which exhibit enhanced cooperation compared to prioritized planning algorithms. However, MARL-based planners inevitably exhibit performance degradation when deviating from their training setup (usually only 8 agents). Moreover, the one-shot Success Rate (SR) of MARL-based planners often drops sharply due to a few agents failing to reach their goals within the timestep limit, thereby wasting high-quality paths for the majority that succeeded.

In this paper, we propose combining large neighborhood search and MARL, introducing a new algorithm called LNS2+RL, which balances the drawbacks of both LNS2 and MARL-based planners. The overall framework of LNS2+RL follows that of LNS2, iteratively updating the overall solution with re-planned paths. However, during the challenging early stages of iterative replanning, it replaces PP+SIPPS with a MARL-based planner, generating paths with fewer collisions at a slower speed. There, we rely on SIPPS to fix the paths planned by agents unable to reach their goals within a limited number of timesteps, thus ensuring the completeness of our MARL-based planner and maximizing the utilization of existing high-quality paths. LNS2+RL then adaptively switches to the original PP+SIPPS to swiftly clean up remaining collisions, naturally trading-off solution quality (number of collisions in the so-

lution) and replanning speed. Moreover, since the number of agents involved in each replanning task (i.e., neighborhood size) is fixed to 8, LNS2+RL can naturally scale the overall planning to thousands of agents while keeping encountered tasks close to MARL’s training distribution.

We trained only a single MARL model but tested LNS2+RL on an extensive set of high-agent-density tasks, where the team sizes and map structures differed from those used during training. Our results indicate that LNS2+RL consistently outperforms LNS2 in SR across nearly all tasks, maintaining a comparable Sum of Costs (SoC), and notably reducing the number of remaining Colliding Pairs (CP) in the overall solution. Additionally, LNS2+RL surpasses the search-based bounded-suboptimal algorithm EECBS (Li, Ruml, and Koenig 2021) and MARL-based planner SCRIMP (Wang et al. 2023) in SR across all tasks. Compared to the search-based unbounded-suboptimal algorithm LaCAM (Okumura 2023), LNS2+RL exhibits a substantial advantage in SR on maps with complex obstacle structure (achieving 100% SR compared to LaCAM’s 0%), while also achieving significantly lower SoC on maps with simple obstacle structures. Finally, we experimentally validated our planner in a hybrid simulation of a warehouse environment under real-world conditions.

Prior Work

MAPF algorithms developed in recent decades fall into three main categories based on solution optimality: optimal, bounded-suboptimal, and unbounded-suboptimal. Common optimal algorithms include Conflict-Based Search (CBS) (Sharon et al. 2015), joint-state A*, and their variants with additional speed-up techniques such as ICBS (Boyerski et al. 2015), M* (Wagner and Choset 2011), EPEA* (Goldenberg et al. 2014), among others. Bounded-suboptimal algorithms trade off solution quality and runtime by finding a solution whose cost is at most w times the optimal cost, where w is a user-specified suboptimality factor. Such MAPF algorithms are usually variants of optimal MAPF algorithms, including EECBS (Li, Ruml, and Koenig 2021), and inflated M* (Wagner and Choset 2015), etc. Both optimal and bounded-suboptimal algorithms are able to ensure high solution quality but face scalability challenges when handling thousands of agents, often encountering timeouts or memory overflows. Conversely, unbounded-suboptimal algorithms (such as prioritized planning (Erdmann and Lozano-Perez 1987) and rule-based algorithm (Luna and Bekris 2011)) prioritize speed over optimality, scaling to thousands of agents easily, but may yield substantially sub-optimal solutions in complex tasks and often do not guarantee completeness.

More recently, learning-involved approaches have begun to emerge. One class of approaches, exemplified by PRIMAL (Sartoretti et al. 2019), leverages MARL or imitation learning to directly train MAPF planners. These approaches often incorporate non-learning algorithms to further improve the performance of the learning-based planners by enriching observations (Damani et al. 2021), designing rewards (Skrynnik et al. 2024), or generating training

data (Li et al. 2021). Although these learning-based planners emphasize their superior scalability compared to optimal or bounded-suboptimal search-based algorithms, nearly all of them still fail to outperform scalable suboptimal non-learning algorithms like LNS2 within the same time constraints. Another class of methods seeks to enhance conventional algorithms by using machine learning techniques to replace manually crafted heuristics or components, exemplified by learning a priority ordering for prioritized planning (Zhang et al. 2022), selecting subsets of agents that need to be replanned (Huang et al. 2022). However, these approaches do not directly optimize the path planner itself.

Background

Multi-Agent Path Finding

In this paper, we consider a common formulation of the MAPF problem, defined within a 2D 4-neighbor grid with a finite set of static obstacles, and encompassing a set of m agents $A = \{a_1, \dots, a_m\}$. Each agent a_i initiates from a distinct start s_i , with the objective of reaching its designated goal g_i . At each discretized timestep, an agent can either move to an adjacent cell not occupied by an obstacle or wait at its current cell. The *path* p_i with length k (where the path length k may vary for different paths) for each agent is a sequence of adjacent or identical cells (c_1^i, \dots, c_k^i) satisfying $c_1^i = s_i$, $c_k^i = g_i$. We assume that the agents wait at their goals until all agents have reached their goals. A *solution* of MAPF $P = \{p_i \mid a_i \in A\}$ is a path set for all agents, ensuring all agents reach their goals without any vertex (agents occupy the same cell at the same timestep) or swapping collisions (agents simultaneously attempting to move into each other’s current cells). This paper adopts the *sum of costs* metric as one of the evaluation criteria, which is defined as the sum of timesteps required by each agent to reach its goal $\sum_{1 \leq i \leq m} |p_i|$.

LNS2

LNS2 (Li et al. 2022) is an unbounded-suboptimal and memory-efficient MAPF algorithm based on large neighborhood search. It begins by planning a set of paths that may contain collisions, and then iteratively selects subsets for replanning using adaptive neighborhood selection based on past experience. A replanning solver then computes new paths for these selected agents and evaluates whether to adopt the new paths depending on whether they reduce or maintain the number of CP within the overall solution. This replanning process continues iteratively until the overall path set is collision-free or the time limit is exceeded.

The low-level replanning solver employed in LNS2 is a prioritized planning algorithm PP+SIPPS, which utilizes SIPPS to sequentially plan paths for each agent based on a randomly assigned priority order, where lower-priority agents need to avoid the paths already planned for higher-priority agents as much as possible. SIPPS is designed to find a path for an agent that avoids collisions with hard obstacles (static obstacles) while minimizing collisions with soft obstacles (pre-existing paths). It ensures completeness and returns the shortest path with zero soft collisions if one

exists. However, in dense-obstacle scenarios, SIPPS may produce a path with significantly more soft collisions than the minimum due to its heuristic function ignoring soft collisions that occur when the agent waits within a safe interval (a contiguous period of time) that contains soft obstacles. More importantly, the prioritized planning algorithm relies on a static-priority-based conflict avoidance mechanism, which significantly reduces the flexibility of the planning process. This limitation results in suboptimal paths and a larger number of collisions for lower-priority agents, making it difficult to achieve global cooperation. Despite the fast speed of PP+SIPPS, this reduced solution quality (more collisions) necessitates a higher number of iterations to fully solve high-difficulty MAPF instances, resulting in substantial cumulative time expenditure. Therefore, there is an urgent need to develop a new replanning solver that can generate high-quality solutions (fewer collisions) while potentially sacrificing speed, to find a better balance between quality and time throughout the overall planning.

Path Finding with Mixed Dynamic Obstacles

The replanning task of LNS2 for selected agents can be defined as the Path Finding with Mixed Dynamic Obstacles (PMDO) problem. In the PMDO problem, two finite sets of obstacles are considered. The first category, hard obstacles, aligns with the immovable obstacles in conventional MAPF. The second category, soft obstacles, refers to other unselected agents and is characterized by the ability to move to any cell not occupied by hard obstacles and remain there for an arbitrary duration. Notably, the complete path of each soft obstacle is fully disclosed and remains constant throughout the replanning task. Any occupied cell at a given timestep t can contain either a single hard obstacle or multiple soft obstacles/agents. The objective of the replanning task is to find a path for each selected agent from its start to its goal that avoids collisions with hard obstacles, prioritizes minimizing collisions with soft obstacles and other selected agents, and then minimizes the path length.

Method: LNS2+RL

In this section, we first present the overall framework of LNS2+RL and then detail the MARL model training.

Overall Framework

The strengths and weaknesses of PP+SIPPS and our MARL-based planner are distinct yet complementary. PP+SIPPS is fast and effective for simple tasks but tends to generate paths with more soft collisions in complex tasks. Conversely, our MARL-based planner, although slower (approximately seven times slower), excels at planning paths with fewer collisions in difficult tasks (about one-fourth of PP+SIPPS, with detailed data provided in ‘‘Performance analysis’’ subsection). Therefore, LNS2+RL algorithm retains the overarching framework of LNS2 but uses the MARL-based planner in the initial, more challenging stages of iterative planning for better solution quality (fewer collisions). It then adaptively switches to PP+SIPPS to quickly clean up the remaining collisions, thus balancing solution quality (number of

Algorithm 1: LNS2+RL

```

1: Input: A MAPF instance  $I$ 
2:  $P = \{p_i \mid a_i \in A\} \leftarrow$  InitialSolver ( $I$ )
3:  $Q = \emptyset$ 
4:  $t_m = \max\{|p_i| : p_i \in P\}$ ,  $t_l = d_l \cdot t_m$ ,  $t_h = d_h \cdot t_m$ 
5: while runtime limit not exceeded and  $CP(P) > 0$  do
6:    $A_s, P^- \leftarrow$  selectAgentSet( $I, P$ )
7:   if  $|Q| < \mu$  or  $\bar{Q} > \rho$  then
8:      $P^{rl}, \hat{P} \leftarrow$  MARLPlanner( $I, P \setminus P^-, A_s, t_l$ )
9:      $A_r \leftarrow \{a_i \in A_s : p_{i,end}^{rl} \neq g_i\}$ 
10:     $P^{rl} \leftarrow$  PrioritizedPlanner( $I, P \setminus P^-, P^{rl}, A_r, t_l$ )
11:     $A_h \leftarrow \{a_i \in A_s : |p_i^{rl}| \geq t_h\}$ 
12:     $P^{rl} \leftarrow$  PrioritizedPlanner( $I, P \setminus P^-, P^{rl}, A_h, 0$ )
13:    Update  $Q$  according to if  $CP(P^{rl}) \geq CP(\hat{P})$ 
14:    if  $CP((P \setminus P^-) \cup P^{rl}) \leq CP(P)$  then
15:       $P \leftarrow (P \setminus P^-) \cup P^{rl}$ 
16:    continue
17:  end if
18:  else
19:     $\hat{P} \leftarrow$  PrioritizedPlanner( $I, P \setminus P^-, \emptyset, A_s, 0$ )
20:  end if
21:  if  $CP((P \setminus P^-) \cup \hat{P}) \leq CP(P)$  then
22:     $P \leftarrow (P \setminus P^-) \cup \hat{P}$ 
23:  end if
24: end while
25: return  $P$ 

```

collisions in the solution) with the speed of the replanning task solvers. Moreover, the one-shot SR of MARL-based planners often drops sharply because a few agents fail to reach their goals within the constrained episode length, even though most agents have already found high-quality paths. To address this, we employ PP+SIPPS to supplement paths for these few agents, ensuring the completeness of the final solution.

The pseudocode for the LNS2+RL algorithm is presented in Algorithm 1. Specifically, for solving a MAPF instance I , LNS2+RL starts by invoking the initial solver (PP+SIPPS) to generate a path set P containing collisions (lines 2). Furthermore, LNS2+RL introduces a queue Q with a maximum capacity of μ , used to store the comparison results of the number of CP in the paths generated by PP+SIPPS and the MARL planner, enabling adaptive switching (line 3). Subsequently, LNS2+RL counts the maximum length of individual paths in the initial solution (t_m), sets the timestep threshold t_l for pausing the MARL-planner, along with the maximum allowable path length t_h , where the constant value $d_h > d_l$ (line 4). If the termination criteria are not met, LNS2+RL iteratively selects a subset of agents $A_s \subseteq A$ using adaptive neighborhood selection (the same as in LNS2) and removes their paths $P^- = \{p_i \in P : a_i \in A_s\}$ from P , treating the remaining part of the path set $P \setminus P^-$ as fixed to construct a PMDO replanning task (lines 6). The neighborhood size (the size of A_s) is fixed at 8 (equivalent to the number of agents used in MARL training). If the size of Q is below μ , or its average \bar{Q} exceeds another threshold ρ , LNS2+RL infers that the benefit of using the MARL-planner exceeds that of PP+SIPPS at this stage (line 7). Therefore, LNS2+RL chooses the trained MARL plan-

ner as the replanning solver. The function in line 8 uses the MARL-planner to generate actions until timestep t_l (or until the replanning task is fully resolved, i.e., all agents have reached their goal without future soft collisions). It yields two path sets: \hat{P} , derived by PP+SIPPS at the onset of the replanning task, acting as the reference paths for the MARL policy, and path set P^{rl} , computed by the MARL policy. For incomplete paths in P^{rl} where agents haven't reached their goals (Line 9), we employ PP+SIPPS to plan the remaining paths from t_l and merge these new paths with the existing ones (Line 10). Subsequently, all path lengths in P^{rl} are reviewed, and those exceeding t_h are directly replaced with paths re-planned by PP+SIPPS from timestep zero (lines 11-12). This operation aims to remove non-contributory segments planned by the MARL planner, thus reducing the SoC of the final solution P and planning time for subsequent iterations. Next, LNS2+RL updates Q by appending 0 if $CP(P^{rl}) \geq CP(\hat{P})$, and 1 otherwise (lines 13). Finally, LNS2+RL first checks whether P^{rl} meets the update condition (line 14), namely if the CP of the repaired path set $(P \setminus P^-) \cup P^{rl}$ is no larger than that of the old path set. If not, LNS2+RL continues to check whether \hat{P} meets the update condition (line 21). The MARL planner continues to re-plan until either the switching condition or termination condition is met. Once the replanning solver switches to PP+SIPPS, it cannot switch back to the MARL-planner and always addresses remaining collisions using the same procedure as LNS2.

Replanning Tasks as a MARL Problem

RL Environment Setup Aligned with previously outlined MAPF and PMDO problems, we consider a discrete 2D 4-neighbor grid environment where the size of each agent, goal, and obstacle is one grid cell. The agents' actions that result in collisions with hard obstacles or exceed world boundaries are classified as invalid. Only valid actions are sampled during training and evaluation, and an additional supervised loss is added to the overall MARL loss to reduce the probability of choosing invalid actions.

We initially establish a MAPF task comprising m agents, and use PP+SIPPS to compute the initial path set P for these agents. Subsequently, a subset of agents is selected, and during the training phase of the MARL model, we keep the subset size n (i.e., the neighborhood size) to 8. The PMDO replanning task is then formulated by designating the chosen subset as the controlled agents of the MARL model while the remaining agents act as soft obstacles. Throughout the remainder of this section, we refer to the selected agents simply as agents and to the others as soft obstacles. A training episode terminates either when all agents are on goal at the end of a timestep (and the timestep is larger than the maximum length of the soft obstacle paths) or when the episode reaches the pre-defined timestep limit.

Observation and Reward Aligned with previous works (Sartoretti et al. 2019), the environment is assumed to be partially observable for each agent, restricting it to only access information within its self-centered field of view (FoV). This assumption enables the trained policy to scale

to arbitrary world sizes while fixing the input dimensions of the neural network.

Each agent's observation is divided into two parts. Here, we take agent i and timestep t as an example. The first part of the observation is an eight-element vector consisting of the normalized distance between c_t^i (the position of agent i at timestep t) and g_i along the x-axis and y-axis, as well as the total Euclidean distance, the ratio of the number of CP in the path planned so far by the MARL-planner compared to that in \hat{p}_i (agent i 's SIPPS path), the ratio of t to the predefined episode length, the ratio of t to the maximum length of \hat{P} , the ratio of agents who have reached goal, and agent i 's action at $t - 1$. The section part of the observation consists of thirty-one 2D matrices (Summary of these metrics is in the supplementary material). These matrices can be classified into four categories.

1. Design from prior work Building on previous research (Wang et al. 2023), we introduce four matrices to represent the positions of hard obstacles, other agents, observable agents' goals, and agent i 's own goal. Additionally, there are four binary matrices (one for each move action) that indicate whether selecting the action will bring agent i closer to its goal, as determined by the backward Dijkstra algorithm.

2. Reference path To better guide agent i closer to its goal, we introduce the *SIPPS map*, which represents all cells traversed by agent i 's SIPPS path \hat{p}_i (computed only at the beginning of the replanning task) during timesteps $t - 15$ to $t + 15$. To motivate agents to refer to the SIPPS paths while allowing necessary detours to avoid collisions, we define an off-route penalty ro_t^i :

$$ro_t^i = -\alpha \min \|c_t^i - c\|_2, c \in \{(\hat{p}_i)_j \mid t - 15 \leq j < t + 15\}, \quad (1)$$

where α is a predefined hyperparameter, c represents all cells contained in the SIPPS map of agent i .

3. Avoid collisions To avoid collisions with hard and soft obstacles (with known paths), we further introduce *future path maps of soft obstacles*, and an *occupancy duration map* along with its complementary *blank duration map*. The future path maps, comprising nine matrices, represent the positions of soft obstacles from t to $t + 8$. The occupancy duration map denotes the ratio of time a cell will be occupied by obstacles to the predefined episode termination length, while the blank duration map reflects the ratio of time a cell will remain obstacle-free.

The future paths of other agents are crucial to enhancing cooperation, but these paths are unknown in advance. Hence, we introduce *predicted path maps of other agents* and employ a straightforward yet efficient method for forecasting them. To predict the future path of agent j , we first calculate the spatial Euclidean distance between c_t^j and all cells on \hat{p}_j , and treat the absolute difference of their timestep as the temporal distance. Next, we compute a weighted sum of the spatial and temporal distances for each cell, with weights of 0.9 for spatial distance and 0.1 for temporal distance. If the cell with the smallest weighted sum can be reached from

Action	Move/Stay (not on goal)	Stay (on goal)	Exceed max length of \hat{P}	Revisit	Collision with obstacles	Avoid congestion	Off-route
Reward	[-0.4,-0.5,-0.6]	0.0	-0.2	-0.4	-1.5	Eq. 2	Eq.1

Table 1: Reward structure.

c_t^j in a single step, it is considered the predicted position at $t+1$; otherwise, agent j is assumed to remain stationary, and its current position c_t^j is regarded as the predicted position at $t+1$. Subsequently, the next four cells on \hat{p}_j starting from the cell with the smallest weighted sum are designated as the predicted positions from $t+2$ to $t+5$.

4. Alleviate congestion In path planning tasks involving a large number of agents, congestion often arises in bottleneck areas of the map, where all agents attempt to traverse the same set of cells, leading to under-utilization of other cells and significantly increased collision risks. To alleviate congestion, we introduce the *cell utilization map* and the *direction utilization maps*, which visualize congestion and assist agents in identifying and avoiding high-density areas, promoting more evenly dispersed paths. Specifically, denote $U(c)$ and $U(c_1, c_2)$ as the usage of cell c and direction (c_1, c_2) during the time period from $t-2$ to $t+15$. They represent, respectively, the total timesteps during that time period in which soft obstacles and agents occupy cell c and traverse direction (c_1, c_2) . The normalized value $10 \frac{U(c)}{m}$ and $10 \frac{U(c_1, c_2)}{m}$ are then used as the value of each grid in the cell utilization map and four direction utilization maps. Moreover, we further encourage agents to avoid these congested areas by giving them a small penalty rc_t^i calculated by

$$rc_t^i = 10(\delta_c \frac{U(c)}{m} + \delta_d \frac{U(c_1, c_2)}{m}). \quad (2)$$

Here $\delta_c = 0.225$ and $\delta_d = 0.075$ are used to balance between cell and direction information.

Our reward structure is shown in Table 1. To expedite task completion, agents currently off-goal are penalized at each step, with its magnitude scaling up according to the task’s difficulty within the curriculum learning framework (refer to “Training process“ subsection for details). Agents receive an extra penalty if the current timestep exceeds the maximum length of the current \hat{P} . Additionally, to foster exploration and learn effective policies, agents are penalized for revisiting their previous timestep’s location. Furthermore, we employ reward shaping based on the distance from the goal to guide learning while preserving the invariance of the optimal policy.

Network Structure Figure 1 illustrates the network architecture of our MARL model. Given our assumption of a homogeneous team, we rely on weight sharing across all agents, enabling the model to scale effectively with the number of agents during evaluation.

Understanding the underlying dynamics of the environment is crucial in PMDO tasks. Therefore, we additionally incorporate observations from the past three time steps as input to provide a more comprehensive understanding of how the environment changes over time. The aggregated observation first passes through a Convolutional Long Short-Term

Memory (ConvLSTM) (Shi et al. 2015) unit capable of capturing spatial and temporal relationships and then concatenates its last hidden state with the current observation o_t^i . The concatenated data and the eight-element vector v_t^i are processed through multiple neural network layers to generate a vector \hat{o}_t^i for imitation learning and a predicted state value. Subsequently, \hat{o}_t^i passes through a softmax activation to derive the agent’s policy and a sigmoid activation to yield a vector \hat{o}_t^i for supervised invalid action loss.

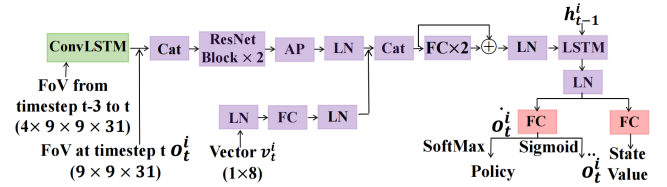


Figure 1: Network structure. AP, Cat, LN, and FC represent average pooling, concatenation, layer normalization, and fully connected layer, respectively. h_{t-1}^i represents the hidden state output by the LSTM unit at the previous time step.

Training Process Our MARL model is trained by Proximal Policy Optimization (PPO) (Schulman et al. 2017). The total loss of the MARL model is $L_{final} = \omega_v \cdot L_v + \omega_\pi \cdot L_\pi + \omega_{invalid} \cdot L_{invalid} - \omega_e \cdot L_e$, where L_v and L_π are the standard critic and actor loss in PPO. L_e is the entropy of the agent’s policy, suppressing its decline shown to encourage exploration and inhibit premature convergence. $L_{invalid}$ is a binary cross-entropy loss (supervised invalid action loss) designed to reduce the logarithmic likelihood of selecting an invalid action. ω_v , ω_π , $\omega_{invalid}$, and ω_e are all manually set hyperparameters. Details regarding the hyperparameters, computational resources, and training duration are provided in the supplementary materials.

The training of the MARL model is divided into two stages. Both stages commence by generating a high-level MAPF task and obtaining the initial path set P through PP+SIPPS. Subsequently, a series of PMDO replanning tasks are generated to train the MARL model. In the first training stage, to reduce training time, the MARL-controlled agents are randomly selected when constructing each PMDO task, bypassing adaptive neighborhood selection. After a MARL episode ends, the path set P is updated with the SIPPS paths \hat{P} calculated at the beginning of the episode, regardless of whether these new paths reduce CP in P . PMDO replanning tasks are continually generated and only switch to the next new MAPF task when the maximum iteration limit is reached. Curriculum learning is adopted during the first stage of training, allowing agents to rapidly acquire fundamental skills from simpler tasks before progressively facing more challenging scenar-

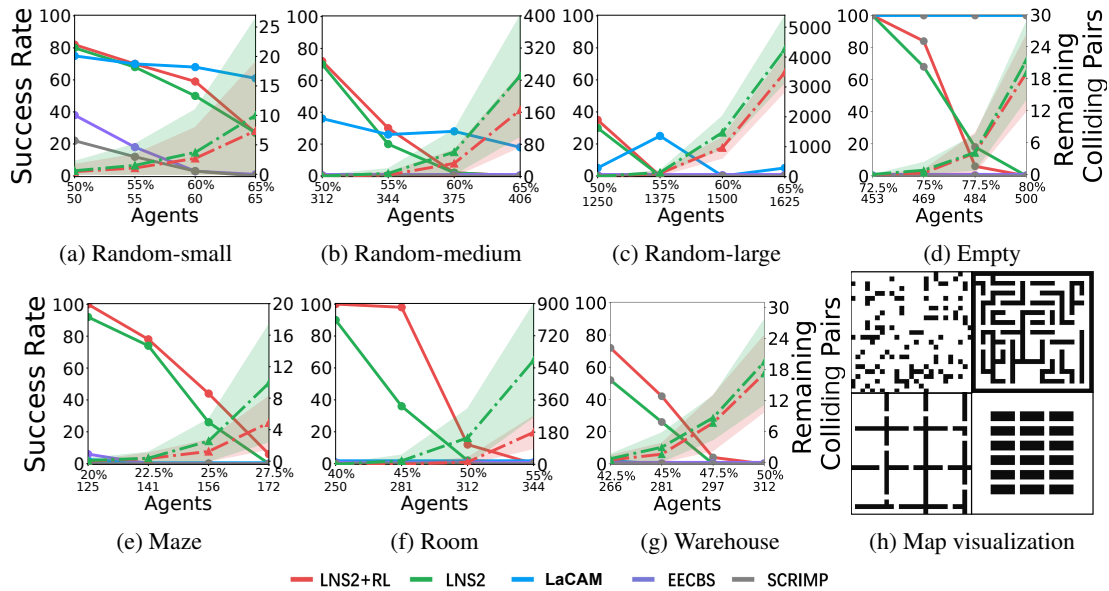


Figure 2: Result on different map types and visualization of representative maps. On the x-axis, the percentage represents agent density, while the number below it represents the number of agents. The success rate is the percentage of tasks that are fully solved before the termination condition is met. The remaining colliding pairs metric represents the number of CP remaining in the overall solution after reaching the time constraint and is only available for LNS2+RL and LNS2. The shaded area represents one standard deviation of the remaining CP. All three random maps have a 17.5% obstacle density. Except for the random-small map with a size of 10×10 and the random-large map with a size of 50×50 , all other five maps have a size of 25×25 . For maps with sizes 10×10 , 25×25 , and 50×50 , the test counts are set to 100, 50, and 20, respectively, and the time constraints for each task, except when using SCRIMP, are set to 100 seconds, 600 seconds, and 5,000 seconds. To ensure consistency with the original paper, SCRIMP uses the maximum timestep as the termination condition, with maximum timesteps of 356, 556, and 1,024 for maps sized 10×10 , 25×25 , and 50×50 , respectively. Figure 2h shows the visualization of the random-medium (top left), maze (top right), room (bottom left), and warehouse (bottom right) map.

ios. Specifically, we increase the task difficulty by increasing the hard obstacles rate and the total number of agents (MARL-controlled agents and soft obstacles) in the high-level MAPF tasks and advancing to greater difficulties once the predetermined volume of training data is accumulated. As task difficulty increases, the advantage of SIPPS weakens, and the agents’ preference for safely wandering in a certain area increases. Therefore, we gradually lower the value of α in the off-route penalty to reduce references to SIPPS paths and gradually increase the action cost to incentivize agents to move toward their goal. Additionally, we employ imitation learning (behavior cloning based on the centralized planner ODrM* (Wagner and Choset 2015)) at the beginning of each task difficulty in the first stage to explore high-quality regions of the agents’ state-action space. Upon reaching the manually predefined number of training steps, the training transitions to the second stage. This stage aims to acclimate the MARL policy to scenarios it will face when inserted into the LNS2+RL algorithm. Therefore, both the selection of agents and the method for updating path set P are exactly the same as the LNS2+RL algorithm’s execution.

Experiments

We compare LNS2+RL with vanilla LNS2¹ as well as a representative set of MAPF algorithms: the bounded-suboptimal algorithm EECBS (with a suboptimality factor $w = 5$)², the complete but unbounded-suboptimal algorithm LaCAM³, and the pure MARL-based algorithm SCRIMP⁴. Code for LNS2, EECBS, LaCAM was written in C++, SCRIMP was written in Python, and LNS2+RL was written in both C++ and Python and connected using pybind. To obtain results faster, LNS2+RL, LNS2, EECBS used Ray to test multiple MAPF tasks simultaneously. However, due to high memory demands (we limited the maximum memory to 100GB per task), LaCAM and SCRIMP were tested using a single process. The same MARL model was used in all evaluations, and it was trained on maps with random obstacles.

Comparison with Baselines

The evaluation results in Figure 2 show that LNS2+RL consistently achieves equal or higher SR compared to LNS2

¹Code from <https://github.com/Jiaoyang-Li/MAPF-LNS2>

²Code from <https://github.com/Jiaoyang-Li/EECBS>

³Code from <https://github.com/Kei18/lacam>

⁴Code from <https://github.com/marmotlab/SCRIMP>

across all tasks, except for those on the empty map with 77.5% agent density (484 agents). It is noteworthy that the improvement of LNS2+RL on CP is more pronounced. In tasks with 65% agent density (65 agents) on the random-small map, despite both algorithms achieving the same SR, LNS2+RL remarkably reduces the remaining CP by 26.3% compared to LNS2. The structural complexity of maps significantly influences algorithm performance. Despite being trained on random maps without prior exposure to environments with complex obstacle structures, such as mazes, rooms, and warehouses, LNS2+RL demonstrates remarkable improvements in SR on these three types of maps. Conversely, LaCAM excels on maps with simple or no obstacle structures, achieving a 100% SR on empty maps, but close to 0% SR on maps with complex obstacle structures. Due to the high difficulty of the selected tasks, the SR of EECBS and SCRIMP remains at 0% for most tasks. Furthermore, taking the tasks in Table 2 as an example, the SoC for LNS2+RL only averages 1.03% higher than that of LNS2. Despite LaCAM displaying higher SR on the random-small map, its SoC notably lags behind. Conversely, EECBS consistently achieves the best SoC but at the expense of obtaining the lower SR. SCRIMP exhibits the poorest performance in both SR and SoC. Moreover, LNS2+RL and LNS2 offer the added benefit of maintaining consistent memory usage, in contrast to the other three baselines, whose memory usage escalates rapidly over time due to the generation of longer paths, larger search frontiers, and more collisions.

Algorithms	Sum of Cost			
	50% agent density	55% agent density	60% agent density	65% agent density
LNS2+RL	852.7	942.6	994.9	1,052.8
LNS2	831.0	939.4	972.4	1,064.2
LaCAM	2,395.2	5,288.3	5,035.1	306,438.3
EECBS	825.1	912.7	929.7	1,030.0
SCRIMP	3,719.4	4,444.3	9,869.7	-

Table 2: Sum of costs on random-small map (accounting only for successful tasks). “-” represents unavailable data due to 0% SR.

Performance Analysis

Table 3 shows the performance ratio of the solution obtained by the MARL policy compared to that of PP+SIPPS in replanning tasks across different numbers of iterations. The MARL policy clearly produces fewer collisions overall, but its advantage in reducing CP is less pronounced, indicating that the MARL policy tends to cause collisions with different agents. With replanning task difficulty decreasing as iterations increase, the MARL policy’s advantage in reducing CP diminishes. Given that the speed of the MARL policy is nearly seven times slower than that of PP+SIPPS, switching to PP+SIPPS becomes a necessary trade-off between time and solution quality (the number of collisions in the solution). Figure 3 illustrates the trend of the number of remaining CP over 2000 iterations in MAPF tasks on random-medium maps with a 65% agent density. As expected, in the early stage, before switching to PP+SIPPS, the number of

CP for LNS2+RL decreases at a faster rate. After switching, the gap between LNS2+RL and LNS2 continues to widen, as the high solution quality achieved by the MARL policy reduces the difficulty of subsequent replanning tasks.

Iterations	Collisions	Colliding Pairs	Sum of Cost	Time
1	-73.7%	-14.6%	+35.0%	+1,175%
50	-72.5%	-1.2%	+19.7%	+737%
100	-73.2%	+2.3%	+15.5%	+633%

Table 3: Performance ratio of the solution obtained by the MARL policy compared to that of PP+SIPPS on low-level replanning tasks in random-medium maps with 65% agent density. The difference between collisions and CP is that for n collisions involving the same agents, the number of collisions is counted as n , but CP is counted as 1.

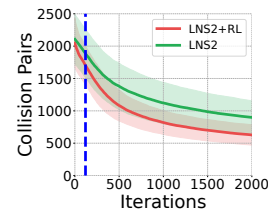


Figure 3: Number of remaining CP through the iterations. The dashed blue line shows the switch from MARL to PP+SIPPS.

Real World Experiment

We implemented our algorithm on a team of 90 virtual and 10 real robots in a warehouse environment. The planned path may not be executed accurately due to disturbances and control inaccuracies. To eliminate these errors, we implement an *Action Dependency Graph* (ADG) (Hönig et al. 2019). From our experiment, we observe that agents can reach their goals quickly without collisions, and errors are eliminated by the ADG, demonstrating the potential of using our method in the real world. More information about our experiment are included in the supplementary material.

Conclusion

This paper introduces LNS2+RL, a novel MAPF algorithm that combines MARL with large neighborhood search. It follows the overall framework of LNS2, beginning with a path set that may contain collisions and then iteratively replanning paths for subsets of agents until the path set becomes collision-free. In early iterations, LNS2+RL addresses these challenging replanning tasks with a MARL-based planner, which offers higher solution quality (fewer collisions) but at a slower pace. As iterations progress, it adaptively switches to a fast, prioritized planning algorithm to quickly resolve remaining collisions. Our experiments on high-agent-density tasks across various world sizes, team sizes, and map structures consistently demonstrate the superior performance of LNS2+RL compared to other state-of-the-art MAPF algorithms. The practicality of LNS2+RL is validated in a hybrid warehouse mockup simulation.

Acknowledgments

This research was supported by the Singapore Ministry of Education (MOE), as well as by an Amazon Research Award.

References

- Boyarski, E.; Felner, A.; Stern, R.; Sharon, G.; Betzalel, O.; Tolpin, D.; and Shimony, E. 2015. Icbs: The improved conflict-based search algorithm for multi-agent pathfinding. In *Proceedings of the International Symposium on Combinatorial Search*, 223–225. AAAI Press.
- Damani, M.; Luo, Z.; Wenzel, E.; and Sartoretti, G. 2021. PRIMAL₂: Pathfinding via reinforcement and imitation multi-agent learning-lifelong. *IEEE Robotics and Automation Letters*, 6: 2666–2673.
- Erdmann, M.; and Lozano-Perez, T. 1987. On multiple moving objects. In *Proceedings of IEEE International Conference on Robotics and Automation*, 477–521. IEEE.
- Goldenberg, M.; Felner, A.; Stern, R.; Sharon, G.; Sturtevant, N.; Holte, R. C.; and Schaeffer, J. 2014. Enhanced partial expansion A. *Journal of Artificial Intelligence Research*, 50: 141–187.
- Hönig, W.; Kiesel, S.; Tinka, A.; Durham, J. W.; and Ayanian, N. 2019. Persistent and robust execution of MAPF schedules in warehouses. *IEEE Robotics and Automation Letters*, 4: 1125–1131.
- Huang, T.; Li, J.; Koenig, S.; and Dilkina, B. 2022. Anytime multi-agent path finding via machine learning-guided large neighborhood search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9368–9376. AAAI Press.
- Li, J.; Chen, Z.; Harabor, D.; Stuckey, P. J.; and Koenig, S. 2022. MAPF-LNS2: Fast repairing for multi-agent path finding via large neighborhood search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 10256–10265. AAAI Press.
- Li, J.; Ruml, W.; and Koenig, S. 2021. Eecbs: A bounded-suboptimal search for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12353–12362. AAAI Press.
- Li, Q.; Lin, W.; Liu, Z.; and Prorok, A. 2021. Message-aware graph attention networks for large-scale multi-robot path planning. *IEEE Robotics and Automation Letters*, 6: 5533–5540.
- Luna, R.; and Bekris, K. E. 2011. Push and swap: Fast cooperative path-finding with completeness guarantees. In *Proceedings of the international Joint Conference on Artificial Intelligence*, 294–300. IJCAI Organization.
- Okumura, K. 2023. Lacam: Search-based algorithm for quick multi-agent pathfinding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 11655–11662. AAAI Press.
- Sartoretti, G.; Kerr, J.; Shi, Y.; Wagner, G.; Kumar, T. S.; Koenig, S.; and Choset, H. 2019. Primal: Pathfinding via reinforcement and imitation multi-agent learning. *IEEE Robotics and Automation Letters*, 4: 2378–2385.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. arXiv:1707.06347.
- Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. R. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219: 40–66.
- Shi, X.; Chen, Z.; Wang, H.; Yeung, D.-Y.; Wong, W.-K.; and Woo, W.-c. 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Proceedings of Advances in Neural Information Processing System*. Curran Associates, Inc.
- Skrynnik, A.; Andreychuk, A.; Nesterova, M.; Yakovlev, K.; and Panov, A. 2024. Learn to Follow: Decentralized Lifelong Multi-Agent Pathfinding via Planning and Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 17541–17549. AAAI Press.
- Wagner, G.; and Choset, H. 2011. M*: A complete multi-robot path planning algorithm with performance bounds. In *2011 IEEE/RSJ international conference on intelligent robots and systems*, 3260–3267. IEEE.
- Wagner, G.; and Choset, H. 2015. Subdimensional expansion for multirobot path planning. *Artificial intelligence*, 219: 1–24.
- Wang, Y.; Xiang, B.; Huang, S.; and Sartoretti, G. 2023. SCRIMP: Scalable communication for reinforcement-and imitation-learning-based multi-agent pathfinding. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 9301–9308. IEEE.
- Zhang, S.; Li, J.; Huang, T.; Koenig, S.; and Dilkina, B. 2022. Learning a priority ordering for prioritized planning in multi-agent path finding. In *Proceedings of the International Symposium on Combinatorial Search*, 208–216. AAAI Press.