

Conflict-Based Search for Multi-Agent Path Finding with Elevators

Haitong He^{1,2}, Xuemian Wu¹, Shizhe Zhao¹, Zhongqiang Ren^{1†}

Abstract—This paper investigates Multi-Agent Path Finding with Elevators (MAPF-E), which seeks conflict-free paths for multiple agents whose start and goal locations may be on different floors, and the agents can use elevators to travel between floors. The existence of elevators complicates the interaction among the agents and introduces new planning challenges. On the one hand, elevators can cause many conflicts among agents due to their relatively long traversal time across floors, especially when many agents need to reach a different floor. On the other hand, the planner has to reason in a larger state space that includes elevator states in addition to agent locations. To address these challenges, this paper proposes CBS-E, an extension of Conflict-Based Search (CBS) to solve MAPF-E optimally. CBS-E introduces new concepts of elevator constraints to incorporate the state of the elevator into the conflict resolution process of CBS. We also extend Multi-Valued Decision Diagrams (MDDs) to CBS-E, referred to as MDD-E. MDD-E enables the intelligent selection of conflicts to resolve in each iteration, thereby improving the runtime efficiency. The results show that our CBS-E often doubles or triples the success rates of the baseline, and the proposed conflict reasoning can reduce the number of iterations of CBS-E by up to an order of magnitude.

I. INTRODUCTION

Multi-Agent Path Finding (MAPF) seeks conflict-free paths for multiple agents from their respective starts to goals within a shared environment, while minimizing the path cost. This paper considers a variant of MAPF where the environment consists of multiple floors connected by elevators, and each elevator can transport one agent at a time. The start and goal locations of the agents may lie on different floors, and the agents use elevators to move between floors. We refer to this variant as Multi-Agent Path Finding with Elevators (MAPF-E). When there is only one floor, MAPF-E becomes the regular MAPF. MAPF is NP-hard [1] and so is MAPF-E. MAPF-E naturally arises in logistics when the environment consists of multiple floors. Consider a 3D warehouse, where mobile robots must use elevators to reach goal locations for retrieving or placing items (Fig. 1).

Although a variety of planners and CBS enhancements for MAPF and its variants have been developed, ranging from optimal planners [3]–[5], bounded sub-optimal planners [6], [7] to unbounded sub-optimal planners [8], [9], we are not aware of any research focusing on MAPF-E. A recent work [10] considers MAPF in a 3D space in which agents can move vertically at arbitrary locations, which is different from the elevator-based setting considered in this work.

¹ Global College, Shanghai Jiao Tong University, China. This work was supported by the Natural Science Foundation of Shanghai under Grant 24ZR1435900.

² Harbin Institute of Technology, China.

† Correspondence: zhongqiang.ren@sjtu.edu.cn

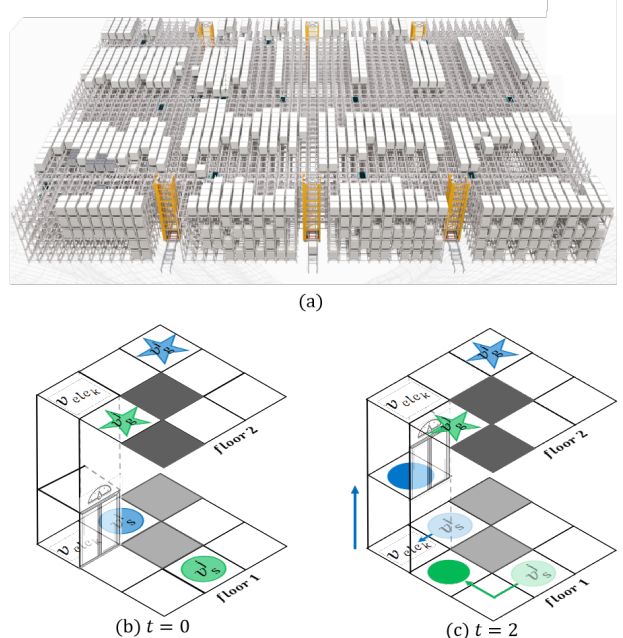


Fig. 1. (a) The motivation of MAPF-E. Robots move in a 3D warehouse, while inter-layer transitions are enabled by the elevators (the orange structures). The picture is from ZS Robotics [2]. (b) A toy example of MAPF-E, which seeks start-goal paths for the agents. All elevator edge costs are 2, and the black squares represent obstacles. (c) At time $t = 2$, agent i is using the elevator, and agent j needs to wait at a vertex outside the elevator.

Existing approaches for MAPF can be adapted to solve MAPF-E, by describing all floors as a graph with special edges representing the elevator between floors and running a MAPF planner on this graph. However, such a naive adaptation can lead to poor runtime efficiency. The existence of elevators complicates the interaction among the agents and introduces new challenges for planning. First, an elevator can couple multiple agents on different floors, even though these agents are spatially far away from each other, thereby complicating the collision resolution process. Second, it often takes much longer for an elevator to traverse between adjacent floors than for an agent to move from one location to another on the same floor, which can lead to frequent collisions and congestion among the agents, especially when many agents seek to take the same elevator. In this case, the planner has to consider the possibility of letting an agent take a long detour to use another elevator to optimize the arrival time. Consequently, the planner must reason over a larger state space that explicitly incorporates elevator states and positions, increasing the computational burden.

As a first attempt to study MAPF-E, this paper focuses on developing exact algorithms that can find optimal solutions

for MAPF-E. Based on Conflict-Based Search (CBS) [11], a popular approach for MAPF, we develop CBS-E and its variants to optimally solve MAPF-E. In particular, we find that, a naive adaptation of CBS for MAPF-E may need an exponentially growing number of iterations to resolve agent-agent conflicts as the traversal duration of the elevator per layer increases. To address this difficulty, our CBS-E can often resolve an elevator conflict in a single iteration, by introducing a new approach for elevator constraints generation based on the state of the elevator. We prove that, with these new constraints, our CBS-E preserves the solution optimality. Additionally, we also develop new conflict reasoning techniques by extending Multi-Valued Decision Diagram (MDD), an important data structure for conflict reasoning in CBS [12], to MDD with elevator (MDD-E), enabling CBS-E to intelligently select conflicts to resolve in each iteration and further enhance runtime efficiency.

We conduct an ablation study of the proposed conflict reasoning techniques in CBS-E, and compare our CBS-E against a naive adaptation of CBS to MAPF-E. Experimental results show that CBS-E often doubles or triples the success rate of the naive CBS adaptation in various maps, and the proposed conflict reasoning techniques can reduce the number of iterations of CBS-E by up to an order of magnitude.

II. PROBLEM FORMULATION

Let the index set $I = \{1, 2, \dots, N\}$ represent a set of N agents. All agents operate within a shared workspace, modeled as a finite graph $G = (V, E)$, which consists of m subgraphs $G = G_1 \cup G_2 \cup \dots \cup G_m$. Each subgraph $G_l, l \in L = \{1, 2, \dots, m\}$ represents the l -th floor (or layer). The subgraphs G_l may differ from one another. The vertex set V represents all possible locations of the agents. The edge set $E \subseteq V \times V$ represents the transitions between vertices in V . Each edge e between $u, v \in V$ has a positive integer cost value $c(e) = c(u, v) \in \mathbb{Z}^+$, indicating the traversal time of any agent along that edge.

The vertices are classified into two categories $V = V_{ele} \cup V_{reg}$, where V_{ele} denotes the set of vertices corresponding to elevators, and V_{reg} are all other regular vertices. The edges are also classified into two categories $E = E_{ele} \cup E_{reg}$ as follows. An edge $e = (u, v) \in E_{ele}$ is an elevator edge if both ends of e are elevator vertices $u, v \in V_{ele}$. Otherwise, an edge is a regular edge. Let the index set $I_K = \{1, 2, \dots, K\}$ denote a set of K elevators. Each elevator $k \in I_K$ corresponds to a set of vertices $V_{ele,k} \subseteq V_{ele}$ and a set of edges $E_{ele,k} \subseteq E_{ele}$. Each vertex $v_{k,l} \in V_{ele,k}$ corresponds to a layer G_l and any two vertices $v_{k,l_1}, v_{k,l_2} \in V_{ele,k}$ cannot share the same layer $l_1 \neq l_2$. Each edge $e \in E_{ele,k}$ connects two adjacent layers. Each regular edge $e \in E_{reg}$ has a unit-cost value $c(e) = 1$, while each elevator edge can have various known and fixed integer costs. An elevator can hold at most one agent at a time, and each agent enters an elevator at most once¹.

¹This implies every agent can take an elevator to directly reach its goal floor, without taking any elevator for a second time.

Let $v_s^i, v_g^i \in V$ denote the start and goal vertices of agent $i \in I$, and l_s^i, l_g^i denote the floors on which v_s^i, v_g^i are located, respectively. All agents share a global clock and start moving along their paths from v_s^i at $t = 0$. Let $\pi^i(v_s^i, v_g^i) = (v_s^i, v_1^i, v_2^i, \dots, v_g^i)$ denote a path for agent i between vertices v_s^i and v_g^i . The time steps $t(v_1^i), t(v_2^i)$ between two adjacent elevator vertices v_1^i and v_2^i in π^i are equal to the elevator edge cost, i.e., $|t(v_2^i) - t(v_1^i)| = c(e), e = (v_1^i, v_2^i) \in E_{ele}$. For all other adjacent vertices v_1^i, v_2^i in π^i , the time steps are one time unit, i.e., the cost of a regular edge $c(v_1^i, v_2^i)$. Let $g(\pi^i(v_s^i, v_g^i))$ denote the cost of a single path, which is defined as the sum of costs of the edges along the path. Let $\pi = (\pi^1, \pi^2, \dots, \pi^N)$ denote a joint path of all the agents, where π^i is the path of agent i . The total cost of the joint path is defined as the sum of the individual costs, i.e., $g(\pi) = \sum_i g(\pi^i)$.

A conflict between two agents $i, j \in I$ occurs if any of the following three cases arises. The first two cases are the same as in MAPF and the third case is new in MAPF-E: (1) A vertex conflict (i, j, v, t) occurs when two agents $i, j \in I$ occupy the same vertex v at the same time step t ; (2) An edge conflict (i, j, e, t) occurs when two agents $i, j \in I$ go through the same edge e from opposite directions between times t and $t + 1$; (3) An *elevator conflict*, as defined below.

Definition 1 (Elevator Conflict). *For an agent i , let t_s^i, t_g^i denote the time steps when i enters and leaves an elevator $k \in I_K$ to travel from v_{k,l_1} to v_{k,l_2} (i.e., v_{k,l_1} and $v_{k,l_2} \in V_{ele,k}$), respectively. Consider another agent j who travels from $v_{k,l_3} \in V_{ele,k}$ after i using the same elevator. We identify three cases for j concerning elevator k :*

- *The elevator is occupied during the interval $[t_s^i, t_g^i]$ because it is transporting agent i ;*
- *The elevator is resetting during the interval $[t_g^i, t_g^i + \Delta]$, where $\Delta = c((v_{k,l_2}, v_{k,l_3}))$, since k needs to move from v_{k,l_2} to v_{k,l_3} before picking up agent j ;*
- *The elevator is available after $t_g^i + \Delta$ because it has moved to the required floor after transporting i ;*

An elevator conflict occurs when agent j arrives at any vertex in the set $V_{ele,k}$ related to the elevator $k \in I_K$ while that elevator is resetting or occupied (Fig. 2). In this work, we assume all elevators are available to all agents at $t = 0$.

The goal of MAPF-E is to find a conflict-free joint path π with the minimal cost.

III. PRELIMINARIES

A. Conflict-Based Search

Conflict-Based Search (CBS) [11] is a two-level search algorithm that finds an optimal joint path for all agents for MAPF. At the high-level, CBS searches by building a binary tree, called the *constraint tree* (CT). Each node in the CT represents a set of constraints Ω , a joint path π satisfying the constraints in Ω , and its g -value, the cost of π .

CBS begins by constructing a root node $R_{\text{root}} = (\pi_0, g_0, \Omega_0)$, where $\pi_0 = (\pi^1, \pi^2, \dots, \pi^N)$ is the shortest path for each agent ignoring any agent-agent conflicts and

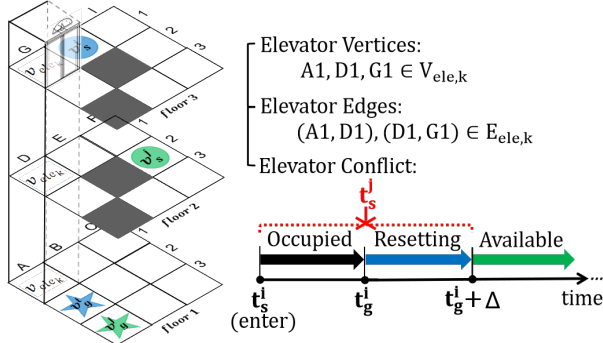


Fig. 2. Agent i (blue) and j (green) need to use the same elevator k to reach their goals. An elevator conflict arises when agent j attempts to enter the elevator with the *occupied* or *resetting* state, i.e., when $t_s^j \in [t_s^i, t_g^i + \Delta]$.

$\Omega_0 = \emptyset$. At any time during the search, all leaf nodes of the CT form an open list, denoted as OPEN, which is a priority queue sorted in ascending order of g -value. CBS repeatedly selects the node with the smallest g -value from OPEN for expansion. CBS expands a node $P = (\pi, g, \Omega)$ by checking for a conflict among the agent paths in π . If no conflict is detected, the current joint path π is returned as an optimal solution. Otherwise, for a detected conflict $C = (i, j, v, t)$, two new constraints, (i, v, t) and (j, v, t) , are generated. Each constraint leads to a new constraint set, $\Omega \cup \{(i, v, t)\}$ and $\Omega \cup \{(j, v, t)\}$, respectively (the same applies to edge conflicts (i, j, e, t)). For each new constraint set Ω' , a low-level search is performed to compute a new optimal path for the affected agent i that satisfies all the constraints in Ω' . The new joint path π' is formed by replacing the old path of agent i while keeping the paths of the other agents unchanged. A corresponding node $P' = (\pi', g', \Omega')$ is created and added to OPEN for future expansion. The search continues until a node with conflict-free joint paths is found. CBS guarantees that the returned solution has the minimum cost.

B. MDD, Bypassing Conflicts and Conflict Selection

At the high-level of CBS, for a selected node $P = (\pi, g, \Omega)$ from OPEN, there may be multiple conflicts in π . Both the choice of which conflict to resolve and the method used to select an appropriate resolution can significantly affect the runtime efficiency of CBS [5]. To optimize this process, we employ two techniques, namely *Bypassing Conflict* and *Conflict Selection*, both of which are based on Multi-Valued Decision Diagrams (MDD) [13].

1) *Multi-Valued Decision Diagrams (MDD)*: The MDD_d^i for agent i in a CT node P is defined as a directed graph representing all optimal paths (each with cost d) that are subject to the set of constraints Ω of P . Each node in the MDD corresponds to a tuple (v, t) , where v is a vertex reached by agent i at time t on an optimal path, and it appears at time t in MDD_d^i . When an MDD contains only a single node (v, t) at time t , and a new constraint prohibits the agent from occupying v at that time, all optimal paths (with cost d) are eliminated. The agent with this new constraint must take

a new path with path cost $d + 1$, thereby increasing the path cost. This property allows MDDs to classify a conflict by checking whether a constraint eliminates all paths in MDD_d^i . Specifically, this check is performed by constructing the joint MDD^{ij} for agents i and j . A node in the joint MDD^{ij} is represented as $((v_i, t), (v_j, t))$, indicating that there exists a pair of conflict-free paths for agents i and j , with the respective optimal paths passing through (v_i, t) and (v_j, t) . Fig. 3 shows an example.

2) *Bypassing Conflicts (BP)*: For a popped node $P = (\pi, g, \Omega)$ containing a conflict C in π , sometimes it is beneficial to directly modify the joint path π in P to avoid branching that would generate two new nodes and increase the size of the CT. This technique is known as bypass conflicts (BP). Specifically, a path π_{new}^i is a *valid bypass* to path π^i for agent i with respect to a conflict C and the CT node P , if the following conditions are satisfied: (i) $g(\pi_{new}^i) = g(\pi^i)$, (ii) both π_{new}^i and π^i satisfy the constraints Ω of P , (iii) π_{new}^i does not violate the conflict C . Since the joint MDD^{ij} contains all optimal conflict-free paths between agents i and j , π_{new}^i can be extracted from the joint MDD^{ij} if it exists. Then, P adopts π_{new}^i directly and is added back to OPEN.

3) *Conflict Selection*: All conflicts can be classified into three types [5]: (1) *Cardinal conflict*: A conflict $C = (i, j, v, t)$ is cardinal for a CT node P if neither agent i nor agent j can find a valid bypass in the joint MDD^{ij} . This implies that each of the two constraints derived from C increases the cost g of the CT node P . (2) *Semi-cardinal conflict*: A conflict C is semi-cardinal if only one of the agents, either i or j , can find a valid bypass in the joint MDD^{ij} . This implies that only one of the constraints derived from C increases the cost g of the CT node P . (3) *Non-cardinal conflict*: A conflict C is non-cardinal if both agents i and j can find a valid bypass in the joint MDD^{ij} . This implies that neither of the two constraints derived from C increases the cost g of the CT node P . Since the overall process of CBS aims to increase g toward its final optimal value, we prefer to resolve conflicts that cause g to increase more rapidly. Therefore, CBS often selects and resolves conflicts in the following order [5]: cardinal conflicts, semi-cardinal conflicts, and then non-cardinal conflicts.

IV. METHOD

CBS-E (Alg. 1) follows a process similar to that of CBS, with the primary differences occurring on Line 7 and 12, which represent our contributions. On Line 7, we introduce MDD-E to classify and select conflicts. On Line 12, we introduce a new type of elevator constraint to efficiently handle elevator conflicts.

A. Elevator Constraints (EC)

The conventional conflict resolution technique in CBS can also be applied to resolve elevator conflicts. For two agents, i and j , that are in conflict, each iteration of CBS forbids either agent i or j from occupying a vertex at a specific time step t . However, this approach is inefficient for resolving an




Fig. 3. Each node (v, t) in the MDD means the agent can optimally reach the goal via vertex v at time t . For instance, in MDD_2^i , at $t = 1$, the agent can reach either $B2$ or $C1$, resulting in two nodes at this layer. In the joint MDD, each node must correspond to a conflict-free partial path. For instance, in MDD_3^{ij} at $t = 1$, a conflict would occur if agents i and j both visited $B2$ simultaneously. So the joint MDD at $t = 1$ contains only the node $((C3, 1), (B2, 1))$.

Algorithm 1 CBS-E

INPUT: $G = (V, E, C)$
OUTPUT: a conflict-free joint path π in G

- 1: $\Omega_0 \leftarrow \emptyset, \pi_0 \leftarrow \text{LOWLEVEL}(i, \Omega), \forall i \in I, g_0 \leftarrow g(\pi_0)$
- 2: Add $R_{root} = (\pi_0, g_0, \Omega_0)$ to OPEN
- 3: **while** OPEN $\neq \emptyset$ **do**
- 4: $P = (\pi, g, \Omega) \leftarrow$ best node from OPEN
- 5: **if** π has no conflict **then**
- 6: **return** π
- 7: $C \leftarrow \text{FINDCONFLICT}(\pi)$
- 8: **if** C is not cardinal **then**
- 9: **if** $\text{FINDBYPASS}(P, C)$ **then**
- 10: **continue**
- 11: **for all** agent i in C **do**
- 12: $\Omega' \leftarrow \Omega \cup \text{NEWCONSTRAINTS}(i, C)$
- 13: $\pi', g' \leftarrow \text{LOWLEVEL}(i, \Omega')$
- 14: Add $P' = (\pi', g', \Omega')$ to OPEN
- 15: **return** failure

elevator conflict since many time steps must be forbidden for an agent. Fig. 4(b) provides such a toy example, where a large number of CT nodes are generated when resolving the elevator conflict using conventional CBS. By contrast, the same elevator conflict can be resolved in a single iteration (Fig. 4(c)) using our new method as explained below.

Let $(i, j, k, t_s^i, l_s^i, t_s^j, l_s^j, t_r^i, l_r^i, t_r^j, l_r^j)$ denote an elevator conflict, indicating that agents i and j encounter a conflict involving elevator k . Here, agent i , with start and goal vertices located at floor l_s^i, l_g^i respectively, starts to use elevator k at time t_s^i , and agent j , with start and goal vertices located at floor l_s^j, l_g^j respectively, starts to use elevator k at time t_s^j .

To simplify the presentation, we assume that the time required for an elevator to traverse any elevator edge between two adjacent layers is constant. Let $t_{floor} = c(e), e \in E_{ele,k}$ denote the cost of any elevator edge. Then, the duration for an elevator to transport agent i from layer l_s^i to l_g^i is $t_o^i = |l_s^i - l_g^i| \cdot t_{floor}$. Before the elevator can serve another agent $j \neq i$, the elevator must move to the start floor of agent j incurring a resetting cost of $t_r^i = |l_g^i - l_s^j| \cdot t_{floor}$. Similarly, the corresponding parameters for agent j , such as t_s^j (the time when agent j enters this elevator), t_o^j (duration that the elevator is occupied by j during transportation) and t_r^j (time for the elevator to return to agent i 's floor after transporting agent j) can be computed. Thus, due to agent

i 's usage of the elevator, the time interval during which the elevator k cannot be used by agent j is at least $IN_i(1)$ and $IN_j(2)$ can be similarly defined.

$$IN_i = [t_s^i, t_s^i + t_o^i + t_r^i] \quad (1)$$

$$IN_j = [t_s^j, t_s^j + t_o^j + t_r^j] \quad (2)$$

When an elevator conflict is first detected, without loss of generality, we consider the case where $t_s^i < t_s^j$, i.e., agent i starts using elevator k earlier than agent j , and that $t_s^j \leq t_s^i + t_o^i + t_r^i$, indicating that the two agents experience an elevator conflict over elevator k .

To resolve such an elevator conflict, the following two constraint sets are generated.

$$\Omega_i = \{\omega^i = (i, v_{k,l_s^i}, t^i) | t^i \in [t_s^i, t_s^j + t_o^j + t_r^j]\} \quad (3)$$

$$\Omega_j = \{\omega^j = (j, v_{k,l_s^j}, t^j) | t^j \in [t_s^j, t_s^i + t_o^i + t_r^i]\}, \quad (4)$$

where v_{k,l_s^i} and v_{k,l_s^j} denote the entrances of elevator k at floors l_s^i and l_s^j , respectively. The first set of constraints prohibits agent i from entering elevator k from vertex v_{k,l_s^i} during the time interval $[t_s^i, t_s^j + t_o^j + t_r^j]$. The second set of constraints prohibits agent j from entering elevator k from vertex v_{k,l_s^j} during the time interval $[t_s^j, t_s^i + t_o^i + t_r^i]$.

As shown in Fig. 4(c), by using EC, the CT needs to branch only once to resolve the elevator conflict, thereby addressing the long-term coupling between agents i and j caused by the elevator. In this example, $t_s^i = 1, t_s^j = 1, t_o^i = 1, t_r^j = 1$. As a result, the set of constraints for agent i is $\Omega_i = \{i, A1, [1, 3]\}$. Similarly the set of constraints for agent j is $\Omega_j = \{j, A1, [1, 3]\}$.

To prove CBS with this new branching rule returns an optimal solution, we need to show the two constraint sets are *mutually disjunctive* [14]. Specifically, two sets Ω_i and Ω_j are said to be mutually disjunctive if no conflict-free joint path exists that can simultaneously violate any pair of constraint sets (ω^i, ω^j) with $\omega^i \in \Omega_i$ and $\omega^j \in \Omega_j$. In other words, if paths π^i and π^j simultaneously violate constraint sets (3) and (4), then there must exist an elevator conflict.

Theorem 1. *Constraints (3) and (4) are mutually disjunctive.*

Proof. Suppose agent i (entering elevator k from vertex v_{k,l_s^i} at t_s^i) and agent j (entering elevator k from vertex v_{k,l_s^j} at t_s^j) encounter their first elevator conflict at elevator k , agents i and j are then assigned two constraints Ω_i and Ω_j . Assume that in the new paths of agents i and j , agent i enters elevator k at time x such that elevator k cannot be used by agent j for at least the interval $I_i = [x, x + t_o^i + t_r^i]$ by Eq.1, and agent j enters elevator k at time y such that elevator k cannot be used by agent i for at least the interval $I_j = [y, y + t_o^j + t_r^j]$ by Eq.2. We need to show that if the new paths of agents i and j violate the constraints Ω_i and Ω_j , for any pair (x, y) with $x \in [t_s^i, t_s^j + t_o^j + t_r^j]$ and $y \in [t_s^j, t_s^i + t_o^i + t_r^i]$, the two corresponding intervals I_i and I_j must intersect, which indicates a conflict. We show the following cases cannot happen by contradiction: case 1: $y > x + t_o^i + t_r^i$, and case 2: $x > y + t_o^j + t_r^j$. For the case 1, since $x \in [t_s^i, t_s^j + t_o^j + t_r^j]$ and $y \in [t_s^j, t_s^i + t_o^i + t_r^i]$,

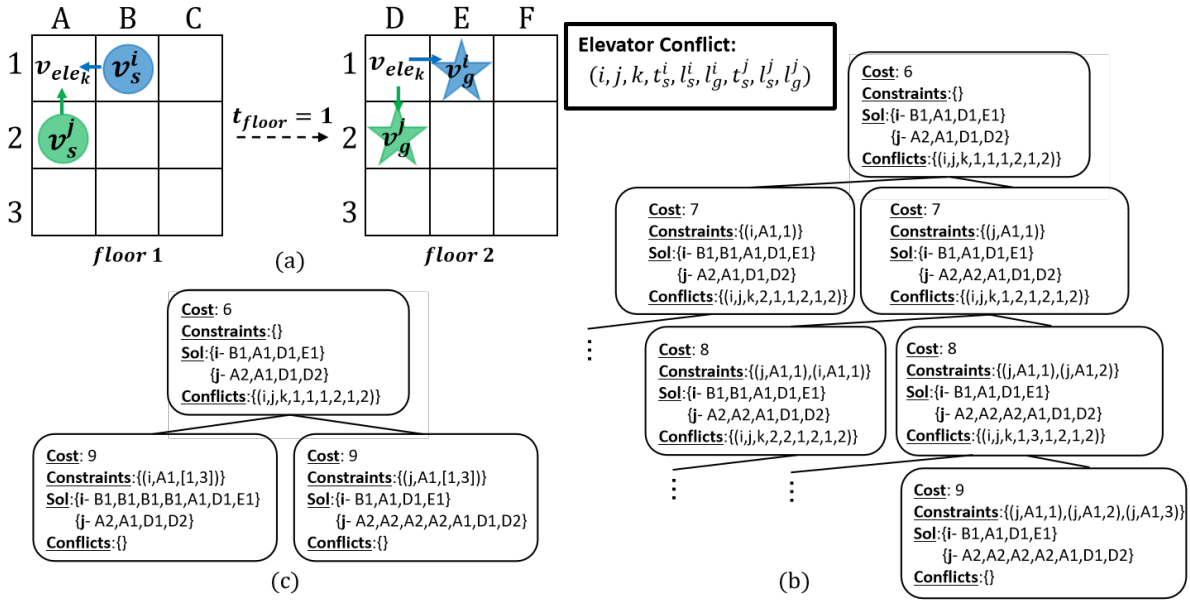


Fig. 4. (a) A toy example of MAPF-E, which seeks start-goal paths for the agents. The colors indicate the agents. The traversal time of any elevator edge is $t_{floor} = 1$. (b) Inefficiency of resolving the elevator conflict by single-timestep constraint in CBS, which requires expanding the CT at least $\min\{t_s^j + t_o^j + t_r^j - t_s^i + 1, t_s^i + t_o^i + t_r^i - t_s^j + 1\}$ times. (c) With EC, the elevator conflict can be resolved with a single expansion.

we have $y \leq t_s^i + t_o^i + t_r^i$ and $x + t_o^i + t_r^i \geq t_s^i + t_o^i + t_r^i$. Thus, $y \leq t_s^i + t_o^i + t_r^i \leq x + t_o^i + t_r^i$, which contradicts the assumption $y > x + t_o^i + t_r^i$. For the case $x > y + t_o^j + t_r^j$, since $x \in [t_s^i, t_s^i + t_o^i + t_r^i]$ and $y \in [t_s^j, t_s^j + t_o^j + t_r^j]$, we have $x \leq t_s^j + t_o^j + t_r^j$, and $y + t_o^j + t_r^j \geq t_s^i + t_o^i + t_r^i$. Thus, $x \leq t_s^j + t_o^j + t_r^j \leq y + t_o^j + t_r^j$, which contradicts the assumption $x > y + t_o^j + t_r^j$. Therefore, if a joint path π simultaneously violates any pair of constraints $(\omega^i, \omega^j), \omega^i \in \Omega_i, \omega^j \in \Omega_j$, π must contain elevator conflicts, and constraints (3) and (4) are mutually disjunctive. \square

Remark 1. *The conventional CBS is known to be inefficient when resolving conflict in long corridor [3], a problem that is addressed by corridor reasoning [15]. The key difference between corridor conflicts and elevator conflicts lies in the agents' behavior: in a feasible solution, agents can wait or avoid other agents by temporarily entering, whereas an elevator can only transport agents between floors and can be unavailable even when no agents are using it (due to resetting). As a result, corridor reasoning is not applicable to elevator conflicts.*

B. MDD with Elevator Details (MDD-E)

Multi-Valued Decision Diagrams (MDD) [12] cannot be directly applied to MAPF-E as it does not consider the state of the elevator and therefore fails to detect elevator-related conflicts. As illustrated in Fig. 5(b), when constructing the joint MDD using the classical method, no edge or vertex conflict is detected between agent i and j at $t = 3$. Consequently, the node $((E1, 3), (A1, 3))$ is added to MDD_5^{ij} . Nevertheless, an elevator conflict actually occurs at $t = 3$: the elevator is in the resetting state due to agent i 's usage, while agent j attempts to use it. Since elevator-related information is not encoded in the nodes, this conflict remains undetected

during the construction of MDD_5^{ij} . We therefore modify MDD to incorporate elevator states.

Let $(v, t, k, t_s^i, l_s^i, l_g^i)$ denote an MDD-E node, where agent i has start and goal vertices located on layers l_s^i and l_g^i , respectively, and enters elevator k at time t_s^i . Let $k, t_s^i, l_s^i, l_g^i = -1$ to indicate that agent i has not yet used any elevator. Each node at time t in $MDD-E_d^i$ corresponds to a possible location of agent i at time t along a path of cost d from v_s^i to v_g^i . Edges between nodes in $MDD-E_d^i$ are directed, representing feasible transitions where agent i can move from one node to another in the next time step.

When constructing the joint $MDD-E_d^{ij}$ for agents i and j , if the corresponding nodes in $MDD-E_d^i$ and $MDD-E_d^j$ have the same k value of elevator and $k \neq -1$, then the other information in the nodes is used to determine whether an elevator conflict occurs, according to the detection method described in Sec. II (Example 1). Specifically, an elevator conflict exists if either (5) or (6) holds:

$$t_s^i \in [t_s^j, t_s^j + t_o^j + t_r^j] \quad (5)$$

$$t_s^j \in [t_s^i, t_s^i + t_o^i + t_r^i] \quad (6)$$

By identifying the types of elevator conflicts using MDD-E and integrating Conflict Selection and Bypass techniques at the high-level, the CBS with MDD-E can compute a lower bound on the joint path cost that is at least as tight as that of CBS without MDD-E, given the same number of CT nodes.

Example 1. *Using our MDD-E, each node is augmented with elevator-related information, allowing the joint MDD-E to correctly detect elevator conflicts. As shown in Fig. 5(c), node $(E1, 3, k, 1, 1, 2)$ of $MDD-E_4^i$ encodes that agent i starts using elevator k at $t = 1$, traveling from floor 1 to floor 2. Here, elevator k is still in the resetting state at $t = 3$. Since agent j is not allowed to use an elevator that is in the*

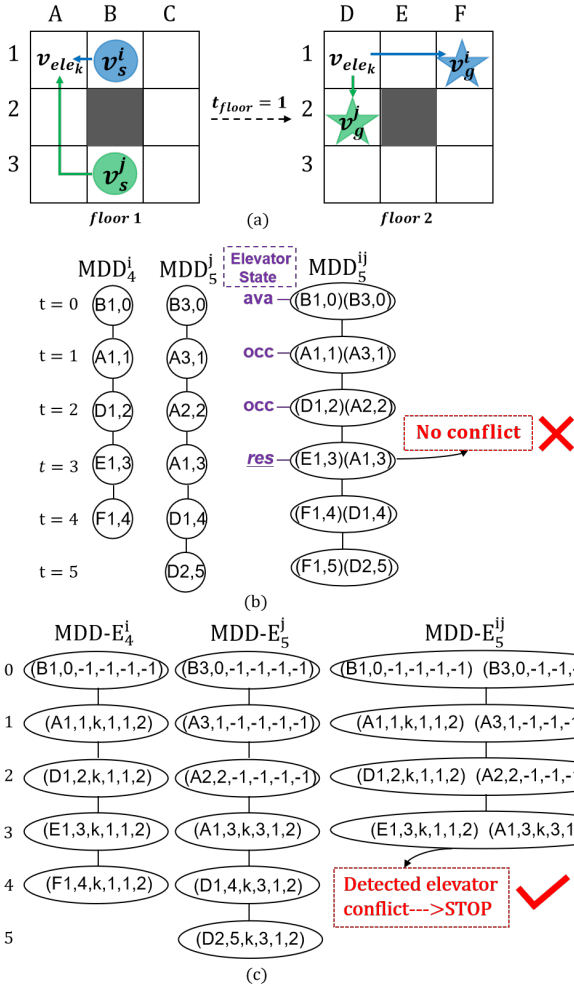


Fig. 5. (a) A toy example of MAPF-E, where agents (green and blue) move from start to goal, black blocks denote obstacles, and each elevator edge has a traversal time of $t_{floor} = 1$. (b) shows that when constructing joint MDD by the classical method, an elevator conflict occurring at $t = 3$ cannot be detected, because the nodes in the traditional MDD do not contain elevator-related information and cannot capture the elevator state; (c) shows that with MDD-E, the elevator conflict can be successfully detected during the construction of the joint MDD-E.

resetting state, the elevator conflict can be correctly detected during the construction of the joint MDD-E.

V. EXPERIMENTAL RESULTS

We use two 4-neighbor grids of sizes 8×8 and 16×16 as the single-floor maps, both with 10% randomly generated obstacles, and vary the number of layers in our experiments.² Each floor shares the same grid map, and all the floors are connected by a certain number of elevators. As shown in Fig. 6, a multi-floor scenario is generated based on maps of size 8×8 . For all the algorithms, we adopt SIPP [16] as the low-level planner. All tests are on a laptop with Intel i7-13700 CPU with 32GB RAM.

²Note that, the overall scale of the workspace quickly grows as the number of layers increases. For example, when each floor is represented by a relatively small map of size 8×8 with 64 vertices, the resulting workspace graph contains 320 vertices when 5 floors are considered.




Fig. 6. A map with 5 floors connected by 3 elevators, where each floor is the random-8-8-10 map.

A. Experiment 1: Varying Types of Map

Our first experiment considers two floors connected by 3 elevators, with $t_{floor} = 3$. For each map, we run 15 instances with varying numbers of agents N , and set a runtime limit of 60 seconds for each instance. In each instance, the start and goal positions of the agents, as well as the locations of the 3 elevators, are randomly generated, ensuring that a valid solution exists. In these experiments, we compare EC and MDD-E against a naive adaptation of CBS to MAPF-E (denoted as CBS).

1) *Success Rates*: Fig. 7 (a), (d) show the success rates (i.e., the rate of instances solved by each algorithm within 60 seconds) for the algorithms. Incorporating EC into the algorithms often improves the success rate within the same time limit. The introduction of MDD-E alone brings limited improvement and may even perform worse than CBS when $N \geq 11$. This is because, as the number of agents increases, conflicts become more frequent, and the computational overhead of MDD-E in the conflict selection process restricts the number of conflicts that can be resolved within the time limit. The computational overhead of MDD-E is often larger for agents with long paths, which is often the case since the agents take the elevator to move across multiple floors.

MDD-E demonstrates its advantages when combined with more efficient constraint-handling techniques (e.g., EC). For instance, on the random-16-16-10 map, CBS+EC+MDD-E achieves a higher success rate than CBS+EC. This is because, with EC efficiently resolving conflicts, MDD-E can further guide the conflict selection process by classifying conflicts, thereby improving the success rate. However, on the random-8-8-10 map, this improvement is less pronounced. In particular, when $N = 14$, the success rate of CBS+EC+MDD-E is lower than that of CBS+EC in some cases. This is because, with the same number of agents, conflicts are denser on the smaller map and the additional overhead of MDD-E leads to more timeouts. Furthermore, compared with CBS+MDD-E, CBS+EC+MDD-E achieves higher success rates due to EC. In fact, on both maps, CBS+EC+MDD-E achieves a success rate that is up to

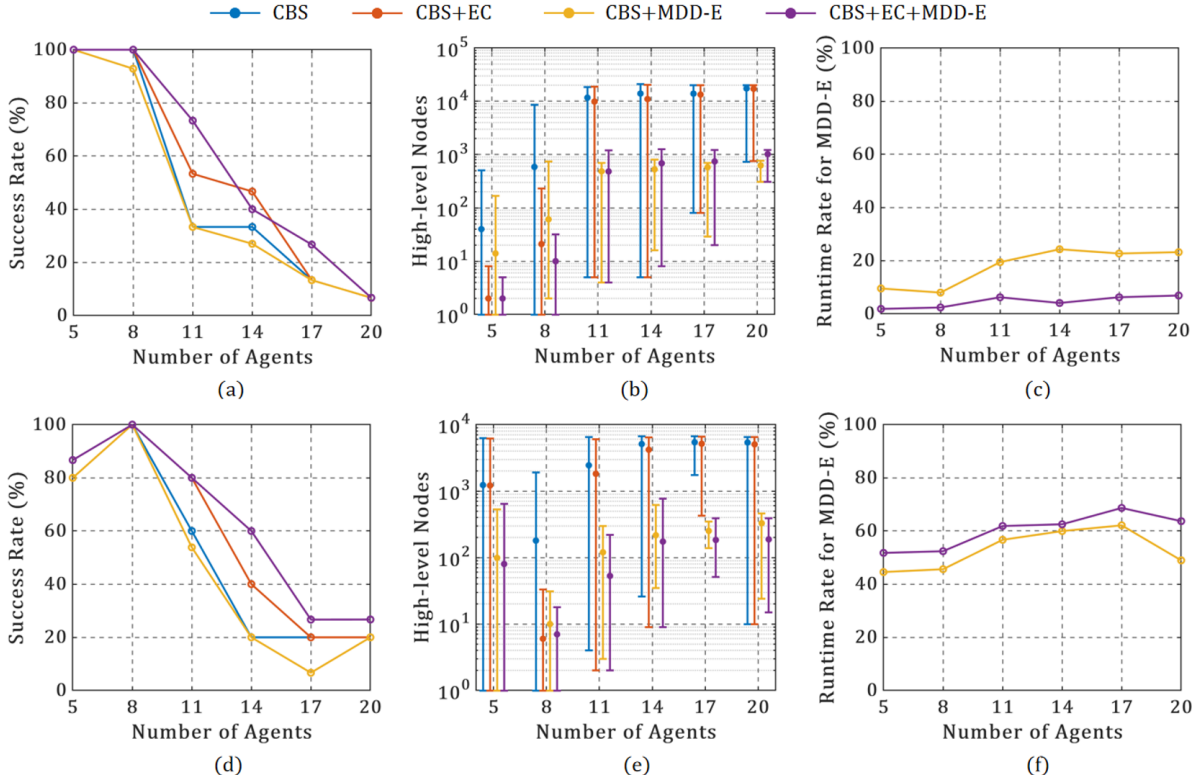


Fig. 7. (a),(d): the success rates of the algorithms when the number of agents $N \in \{5, 8, 11, 14, 17, 20\}$ on random-8-8-10 and random-16-16-10 maps respectively. (b),(e): the minimum, average and maximum number of high-level nodes expanded over all instances with varying numbers of agents N . (c),(f): the fractions of total search time spent on MDD-E and related techniques in CBS+MDD-E and CBS+EC+MDD-E within the time limit over all instances with varying numbers of agents N .

around 40% higher than CBS+MDD-E.

2) *Number of Expansions*: Fig. 7 (b), (e) show the average number of resolved conflicts (i.e., expanded high-level nodes). Compared with CBS, CBS+EC resolves elevator conflicts more efficiently, leading to fewer average CT nodes and a higher success rate, as observed in both the random-8-8-10 and random-16-16-10 maps. Furthermore, using MDD-E with Conflict Selection and BP techniques at the high level reduces expansions by prioritizing more critical conflicts. The CBS+EC+MDD-E algorithm can solve the same instances with fewer high-level node expansions compared to the CBS algorithm. For example, on the random-16-16-10 map with $N \geq 11$, the number of expanded CT nodes in CBS+EC+MDD-E is much lower than that in CBS, and its success rate is also higher.

3) *MDD-E Runtime Overhead*: Fig. 7 (c), (f) show the fraction of total search time spent on MDD-E and related techniques within the 60-second runtime limit. In both maps, as N increases, more conflicts require repeated construction and querying of MDD-E and joint MDD-E for conflict classification, thus increasing the runtime overhead. As larger maps often lead to longer paths, the MDD-E overhead increases accordingly, resulting in a higher fraction in Fig. 7 (f) than in Fig. 7 (c).

B. Experiment 2: Varying Value of t_{floor}

We then investigate the effect of increasing t_{floor} . We use five floors connected by three elevators, where each floor is

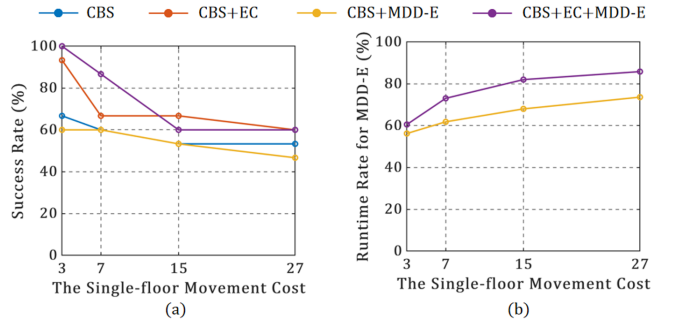


Fig. 8. (a): the success rates of CBS, CBS+EC, CBS+MDD-E, and CBS+EC+MDD-E as t_{floor} increases, on a 5-floor map where each floor is the random-16-16-10 map. (b): the fractions of total search time spent on MDD-E and related techniques in CBS+MDD-E and CBS+EC+MDD-E within the time limit over all instances in each group with varying values of t_{floor} .

based on the random-16-16-10 map and the number of agents is fixed at $N = 8$. Fig. 8 (a) presents the success rates. As t_{floor} increases, both CBS+EC and CBS+EC+MDD-E maintain higher success rates due to the advantages of EC compared with CBS. Combined with Fig. 8 (b), which shows the fraction of total search time spent on MDD-E and related techniques, the overhead of MDD-E increases with t_{floor} . This, in turn, leads to a decrease in the success rates of MDD-E based algorithms, since longer agent paths result in larger and more costly MDD-E and joint MDD-E structures.

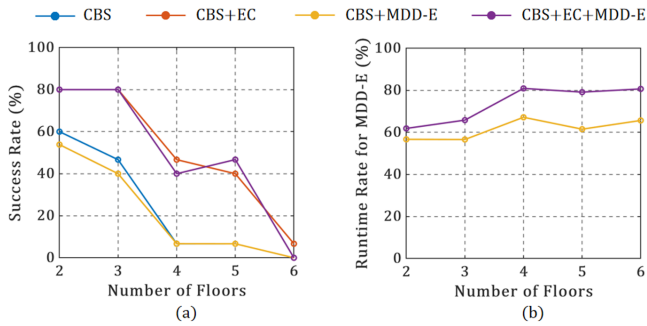


Fig. 9. (a): the success rates of CBS, CBS+EC, CBS+MDD-E, and CBS+EC+MDD-E as the number of floors increases, with each floor being the `random-16-16-10` map. (b): the fractions of total search time spent on MDD-E and related techniques in CBS+MDD-E and CBS+EC+MDD-E within the time limit over all instances in each group with varying numbers of floors.

For $t_{floor} \geq 15$, elevator conflicts in the successfully solved cases tend to be less severe and less sensitive to increases in t_{floor} , resulting in stable success rates.

C. Experiment 3: Varying Number of Floors

The third experiment investigates the effect of varying the number of floors with 3 elevators, where each floor is based on the `random-16-16-10` map and the number of agents is fixed at $N = 11$. Fig. 9 (a) presents the success rates, showing that both CBS+EC and CBS+EC+MDD-E maintain higher success rates than CBS, and this advantage persists even as the number of floors increases due to the efficient conflict resolution provided by EC. However, as the number of floors increases, the average elevator occupied time per agent also increases, resulting in stronger coupling among agents and a gradual decline in success rates. Intuitively, the shared use of elevators introduces a combinatorial growth in potential conflicts as the number of floors increases. Fig. 9 (b) shows that the runtime ratio spent on MDD-E and related techniques grows significantly with the number of floors. When the number of floors exceeds 4, this rate in CBS+EC+MDD-E already surpasses 80%, indicating that constructing and maintaining MDD-E imposes substantial overhead in such scenarios.

VI. CONCLUSIONS AND FUTURE WORK

This paper investigates a new problem, MAPF-E, and develops efficient methods for resolving elevator conflicts, addressing the coupling of multiple agents introduced by elevators. By using elevator constraints, elevator conflicts can be resolved with only a single expansion while still guaranteeing optimal solutions. By extending the information stored in MDD nodes, we construct a new structure called MDD-E, which serves as the foundation for applying PC and BP techniques. Experimental results demonstrate the advantages of our new approaches under different settings.

Future work includes fast construction of MDD-E, considering elevators holding more than one agents, extension to agents with asynchronous/continuous-time actions [17]–[19], and combination with task planning [20], [21].

REFERENCES

- [1] J. Yu and S. M. LaValle, “Structure and intractability of optimal multi-robot path planning on graphs,” in *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [2] Z. Robotics. (2025) Wcs system. Accessed: 2025-08-16. [Online]. Available: <https://www.zsrobotics.com/en/zsmart.html>
- [3] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, “Conflict-based search for optimal multi-agent pathfinding,” *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [4] G. Wagner and H. Choset, “Subdimensional expansion for multirobot path planning,” *Artificial Intelligence*, vol. 219, pp. 1–24, 2015.
- [5] E. Boyarski, A. Felner, R. Stern, G. Sharon, D. Tolpin, O. Betzalel, and E. Shimony, “Icbs: improved conflict-based search algorithm for multi-agent pathfinding,” in *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [6] M. Barer, G. Sharon, R. Stern, and A. Felner, “Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem,” in *Seventh Annual Symposium on Combinatorial Search*, 2014.
- [7] J. Li, W. Ruml, and S. Koenig, “Eecbs: A bounded-suboptimal search for multi-agent path finding,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 14, 2021, pp. 12 353–12 362.
- [8] K. Okumura, M. Machida, X. Défago, and Y. Tamura, “Priority inheritance with backtracking for iterative multi-agent path finding,” *Artificial Intelligence*, p. 103752, 2022.
- [9] B. De Wilde, A. W. Ter Mors, and C. Witteveen, “Push and rotate: cooperative multi-agent path planning,” in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, 2013, pp. 87–94.
- [10] Q. Wang, R. Veerapaneni, Y. Wu, J. Li, and M. Likhachev, “Mapf in 3d warehouses: Dataset and analysis,” in *Proceedings of the Thirty-Fourth International Conference on Automated Planning and Scheduling (ICAPS 2024)*. IEEE, 2024, pp. 623–632.
- [11] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, “Conflict-based search for optimal multi-agent pathfinding,” *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [12] —, “Meta-agent conflict-based search for optimal multi-agent path finding,” *SoCS*, vol. 1, pp. 39–40, 2012.
- [13] G. Sharon, R. Stern, M. Goldenberg, and A. Felner, “The increasing cost tree search for optimal multi-agent pathfinding,” vol. 195, 2 2013, pp. 470–495.
- [14] J. Li, P. Surynek, A. Felner, H. Ma, T. S. Kumar, and S. Koenig, “Multi-agent path finding for large agents,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, 2019, pp. 7627–7634.
- [15] J. Li, D. Harabor, P. J. Stuckey, H. Ma, G. Gange, and S. Koenig, “Pairwise symmetry reasoning for multi-agent path finding search,” *Artificial Intelligence*, vol. 301, p. 103574, 2021.
- [16] M. Phillips and M. Likhachev, “Sipp: Safe interval path planning for dynamic environments,” in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 5628–5635.
- [17] Z. Ren, S. Rathinam, and H. Choset, “Loosely synchronized search for multi-agent path finding with asynchronous actions,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021.
- [18] S. Zhou, S. Zhao, and Z. Ren, “Loosely synchronized rule-based planning for multi-agent path finding with asynchronous actions,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 39, Apr. 2025, pp. 14 763–14 770.
- [19] X. Wu, S. Zhao, and Z. Ren, “Conflict-based search for multi agent path finding with asynchronous actions,” in *Proceedings of the 25th International Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS ’26. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2026, p. 1174–1182.
- [20] Z. Ren, S. Rathinam, and H. Choset, “CBSS: A New Approach for Multiagent Combinatorial Path Finding,” *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2669–2683, 2023.
- [21] —, “Ms*: A new exact algorithm for multi-agent simultaneous multi-goal sequencing and path finding,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 11 560–11 565.