

# Leveraging Experience in Lifelong Multi-Agent Pathfinding

Nitzan Madar, Kiril Solovey, Oren Salzman

Technion, Israel Institute of Technology  
nitzanmadar@campus.technion.ac.il, kirilsol@technion.ac.il, osalzman@cs.technion.ac.il

## Abstract

In Lifelong Multi-Agent Path Finding (L-MAPF) a team of agents performs a stream of tasks consisting of multiple locations to be visited by the agents on a shared graph while avoiding collisions with one another. L-MAPF is typically tackled by partitioning it into *multiple consecutive*, and hence *similar*, “one-shot” MAPF queries, as in the Rolling-Horizon Collision Resolution (RHCR) algorithm. Therefore, a solution to one query informs the next query, which leads to similarity with respect to the agents’ start and goal positions, and how collisions need to be resolved from one query to the next. Thus, *experience* from solving one MAPF query can potentially be used to speedup solving the next one. Despite this intuition, current L-MAPF planners solve consecutive MAPF queries *from scratch*. In this paper, we introduce a new RHCR-inspired approach called exRHCR, which exploits experience in its constituent MAPF queries. In particular, exRHCR employs an extension of Priority-Based Search (PBS), a state-of-the-art MAPF solver. The extension, which we call exPBS, allows to *warm-start* the search with the priorities between agents used by PBS in the previous MAPF instances. We demonstrate empirically that exRHCR solves L-MAPF instances up to 39% faster than RHCR, and has the potential to increase system throughput for given task streams by increasing the number of agents a planner can cope with for a given time budget.

## 1 Introduction and Related Work

*Multi-Agent Path Finding* (MAPF) is the problem of finding collision-free paths for a fleet of agents operating on a shared graph (Stern et al. 2019). MAPF has been widely used in modeling a variety of applications including autonomous warehouse management (Hönig et al. 2019), multi-robot motion planning (Dayan et al. 2021), and multi-drone delivery (Choudhury et al. 2021).

A significant body of work is devoted to the study of the one-shot version of MAPF wherein each agent needs to reach a specific goal location from a given start. It is usually desirable to find *high-quality solutions*, which minimize a given objective function such as the total travel time of all the agents. Finding optimal solutions to MAPF is typically computationally hard (Yu 2016; Ma et al. 2016). Nevertheless, numerous approaches have been developed which

strive to return optimal or bounded-suboptimal solutions, including SAT solvers (Surynek et al. 2016), graph-theoretic approximation algorithms (Yu 2016; Demaine et al. 2019), and search-based approaches (Sharon et al. 2013; Wagner and Choset 2015; Li et al. 2019a).

One of the most popular approaches for MAPF is Conflict-Based Search (CBS) (Sharon et al. 2015), which is a two-level optimal solver. The top level performs a best-first search on a binary *constraint tree* (CT), where a given tree node specifies temporal and spatial constraints between agents in order to avoid collisions. The bottom level searches for single-agent paths that abide to the constraints in the given CT node. The high-level search continues until a feasible (collision-free) plan is found. Due to the possibly long runtime of the approach (Gordon, Filmus, and Salzman 2021), subsequent work introduced various suboptimal CBS extensions to improve runtime (Barer et al. 2014; Li, Ruml, and Koenig 2021). A recent approach called Priority-Based Search (PBS) trades off the completeness guarantees of CBS with improved efficiency, by exploring a binary *priority tree* (PT), whose nodes specify priorities between agents (Ma et al. 2019a). Priorities can be viewed as a coarser and more computationally-efficient alternative to CBS’s constraints (see more details in Section 3).

The efficiency of PBS makes it particularly suitable for solving large problem instances, or in cases where multiple MAPF instances need to be solved rapidly. For example, PBS has recently been used as a building block within an approach for tackling the Lifelong MAPF (L-MAPF) problem (Ma et al. 2017, 2019b; Liu et al. 2019; Salzman and Stern 2020). In L-MAPF the agents need to execute a stream of tasks consisting of multiple locations to be visited by each agent (rather than moving to a specific goal location as in the one-shot version). A recent work (Li et al. 2021) proposed an effective approach called Rolling-Horizon Collision Resolution (RHCR) to solve L-MAPF, by breaking the L-MAPF problem into a *sequence* of one-shot MAPF problems (see Section 3). Clearly, by improving the efficiency of the internal MAPF solver (e.g., PBS) we can improve the efficiency of RHCR overall, as it solves *multiple* MAPF queries.

In this work we explore the use of *experience* gained from solving previous MAPF queries to speed up the solution of the entire L-MAPF problem. Using experience has been considered in various domains and properly using experi-

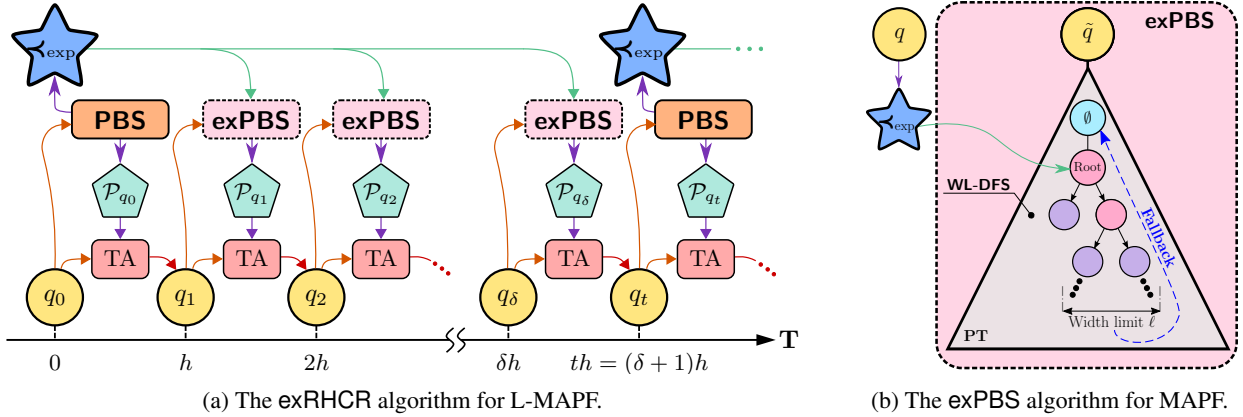


Figure 1: Visualization of our algorithmic framework. (a) To solve the L-MAPF problem, **exRHCR** solves a sequence of (bounded horizon) MAPF queries  $q_0, q_1, \dots$  where a new query is solved every  $h$  timesteps. The task assigner (TA) generates the next query  $q_{i+1}$  from a given query  $q_i$  and the solution  $\mathcal{P}_{q_i}$  computed for  $q_i$ . In **exRHCR** queries are grouped into batches of  $\delta + 1$  queries, where the first query in each batch is solved using **PBS** whereas the next  $\delta$  queries are solved using **exPBS**. This is in contrast to **RHCR**, which uses only **PBS**. The **PBS** priority set  $\prec_{\text{exp}}$  that was used to solve  $q_0$  is extracted and used to initialize **exPBS** in the next  $\delta$  queries. This process repeats iteratively. (b) The internal MAPF solver **exPBS** leverages experience  $\prec_{\text{exp}}$  by using it in the root node of the priority tree (PT) to constrain and guide the search. **exPBS** uses **WL-DFS**, which is a **DFS** variant that limits the search tree width to avoid searching in over-constrained trees. When the width limit is violated or the root PT node contains an infeasible seed priority, the original **PBS** is used as a fallback.

ence to improve MAPF solvers has been identified as a key challenge and opportunity (Salzman and Stern 2020). For instance, in robot motion planning previous solutions can be reused for a new query (Coleman et al. 2015). Similar ideas were employed within **CBS**-based approaches to incentivize agents to traverse certain regions, e.g., corridors, in a specific manner to resolve conflicts (Cohen and Koenig 2016; Li et al. 2019b, 2020). Learning-based approaches have been used to implicitly encode experience for conflict resolution in **CBS** (Huang, Koenig, and Dilkina 2021) and MAPF algorithm selection (Kaduri, Boyarski, and Stern 2020). With that said, we are not familiar with systematic approaches that exploit experience within MAPF queries in L-MAPF.

**Contribution.** In this paper we develop the Experienced **RHCR** (**exRHCR**) approach for L-MAPF, which allows to transfer *experience* gained from solving one MAPF query to the next one in order to improve planning times (Figure 1a and Section 4). In particular, **exRHCR** solves constituent MAPF queries by interleaving between calls to the vanilla **PBS** and an extension of **PBS** which we call Experienced **PBS** (**exPBS**) (Figure 1b). Unlike the standard **PBS**, which begins the high-level search on a given MAPF instance from an empty priority set at the PT root, **exPBS** is initialized with a specific priority set for the root node that encodes the experience gained from the solution of a previous MAPF query using **PBS**. This allows to reduce the depth of the search tree and speed up the solution of MAPF queries. We also consider a lightweight version of **exPBS** that uses total priorities as experience, which leads to good performance in easier instances. We demonstrate that **exRHCR** solves L-MAPF instances up to 39% faster than **RHCR**, which allows to potentially increase the throughput of a given stream of

task streams by increasing the number of agents we can cope with for a given time budget (Section 5).

## 2 Problem Definition

In this section we provide a definition of Lifelong MAPF (L-MAPF), for which we design an effective algorithmic approach in Section 4. Before describing L-MAPF, we first define the single-query, or one-shot, setting termed MAPF.

### 2.1 Multi Agent Path Finding (MAPF)

In the MAPF problem, we are given a graph  $G = (V, E)$  and a set of  $k$  agents  $\mathcal{A} = \{a_1, \dots, a_k\}$ , where an agent  $a_i \in \mathcal{A}$  starts at a vertex  $s_i \in V$  and needs to reach a goal vertex  $g_i \in V$ . Time is discretized such that at each timestep, each agent occupies exactly one vertex. Between any two consecutive timesteps, each agent can either *move* to an adjacent vertex (i.e., along an edge connecting its current vertex and the destination vertex) or *wait* in its current vertex, where each action is assigned a unit cost.

The goal of MAPF is to find a *feasible* solution plan, which consists of a set of paths  $\mathcal{P} = \{p_1, \dots, p_k\}$ , where  $p_i$  is a path for agent  $a_i$  from  $s_i$  to  $g_i$ , such that no *conflicts* arise between the different agents. In particular, for a given agent  $a_i \in \mathcal{A}$ , a path  $p_i = (v_0^i, v_1^i, \dots, v_{T_i}^i)$  is a sequence of  $T_i \geq 1$  vertices, such that  $(v_t^i, v_{t+1}^i) \in E$  or  $v_t^i = v_{t+1}^i$ , and  $v_0^i = s_i, v_{T_i}^i = g_i$ . There are two types of conflicts between any two given agents  $a_i \neq a_j$  to avoid: In a *vertex conflict* the agents occupy the same vertex at the same timestep, i.e.,  $v_t^i = v_t^j$  for some  $1 \leq t \leq \min\{T_i, T_j\}$ . In an *edge conflict*, the agents cross the same edge at the same timestep, i.e.,  $v_t^i = v_{t+1}^j$  and  $v_{t+1}^i = v_t^j$  for timestep  $1 \leq t < \min\{T_i, T_j\}$ .

We will also consider a relaxed version of MAPF called Windowed-MAPF (W-MAPF), which is defined for a constant parameter  $w \geq 1$  indicating the window size wherein conflicts should be resolved (Silver 2005). I.e., a solver for W-MAPF needs to avoid conflicts only up to timestep  $w$ , after which conflicts are allowed. This relaxation reduces the size of the search tree, in comparison to MAPF, which leads to faster computation times. Most solvers for vanilla MAPF can be adapted to work for the windowed case.

## 2.2 Lifelong Multi Agent Path Finding (L-MAPF)

In L-MAPF, agents need to execute a stream of tasks, where each task consists of moving between two specific vertices, while avoiding conflicts. In this setting the objective can be completing all tasks as quickly as possible, or maximizing throughput, which is the average number of tasks completed per unit of time.

A critical component within L-MAPF solvers is a *task assigner*, which specifies for each agent the next task it performs. There are various approaches to design task assigners (Ma et al. 2017; Liu et al. 2019), which are outside the scope of this work. One common approach for tackling L-MAPF problems is partitioning it into a sequence of W-MAPF problems, where the task assigner is invoked after every MAPF query is completed to generate start and goal locations for the next MAPF instance (Ma et al. 2017; Liu et al. 2019; Li et al. 2021). We describe one such approach called RHCR in the next section.

## 3 Algorithmic Background

We describe two algorithmic components, which we will build upon to design our approach for L-MAPF (Section 4). In particular, we describe Priority-Based Search (PBS) for MAPF, and Rolling-Horizon Collision Resolution (RHCR) for L-MAPF.

### 3.1 Priority-Based Search (PBS)

PBS is a recent approach for MAPF that can be used to solve instances with a large number of agents, or situations where multiple MAPF instances need to be solved rapidly (as in L-MAPF). We now provide an overview of PBS, and refer the reader to (Ma et al. 2019a) for the full description.

PBS maintains priorities between agents which are used to resolve conflicts. On the high-level search PBS explores a *priority tree* (PT), where a given node  $N$  of PT encodes a (partial) priority set  $\prec_N = \{a_h \prec a_i, a_j \prec a_l, \dots\}$ . A priority  $a_i \prec a_j$  means that agent  $a_i$  has precedence over agent  $a_j$  whenever a low-level search is invoked (see below). In addition to the ordering, each PT node maintains single-agent plans that represent the current MAPF solution (possibly containing conflicts). PBS starts the high-level search with the tree root whose priority set is empty, and assigns to each agent its shortest path. We mention that PBS can be initialized with a non-empty priority set in its root node, albeit previous work has not specified an effective method to specify this set. Whenever PBS expands a node  $N$ , it invokes a low-level search to compute a new set of plans which abide to the priority set  $\prec_N$ . If a collision between agents, e.g.,

$a_i$  and  $a_j$ , is encountered in the new plans, PBS generates two child PT nodes  $N_1, N_2$  with the updated priority sets  $\prec_{N_1} = \prec_N \cup \{a_i \prec a_j\}$ ,  $\prec_{N_2} = \prec_N \cup \{a_j \prec a_i\}$ , respectively. The high-level search chooses to expand at each step a PT node in a depth-first search (DFS) manner. The high-level search terminates when a valid solution is found at some node  $N$ , or when no more nodes for expansion remain, in which case PBS declares failure.

The low-level search of PBS proceeds in the following manner. For a given PT node  $N$ , PBS performs a topological sort of the agents according to  $\prec_N$  from high priority to low, and plans individual-agent paths based on the sorting. Given a topological sort  $(a'_1, \dots, a'_{k'}) \subset \mathcal{A}$ , for some  $1 \leq k' \leq k$ , the low-level search iterates over the  $k'$  agents in the topological sort, and updates their plans such that they do not collide with any higher-priority agents (agents that do not appear on this list maintain their original plans). It then checks whether collisions occur between all the agents combined.

### 3.2 Rolling-Horizon Collision Resolution (RHCR)

RHCR (Li et al. 2021), is a state-of-the-art framework for solving L-MAPF. RHCR accepts as parameters a time window  $w \geq 1$ , and a replanning rate  $h \leq w$ , and decomposes the L-MAPF problem into a sequence of W-MAPF queries  $q_0, q_1, \dots$  that are solved one by one. In particular, after obtaining a W-MAPF plan for a query  $q_i$  with time window  $w$ , the next query  $q_{i+1}$  is generated by executing the solution plan for  $q_i$  for  $h$  steps. This yields the agent’s start locations for query  $q_{i+1}$ . The goal locations remain the same for agents that did not complete their task and new goal locations are assigned by the task assigner for agents that did. For a concrete example of a task assigner, see Section 5.

RHCR requires a W-MAPF solver as an internal submodule. Accordingly, several bounded-horizon versions of state-of-the-art MAPF algorithms were used, among which PBS proved to be the most effective as a W-MAPF solver within RHCR. Overall, it was observed that by using a small time window  $w$ , and consequently a small replanning rate  $h$ , RHCR can obtain faster solutions than alternative approaches. This, in turn, allows to solve queries containing more agents and potentially improve the system’s throughput (Li et al. 2021).

## 4 Leveraging Experience in Lifelong MAPF

In this section we present our algorithmic approach for leveraging experience in L-MAPF, which we call Experienced RHCR (exRHCR). First we discuss some properties of the original RHCR framework, which will be instrumental in the development of exRHCR.

Clearly, by improving the efficiency of the internal MAPF solver (e.g., PBS) we can improve the efficiency of RHCR overall, as it solves *multiple* W-MAPF queries. However, speeding up MAPF solvers without leveraging additional structure within the RHCR approach can be difficult. Next, we observe that there is an underlying structure emerging from the fact that the W-MAPF queries are *consecutive*, i.e., a solution plan to one query  $q_i$  defines the next query  $q_{i+1}$ . In

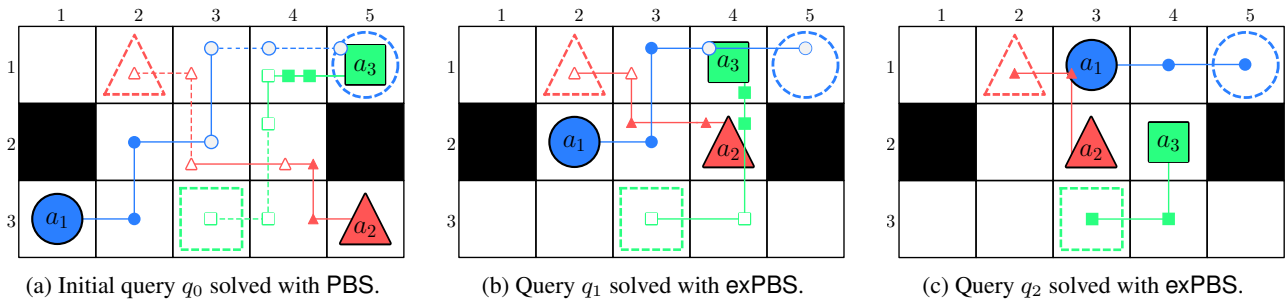


Figure 2: A toy example visualizing an execution of RHCR and  $\text{exRHCR}$  for three agents, whose current start and goal positions are illustrated as solid and dashed shapes, respectively, with parameters  $w = 4$ ,  $h = 2$ , and  $\delta = 2$  for  $\text{exRHCR}$ . Solid lines represent the path’s portion up to the  $w$  first timesteps (where collisions are avoided) and dashed line represent the rest (where collisions are ignored). Solid markers on the paths represent the first  $h$  steps and the rest are represented by hollow markers. Here, we chose  $\text{exPBS}$  and  $\text{PBS}$  to break ties in the low and high-level search in a specific manner, but other valid execution outcomes are possible. (a) The W-MAPF instance  $q_0$  with its solution obtained by  $\text{PBS}$ , where the experience gained is  $\prec_{\text{exp}} = \{a_2 \prec a_3, a_1 \prec a_2\}$ , which forces  $a_2$  to avoid  $a_1$ ’s path, and for  $a_3$  to avoid the paths of  $a_1$  and  $a_2$  (which causes  $a_3$  to wait). See full details of  $\text{PBS}$ ’s execution in the extended version (Madar, Solovey, and Salzman 2022). (b) and (c) show the next two queries  $q_1$  and  $q_2$  where they both use  $\text{exPBS}(\prec_{\text{exp}})$ , which allows to obtain a feasible solution already in the tree root, i.e., without further expansion of PT nodes. In contrast, to solve  $q_1$  RHCR would expand two PT nodes representing the same conflicts in  $q_0$  (to see this, note that for some shortest paths the agents do collide at the same positions as in  $q_0$ ). In both (b) and (c)  $\text{exRHCR}$  finds the continuation of the solution found in (a). In (b)  $a_2$  waits in place to avoid  $a_1$ ’s path, and  $a_3$  waits in its start position to avoid  $a_2$ . In (c), query  $q_2$ , the agents can follow their shortest paths without colliding, thus RHCR finds a solution in the first PT node without expanding further nodes.  $\text{exRHCR}$  also finds a solution from the first node although an experience is given.

particular, considering that a short replan rate  $h$  is typically used by RHCR, consecutive W-MAPF queries can be quite similar to one another in terms of initial and goal agent locations. This, in turn, can lead to *similar* queries, in terms of the conflicts that need to be avoided and the agents between which they arise. Despite this, current RHCR implementations solve every subsequent W-MAPF query *from scratch*.

In contrast to RHCR, our approach allows to transfer *experience* gained from solving one W-MAPF query to the next one in order to improve planning times. In particular,  $\text{exRHCR}$  solves W-MAPF queries by interleaving between calls to the vanilla  $\text{PBS}$  and an extension of  $\text{PBS}$  which we call Experienced  $\text{PBS}$  ( $\text{exPBS}$ ). Unlike the standard  $\text{PBS}$ , which begins the high-level search on a given W-MAPF instance from an empty priority set  $\emptyset$  at the root of the PT,  $\text{exPBS}$  is initialized with a specific priority set  $\prec_{\text{exp}}$  for the root of the PT. This priority  $\prec_{\text{exp}}$  encodes the experience gained from the solution of a previous W-MAPF query obtained via  $\text{PBS}$ . See illustration of  $\text{exRHCR}$  and  $\text{exPBS}$  in Figure 1. Additionally, in Figure 2 we illustrate two consecutive W-MAPF sub-queries ( $q_0$  and  $q_1$ ) within an L-MAPF problem, that have the same conflicts in terms of types of agents and positions in case that  $\text{PBS}$  is used for solving both queries. However, the conflicts in  $q_1$  can be avoided by reusing the experience obtained in  $q_0$  to solve  $q_1$  using  $\text{exPBS}$ .

Encoding experience as priorities has the desirable property of being *generalizable* in the sense that a priority used for solving one query can be applied to a similar query without necessarily over-constraining the solution. This is because priorities provide high-level specification that are

likely to apply in a variety of queries, rather than hard constraints which need to be followed exactly (e.g., the constraints used in CBS that specify for each agent exact locations that should be avoided in specific points in time). Experience encoded in the form of a priority set can thus be thought as a way to warm-start consecutive MAPF queries.

In some cases previous experience might not be useful for the given query, in which case we want to ensure the *robustness* of our MAPF solver to uninformative experience. In  $\text{exPBS}$  this is achieved by restarting the search with an empty priority set in case that experience causes the search to diverge. In the remainder of this section we detail  $\text{exRHCR}$  and the internal MAPF solver  $\text{exPBS}$ .

#### 4.1 Experienced RHCR ( $\text{exRHCR}$ )

The  $\text{exRHCR}$  algorithm accepts as inputs a graph  $G$  representing the environment, and a task assigner  $\mathcal{T}$ , which implicitly maintains the agent’s current start and goal positions as well as the remaining tasks. Using the task assigner as a black box helps keeping our description below generic. Similarly to RHCR,  $\text{exRHCR}$  has the parameters of time window  $w$ , and replanning rate  $h$ . Two additional parameters that we introduce for  $\text{exRHCR}$  are the *experience lookahead*  $\delta$ , which determines the number of  $\text{exPBS}$  calls after every  $\text{PBS}$  call (see details below), and the  $\text{exPBS}$  high-level PT *width limit*  $\ell > 1$ . The latter parameter helps to identify situations where previous experience turns out to be uninformative for the current query, and to restart the high-level search with an empty priority set (see Section 4.2).

The motivation behind using the lookahead parameter  $\delta$  is keeping the experience up-to-date, and employing it only

---

**Algorithm 1: exRHCR**

---

**Inputs:** L-MAPF query, graph  $G$ , task assigner  $\mathcal{T}$ **Parameters:** Window size  $w$ , replanning rate  $h$ , experience lookahead  $\delta$ , width limit  $\ell$ **Output:** Paths for all agents

```

1:  $q_0 \leftarrow \mathcal{T}.\text{INITIALQUERY}()$ ;  $i = 0$ ;
2: while not  $\mathcal{T}.\text{EMPTY}()$  do
3:    $(\mathcal{P}_{q_i}, \prec_{\text{exp}}) \leftarrow \text{PBS}(q_i, w)$ ;
4:    $q_{i+1} \leftarrow \mathcal{T}.\text{NEXTQUERY}(q_i, \mathcal{P}_{q_i}, h)$ ;  $i++$ ;
5:   repeat  $\delta$  times // run exPBS with PBS as fallback
6:     if  $\mathcal{T}.\text{EMPTY}()$  return // tasks finished
7:      $\mathcal{P}_{q_i} \leftarrow \text{exPBS}(q_i, \prec_{\text{exp}}, w, \ell)$ ;
8:      $q_{i+1} \leftarrow \mathcal{T}.\text{NEXTQUERY}(q, \mathcal{P}_{q_i}, h)$ ;  $i++$ ;
9: end while

```

---

when it is likely to be relevant to the current query, which should not differ vastly from the query from which the experience was extracted. We also ensure the experience does not become overconstrained by generating a brand new experience every few runs (rather than, e.g., constantly passing the priority set obtained from solving the current query as experience to the next query, and so on, where priorities are accumulated from each run).

We suggest to set  $\delta$  around the value  $\lfloor \frac{w}{h} \rfloor - 1$ , which means the experience is used in all the queries of the planning horizon it was created in to maximize the experience utilization. See more details in Section 5.2.

exRHCR is detailed in Algorithm 1. exRHCR first uses the task assigner to obtain the initial W-MAPF instance  $q_0$  and initializes to zero the counter  $i$  which represent the current W-MAPF query index [Line 1]. It then iteratively solves  $q_i$  and all subsequent W-MAPF instances until all tasks are completed [Line 2]. It runs PBS with the time window  $w$  on query  $q_i$  to obtain a plan  $\mathcal{P}_{q_i}$ , which specifies agent paths for the current W-MAPF instance [Line 3]. In contrast to RHCR, the priority set of the PT node for which  $\mathcal{P}_{q_i}$  was obtained is stored in memory. We term this the *seed priority set* and denote it by  $\prec_{\text{exp}}$ . Subsequently, and similar to RHCR, a new W-MAPF query  $q_{i+1}$  is generated by the task assigner, and the counter  $i$  is updated [Line 4]. This is done by executing the plan  $\mathcal{P}_{q_i}$  for  $h$  steps to update the agents’ start locations, and potentially updating the agents’ goals and tasks. At this point, exRHCR differs from RHCR: for the next  $\delta$  W-MAPF instances [Line 5], unless all tasks are finished [Line 6], the algorithm invokes exPBS with the seed priority set  $\prec_{\text{exp}}$  (using the same window size  $w$  and with a parameter  $\ell$ , which will be explained shortly) [Line 7]. Given the new solution plan  $\mathcal{P}_{q_i}$  and replanning rate  $h$ , the query and the counter are updated as before [Line 8] and this inner loop is repeated.

## 4.2 Experienced PBS (exPBS)

exPBS utilizes PBS’s option of starting the high-level search from a given seed priority set which we denote as an experience rather than from the typically-used empty priority set. An additional difference is that exPBS avoids overexploring the PT by limiting the width of the explored PT

rooted in  $\prec_{\text{exp}}$ , in case that  $\prec_{\text{exp}}$  does not lead to a solution fast enough (or does not find a solution at all), and restarts the search with a priority  $\emptyset$  by calling the “vanilla” PBS. We call the usage of PBS after exPBS terminates with no solution the “fallback”. To limit the tree width, exPBS explores the high-level PT using a *Width-Limited Depth-First Search* (WL-DFS), which we describe below, rather than the standard DFS used by PBS. Thus, exPBS accepts two additional parameters when compared to PBS: the seed priority set  $\prec_{\text{exp}}$  and the width-limit parameter  $\ell$ .

We provide additional details on WL-DFS. Given a tree graph, let its *width* denote the maximal number of nodes across all levels (where two nodes are on the same tree level if their distance from the root, or depth, is the same). The parameter  $\ell$  specifies the maximal width allowed when exploring a PT using WL-DFS. To keep track of the current width of the PT we maintain for each level a counter representing the number of nodes in the level, and increment it whenever new nodes are added. When the width of the PT exceeds  $\ell$  (or when no solution exists), WL-DFS aborts the search of the PT rooted in  $\prec_{\text{exp}}$ , and invokes the vanilla PBS solver without width limitation.

We chose to limit the search efforts by limiting the PT width as we found that it is a good indicator whether the experience is over-constrained or not (over-constrained experience tend to force the search to expand entire, wide subtrees). We found (see Section 5) that alternative measures such as the number of expanded nodes require per-instance tuning as exPBS uses a Depth First Search (DFS) on the PT and thus we need to account for the scenario and the number of agents. In contrast, we empirically found that the width is a robust parameter that does not require tuning. Indeed, in all our experiments (Section 5) we used the same width value. Finally, in Section 5.3 we study the effect of WL-DFS with different values of  $\ell$  on the overall performance and show the robustness of the method to the specific choice of  $\ell$ .

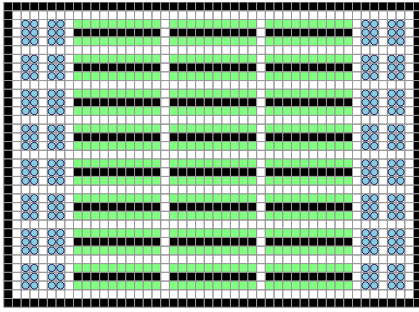
## 4.3 Alternative Instantiation with Total Priority

The key ingredients of our algorithmic framework are (i) the notion of experience derived from a non-informed planner (partial priority and PBS in the method described) and (ii) how the experience is used in the W-MAPF planner (exPBS in the method described). Here, we suggest an alternative instantiation and discuss its merits.

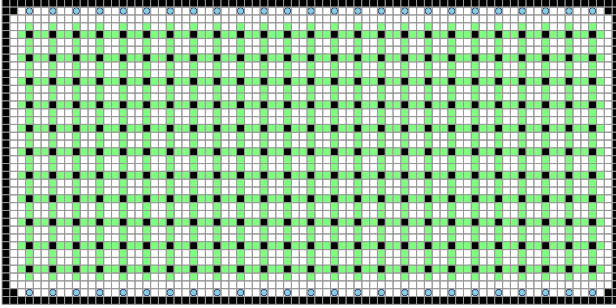
Recall that after running PBS, we obtain a partial priority  $\prec_{\text{exp}}$ , which is used within exPBS. Here we introduce a lightweight alternative which computes a *total priority*<sup>1</sup>  $\prec_{\text{exp-tot}}$  that is consistent with  $\prec_{\text{exp}}$  (namely, if  $a_i \prec a_j$  in  $\prec_{\text{exp}}$  then  $a_i \prec a_j$  in  $\prec_{\text{exp-tot}}$ ). Note that such a consistent total priority always exists and is easy to compute. Now, we run exPBS with the seed  $\prec_{\text{exp-tot}}$ , which boils down to running a prioritized planner (Silver 2005) and running PBS in case of failure.

Using a total priority is less generalizable than a partial priority and thus the planner is more likely to fall back to PBS. On the other hand, running a prioritized planner is ex-

<sup>1</sup>A total priority specifies priorities between all the agents of the form  $a'_1 \prec a'_2 \prec \dots \prec a'_n$ .



(a) WAREHOUSE environment.



(b) SORTING environment.

Figure 3: Benchmark environments for L-MAPF problems. (a) A  $33 \times 46$  WAREHOUSE domain with  $\sim 16\%$  obstacles (black) representing 240 pods (Liu et al. 2019). Working stations are drawn in blue and task locations around the inventory pods are drawn in green. (b) A  $37 \times 77$  SORTING center with  $\sim 10\%$  obstacles representing chutes (Li et al. 2021). Working stations are marked in blue, and task locations, which represent drop-off locations around chutes, are shown in green. Note that we used the undirected versions of these benchmarks.

tremely fast and when the problem is “easy” it may often succeed even with this more constrained notion of experience. As we will see in our experiments, this approach is advantageous in easier settings due to its simplicity.

## 5 Experimental Results

We provide an empirical evaluation of our exRHCR approach and compare it to RHCR. We implemented the algorithms in C++ and tested them on an Ubuntu machine with 4 GB RAM and a 2.7 GHz Intel i7 CPU. We used benchmarks simulating a warehouse (Figure 3a) and a sorting center (Figure 3b)<sup>2</sup>.

For both environments we randomly initialize the start location for each agent from all possible locations, and the task assigner samples a goal location randomly from the blue and green locations depicted in Figure 3. Each time an agent reaches a goal location of a specific color the task assigner specifies a new goal location uniformly at random from the other color that is not currently assigned to another agent.

<sup>2</sup>Code and benchmarks are available at <https://github.com/NitzanMadar/exPBS-exRHCR>.

In Section 5.1 we compare our exRHCR framework with RHCR. Next, in Sections 5.2 and 5.3 we focus on exRHCR with partial priorities, which has more promise in tackling hard L-MAPF instances, and discuss the effect that the lookahead  $\delta$  and the width limit  $\ell$  parameters have on exRHCR and exPBS, respectively.

### 5.1 L-MAPF Experiments

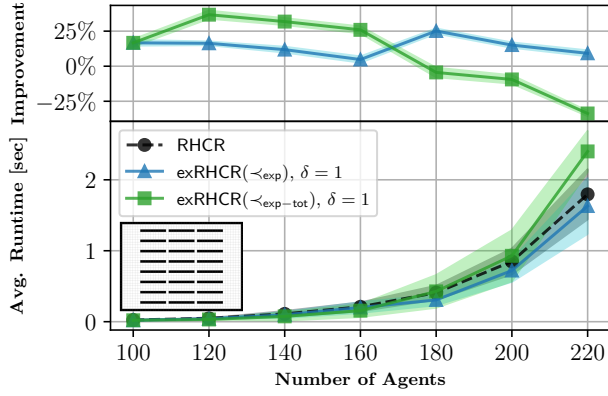
For each of the two environments (i.e., WAREHOUSE and SORTING) we generated multiple L-MAPF instances and tested them on a varying numbers of agents  $k$ . In particular, we set  $k \in \{100, 120, \dots, 220\}$  and  $k \in \{100, 150, \dots, 450\}$  for WAREHOUSE and SORTING, respectively. For each combination of an environment and  $k$ , we randomly generated 50 L-MAPF instances.

We consider three planners, RHCR, and our two planners: (i) using a partial priority with exPBS, denoted by  $\text{exRHCR}(\leftarrow_{\text{exp}})$  and (ii) the alternative instantiation described in Section 4.3 using a total priority with prioritized planning, denoted by  $\text{exRHCR}(\leftarrow_{\text{exp-tot}})$ . We set the replanning rate  $h$  and the time window  $w$  to be  $h = 5$  and  $w = 10$  for both WAREHOUSE and SORTING environments. For exRHCR we used an experience lookahead of  $\delta = \lfloor \frac{w}{h} \rfloor - 1 = 1$ , and width limit  $\ell = 10$  when using partial priority experience.

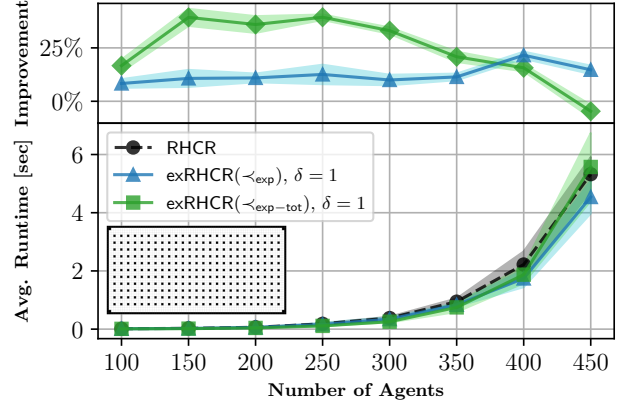
In Figure 4, we report the solvers’ average MAPF query runtime for a total L-MAPF execution of 250 timesteps. In particular, each of the 50 L-MAPF instances induces 50 W-MAPF instances (a total of 250 timesteps divided by the replan rate of  $h = 5$ ). We report the average runtime and standard deviation across all W-MAPF instances solved, as well as success rates. We limit the runtime for a W-MAPF query to 30 seconds, after which we declare failure. In such a case the runtime of a failed W-MAPF instance is 30 seconds.

exRHCR with  $\delta = 1$  improved the average runtime (over RHCR) up to 37% in WAREHOUSE and 39% in SORTING. When considering  $\text{exRHCR}(\leftarrow_{\text{exp-tot}})$ , we can see (as in Ma et al. (2019a)) that the approach is highly effective when the number of agents is small. However, given a large number of agents, this method has a high failure rate (fallback is used roughly 86% of the time for the largest number of agents). For a large number of agents  $\text{exRHCR}(\leftarrow_{\text{exp}})$  achieves the best performance compared to RHCR and  $\text{exRHCR}(\leftarrow_{\text{exp-tot}})$ .

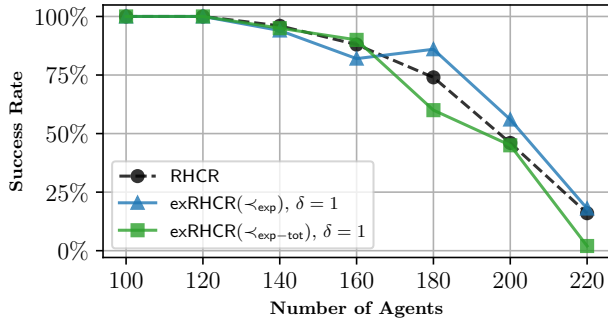
We provide a few more observations. (i) In all the L-MAPF experiments we performed, the difference in the average solution cost between PBS and exPBS was negligible (roughly  $\pm 1\%$ ), which suggests that reusing experience does not hinder solution quality. (ii) As a result, the throughput difference per instance and number of agents is also negligible (roughly  $\pm 2\%$ ). (iii) The improved runtime we described above (Figure 4) suggests that for a given time budget and an average W-MAPF query, exRHCR can accommodate more agents than RHCR. This suggests that the improved efficiency of our approach can improve the overall throughput in automated logistic domains.



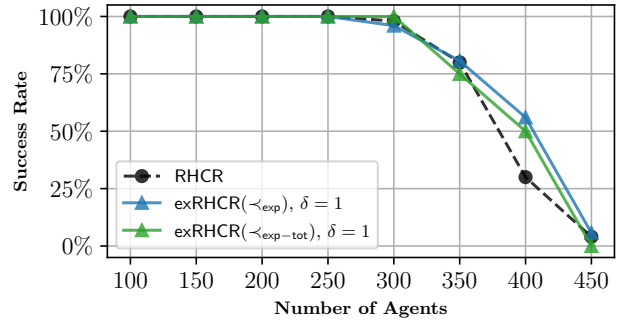
(a) Runtime for WAREHOUSE.



(b) Runtime for SORTING.



(c) Success Rate for WAREHOUSE.



(d) Success Rate for SORTING.

Figure 4: L-MAPF experiments results. (a) and (b) depict the average runtimes and standard deviation of a W-MAPF query for the different L-MAPF solvers. For each environment, the top plot shows the relative improvement of each method compared to RHCR. (c) and (d) depict the success rate obtained by all planners.

## 5.2 Effect of Lookahead Parameter $\delta$ on RHCR

We consider the effect that the lookahead parameter  $\delta$  has on the performance of  $\text{exRHCR}(\prec_{\text{exp}})$ . Specifically, we fixed  $w = 10$  and  $w = 20$  for the SORTING and WAREHOUSE environments, respectively and evaluate the average PT depth as a function of  $\delta$  for different values of  $h$  (Figure 5)

First, we observe that our method improves over RHCR for every value of  $\delta$  considered. Additionally, our suggested value of  $\delta = \lfloor \frac{w}{h} \rfloor - 1$  yields close-to-optimal values and is a good rule of thumb in the absence of any other information. Not surprisingly, selecting  $\delta > \lfloor \frac{w}{h} \rfloor - 1$  reduces performance in all of the cases tested, as it uses experience outside the planning horizon for which it was created. Finally, selecting  $\delta < \lfloor \frac{w}{h} \rfloor - 1$  may be beneficial when the replanning rate is smaller and the time window is large as the relevance of the experience seems to diminish as  $\delta$  increases. For example, if  $w = 20$ ,  $h = 2$  and  $\delta = \lfloor \frac{w}{h} \rfloor - 1 = 9$ , in the last query the experience used, it's relevant for only 2 out of the 20 steps.

We also compare our framework to two alternatives that do not use the lookahead parameter  $\delta$ , as mentioned in Section 4.1: (i) running PBS once and using the same experience for all subsequent queries within  $\text{exPBS}$  and (ii) using the priority set from a solution for a given query

as an experience in the next query. We used the WAREHOUSE domain with the same parameter as in Section 5.1 and  $k \in \{180, 220\}$ . We found that the average runtime decreases by 8% and 21%, respectively, using alternative (i) compared to  $\text{exRHCR}$ , and by 9% and 15%, respectively using alternative (ii). Additionally, the success rate dropped by 15.2% and 11.1% using (i), and by 10.6% and 11.1% using (ii).

## 5.3 Effect of Width Limit $\ell$ on WL-DFS

We assess how the width limit parameter  $\ell$  affects the performance of  $\text{exPBS}$  and consequently  $\text{exRHCR}(\prec_{\text{exp}})$ . We use the WAREHOUSE environment with  $w = 20$ ,  $h = 5$  and  $\delta = 3$  and fixed the number of agents to be  $k = 150$ . We report average metrics (detailed shortly) for MAPF queries for different values of  $\ell \in \{2, 5, 10, 20, 50, 1,000\}$  in Figure 6. Note that for  $\ell = 2$   $\text{exPBS}$  behaves very similarly to PBS as it reverts to PBS after the first backtrack. The setting of  $\ell = 1,000$  simulates a version wherein no fallback is taken and WL-DFS is equivalent to DFS.

For every value of  $\ell$  we logged the following attributes and report averaged values across all MAPF queries (starting with the top attribute and going in a clockwise manner with respect to Figure 6): (i) average runtime; (ii) average

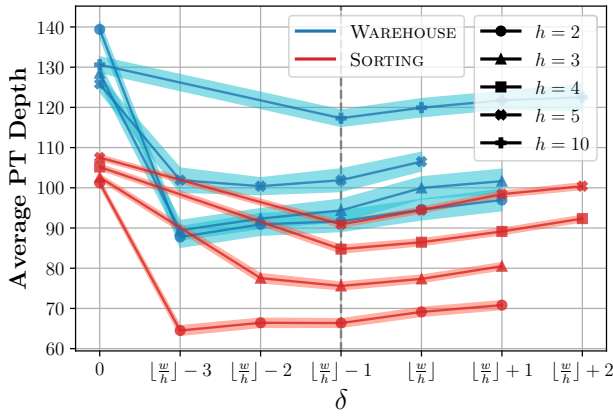


Figure 5: Average PT depth of  $\text{exRHCR}(\prec_{\text{exp}})$  for the WAREHOUSE and SORTING domains using 150 and 300 agents, respectively for different values of  $\delta$ . Here,  $\delta = 0$  corresponds to running RHCR without using experience.

PT width; (iii) average number of  $A^*$  node expansions in the low-level  $\text{exPBS}$  and  $\text{PBS}$  search; (iv) average number of PT node expansions in the high-level  $\text{exPBS}$  and  $\text{PBS}$  search; (v) average depth of PT (in case of a fallback, it consists of the sum between the depth of the  $\text{exPBS}$  and  $\text{PBS}$  trees). The reported values in the plot are normalized by the maximal value per attribute.

Having a small opportunity to reuse experience ( $\ell = 2$ ), or allowing the search to explore the PT rooted in a seed experience indefinitely ( $\ell = 1,000$ ), significantly reduces the performance of  $\text{exRHCR}$  in comparison with the other values of  $\ell$  for all attributes indicating the benefits of WL-DFS.

However, the reason why each one of these extreme values is not useful is different: When  $\ell = 2$ ,  $\text{exPBS}$  starts to run and in many cases (roughly, 65% of the times) falls back to  $\text{PBS}$ . This implies a small overhead caused by unnecessarily running  $\text{exPBS}$  many times. In contrast, when  $\ell = 1,000$ ,  $\text{exPBS}$  rarely falls back to  $\text{PBS}$  which may imply PT overexploration the PT when the  $\text{PBS}$  fallback could have found a solution quicker. This implies a large overhead of running  $\text{exPBS}$  incurred a small number of times.

Finally, we consider an alternative implementation that uses a limit on the number of node expansions (rather than width). We used the WAREHOUSE domain with the same parameter as used in Section 5.1 and  $k \in \{180, 220\}$ . To estimate the number of nodes to be used as a parameter, we empirically found from previous experiments the average number of nodes expanded when  $\text{exPBS}$  succeeds in a given scenario. In our case, this was 99 and 209 for 180 and 220 agents, respectively. We evaluated a node limit of 80, 100 and 150 for 180 agents and a node limit of 150, 200 and 250 for 220 agents. We found that the success rate dropped by 0%–2.6% for 180 agents and by 5.6%–56% for 220 agents. Moreover, the runtime increased by roughly 1%–18% for 180 agents and by 0.7%–10% for 220 agents. We note the parameter selection is more sensitive and domain-dependent compared to width limit, as mentioned in Section 4.2.

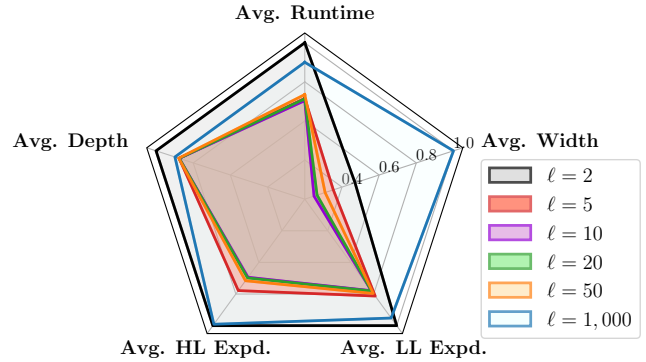


Figure 6: Radar plot depicting the effect of  $\ell$  on various attributes of  $\text{exRHCR}(\prec_{\text{exp}})$ . All values are normalized according to the maximal value attained on the axis they belong to.

## 6 Conclusions and Future Work

In this paper we described  $\text{exRHCR}$ , a new approach for leveraging experience in L-MAPF instances, which allows to reduce computational effort in constituent MAPF queries by reusing priority sets from previous queries within W-MAPF solvers. We demonstrated empirically that our approach can substantially improve runtime and has the potential to increase system throughput by incorporating additional agents.

Our work introduces various directions for future research. In the short run, we plan to explore approaches for systematic selection of the parameter  $\ell$ , and additional heuristics for terminating and restarting the  $\text{exPBS}$  search.

In the long run, it would be beneficial to consider advanced experience-retrieval strategies. For example, can we design an “experience database” that contains queries and their solution priorities? Here one can potentially retrieve the nearest-neighbor query under some metric to be used in  $\text{exPBS}$ . Learning-based approaches can come in handy as well by, e.g., identifying similar queries and generating experience artificially.

Finally, the algorithms we introduced can be used to create a *hierarchical approach* for solving W-MAPF problems in L-MAPF: after running  $\text{PBS}$  we obtain a partial priority  $\prec_{\text{exp}}$  and then compute a consistent total priority  $\prec_{\text{exp-tot}}$ . We then start by running a prioritized planner using  $\prec_{\text{exp-tot}}$ . In the case of failure, we fall back to  $\text{exPBS}$  which uses  $\prec_{\text{exp}}$  to warm-start the search and if this planner fails, we fall back to  $\text{PBS}$ .

## Acknowledgments

This research was supported in part by the Israeli Ministry of Science & Technology grants no. 3-16079 and 3-17385, by the United States-Israel Binational Science Foundation (BSF) grants no. 2019703 and 2021643, and the Ravitz Fellowship. The authors also thank the anonymous reviewers for their insightful comments and suggestions.



## References

- Barer, M.; Sharon, G.; Stern, R.; and Felner, A. 2014. Sub-optimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *SoCS*, 19–27.
- Choudhury, S.; Solovey, K.; Kochenderfer, M. J.; and Pavone, M. 2021. Efficient Large-Scale Multi-Drone Delivery using Transit Networks. *J. Artif. Intell. Res.*, 70: 757–788.
- Cohen, L.; and Koenig, S. 2016. Bounded suboptimal multi-agent path finding using highways. In *IJCAI*, 3978–3979.
- Coleman, D.; Şucan, I. A.; Moll, M.; Okada, K.; and Correll, N. 2015. Experience-based planning with sparse roadmap spanners. In *ICRA*, 900–905.
- Dayan, D.; Solovey, K.; Pavone, M.; and Halperin, D. 2021. Near-Optimal Multi-Robot Motion Planning with Finite Sampling. In *ICRA*, 9190–9196.
- Demaine, E. D.; Fekete, S. P.; Keldenich, P.; Meijer, H.; and Scheffer, C. 2019. Coordinated Motion Planning: Reconfiguring a Swarm of Labeled Robots with Bounded Stretch. *SIAM J. on Comput.*, 48(6): 1727–1762.
- Gordon, O.; Filmus, Y.; and Salzman, O. 2021. Revisiting the Complexity Analysis of Conflict-Based Search: New Computational Techniques and Improved Bounds. In *SoCS*, 64–72.
- Hönig, W.; Kiesel, S.; Tinka, A.; Durham, J. W.; and Ayanian, N. 2019. Persistent and robust execution of MAPF schedules in warehouses. *RA-L*, 4(2): 1125–1131.
- Huang, T.; Koenig, S.; and Dilkina, B. 2021. Learning to Resolve Conflicts for Multi-Agent Path Finding with Conflict-Based Search. In *AAAI*, 11246–11253. AAAI Press.
- Kaduri, O.; Boyarski, E.; and Stern, R. 2020. Algorithm Selection for Optimal Multi-Agent Pathfinding. In *ICAPS*, 161–165.
- Li, J.; Gange, G.; Harabor, D.; Stuckey, P. J.; Ma, H.; and Koenig, S. 2020. New techniques for pairwise symmetry breaking in multi-agent path finding. In *ICAPS*, volume 30, 193–201.
- Li, J.; Harabor, D.; Stuckey, P. J.; Felner, A.; Ma, H.; and Koenig, S. 2019a. Disjoint splitting for multi-agent path finding with conflict-based search. In *ICAPS*, volume 29, 279–283.
- Li, J.; Harabor, D.; Stuckey, P. J.; Ma, H.; and Koenig, S. 2019b. Symmetry-breaking constraints for grid-based multi-agent path finding. In *AAAI*, volume 33, 6087–6095.
- Li, J.; Ruml, W.; and Koenig, S. 2021. EECBS: A bounded-suboptimal search for multi-agent path finding. In *AAAI*.
- Li, J.; Tinka, A.; Kiesel, S.; Durham, J. W.; Kumar, T. K. S.; and Koenig, S. 2021. Lifelong Multi-Agent Path Finding in Large-Scale Warehouses. In *AAAI*, 11272–11281.
- Liu, M.; Ma, H.; Li, J.; and Koenig, S. 2019. Task and path planning for multi-agent pickup and delivery. In *AAMAS*, 1152–1160.
- Ma, H.; Harabor, D.; Stuckey, P. J.; Li, J.; and Koenig, S. 2019a. Searching with consistent prioritization for multi-agent path finding. In *AAAI*, volume 33, 7643–7650.
- Ma, H.; Hönig, W.; Kumar, T. K. S.; Ayanian, N.; and Koenig, S. 2019b. Lifelong Path Planning with Kinematic Constraints for Multi-Agent Pickup and Delivery. In *AAAI*, 7651–7658.
- Ma, H.; Li, J.; Kumar, T. K. S.; and Koenig, S. 2017. Lifelong multi-agent path finding for online pickup and delivery tasks. In *AAMAS*, 837–845.
- Ma, H.; Tovey, C. A.; Sharon, G.; Kumar, T. K. S.; and Koenig, S. 2016. Multi-Agent Path Finding with Payload Transfers and the Package-Exchange Robot-Routing Problem. In *AAAI*, 3166–3173.
- Madar, N.; Solovey, K.; and Salzman, O. 2022. Leveraging Experience in Lifelong Multi-Agent Pathfinding. *CoRR*, abs/2202.04382.
- Salzman, O.; and Stern, R. 2020. Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems. In *AAMAS*, 1711–1715.
- Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. R. 2015. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219: 40–66.
- Sharon, G.; Stern, R.; Goldenberg, M.; and Felner, A. 2013. The increasing cost tree search for optimal multi-agent pathfinding. *Artificial Intelligence*, 195: 470–495.
- Silver, D. 2005. Cooperative pathfinding. *Artificial Intelligence and Interactive Digital Entertainment*, 1: 117–122.
- Stern, R.; Sturtevant, N. R.; Felner, A.; Koenig, S.; Ma, H.; Walker, T. T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T. S.; et al. 2019. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *SoCS*, 151–158.
- Surynek, P.; Felner, A.; Stern, R.; and Boyarski, E. 2016. Efficient SAT approach to multi-agent path finding under the sum of costs objective. In *ECAI*, 810–818. IOS Press.
- Wagner, G.; and Choset, H. 2015. Subdimensional expansion for multirobot path planning. *Artificial Intelligence*, 219: 1–24.
- Yu, J. 2016. Intractability of Optimal Multirobot Path Planning on Planar Graphs. *RA-L*, 1(1): 33–40.