

A Real-Time Rescheduling Algorithm for Multi-robot Plan Execution

Ying Feng¹, Adittyo Paul¹, Zhe Chen², Jiaoyang Li¹

¹Carnegie Mellon University, USA

²Monash University, Australia

{yingfeng,adittyop}@andrew.cmu.edu, zhe.chen@monash.edu, jiaoyangli@cmu.edu

Abstract

One area of research in multi-agent path finding is to determine how replanning can be efficiently achieved in the case of agents being delayed during execution. One option is to reschedule the passing order of agents, i.e., the sequence in which agents visit the same location. In response, we propose Switchable-Edge Search (SES), an A*-style algorithm designed to find optimal passing orders. We prove the optimality of SES and evaluate its efficiency via simulations. The best variant of SES takes less than 1 second for small- and medium-sized problems and runs up to 4 times faster than baselines for large-sized problems.

Introduction

Multi-Agent Path Finding (MAPF) is the problem of finding collision-free paths that move a team of agents from their start to goal locations. MAPF is fundamental to numerous applications, such as automated warehouses (Wurman, D’Andrea, and Mountz 2007; Kou et al. 2020), computer games (Silver 2005), and drone swarms (Hönig et al. 2018).

Classic MAPF models assume flawless execution. However, in real-world scenarios, agents may encounter unexpected delays due to mechanical differences, unforeseen events, localization errors, and so on. To accommodate such delays, existing research suggests the use of a Temporal Plan Graph (TPG) (Hönig et al. 2016). The TPG captures the precedence relationships within a MAPF solution and maintains them during execution. Each precedence relationship specifies an order for two agents to visit the same location. An agent advances to the next location in its path only if the corresponding precedence conditions are met. Consequently, if an agent experiences a delay, all other agents whose actions depend on this agent will pause. Despite its advantages of providing rigorous guarantees on collision-freeness and deadlock-freeness, the use of TPG can introduce a significant number of waits into the execution results due to the knock-on effect in the precedence relationship.

In this paper, we adopt a variant of TPG, named *Switchable TPG* (STPG) (Berndt et al. 2020). STPG allows for the modification of some precedence relationships in a TPG, resulting in new TPGs. To address delays, we propose an A*-style algorithm called *Switchable-Edge Search* (SES) to find

the new TPG based on a given STPG that minimizes the travel times of all agents to reach their goal locations. We prove the optimality of SES and introduce two variants: an execution-based variant (ESES) and a graph-based variant (GSES). Experimental results show that GSES finds the optimal TPG with an average runtime of less than 1 second for various numbers of agents on small- and medium-sized maps. On larger maps, GSES runs up to 4 times faster than existing replanning algorithms.

Preliminaries

Definition 1 (MAPF). Multi-Agent Path Finding (MAPF) aims to find collision-free paths for a team of agents \mathcal{A} on a given graph. Each agent $i \in \mathcal{A}$ has a unique start location and a unique goal location. In each discrete timestep, every agent either moves to an adjacent location or waits at its current location. A path for an agent specifies its action at each timestep from its start to goal locations. A collision occurs if either of the following happens:

1. Two agents are at the same location at the same timestep.
2. One agent leaves a location at the same timestep when another agent enters the same location.

A MAPF solution is a set of collision-free paths of all agents.

Remark 1. The above definition of collision coincides with that in the setting of k-robust MAPF (Atzmon et al. 2018) with $k = 1$. We disallow the second type of collision because, if agents follow each other and the front agent suddenly stops, the following agents may collide with the front agent. Thus, this restriction ensures better robustness when agents are subject to delays. Note that the swapping collision, where two agents swap their locations simultaneously, is a special case of the second type of collision.

A MAPF solution can be represented in different formats. We stick to the following format for our discussion, though our algorithms do not depend on specific formats.

Definition 2 (MAPF Solution). A MAPF solution takes the form of a set of collision-free paths $\mathcal{P} = \{p_i : i \in \mathcal{A}\}$. Each path p_i is a sequence of location-timestep tuples $(l_0^i, t_0^i) \rightarrow (l_1^i, t_1^i) \rightarrow \dots \rightarrow (l_{z_i}^i, t_{z_i}^i)$ with the following properties: (1) The sequence follows a strict temporal ordering: $0 = t_0^i < t_1^i < \dots < t_{z_i}^i$. (2) l_0^i and $l_{z_i}^i$ are the start and goal locations of agent i , respectively. (3) Each tuple (l_k^i, t_k^i) with $k > 0$ specifies a move action of i from l_{k-1}^i to l_k^i at timestep t_k^i .

These properties force all consecutive pairs of locations l_k^i and l_{k+1}^i to be adjacent on the graph. A wait action is implicitly defined between two consecutive tuples. Namely, if $t_{k+1}^i - t_k^i = \Delta > 1$, then i is planned to wait at l_k^i for $\Delta - 1$ timesteps before moving to l_{k+1}^i . Additionally, t_{zi}^i records the time when i reaches its goal location, called the *travel time* of i . The cost of a MAPF solution \mathcal{P} is $cost(\mathcal{P}) = \sum_{i \in \mathcal{A}} t_{zi}^i$, and \mathcal{P} is *optimal* if its cost is minimum.

Remark 2. Definition 2 discards the explicit representation of wait actions. This is because, when executing \mathcal{P} as a TPG (specified in the next section), it may reduce the travel time if \mathcal{P} has unnecessary wait actions, e.g., when \mathcal{P} is suboptimal.

Related Works Numerous recent studies on MAPF have explored strategies for managing unexpected delays during execution. A simple strategy is to re-solve the MAPF problem when a delay occurs. However, this strategy is computationally intensive, leading to prolonged agent waiting time. To avoid the need for replanning, Atzmon et al. (2018) suggested the creation of a k -robust MAPF solution, allowing agents to adhere to their planned paths even if each agent is delayed by up to k timesteps. However, replanning is still required if an agent’s delay exceeds k timesteps. Atzmon et al. (2020) then proposed a different model, called p -robust MAPF solutions, that ensures execution success with a probability of at least p , given an agent delay probability model. Nevertheless, planning a k -robust or p -robust MAPF solution is considerably more computational-intensive than computing a standard MAPF solution. Another strategy for managing delays involves the use of an execution policy that preserves the precedence relationships of a MAPF solution during execution (Hönig et al. 2016; Ma, Kumar, and Koenig 2017; Hönig et al. 2019). This strategy is quick and eliminates the need for replanning paths. However, the execution results often leave room for improvement, as many unnecessary waits are introduced. Our work aims to address this limitation by formally exploring the concept of optimizing precedence relationships online (Berndt et al. 2020; Manucci, Pallottino, and Pecora 2021).

Temporal Plan Graph (TPG)

In essence, we aim to optimize the passing order for multiple agents to visit the same location. This is achieved using a graph-based abstraction known as the TPG.

Definition 3 (TPG). A Temporal Plan Graph (TPG) (Hönig et al. 2016) is a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ that represents the precedence relationships of a MAPF solution \mathcal{P} . The set of vertices is $\mathcal{V} = \{v_k^i : i \in \mathcal{A}, k \in [0, z_i]\}$, where each vertex v_k^i corresponds to (l_k^i, t_k^i) , namely the k^{th} move action in path p_i . There are two types of edges \mathcal{E}_1 and \mathcal{E}_2 , where each directed edge $(u, v) \in \mathcal{E}_1 \cup \mathcal{E}_2$ encodes a precedence relationship between a pair of move actions, namely movement u is planned to happen *before* movement v .

- A *Type 1 edge* connects two vertices of the same agent, specifying its path. Specifically, $\mathcal{E}_1 = \{(v_k^i, v_{k+1}^i) : \forall i \in \mathcal{A}, k \in [0, z_i]\}$.
- A *Type 2 edge* connects two vertices of distinct agents, specifying their ordering of visiting the same location.




Figure 1: Example of converting a MAPF solution to a TPG. The solid arrows in the TPG represent Type 1 edges, and the dashed arrow represents a Type 2 edge.

Procedure 1: TPG Execution

```

1 Define a counter cost;
2 Function INITEXEC( $\mathcal{V}$ )
3   cost  $\leftarrow$  0;
4   Mark vertices in  $\mathcal{V}_0 = \{v_0^i : i \in \mathcal{A}\}$  as satisfied;
5   Mark vertices in  $(\mathcal{V} \setminus \mathcal{V}_0)$  as unsatisfied;
6 Function STEPEXEC( $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ )
7    $S \leftarrow \emptyset$ ; // Vertices to be marked as satisfied
8   forall  $i \in \mathcal{A} : v_{zi}^i$  unsatisfied do
9     cost  $\leftarrow$  cost + 1;
10     $k \leftarrow \min\{k : v_k^i \text{ unsatisfied}\}$ ;
11    if  $\forall (u, v_k^i) \in \mathcal{E}_2 : u$  satisfied then
12      Add  $v_k^i$  into  $S$ ;
13  return  $S$ ;
14 Function EXEC( $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$ )
15  INITEXEC( $\mathcal{V}$ );
16  while  $\exists v \in \mathcal{V} : v$  unsatisfied do
17     $S \leftarrow$  STEPEXEC( $\mathcal{G}$ );
18    forall  $v_k^i \in S$  do Mark  $v_k^i$  as satisfied;
19  return cost;

```

Specifically, $\mathcal{E}_2 = \{(v_{s+1}^j, v_k^i) : \forall i \neq j \in \mathcal{A}, s \in [0, z_j], k \in [0, z_i] \text{ satisfying } l_s^j = l_k^i \text{ and } t_{s+1}^j < t_k^i\}$.

Example 1. Figure 1 shows an example of converting a MAPF solution into a TPG. Both agents are planned to visit location D, and the red agent is planned to visit D earlier than the blue agent. Consequently, there is a Type 2 edge from v_3^{red} to v_2^{blue} , signifying that the blue agent can move to D only after the red agent has reached F. Note that we define Type 2 edge as $(v_3^{\text{red}}, v_2^{\text{blue}})$ instead of $(v_2^{\text{red}}, v_2^{\text{blue}})$ to avoid the second type of collision in Definition 1.

Executing a TPG Procedure 1 describes how to execute a TPG \mathcal{G} , which includes a main function EXEC and two helper functions INIT_{EXEC} and STEP_{EXEC}, along with two marks “satisfied” and “unsatisfied” for vertices. Marking a vertex as satisfied corresponds to moving an agent to the corresponding location, and we do so if and only if all in-neighbors of this vertex have been satisfied. The execution terminates when all vertices are satisfied, i.e., all agents have reached their goal locations. The *cost* of executing \mathcal{G} , namely $cost(\mathcal{G}) = \text{EXEC}(\mathcal{G})$, is the sum of travel time for

agents following \mathcal{G} (while assuming no delays happen).

We now introduce some known properties of TPGs. All proofs are delayed to the appendix. We use \mathcal{G} to denote a TPG constructed from a MAPF solution \mathcal{P} as in Definition 3.

Proposition 1 (Cost). $cost(\mathcal{G}) \leq cost(\mathcal{P})$.

Intuitively, $cost(\mathcal{G}) < cost(\mathcal{P})$ if \mathcal{P} has unnecessary wait actions and $cost(\mathcal{G}) = cost(\mathcal{P})$ otherwise.

Proposition 2 (Collision-Free). *Executing a TPG with Procedure 1 does not lead to collisions among agents.*

Next, we present two lemmas regarding *deadlocks* of executing a TPG, which were used in previous work (Berndt et al. 2020; Su, Veerapaneni, and Li 2024) and are helpful for our discussion of switchable TPGs in the next section.

Definition 4 (Deadlock). When executing a TPG, a *deadlock* is encountered iff, in an iteration of the while-loop of EXEC(\mathcal{G}), \mathcal{V} contains unsatisfied vertices but $\mathcal{S} = \emptyset$.

Lemma 3 (Deadlock \iff Cycle). *Executing a TPG encounters a deadlock iff the TPG contains cycles.*

Lemma 4 (Deadlock-Free). *If a TPG is constructed from a MAPF solution, then executing it is deadlock-free.*

Switchable TPG (STPG)

TPG is a handy representation for precedence relationships. Yet, a TPG constructed as in Definition 3 is fixed and bound to a given set of paths. In contrast, our optimization algorithm will use the following extended notion of TPGs, which enables flexible modifications of precedence relationships.

Definition 5 (STPG). Given a TPG $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$, a *Switchable TPG* (STPG) $\mathcal{G}^S = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{\mathcal{E}_2}, \mathcal{N}_{\mathcal{E}_2}))$ partitions Type 2 edges \mathcal{E}_2 into two disjoint subsets $\mathcal{S}_{\mathcal{E}_2}$ (switchable edges) and $\mathcal{N}_{\mathcal{E}_2}$ (non-switchable edges) and allows two operations on any switchable edge $(v_{s+1}^j, v_k^i) \in \mathcal{S}_{\mathcal{E}_2}$:

- *fix* (v_{s+1}^j, v_k^i) removes (v_{s+1}^j, v_k^i) from $\mathcal{S}_{\mathcal{E}_2}$ and add it into $\mathcal{N}_{\mathcal{E}_2}$. It fixes a switchable edge to be non-switchable.
- *reverse* (v_{s+1}^j, v_k^i) removes (v_{s+1}^j, v_k^i) from $\mathcal{S}_{\mathcal{E}_2}$ and add (v_{k+1}^i, v_s^j) into $\mathcal{N}_{\mathcal{E}_2}$. It switches the precedence relationship and then fixes it to be non-switchable.

Remark 3. Reversing the precedence relationship represented by (v_{s+1}^j, v_k^i) produces (v_{k+1}^i, v_s^j) because, based on Definition 3, Type 2 edge (v_{s+1}^j, v_k^i) indicates locations l_s^j and l_k^i are the same. Thus, after reversing, vertex v_{k+1}^i needs to be satisfied before v_s^j can be marked as satisfied.

Example 2. Figure 2 shows an example of reversing an edge. After the *reverse* operation, edge $(v_1^{\text{red}}, v_2^{\text{blue}})$ in the left TPG is replaced with edge $(v_3^{\text{blue}}, v_0^{\text{red}})$ in the right TPG.

Definition 5 defines a strict superclass of Definition 3. A STPG degenerates into a TPG if $\mathcal{S}_{\mathcal{E}_2}$ is empty.

Definition 6 (\mathcal{G}^S -producible TPG). Given a STPG \mathcal{G}^S , a TPG is \mathcal{G}^S -producible if it can be generated through a sequence of *fix* and *reverse* operations on \mathcal{G}^S .




Figure 2: Example of reversing an edge in a TPG.

We now show the roadmap of our algorithm. Given a MAPF solution \mathcal{P} , we construct TPG \mathcal{G}_0 from \mathcal{P} as in Definition 3 and then run Procedure 1. When a delay happens, we (1) construct a STPG \mathcal{G}^S based on \mathcal{G}_0 and (2) find a TPG \mathcal{G}^* with $cost(\mathcal{G}^*) = \min\{cost(\mathcal{G}) : \mathcal{G} \text{ is } \mathcal{G}^S\text{-producible}\}$, representing an optimal ordering of agents visiting each location, upon sticking to the original location-wise paths. We describe Step (1) below and Step (2) in the next section.

Construction 1. Assume that, during the execution of $\mathcal{G}_0 = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$, agent d is forced to delay at its current location l_c^d for Δ timesteps. We construct STPG \mathcal{G}^S as follows:

1. Construct STPG $\mathcal{G}^S = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{\mathcal{E}_2}, \mathcal{N}_{\mathcal{E}_2}))$ with $\mathcal{S}_{\mathcal{E}_2} = \{(v_{s+1}^j, v_k^i) \in \mathcal{E}_2 : v_{s+1}^j \text{ is unsatisfied and } k < zi\}$ and $\mathcal{N}_{\mathcal{E}_2} = \{(v_{s+1}^j, v_k^i) \in \mathcal{E}_2 : v_{s+1}^j \text{ is satisfied or } k = zi\}$.
2. Create Δ new dummy vertices $\mathcal{V}_{\text{new}} = \{v_1, \dots, v_\Delta\}$ and $(\Delta + 1)$ new Type-1 edges $\mathcal{E}_{\text{new}} = \{(v_c^d, v_1), (v_1, v_2), \dots, (v_{\Delta-1}, v_\Delta), (v_\Delta, v_{c+1}^d)\}$ and modify \mathcal{G}^S with $\mathcal{V} \leftarrow \mathcal{V} \cup \mathcal{V}_{\text{new}}$ and $\mathcal{E}_1 \leftarrow (\mathcal{E}_1 \cup \mathcal{E}_{\text{new}}) \setminus \{(v_c^d, v_{c+1}^d)\}$.

If there are multiple agents delayed at the same timestep, we repeat Step 2 for each delayed agent.

Remark 4. In Step 1, (v_{s+1}^j, v_k^i) is non-switchable when v_{s+1}^j is satisfied because agent j has already visited l_s^j . (v_{s+1}^j, v_{zi}^i) is non-switchable because agent i must be the last one to visit its goal location. The dummy vertices added in Step 2 are used to account for the delays in Procedure 1.

We now show an intuitive yet crucial theorem.

Theorem 5. *If STPG \mathcal{G}^S is constructed by Construction 1, then there is at least one deadlock-free \mathcal{G}^S -producible TPG.*

Proof. We generate a naïve solution $\mathcal{G}_{\text{naive}}$ by *fixing* all switchable edges in \mathcal{G}^S . Lemma 4 ensures that \mathcal{G}_0 is deadlock-free. \mathcal{G}^S constructed in Step 1 is identical to \mathcal{G}_0 if we *fix* all switchable edges. Step 2 behaves as expanding a pre-existing edge (v_{k-1}^i, v_k^i) into a line of connecting edges, which does not create any new cycles. Therefore, by Lemma 3, $\mathcal{G}_{\text{naive}}$ is deadlock-free. \square

Switchable Edge Search (SES) Framework

We describe our algorithm, *Switchable Edge Search* (SES), in a top-down modular manner, starting with a high-level heuristic search framework in Algorithm 2. We define the *partial cost* of a STPG as the cost of its *reduced* TPG, which is defined as follows.

Definition 7 (Reduced TPG). The reduced TPG of a STPG $\mathcal{G}^S = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{\mathcal{E}_2}, \mathcal{N}_{\mathcal{E}_2}))$ is the TPG that omits all switchable edges, denoted as $red(\mathcal{G}^S) = (\mathcal{V}, \mathcal{E}_1, \mathcal{N}_{\mathcal{E}_2})$.

Algorithm 2: Switchable Edge Search (SES)
HEURISTIC and BRANCH are modules to be specified later. \mathcal{X} stores auxiliary information accompanying a STPG and will be specified later.

Input: STPG $\mathcal{G}_{\text{root}}^S$
Output: TPG \mathcal{G}

- 1 $(h_{\text{root}}, \mathcal{X}_{\text{init}}) \leftarrow \text{HEURISTIC}(\mathcal{G}_{\text{root}}^S, \mathcal{X}_{\text{init}})$;
- 2 $\mathcal{Q} \leftarrow \{(\mathcal{G}_{\text{root}}^S, \mathcal{X}_{\text{init}}, 0, h_{\text{root}})\}$; // A priority queue
- 3 **while** \mathcal{Q} is not empty **do**
- 4 $(\mathcal{G}^S, \mathcal{X}, g, h) \leftarrow \mathcal{Q}.\text{pop}()$;
- 5 $(\mathcal{X}', g_{\Delta}, (v_{k+1}^i, v_s^j)) \leftarrow \text{BRANCH}(\mathcal{G}^S, \mathcal{X})$;
- 6 **if** $(v_{k+1}^i, v_s^j) = \text{NULL}$ **then return** \mathcal{G}^S ;
- 7 $\mathcal{G}_f^S \leftarrow \text{fix}(\mathcal{G}^S, (v_{k+1}^i, v_s^j))$;
- 8 **if not** $\text{CYCLEDETECTION}(\mathcal{G}_f^S, (v_{k+1}^i, v_s^j))$ **then**
- 9 $(h_f, \mathcal{X}_f) \leftarrow \text{HEURISTIC}(\mathcal{G}_f^S, \mathcal{X}')$;
- 10 $\mathcal{Q}.\text{push}((\mathcal{G}_f^S, \mathcal{X}_f, g + g_{\Delta}, h_f))$;
- 11 $\mathcal{G}_r^S \leftarrow \text{reverse}(\mathcal{G}^S, (v_{k+1}^i, v_s^j))$;
- 12 **if not** $\text{CYCLEDETECTION}(\mathcal{G}_r^S, (v_{s+1}^j, v_k^i))$ **then**
- 13 $(h_r, \mathcal{X}_r) \leftarrow \text{HEURISTIC}(\mathcal{G}_r^S, \mathcal{X}')$;
- 14 $\mathcal{Q}.\text{push}((\mathcal{G}_r^S, \mathcal{X}_r, g + g_{\Delta}, h_r))$;
- 15 **throw exception** “No solution found”;
- 16 **Function** $\text{CYCLEDETECTION}(\mathcal{G}^S, (u, v))$
- 17 Run depth-first search (DFS) from v on $\text{red}(\mathcal{G}^S)$;
- 18 **if** DFS visits vertex u **then return** true;
- 19 **return** false;

Lemma 6. *The partial cost of a STPG \mathcal{G}^S is no greater than the cost of any \mathcal{G}^S -producible TPG.*

Proof. Let \mathcal{G} be a \mathcal{G}^S -producible TPG. Consider running Procedure 1 on \mathcal{G} and $\text{red}(\mathcal{G}^S)$, respectively. Since an edge appears in $\text{red}(\mathcal{G}^S)$ must appear in \mathcal{G} , we can inductively show that, in any call to $\text{STEP}_{\text{EXEC}}$, if a vertex v can be marked as satisfied in \mathcal{G} , then it can be marked as satisfied in $\text{red}(\mathcal{G}^S)$. Thus, the total timesteps to satisfy all vertices in $\text{red}(\mathcal{G}^S)$ cannot exceed that in \mathcal{G} . \square

As shown in Algorithm 2, SES runs A* in the space of STPGs with a root node corresponding to the STPG $\mathcal{G}_{\text{root}}^S$ constructed as in Construction 1. The priority queue \mathcal{Q} sorts its nodes by their f -values (namely $g + h$). The f -value of a node is defined as the partial cost of its STPG. When expanding a node, SES selects one switchable edge in the STPG by module BRANCH and generates two child nodes with the selected edge being *fixed* or *reversed*. We abuse the operators *fix* and *reverse* on Lines 7 and 11 to take a STPG and a switchable edge as input and return a new STPG.

SES uses function CYCLEDETECTION to prune child nodes with STPGs that definitely produce cyclic TPGs, namely STPGs whose reduced TPGs are cyclic. Specifically, $\text{CYCLEDETECTION}(\mathcal{G}^S, (u, v))$ returns true iff $\text{red}(\mathcal{G}^S)$ contains a cycle involving edge (u, v) . As $\mathcal{G}_{\text{root}}^S$ is acyclic, it holds inductively that $\text{CYCLEDETECTION}(\mathcal{G}, (u, v))$ returns true iff $\text{red}(\mathcal{G}^S)$ contains *any* cycle. This is because, when

we generate a node, we add only one new non-switchable edge, so any cycle formed must contain the new edge.

Assumption 1. The modules in SES satisfy:

- A1 BRANCH($\mathcal{G}^S, \mathcal{X}$) outputs an updated auxiliary information \mathcal{X}' , a value g_{Δ} , and a switchable edge of \mathcal{G}^S if exists or NULL otherwise.
- A2 HEURISTIC($\mathcal{G}^S, \mathcal{X}$) computes a value h such that $g + h$ is the partial cost of \mathcal{G}^S for every node $(\mathcal{G}^S, \mathcal{X}, g, h) \in \mathcal{Q}$.

Theorem 7 (Completeness and Optimality). *Under Assumption 1, SES always finds a deadlock-free TPG \mathcal{G} with $\text{cost}(\mathcal{G}) = \min\{\text{cost}(\mathcal{G}) : \mathcal{G} \text{ is } \mathcal{G}_{\text{root}}^S\text{-producible}\}$.*

Proof. First, SES always terminates within a finite time because there are only finitely many possible operation sequences from $\mathcal{G}_{\text{root}}^S$ to any TPG, each corresponding to a node that can possibly be added to \mathcal{Q} . Second, Theorem 5 ensures that there always exist solutions for SES since $\mathcal{G}_{\text{root}}^S$ is constructed as in Construction 1. Therefore, to prove the completeness of SES, we just need to prove the following claim: At the beginning of each while-loop iteration, for any deadlock-free $\mathcal{G}_{\text{root}}^S$ -producible TPG \mathcal{G} , there exists $\mathcal{G}^S \in \mathcal{Q}$ such that \mathcal{G} is \mathcal{G}^S -producible. Here, we abuse the notation $\mathcal{G}^S \in \mathcal{Q}$ to denote a node in \mathcal{Q} with STPG \mathcal{G}^S . This claim holds inductively: At the first iteration, $\mathcal{G}_{\text{root}}^S \in \mathcal{Q}$. During any iteration, if some $\mathcal{G}^S \in \mathcal{Q}$ such that \mathcal{G} is \mathcal{G}^S -producible is popped on Line 4, then one of the following must hold:

- \mathcal{G}^S contains no switchable edge, i.e., $\mathcal{G}^S = \mathcal{G}$: SES terminates, and the inductive step holds vacuously.
- \mathcal{G} is \mathcal{G}_f^S -producible: Since \mathcal{G} is acyclic, so is $\text{red}(\mathcal{G}_f^S)$. Thus, \mathcal{G}_f^S is added into \mathcal{Q} .
- \mathcal{G} is \mathcal{G}_r^S -producible: This is symmetric to the above case.

In any case, the claim remains true after this iteration. Therefore, SES always outputs a solution within a finite time.

Finally, we prove that the output TPG \mathcal{G} has the minimum cost. Assume towards contradiction that when \mathcal{G} is returned, there exists $\mathcal{G}_0^S \in \mathcal{Q}$ that can produce a better TPG $\mathcal{G}_{\text{better}}$ with $\text{cost}(\mathcal{G}_{\text{better}}) < \text{cost}(\mathcal{G})$. Yet this is impossible since Lemma 6 implies that such \mathcal{G}_0^S must have a smaller $g + h$ value and thus would be popped from \mathcal{Q} before \mathcal{G} . \square

Execution-based Modules

In this and the next sections, we describe two sets of modules and prove that they satisfy Assumption 1. We start with describing a set of “execution-based” modules in Module 3 and refer to it as *Execution-based SES* (ESES).

In essence, ESES simulates the execution of the STPG and branches when encountering a switchable edge. It uses \mathcal{X} to record the index of the most recently satisfied vertex for every agent, indicating their current locations. \mathcal{X} is updated by the BRANCH module, which largely ensembles EXEC in Procedure 1. At the beginning of each while-loop iteration of BRANCH, ESES first checks whether the next vertex of any agent is involved in a switchable edge and, if so, returns that edge together with the updated \mathcal{X}' and the cost of moving agents from the old \mathcal{X} to the new \mathcal{X}' [Lines 8 to 11],

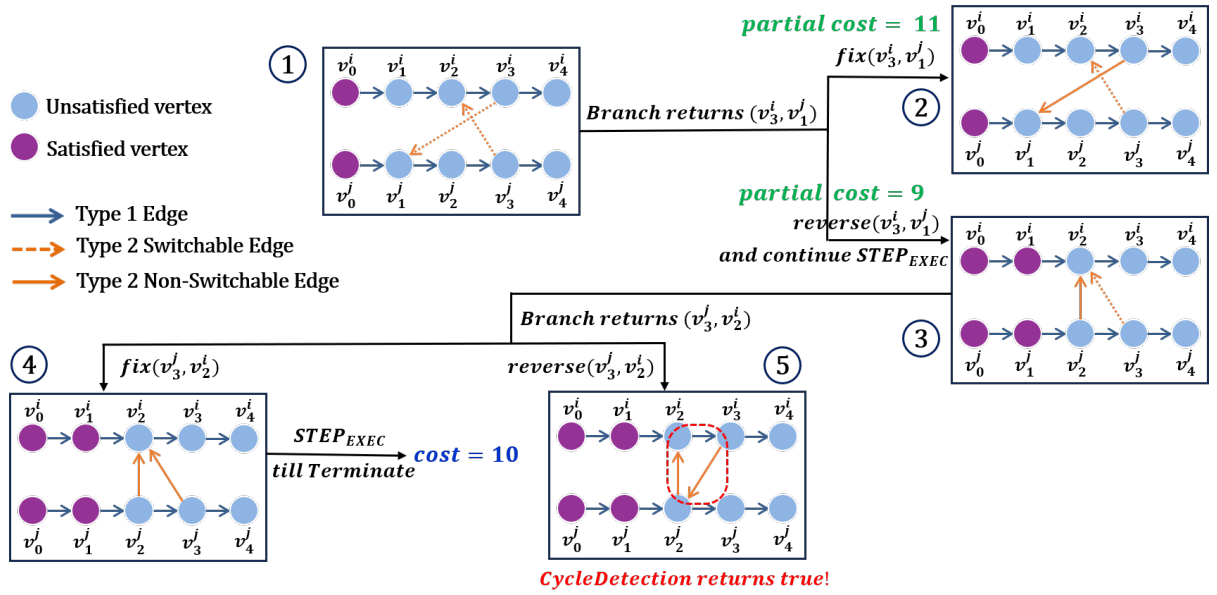


Figure 3: Example of running ESES on the top-left STPG. The circled numbers denote the order of generating these STPGs.

Module 3: Execution-based Modules for ESES

- 1 Auxillary information \mathcal{X} is a map $\mathcal{X} : \mathcal{A} \rightarrow [0 : z_i]$, where $\mathcal{X}[i]$ records the index of the most recently satisfied vertex for agent i ;
- 2 \mathcal{X}_{init} on Line 1 of Algorithm 2 maps all agents to 0;
- 3 $INIT_{EXEC}$ in Procedure 1 is modified by setting \mathcal{V}_0 to $\{v_k^i : i \in \mathcal{A}, k \leq \mathcal{X}[i]\}$ instead of $\{v_0^i : i \in \mathcal{A}\}$;
- 4 **Function** $BRANCH(\mathcal{G}^S = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{E_2}, \mathcal{N}_{E_2})), \mathcal{X})$

```

5    $INIT_{EXEC}(\mathcal{V}, \mathcal{X});$ 
6    $\mathcal{X}' \leftarrow \mathcal{X};$ 
7   while  $\exists v \in \mathcal{V} : v$  unsatisfied do
8     forall  $i \in \mathcal{A} : \mathcal{X}'[i] < z_i$  do
9        $v \leftarrow v_{\mathcal{X}'[i]+1}^i$ ; // First unsatisfied vertex
10      if  $\exists e \in \mathcal{S}_{E_2}$  with  $e = (v, u)$  or  $(u, v)$  then
11        return  $(\mathcal{X}', cost, e)$ ;
12       $\mathcal{S} \leftarrow STEP_{EXEC}(red(\mathcal{G}^S));$ 
13      forall  $v_k^i \in \mathcal{S}$  do
14        Mark  $v_k^i$  as satisfied;
15         $\mathcal{X}'[i] \leftarrow k$ ;
16    return  $(\mathcal{X}', cost, NULL)$ ;

```
- 17 **Function** $HEURISTIC(\mathcal{G}^S, \mathcal{X})$
- 18 **return** $(EXEC(red(\mathcal{G}^S)), \mathcal{X})$;

where the cost is updated inside function $STEP_{EXEC}$. If no such edge is found, it runs $STEP_{EXEC}$ on the reduced TPG to move agents forward by one timestep and repeat the process.

Example 3. Figure 3 shows an example of ESES.¹ We

¹We note that Figure 3 also works as an example for the GSES implementation in the following section. The only difference is that

start with the top-left STPG ① containing two switchable edges. ESES looks at a “horizon” containing the first unsatisfied vertices v_1^i and v_1^j , and then picks the adjacent switchable edge (v_3^i, v_1^j) to branch on. This leads to two copies of STPGs ② and ③, containing non-switchable edge (v_3^i, v_1^j) or (v_2^j, v_2^i) , respectively. ESES expands on STPG ③ first as it has a smaller $g + h$ value. The next switchable edge it encounters is (v_3^j, v_2^i) . ESES first *fixes* it and generates TPG ④ with $cost = 10$, which is the optimal solution. When ESES *reverses* the edge, the resulting TPG ⑤ is pruned as it contains a cycle. Note that STPG ② will not be expanded since it has a partial cost greater than the cost of TPG ④.

Proposition 8. *Module 3 satisfies Assumption 1.*

Proof. Assumption A1 holds by design. To prove Assumption A2, we first prove the following claim by induction: for every node $(\mathcal{G}^S, \mathcal{X}, g, h) \in \mathcal{Q}$, g is the cost of moving agents from their start locations to \mathcal{X} . This holds for the root node with $g = 0$ and $\mathcal{X}[i] = 0, \forall i \in \mathcal{A}$. When we expand a node $(\mathcal{G}, \mathcal{X}, g, h) \in \mathcal{Q}$, g_Δ returned by $BRANCH$ is the cost of moving agents from \mathcal{X} to \mathcal{X}' . Thus, on Lines 10 and 14 of Algorithm 2, the g value of the child nodes are $g + g'$, which is the cost of moving agents from their start locations to \mathcal{X}' ($= \mathcal{X}_f = \mathcal{X}_r$). So our claim holds. Module $HEURISTIC$ runs function $EXEC$ to compute the cost of moving agents from \mathcal{X}' to their goal locations on the reduced TPG, making $g + h$ the partial cost of \mathcal{G}^S for every node $(\mathcal{G}^S, \mathcal{X}, g, h) \in \mathcal{Q}$. \square

Graph-based Modules

We now introduce an alternative set of modules that focus on the graph properties of a TPG. We refer to this implementation as *Graph-based SES* (GSES). We will see later in our

GSES does not use the notion of “(un)satisfied vertex” or $STEP_{EXEC}$.

experiment that this shift of focus significantly improves the efficiency of SES. We start by presenting the following crucial theorem that provides a graph-based approach to computing the cost of a TPG.

Given a TPG \mathcal{G} and a vertex $v \in \mathcal{V}$, let $lp(v)$ denote the longest path among the longest paths from every vertex $v_0^i, i \in \mathcal{A}$ to vertex v on \mathcal{G} and $|lp(v)|$ denote its length.

Theorem 9. *When we execute a TPG, every vertex v is marked as satisfied in the $|lp(v)|^{\text{th}}$ iteration of the while-loop of EXEC in Procedure 1.*

Proof. We induct on iteration t and prove that all vertices v with $|lp(v)| = t$ are marked as satisfied in the t^{th} iteration. In the base case, $\{v_0^i, i \in \mathcal{A}\}$ are the vertices with $|lp(v)| = 0$ and are marked as satisfied in the 0^{th} iteration. In the inductive step, we assume that, by the end of the $(t-1)^{\text{th}}$ iteration, all vertices v with $|lp(v)| < t$ are satisfied, and all vertices v with $|lp(v)| \geq t$ are unsatisfied. Then, in the t^{th} iteration, every vertex v with $|lp(v)| = t$ is marked as satisfied because all of its in-neighbors v' have $|lp(v')| < |lp(v)| = t$ and are thus satisfied. For every vertex v with $|lp(v)| > t$, the vertex right before v on $lp(v)$, denoted as v' , has $|lp(v')| = |lp(v)| - 1 \geq t$ and is thus unsatisfied on Line 17. Thus, every vertex v with $|lp(v)| > t$ has at least one in-neighbor unsatisfied and thus remains unsatisfied. Therefore, the theorem holds. \square

Hence, the last vertex $v_{z_i}^i$ of every agent $i \in \mathcal{A}$ is marked as satisfied in the $|lp(v_{z_i}^i)|^{\text{th}}$ iteration, namely the travel time of agent i is $|lp(v_{z_i}^i)|$. We thus get the following corollary.

Corollary 10. *Given a TPG \mathcal{G} , $cost(\mathcal{G}) = \sum_{i \in \mathcal{A}} |lp(v_{z_i}^i)|$.*

An interesting observation is that, if $lp(u) < lp(v)$ for a given TPG \mathcal{G} , then adding edge (u, v) to \mathcal{G} does not change its cost since adding (u, v) does not change any longest paths. We thus get the following corollary that is useful later.

Corollary 11. *Given a STPG \mathcal{G}^S , we compute $lp(v)$ on $red(\mathcal{G}^S)$. For any switchable edge with $lp(u) < lp(v)$, fixing it does not change the partial cost of \mathcal{G}^S .*

We adopt the following well-known algorithm to compute $lp(v)$ on a given deadlock-free TPG $\mathcal{G} = (\mathcal{V}, \mathcal{E}_1, \mathcal{E}_2)$: (1) Set $lp(v) = 0, \forall v \in \mathcal{V}$. (2) Compute a topological sort of all vertices in \mathcal{V} . (3) For each vertex v in the topological order, we set $lp(v) = \max\{lp(u), lp(v) + 1\}$ for every out-neighbor u (namely $(v, u) \in \mathcal{E}_1 \cup \mathcal{E}_2$). The time complexity of this longest-path algorithm is $\mathcal{O}(|\mathcal{V}| + |\mathcal{E}_1 \cup \mathcal{E}_2|)$.

With this algorithm, we specify the graph-based modules in Module 4. In GSES, \mathcal{X} records $lp(v)$ for every vertex $v \in \mathcal{V}$ and is updated by the HEURISTIC module. Since, with \mathcal{X} , HEURISTIC can directly compute the partial cost of a given STPG, GSES does not use any g values. The BRANCH module chooses a switchable edge (u, v) with $\mathcal{X}[u] \geq \mathcal{X}[v]$ to branch on. If no such edge exists, then, by Corollary 11, fixing all switchable edges produces a TPG with the same cost as the current partial cost. Thus, GSES *fix* all such edges and terminates in this case.

Proposition 12. *Module 4 satisfies Assumption 1.*

Module 4: Graph-based Modules for GSES

- 1 Auxillary information \mathcal{X} is a map $\mathcal{X} : \mathcal{V} \rightarrow [0, |\mathcal{V}|]$, which records $lp(v)$ for every vertex v ;
 - 2 $\mathcal{X}_{\text{init}}$ on Line 1 of Algorithm 2 is empty;
 - 3 **Function** BRANCH($\mathcal{G}^S = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{\mathcal{E}_2}, \mathcal{N}_{\mathcal{E}_2}), \mathcal{X})$)
 - 4 **if** $\exists (u, v) \in \mathcal{S}_{\mathcal{E}_2} : \mathcal{X}[u] \geq \mathcal{X}[v]$ **then**
 - 5 **return** $(\mathcal{X}, 0, (u, v))$;
 - 6 *fix* all switchable edges in \mathcal{G}^S ;
 - 7 **return** $(\mathcal{X}, 0, \text{NULL})$;
 - 8 **Function** HEURISTIC($\mathcal{G}^S, \mathcal{X}$)
 - 9 $\mathcal{X}' \leftarrow lp$ -values of all vertices in \mathcal{V} on $red(\mathcal{G}^S)$;
 - 10 **return** $(\sum_{i \in \mathcal{A}} \mathcal{X}'[v_{z_i}^i], \mathcal{X}')$;
-

Proof. Assumption A1 holds by design. Assumption A2 holds because of Corollary 10. \square

Remark 5. ESES terminates when all vertices are satisfied, which is possible only when all Type 2 edges are non-switchable. This means that ESES has to expand on all switchable edges before getting a solution. In contrast, GSES can have an *early termination* when *fixing* all switchable edges does not change any longest paths.

Experiment

We use 4 maps from the MAPF benchmark suite (Stern et al. 2019), shown in Figure 4, with 6 agent group sizes per map. Regarding each map and group size configuration, we run the algorithms on 25 different, evenly distributed scenarios (start/goal locations) with 6 trials per scenario. We set a runtime limit of 90 seconds for each trial. In each trial, we execute the TPG constructed from an optimal MAPF solution planned by a k-Robust MAPF solver k-Robust CBS (Chen et al. 2021) with $k = 1$. At each timestep of the execution, each agent that has not reached its goal location is subject to a constant probability p of delay. When a delay happens, we draw a delay length Δ uniformly random from a range $[10, 20]$, construct a STPG as in Construction 1, and run our replanning algorithms. We also develop a baseline that uses k-Robust CBS to find the new optimal solution (that takes into account the delay length Δ) when the delay happens.

We implement all algorithms in C++² and run experiments on a server with a 64-core AMD Ryzen Threadripper 3990X and 192 GB of RAM.

Efficiency Figure 4 compares the runtime of ESES and GSES with replanning using k-Robust CBS. In all cases, GSES runs the fastest. Most remarkably, on the random and warehouse maps, the runtime of GSES is consistently below 1 second and does not increase significantly when the number of agents increases, suggesting the potential of GSES for real-time replanning applications.

²Our SES code is available at <https://github.com/YinggggFeng/Switchable-Edge-Search>. The modified k-Robust CBS code that considers delays is available at <https://github.com/nobodyczcz/Lazy-Train-and-K-CBS/tree/wait-on-start>.

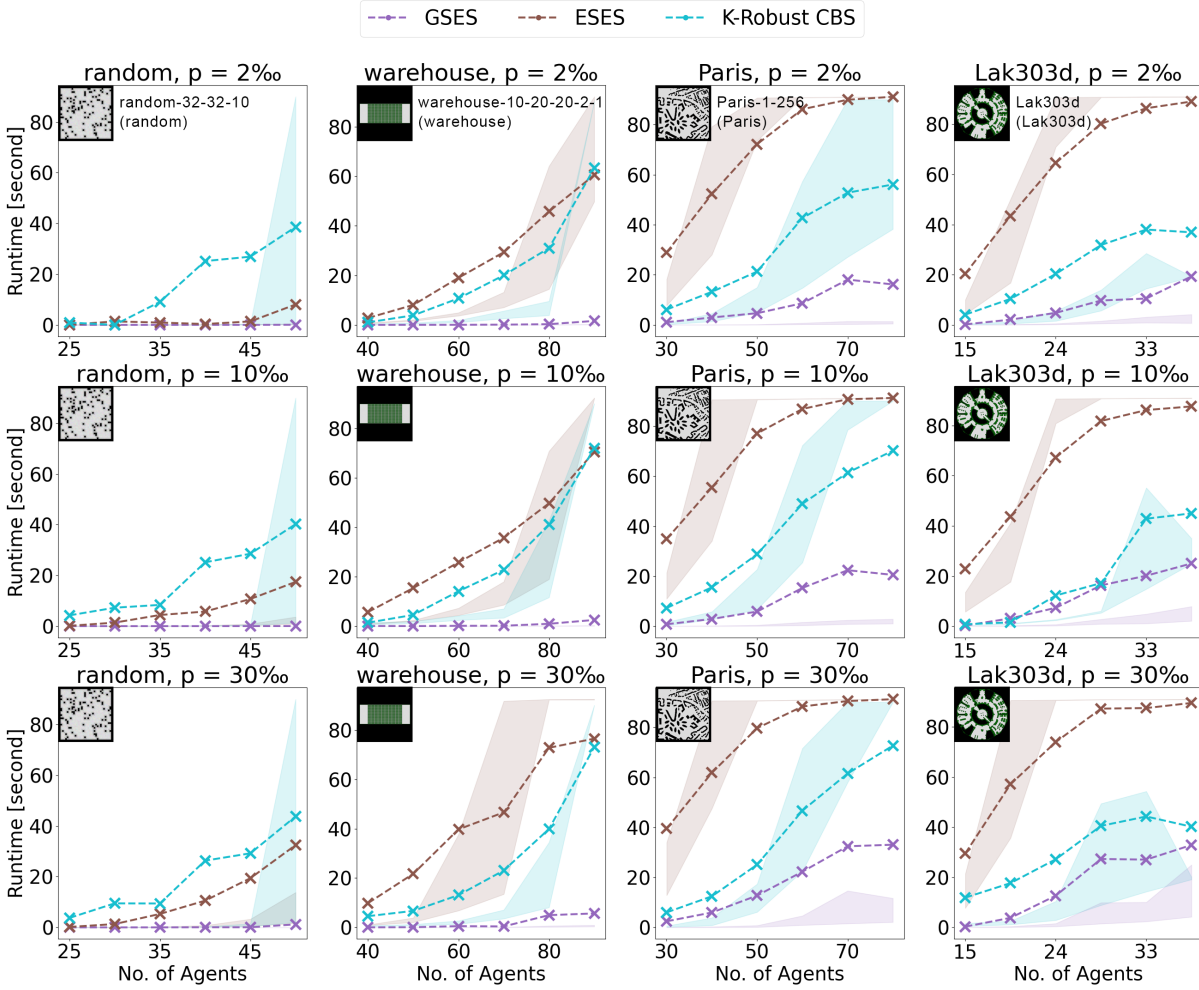


Figure 4: Runtime of ESES, GSES, and k-robust CBS. The dashed lines represent the mean of runtime, and the shaded areas denote the 0.4 to 0.6 quantile range. For trials that exceed the 90-second time limit, we count it as 90 seconds.

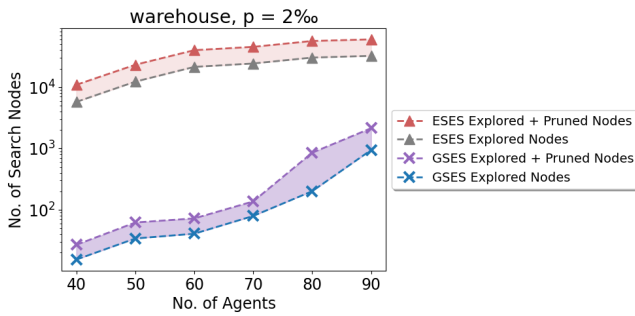


Figure 5: Numbers of nodes explored and pruned by ESES and GSES on the warehouse map. Dashed lines represent the mean values. Shaded area between two lines for the same algorithm indicates the portion of pruned nodes.

Comparing ESES and GSES We observe from Figure 4 that, although ESES and GSES adopt the same framework,

GSES runs significantly faster than ESES. This is because the longest paths used in GSES defines a simple but extremely powerful early termination condition (see Remark 5), which enables GSES to find an optimal solution after a very small number of node expansions. Figure 5 compares the number of search nodes of ESES and GSES, where *explored nodes* are nodes popped from the priority queue, and *pruned nodes* are nodes pruned by CYCLEDETECTION. The gap between the red and grey lines (and the gap between the purple and blue lines) indicates the effectiveness of cycle detection for pruning unnecessary search nodes. The gap between the grey and blue lines indicates the effectiveness of early termination as described in Remark 5.

Improvement of Solution Cost Figure 6 measures the cost of our replanned solution, in comparison to the non-replanned solution produced by the original TPG and the replanned solution produced by k-Robust CBS. We stress that our solution is guaranteed to be optimal, as proven in previous sections, upon sticking to the original location-wise




Figure 6: Mean costs (from the locations where the delays happened to the goal locations) of the non-replanned, GSES-replanned, and K-Robust-CBS-replanned solutions. The means are taken across all trials for all different numbers of agents.

paths, while k-Robust CBS finds an optimal solution that is independent of the original paths. Therefore, the two algorithms solve intrinsically different problems, and the results here serve primarily for a quantitative understanding of how much improvement we can get by changing only the passing orders of agents at different locations. Figure 6 shows that the cost improvement depends heavily on the maps. For example, our solutions have costs similar to the globally optimal solutions on the random map, while the difference is larger on the Lak303d map.

Conclusion

We proposed Switchable Edge Search to find the optimal passing orders for agents that are planned to visit the same location. We developed two implementations based on either execution (ESES) or graph (GESE) presentations. On the random and warehouse maps, the average runtime of GSES is faster than 1 second for various numbers of agents. On harder maps (Paris and game maps), it also runs faster than replanning with a k-Robust CBS algorithm.

Appendix: Proofs for Section 3

We rely on the following lemma to prove Proposition 1.

Lemma A. *For every tuple (l_k^i, t_k^i) in every path $p_i \in \mathcal{P}$, the corresponding vertex v_k^i in \mathcal{G} is satisfied after the $(t_k^i)^{\text{th}}$ iteration of the while-loop of EXEC in Procedure 1.*

Proof. We induct on the while-loop iteration t and prove that any vertex v_k^i with $t_k^i \leq t$ is satisfied after iteration t . When $t = 0$, this holds because of INIT_{EXEC}. Assume that all vertices v_k^i with $t_k^i \leq t - 1$ are satisfied after iteration $t - 1$. At iteration t , for Type 1 edge (v_{k-1}^i, v_k^i) , v_{k-1}^i is satisfied as $t_{k-1}^i < t_k^i \leq t$. For any Type 2 edge (v_s^j, v_k^i) , v_s^j is also satisfied as $t_s^j < t_k^i \leq t$ by Definition 3. This shows that all in-neighbors of v_k^i are satisfied after iteration $t - 1$, thus v_k^i is satisfied after iteration t . \square

Proposition 1 (Cost). $cost(\mathcal{G}) \leq cost(\mathcal{P})$.

Proof. Lemma A implies that the last vertex $v_{z_i}^i$ of every agent i is satisfied after the $(t_{z_i}^i)^{\text{th}}$ iteration, i.e., when executing \mathcal{G} , the travel time of every agent i is no greater than $t_{z_i}^i$. Thus, the proposition holds. \square

Proposition 2 is similar to Lemma 4 in (Hönig et al. 2016) with different terms. We include a proof for completeness.

Proposition 2 (Collision-Free). *Executing a TPG does not lead to collisions among agents.*

Proof. According to Procedure 1, we need to show that, when a vertex v_k^i is marked as satisfied [Line 18], moving agent i to its k^{th} location l_k^i does not lead to collisions. Assume towards contradiction that agent i indeed collides with another agent j at timestep t . Let v_k^i and v_s^j be the latest satisfied vertices for agents i and j , respectively, after the t^{th} iteration of the while-loop. If i and j collide because they are at the same location, then $l_k^i = l_s^j$, indicating that either edge (v_{k+1}^i, v_s^j) or edge (v_{s+1}^j, v_k^i) is in \mathcal{E}_2 . But this is impossible as neither v_{k+1}^i nor v_{s+1}^j is satisfied.

If they collide because j leaves a location at timestep t , and i enters the same location at timestep t , then v_k^i and v_s^j are marked as satisfied exactly at the t^{th} iteration with $l_k^i = l_{s-1}^j$, indicating that either (v_k^i, v_s^j) or (v_{s+1}^j, v_{k-1}^i) is in \mathcal{E}_2 . But this is also impossible as neither v_k^i nor v_{s+1}^j was satisfied before the t^{th} iteration. \square

Lemma 3 (Deadlock \iff Cycle). *Executing a TPG encounters a deadlock iff the TPG contains cycles.*

Proof. If a TPG \mathcal{G} has a cycle, executing it will encounter a deadlock as no vertices in the cycle can be marked as satisfied. If executing \mathcal{G} encounters a deadlock in the t^{th} iteration of the while-loop, we prove that \mathcal{G} has a cycle by contradiction. Let \mathcal{V}' denote the set of unsatisfied vertices, which is non-empty by Definition 4. If \mathcal{G} is acyclic, then there exists a topological ordering of \mathcal{V}' , and \mathcal{S} must contain the first vertex in the topological ordering as all of its in-neighbors are satisfied, contradicting the deadlock condition of $\mathcal{S} = \emptyset$. \square

Lemma 4 (Deadlock-Free). *If a TPG is constructed from a MAPF solution, then executing it is deadlock-free.*

Proof. If a deadlock is encountered, then the execution would enter the while-loop for infinitely many iterations, and $cost$ strictly increases in each iteration. Thus, $cost(\mathcal{G}) = \infty$. Yet, $cost(\mathcal{P})$ is finite, contradicting Proposition 1. \square

Acknowledgements

The research at Carnegie Mellon University was supported by the National Science Foundation (NSF) under Grant 2328671. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the sponsoring organizations, agencies, or the U.S. government.

References

- Atzmon, D.; Stern, R.; Felner, A.; Sturtevant, N. R.; and Koenig, S. 2020. Probabilistic Robust Multi-Agent Path Finding. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 29–37.
- Atzmon, D.; Stern, R.; Felner, A.; Wagner, G.; Barták, R.; and Zhou, N. 2018. Robust Multi-Agent Path Finding. In *Proceedings of the International Symposium on Combinatorial Search*, 2–9.
- Berndt, A.; van Duijkeren, N.; Palmieri, L.; and Keviczky, T. 2020. A Feedback Scheme to Reorder A Multi-Agent Execution Schedule by Persistently Optimizing a Switchable Action Dependency Graph. *ArXiv*.
- Chen, Z.; Harabor, D. D.; Li, J.; and Stuckey, P. J. 2021. Symmetry Breaking for k-Robust Multi-Agent Path Finding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 12267–12274.
- Hönig, W.; Kiesel, S.; Tinka, A.; Durham, J. W.; and Ayanian, N. 2019. Persistent and Robust Execution of MAPF Schedules in Warehouses. *IEEE Robotics and Automation Letters*, 1125–1131.
- Hönig, W.; Kumar, T. K. S.; Cohen, L.; Ma, H.; Xu, H.; Ayanian, N.; and Koenig, S. 2016. Multi-Agent Path Finding with Kinematic Constraints. In *Proceedings of the International Conference on Automated Planning and Scheduling*, 477–485.
- Hönig, W.; Preiss, J. A.; Kumar, T. K. S.; Sukhatme, G. S.; and Ayanian, N. 2018. Trajectory Planning for Quadrotor Swarms. *IEEE Transactions on Robotics*, 856–869.
- Kou, N. M.; Peng, C.; Ma, H.; Kumar, T. K. S.; and Koenig, S. 2020. Idle Time Optimization for Target Assignment and Path Finding in Sortation Centers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 9925–9932.
- Ma, H.; Kumar, T. S.; and Koenig, S. 2017. Multi-Agent Path Finding with Delay Probabilities. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3605–3612.
- Mannucci, A.; Pallottino, L.; and Pecora, F. 2021. On Provably Safe and Live Multirobot Coordination With Online Goal Posting. *IEEE Transactions on Robotics*, 37(6): 1973–1991.
- Silver, D. 2005. Cooperative Pathfinding. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 117–122.
- Stern, R.; Sturtevant, N. R.; Felner, A.; Koenig, S.; Ma, H.; Walker, T. T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T. K. S.; Boyarski, E.; and Bartak, R. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. In *Proceedings of the International Symposium on Combinatorial Search*, 151–159.
- Su, Y.; Veerapaneni, R.; and Li, J. 2024. Bidirectional Temporal Plan Graph: Enabling Switchable Passing Orders for More Efficient Multi-Agent Path Finding Plan Execution. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Wurman, P. R.; D’Andrea, R.; and Mountz, M. 2007. Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1752–1759.