

퍼즐 휴리스틱스: 대규모 환경을 위한 효율적인 다중 에이전트 경로 탐색 알고리즘

Puzzle Heuristics: Efficient Lifelong Multi-Agent Pathfinding Algorithm for Large-scale Challenging Environments

이원종^{1*} · 심준열^{2*} · 남창주[†]

Wonjong Lee^{1*}, Joonyeol Sim^{2*}, Changjoo Nam[†]

Abstract: This paper describes the solution method of Team AIRLAB used to participate in the League of Robot Runners Competition which tackles the problem of Lifelong Multi-agent Pathfinding (MAPF). In lifelong MAPF, multiple agents are tasked to navigate to their respective goal locations where new goals are consecutively revealed once they reach initial goals. The agents need to avoid collisions and deadlock situations while they navigate to perform tasks. Our method consists of (i) Puzzle Heuristics, (ii) MAPF-LNS2, and (iii) RHCR. The Puzzle Heuristics is our own algorithm that generates a compact heuristic table contributing to reduce memory consumption and computation time. MAPF-LNS2 and RHCR are state-of-the-art algorithms for MAPF. By combining these three algorithms, our method can improve the efficiency of paths for all agents significantly.

Keywords: Multi-AMR Control, Path Planning, Search Algorithm

1. Introduction

Multi-agent Pathfinding (MAPF) is the problem that finds collision-free paths for multiple agents from their start locations to their goal locations while optimizing the sum of total path lengths or the time taken for all tasks to complete^[1]. It has numerous applications in various domains, such as logistics, aircraft, and urban traffic systems. In many such applications, the agents are tasked to visit multiple destinations consecutively.

Lifelong MAPF^[2] is an extended version of the MAPF problem. In contrast to the canonical form of MAPF, where each agent has

only one task and remains at its goal location, lifelong MAPF requires agents to generate paths to their respective new goal locations after reaching their current goals. The challenges lie in the fact that the next goals are not known to the agents beforehand and the time that the agents finish their current tasks are not synchronized. Thus, it is paramount to have an efficient method that can find collision-free paths promptly for a dynamically varying task set given a limited time budget.

The League of Robot Runners^[3] is a competition to tackle the lifelong MAPF in challenging environments in terms of the complexity of the environments as well as the scale of the agent team. Our team AIRLAB (ranked 8th) developed a method employing one of the state-of-the-art algorithms MAPF-LNS2^[4] and Rolling Horizon Collision Resolution (RHCR) framework^[5] as the MAPF solver. To enhance this approach, we develop Puzzle Heuristics that generates a compact heuristic table to reduce memory consumption and computation, and successfully integrated it with these existing methods.

Received : Apr. 19. 2024; Revised : Jun. 10. 2024; Accepted : Jul. 23. 2024

※ This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2022R1C1C1008476).

* Wonjong Lee and Joonyeol Sim contributed equally to this work.

1. Graduate Student, Dept. of Artificial Intelligence, Sogang University, Seoul, Korea (leewj1208@sogang.ac.kr)

2. Graduate Student, Dept. of Electronic Engineering, Sogang University, Seoul, Korea (jysim@u.sogang.ac.kr)

† Assistant Professor, Corresponding author: Dept. of Electronic Engineering, Sogang University, Seoul, Korea (cnam@sogang.ac.kr)

2. Problem Definition

Our goal is to solve the lifelong MAPF instances presented in the League of Robot Runners competition. Therefore, the problem definition is dictated by the goal and rules of the competition.

A team of k agents $\mathcal{A} = \{a_1, a_2, \dots, a_k\}$ perform errands (i.e., tasks) that appear in an online manner in a deterministic and a fully observable environment. The environment is represented by a grid map $G = (\mathcal{V}, \mathcal{E})$ where \mathcal{V} is the set of vertices (representing grid cells) and \mathcal{E} is the set of edges representing the adjacency of the cells (i.e., G is a 4-connected graph). The agents operate on discrete time steps, denoted as $t = 0, 1, 2, \dots, \tau$ where a constant τ is the final time step of an instance. At each time step, each agent can perform one of the four actions: move forward, rotate left or right 90 degrees, or wait.

The position $v_i^t \in V$ is the position of agent a_i on the grid map at time step t . Let $P_i^t = \{v_i^t, v_i^{t+1}, \dots, v_i^{t+T}\}$ be a path of agent a_i from t to $t+T$ that the agent plans to move. Along the path, an agent must avoid conflicts with other agents. Two types of conflicts are defined: (i) the vertex conflict, where more than one agent moves to the same position, and (ii) the edge conflict, where two agents swap their positions simultaneously. Specifically, a vertex conflict occurs at t if $\exists i, j$ such that $i \neq j$ and $v_i^t = v_j^t$.

And edge conflict occurs at t if $\exists i, j$ such that $i \neq j$, $v_i^t = v_j^{t-1}$ and $v_j^t = v_i^{t-1}$. An errand is located on a position in the map. It can be completed only if the agent assigned to it arrives at the position of the errand. An agent becomes to know the next errand only after it finishes the current one.

Given the assumptions and definitions, the objective of the lifelong MAPF is to maximize the number of errands to be completed by \mathcal{A} within $t = 0, \dots, \tau$.

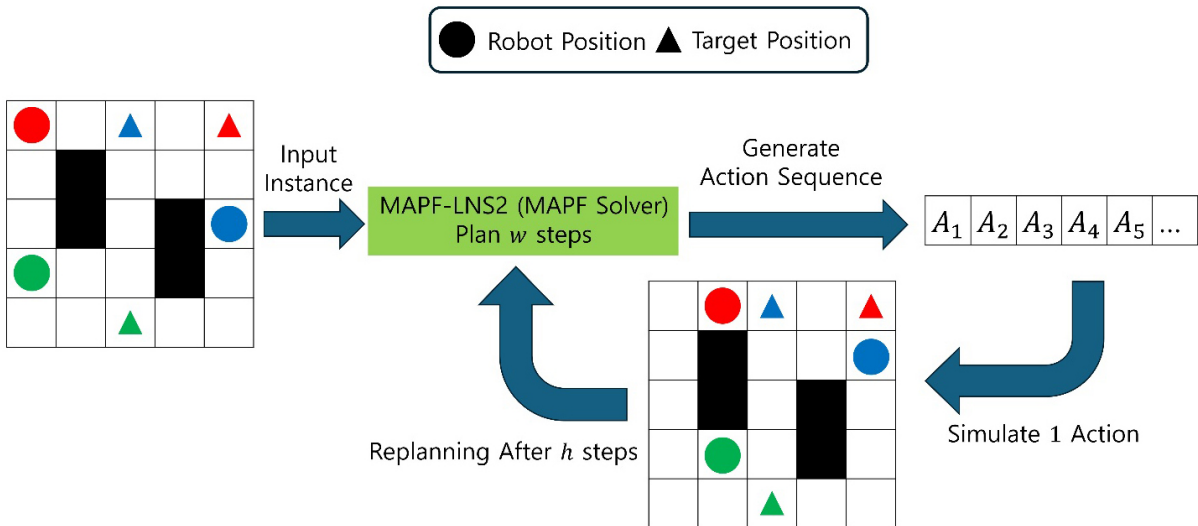
3. Method

By integrating MAPF-LNS2, RHCR, and Puzzle Heuristics, our method can handle large-scale lifelong MAPF instances in challenging environments. [Fig. 1] represents an overview of our algorithm procedure.

We provide a detailed description of each component and how they contribute to the overall effectiveness and efficiency of our solution method.

3.1 Puzzle Heuristics

The shortest distance between a pair of cells is frequently queried in an MAPF instance. In a fast-paced situation where the



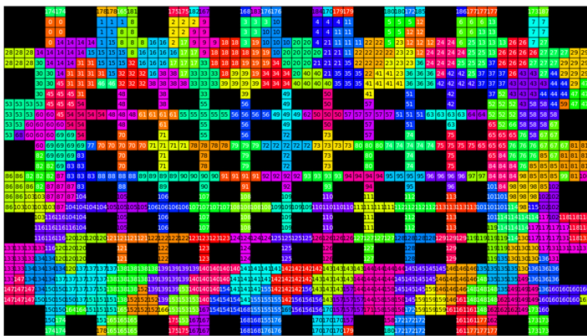
[Fig. 1] Illustration of the process of the MAPF-LNS2 and RHCR for Lifelong Multi-Agent Pathfinding (MAPF). Starting from the input MAPF instance, the solver plans w steps, which is the window size representing the time horizon for conflict resolutions, and generates an action sequence. After each action is simulated, the system replans after h steps, where h is the replanning period indicating the frequency of replanning. This windowed approach represents the RHCR method. The diagram shows the initial state, intermediate steps, and the continuous loop of planning and replanning to achieve collision-free paths for agents in a grid map. The circles represent the positions of the robots, while the triangles indicate their target positions

quick online computation of paths is necessary, it is a common approach to precalculate the distances of all pairs. However, as the size of the map grows, the memory required to store the (Manhattan) distances increases significantly. A naïve method has the space complexity $O(N^2)$ for a grid map with N cells. This is particularly problematic in the competition since the memory size allocated to our method is limited. If we reduce memory usage by storing only part of the distances, some of the computation time needed to resolve inter-agent conflicts at runtime needs to be allocated to calculate the distances. Therefore, the quality of the solution should be compromised to balance the trade-off between storage space and computational overhead.

The core idea of Puzzle Heuristics is to divide the grid map into areas such that each area contains contiguous cells within a depth d in a tree search from its central cell called the pivot point. The algorithm begins by selecting the lowest indexed cell (the index increments from top to bottom and left to right like the pixel coordinates) in the grid map as the first pivot point. Using Breadth-First Search (BFS), the area expands over non-obstacle cells from the pivot point, ensuring that the frontier of the BFS does not exceed a depth of d .

Cells within the area are labeled with the same unique puzzle index. The process selects the next unlabeled cell with the lowest index as the new pivot point to repeat the area expansion. This procedure iterates until all cells in the grid map are labeled. The same-colored neighboring cells shown in the map of [Fig. 2] belong to the same area.

Once all cells are formed into areas, the shortest distances between all pairs of areas (i.e., pivot points) are calculated and stored in the heuristic table. By storing only the distances bet-



[Fig. 2] An example result of the Puzzle Heuristics method for generating a compact heuristic table. The same-colored neighboring cells belong to the same area

ween pivot points instead of all cells, we could meet the memory constraint. When the distance between two cells is queried, the heuristic value in the table effectively approximates the actual distance.

The number of cells in an area can reach up to $2d^2 + 2d + 1$ so at most $2d^2 + 2d + 1$ cells can be represented by a single value. As a result, the size of the heuristic table using Puzzle Heuristics could be reduced to $(N / (2d^2 + 2d + 1))^2$, leading to significant memory savings compared to the naïve approach.

3.2 MAPF-LNS2

MAPF-LNS2 is one of the state-of-the-art algorithms for solving large MAPF instances by leveraging the powerful meta-heuristic technique Large Neighborhood Search^[6]. A key advantage of MAPF-LNS2 is its ability to start from any of feasible or infeasible initial solutions. The algorithm iteratively selects a subset of agents using a neighborhood selection method. The method destroys the current paths of agents and repairs them using an efficient single-agent path planner that minimizes the number of collisions. If the repaired solution reduces the number of collisions, the new paths replace the old ones. This destroy-and-repair procedure continues until a stopping criterion is satisfied.

To achieve high efficiency, MAPF-LNS2 employs Safe Interval Path Planning with Soft Constraints (SIPPS), a variant of the Safe Interval Path Planning algorithm^[7] that can handle hard and soft constraints. Integrating SIPPS has shown significant improvements of MAPF-LNS2. While MAPF-LNS2 can be used with many existing MAPF algorithms, we choose to use Prioritized Planning (PP)^[8] which is simple to implement and fast.

3.3 RHCR

The RHCR is a framework for solving the lifelong MAPF by decomposing an instance into a sequence of Windowed MAPF instances. RHCR utilizes two user-specified parameters: the time horizon w and the replanning period h . The time horizon w specifies the time window such that the MAPF solver must resolve conflicts within the next w steps (i.e., windowed MAPF solver). The replanning period h determines the frequency of the solver to replan paths. To avoid collisions, w should be larger than or equal to h . The values of these parameters are critical for

the performance of RHCR. Too small values of w would lead to deadlocks whereas too large values could result in inefficient solutions and increased computation time.


3.4 Greedy Randomized Search

Although MAPF-LNS2 and RHCR can solve large instances quickly, they would lead to deadlocks (i.e., no solution found) as PP is not proven to be complete.

Consider a scenario where a_i and a_j face each other at t and attempt to move forward. Here a_i has a higher priority. During the planning phase, a_i plans its path P_i^t while ignoring the presence of a_j as a_j has a lower priority. As a result, the planned path of a_i goes through the current location of a_j at $t+1$ (i.e., $v_i^{t+1} = v_j^t$). Given this path P_i^t , a_j begins to plan P_j^t .

In this situation, all possible actions for a_j at $t+1$ result in conflicts because P_i^t passes through v_j^{t+1} in any case. Specifically, if a_j chooses to move forward, then an edge conflict occurs (i.e., $v_i^{t+1} = v_j^t = v_j^{t+1}$ or $v_j^{t+1} = v_j^t = v_i^{t+1}$), respectively). Thus, a_j cannot find a conflict-free action at $t+1$. [Fig. 3] illustrates the three conflict situations incurring deadlocks.

We develop a greedy randomized search that combines greediness and randomness to deal with deadlocks. We probabilistically multiply the number of edge collisions and vertex collisions by a factor of two. By doing so, our method could have an opportunity to explore alternative paths rather than relying on routes determined by the heuristic values. This simple yet effective approach enables the algorithm to explore a broader range of




[Fig. 3] An illustration of deadlock situations where a_i and a_j face each other and attempt to move forward as shown in the grid map in the left. Blue (a_i) has a higher priority so plans P_i^t which passes through the current position of a_j at $t+1$. If a_j moves forward, an edge conflict occurs. If a_j chooses to rotate or wait, a vertex conflict occurs


possibilities and escape from deadlocks. In our experiments, the version with the greedy randomized search recorded of 31 times higher number of errands than the version without the greedy randomized search.

4. Experiments

We evaluate our method in five environments provided by the competition: Sortation, Warehouse, Game, City, and Random map, as shown in [Fig. 4]. Each instance has predetermined agent-task pairs. We vary the number of agents from 50 to 3000 agents. The evaluation metric is the total number of errands completed in 5,000 seconds. The test system is with AMD Ryzen



[Fig. 4] The test environments in the competition. White cells represent spaces that robots can occupy, while black cells indicate the obstacles that robots cannot occupy. From the top and left, (1) Random, (2) City, and (3) Game, (4) Warehouse, and (5) Sortation maps



[Fig. 5] The number of completed errands by the proposed method in five challenging environments: Sortation, Warehouse, Game, City, and Random, with the number of agents ranging from 50 to 3000

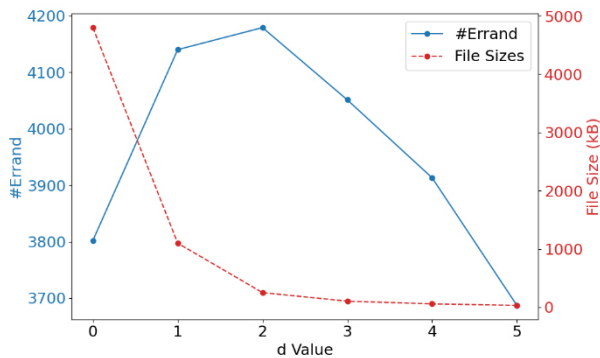
5800X 3.8 GHz CPU and 32G RAM. The source code is written in C++17.

As shown in [Fig. 5], our method effectively solves instances in the random map environment up to 200 agents.


In other environments, the solution quality remains reasonable for instances with up to 1,200 agents but declines as the number of agents increases.

The ability of MAPF-LNS2 and RHCR to solve large instances allows our algorithm to handle scenarios with a significantly large number of agents. The Puzzle Heuristics plays a crucial role in reducing the memory consumption and computation time for large-sized maps like Sortation, Warehouse, and Game. In the greedy randomized search allows us to escape deadlocks effectively even in dense environments like Sortation, Warehouse, and Random.

The algorithm struggles to find individual paths of all agents within the time budget, particularly with a large number of



[Fig. 6] The impact of varying the d value on the number of completed errands (blue line) and file size (red dashed line). As the d value increases from 0 to 5, the number of errands initially increases, reaching a peak at $d = 2$, before declining. Conversely, the file size decreases with increasing d values



[Fig. 7] The impact of varying the window size w on the number of completed errands. As the w size increases from 5 to 25, the number of completed errands initially rises, reaching a peak at $w = 20$, before showing a decline

agents owing to the characteristics of PP, which plans paths sequentially. Since higher-rank agents are considered fixed obstacles to lower-rank agents, the feasible positions become scarce.

In the context of discussing the impact of the d value on the performance of our method, we analyzed how varying d affects both the number of completed errands and the associated file size in Small Warehouse instance. [Fig. 6] illustrates these effects. As depicted, increasing the d value from 0 to 5 leads to a notable pattern: the number of completed errands rises initially, peaking at $d = 2$, before declining as d continues to increase. This indicates the best d value for maximizing the number of errands completed. Concurrently, the file size decreases as d increases, suggesting that higher d values lead to more efficient memory usage. This trade-off between errand completion and file size is critical for optimizing the performance of our method in various environments.

In [Fig. 7], we observe a trade-off in choosing the value of w , which is the planning horizon w in RHCR. A low value of w leads to efficient planning but could cause deadlocks because agents may not be able to plan far enough ahead to avoid conflicts. If we choose a high value, the ability to avoid deadlocks is improved. Balancing between these two cases is critical for the performance of algorithm.


5. Conclusion

In this paper, we presented a novel approach for solving the challenging lifelong MAPF problem in the context of the League of Robot Runners competition. Our method combines three key techniques: Puzzle Heuristics, MAPF-LNS2, and RHCR. Additionally, we introduced a greedy randomized search to escape from deadlock. Experimental results demonstrated the effectiveness of our approach in various environments, particularly in random maps with a moderate number of agents. Through participation in the competition, we found our future directions that include improving the scalability of the method and implementing methods to avoid deadlock situations.

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," *International Symposium on Combinatorial Search*, vol. 10, no. 1, pp. 151-158, 2019, DOI: 10.1609/socs.v10i1.18510.

- [2] H. Ma, J. Li, T. K. S. Kumar, and S. Koenig, "Lifelong multi-agent path finding for online pickup and delivery tasks," *arXiv:1705.10868*, 2017, DOI: 10.48550/arXiv.1705.10868.
- [3] *The League of Robot Runners*, [Online], <https://www.leagueofrobotrunners.org>, Accessed: Aug. 30, 2023.
- [4] J. Li, Z. Chen, D. Harabor, P. J. Stuckey, and S. Koenig, "MAPF-LNS2: Fast repairing for multi-agent pathfinding via large neighborhood search," *AAAI Conference on Artificial Intelligence*, vol. 36, no. 9, pp. 10256-10265, 2022, DOI: 10.1609/aaai.v36i9.21266.
- [5] J. Li, A. Tinka, S. Kiesel, J. W. Durham, T. K. S. Kumar, and S. Koenig, "Lifelong multi-agent path finding in large-scale warehouses," *AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, pp. 11272-11281, 2021, DOI: 10.1609/aaai.v35i13.17344.
- [6] P. Shaw, "Using constraint programming and local search methods to solve vehicle routing problems," *International Conference on Principles and Practice of Constraint Programming*, pp. 417-431, 1998, DOI: 10.1007/3-540-49481-2_30.
- [7] M. Phillips and M. Likhachev, "SIPP: Safe interval path planning for dynamic environments," *2011 IEEE International Conference on Robotics and Automation*, Shanghai, China, pp. 5628-5635, 2011, DOI: 10.1109/ICRA.2011.5980306.
- [8] D. Silver, "Cooperative pathfinding," *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 1, no. 1, pp. 117-122, 2005, DOI: 10.1609/aiide.v1i1.18726.



이 원 종

2024 서강대학교 기계공학과(공학사)
2024~현재 서강대학교 인공지능학과 석사과정

관심분야: Multi-agent path finding



심 준 열

2023 인하대학교 컴퓨터공학과(공학사)
2023~현재 서강대학교 전자공학과 석사과정

관심분야: Multi-agent path finding



남 창 주

2009 고려대학교 전기전자전파공학부(공학사)
2011 고려대학교 전기전자전파공학과(공학 석사)
2016 Texas A&M University(공학박사)
2018~2021 한국과학기술연구원 지능로봇 연구단 선임연구원

2021~2022 인하대학교 정보통신공학과 조교수

2022~현재 서강대학교 전자공학과 조교수

관심분야: Multi-robot systems, robotic manipulation, multi-robot navigation