

# BTPG-max: Achieving Local Maximal Bidirectional Pairs for Bidirectional Temporal Plan Graphs

Yifan Su, Rishi Veerapaneni, Jiaoyang Li

Carnegie Mellon University  
yifansu2003@gmail.com, {vrishi, jiaoyangli}@cmu.edu

## Abstract

Multi-Agent Path Finding (MAPF) requires computing collision-free paths for multiple agents in a shared environment. Most MAPF planners assume that each agent reaches a specific location at a specific timestep, but this is infeasible to directly follow on real systems where delays often occur. To address collisions caused by agents deviating due to delays, the Temporal Plan Graph (TPG) was proposed, which converts a MAPF time-dependent solution into a time-independent solution with a set of inter-agent dependencies. Recently, a Bidirectional TPG (BTPG) was proposed which relaxed some dependencies into “bidirectional pairs” and improved efficiency of agents executing their MAPF solution with delays. Our work improves upon this prior work by designing an algorithm, BTPG-max, that finds more bidirectional pairs. Our main theoretical contribution is in designing the BTPG-max algorithm that is locally maximal, i.e., it constructs a BTPG where no additional bidirectional pairs can be added. We also show how, in practice, BTPG-max leads to BTPGs with significantly more bidirectional edges, superior anytime behavior, and improved robustness to delays.

## 1 Introduction

Multi-Agent Path Finding (MAPF) is the problem of finding collision-free paths for a set of agents in a shared workspace. MAPF is one core component of intelligent multi-agent teams and has direct applications in systems with many robots like warehouse management.

Solving an MAPF instance, in theory, gives a MAPF solution which is a set of paths for robots to follow. This assumes that agents can strictly follow the paths, requiring that they arrive at specific locations at specific timesteps. However, in practice, this is impossible as robots may be delayed during execution (e.g., due to kinematic constraints or wheel slippage). If agents naively follow their original paths without accounting for delays and other temporal differences, they may collide with other agents or get stuck in deadlocks.

To handle temporal differences in execution, Höning et al. (2016) introduce a Temporal Plan Graph (TPG). The main innovation is that a MAPF solution can be processed into a set of dependencies. In particular, agents are only dependent on each other when their paths spatially intersect (e.g.,




Figure 1: Example of TPG and BTPG for a MAPF solution that requires the green agent to pass through E first. The TPG requires blue to wait for green (even if green gets delays) while the BTPG allows either agent to cross the intersection (thus leading to more robustness to delays).

location E in Figure 1). At these intersections, denoted as conflict locations, the inter-agent dependency is that agents must maintain their relative passing order specified in the MAPF solution. For example, in Figure 1, the green agent should pass location E first. During execution, agents can go as fast/slow as long as they satisfy this inter-agent dependency. The TPG paper proves how following these dependencies, regardless of time, will still lead to collision-free and deadlock-free execution.

However, a TPG can be overly restrictive. In Figure 1, during execution, if the green agent encounters a delay, the TPG would make the blue agent wait for the green agent to pass through location E first. Consequently, a delay of the green agent also leads to a delay for the blue agent. To address this issue, the Bidirectional Temporal Plan Graph (BTPG) (Su, Veerapaneni, and Li 2024) is introduced, allowing the blue agent to choose to pass through E first in such situations to avoid unnecessary waiting.

BTPG introduces *bidirectional pairs* (or “bi-pairs”), which are pairs of opposite dependency edges at a conflict location that allow either agent to go first without causing deadlock. In the original BTPG construction, each dependency in the TPG is tested for reversibility: if reversing an edge does not introduce a deadlock, then the agent can safely pass through the conflict point first—even if the MAPF solution originally scheduled it to go later. These reversible edges form bi-pairs. Although any cycle in a TPG indicates deadlock, the BTPG paper observed that some cycles in the BTPG graph would not lead to deadlock. Leveraging this, they proposed *BTPG-o*, which applies hand-crafted rules to detect and exclude such non-deadlocking cycles to find more

bi-pairs.

In this work, we point out that BTPG-o still misses certain non-deadlock cycle types. To address this, we propose *BTPG-max* identifies non-deadlock cycles to construct a locally maximally BTPG (a BTPG where no additional bidirectional pairs can be added). Our experiments show that BTPG-max finds more bi-pairs, yielding faster execution times and markedly better anytime performance.

## 2 Preliminary

### 2.1 Problem Formulation

The MAPF problem takes as input a graph and a set of start and goal locations for a group of agents. The MAPF solver provides each agent with a collision-free path, allowing them to move simultaneously to their goal locations on the graph. In the standard MAPF set-up, time is discretized into timesteps, with solution paths dictating that agents need to be at specific locations at specific timesteps. However, in actual execution, robots cannot meet this requirement due to hardware or communication issues that can cause delays. If agents followed the paths without adjusting for delays, they could collide with each other. We are thus interested in determining how agents should handle delays during execution.

Our work directly builds on top of TPG (Hoenig et al. 2016) and BTPG (Su, Veerapaneni, and Li 2024). Constructing a TPG or BTPG to handle delays is unique in that it is a post-processing technique on top of the MAPF solution that is done *before* execution. This eliminates the need to replan during execution (avoiding additional delays caused by extra computation during execution), as agents only need to follow the dependencies specified by the (B)TPG.

### 2.2 Temporal Plan Graph (TPG)

As discussed in the introduction, TPG is a way to interpret a MAPF solution as a set of dependencies (Hoenig et al. 2016). We now formally define it.

**Definition 1 (TPG).** TPG is a directed acyclic graph  $G = (V, E)$ , where each vertex  $v_i^m \in V$  corresponds to a state of agent  $m$  being at location  $\text{loc}(v_i^m)$ . The index  $i$  indicates that  $\text{loc}(v_i^m)$  is the  $i$ -th location on the path of agent  $m$ . The edges in  $E$  define the precedence dependencies between states and are categorized into two types: type-1 edges  $E_1$  and type-2 edges  $E_2$ . A type-1 edge  $(v_i^m, v_{i+1}^m) \in E_1$  enforces that agent  $m$  must reach  $\text{loc}(v_i^m)$  before moving to  $\text{loc}(v_{i+1}^m)$ . A type-2 edge  $(v_i^m, v_j^n) \in E_2$  enforces that agent  $n$  can only reach  $\text{loc}(v_j^n)$  after or at the same time as agent  $m$  reaches  $\text{loc}(v_i^m)$ .

Conceptually, if all agents had paths that never intersected, the TPG would consist of only type-1 edges. A type-2 edge  $(v_i^m, v_j^n)$  occurs when the paths of both agents  $n$  and  $m$  need to pass through conflict location  $\text{loc}(v_j^n) = \text{loc}(v_{i-1}^m)$  and  $m$  should pass it first.

**TPG execution policy** To move to the next state, an agent must satisfy all dependencies represented by all edges pointing to that state.

For example, the TPG in Figure 1 has a type-2 edge from  $F$  to  $E$ . During execution, the blue agent can only enter  $E$

once the  $F \rightarrow E$  dependency is satisfied. This means that blue can only enter  $E$  after green is entering or has already entered  $F$ .

By following a TPG during execution, we allow agents to travel more flexibly, as they are not tied to specific timesteps. Even if agents experience delays, following the TPG ensures that the passing order at each conflict location remains consistent with the MAPF solution. This enables them to execute the MAPF solution without collisions, deadlocks, or replanning when delays occur.

### 2.3 Bidirectional Temporal Plan Graph (BTPG)

A TPG requires agents to strictly follow the passing order at each location as specified by the MAPF solution. When delays occur, this can cause agents to unnecessarily wait and reduce efficiency. Bidirectional Temporal Plan Graphs (BTPGs) are proposed to address this issue (Su, Veerapaneni, and Li 2024). Conceptually, at certain intersections where it is safe, e.g., Figure 1, BTPG replaces the fixed ordering and allows either of the agents to cross first. In the example, this means that either blue or green can cross the intersection first. This involves changing certain type-2 edges into bidirectional pairs and ensuring that the resultant BTPG is deadlock-free.

**Definition 2 (BTPG).** Building on TPG, BTPG introduces the concept of bidirectional pairs  $E_{pair}$ . A bidirectional pair (bi-pair) consists of a pair of type-2 edges  $\{(v_i^m, v_j^n), (v_{j+1}^n, v_{i-1}^m)\}$  with  $\text{loc}(v_j^n) = \text{loc}(v_{i-1}^m)$  (denoted as conflict location). These two edges represent both options of agent  $m$  going before  $n$  and  $n$  going before  $m$  at the conflict location, respectively.

**BTPG execution policy** The BTPG execution policy is similar to the TPG execution policy but with a key difference. When agents encounter a bidirectional pair, only one edge is selected based on a “first-come, first-served” rule. This means that either agent in the pair can be the first to enter the conflict location. When the first agent arrives, the edge that allows this agent to enter first is selected and the other edge is deleted. For example, in the BTPG of Figure 1, the red type-2 edges form a bi-pair. Now during execution, if green gets delayed, blue can select the  $H \rightarrow E$  edge and enter  $E$  first, and the edge  $F \rightarrow E$  is discarded. This avoids the unnecessary waiting that would have been required if we had followed the original TPG.

Importantly, an edge  $(v_i^m, v_j^n)$  from a bi-pair will only be selected during execution when agent  $m$  reaches  $v_{i-1}^m$  first. This is crucial for additional BTPG optimizations.

**BTPG construction** The method for building a BTPG starts with a TPG and is described in Algorithm 1 without the underlined lines. Each type-2 edge is checked one by one to determine whether it can be reversed to form a bi-pair (Line 6, Algorithm 1). If reversing the edge does not cause a deadlock, the agents could reverse passing orders during execution. In a TPG, the presence of a cycle leads to a deadlock (Berndt et al. 2023; Coskun, O’Kane, and Val-torta 2021). Thus the way to check for a deadlock is to see if the reversed edge forms a cycle (Line 10, Algorithm 1). If not, the type-2 edge is converted into a bi-pair.

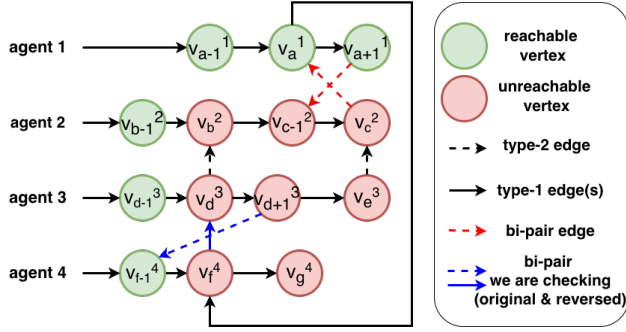


Figure 2: Example of new types of non-deadlock cycle discussed in Section 3.

## 2.4 BTPG-optimized (BTPG-o)

However, due to the existence of bi-pairs in a BTPG, agents choose certain type-2 edges during execution, so some bi-pair edges in the cycle may provably never be selected to execute. This means that certain cycles will not be encountered during execution and will not lead to a deadlock, and are thus called “non-deadlock cycles”.

**Property 1** (Non-deadlock cycles in BTPG-o). If a cycle contains a vertex  $v_i^n$  and a bi-pair edge that points from  $v_j^n$ ,  $j > i$ , then this cycle will not lead to a deadlock, and we call it a non-deadlock cycle.

Conceptually, if the cycle involving  $v_i^n$  can cause a deadlock, agent  $n$  will get stuck at  $v_{i-1}^n$  (or even earlier), preventing it from reaching all subsequent vertices  $v_j^n$ ,  $j \geq i$ . So, according to our BTPG execution policy, the bi-pair edge that originates from  $v_j^n$ ,  $j > i$  can not be selected as  $n$  cannot reach  $v_{j-1}^n$ . Hence, this cycle will never be encountered and will not cause a deadlock during execution. For example, in Figure 2, the cycle  $v_d^3 \rightarrow v_b^2 \rightarrow \dots \rightarrow v_{c-1}^2 \rightarrow v_c^2 \rightarrow v_a^1 \rightarrow v_f^4 \rightarrow v_d^3$  is a non-deadlock cycle, because the edge  $(v_c^2, v_a^1)$  is a bi-pair edge and  $v_b^2$  ( $b < c$ ) is in the cycle. The BTPG-o algorithm uses this property and excludes these non-deadlock cycles, allowing it to find more bi-pairs that make the result BTPG more flexible.

In this work, we identify more types of cycles that do not lead to deadlocks. Section 3 discusses the complete set of *all* types of non-deadlock cycles. Section 4 then describes our BTPG-max algorithm, which is based on this insight and finds the locally maximal number of bi-pairs by ignoring all non-deadlock cycles.

## 2.5 Other Related Works

Recent works have explored online switching of type-2 edges to handle delays, leveraging the key fact that cycles in a TPG imply deadlocks (Berndt et al. 2023; Coskun, O’Kane, and Valtorta 2021). These methods aim to switch type-2 edges without introducing cycles. Switchable-Edge Search (Feng et al. 2024; Jiang, Lin, and Li 2025) uses heuristic search over type-2 edges to minimize expected cost based on current agent positions and delays. Berndt et al. (2023) propose a similar Switchable Action Dependency Graph and solve for switchable edges using receding

horizon planning formulated as a Mixed-Integer Linear Programming problem. Liu et al. (2024) greedily adjust edges for blocked agents to maximize agent progress, ensuring no cycles via DFS checks. Another approach re-solves the MAPF problem online while constraining agents to spatial paths of the original paths (Kottinger et al. 2024).

Separately, some methods plan delay-robust MAPF solutions, either explicitly allowing up to  $K$  delays (Atzmon et al. 2018; Chen et al. 2021) or incorporating delay probabilities (Ma, Kumar, and Koenig 2017), though these require prior delay models. Space-Level CBS (Wagner, Veerapaneni, and Likhachev 2022) minimizes agent wait events, while Space-Order CBS (Wu et al. 2024) directly computes TPGs to reduce type-2 edges.

Our method is orthogonal to these methods and could be used on top of them.

## 3 Non-deadlock Cycles in BTPG

Our objective in this section is to find all *non-deadlock cycles* in BTPG. These are a cycle of vertices and edges in a BTPG that will never lead to deadlocks during execution under the BTPG execution policy. Analogously, a *deadlock cycle* in a BTPG is a cycle of vertices and edges that has a non-zero probability to lead to a deadlock during execution.

In the analysis of Property 1, we claimed a cycle to be non-deadlock if a bi-pair edge in the cycle cannot be selected because the corresponding conflict location is “unreachable” by the deadlocked agent. Our work expands this idea: we will classify all vertices in BTPG as reachable and unreachable and denote those cycles that have bi-pair edges with unreachable conflict locations as non-deadlock.

**Definition 3** (Reachable and unreachable vertex). When a deadlock occurs, all vertices where agents might be located and visited are reachable vertices, and all vertices where they cannot be located are unreachable vertices.

Conceptually, reachable vertices are those before deadlocks. Unreachable vertices are those after deadlocks (since the agent is in a deadlock, it cannot reach future vertices) and, importantly, vertices that are dependent via type-1 and type-2 edges on other unreachable vertices. For example, consider the cycle  $v_d^3 \rightarrow v_b^2 \rightarrow \dots \rightarrow v_{c-1}^2 \rightarrow v_c^2 \rightarrow v_a^1 \rightarrow v_f^4 \rightarrow v_d^3$  in Figure 2. If this cycle leads to a deadlock, then agent 2 can only reach up to vertex  $v_{b-1}^2$ . Agent 2 cannot reach future vertices like  $v_{c-1}^2$ . So,  $v_{c-1}^2$  is unreachable.

It is evident that if a vertex is in a cycle that constitutes a deadlock, then this vertex is unreachable. Consequently, all subsequent type-1 edges and normal type-2 edges cannot be satisfied. This will in turn cause those vertices to be unreachable, and we can cascade the logic down the graph.

**Lemma 1.** During execution, if a cycle leads to a deadlock, vertex  $v_i^n$  is unreachable if and only if it is in the cycle or there is a path from a vertex  $v_j^m$  in the cycle to it only through type-1 edges and normal type-2 edges (i.e., not type-2 edges in bi-pairs).

*Proof.* For the first case, because  $v_i^n$  is in the deadlock cycle, it is definitely unreachable (otherwise the agent would

not be in deadlock there). For the second case, since  $v_j^m$  is in the cycle and thus unreachable, all type-1 edges and normal type-2 edges along the path from  $v_j^m$  to  $v_i^n$  cannot be satisfied, indicating that  $v_i^n$  is also unreachable.

We prove the other direction by contradiction: Assume that the unreachable vertex  $v_i^n$  is not in the cycle and that there is no such path. Then, either (1) there is no path from a vertex in the cycle to  $v_i^n$ , or (2) there is a path passing through some bi-pair edges from a vertex in the cycle to  $v_i^n$ . In case 1, all edges pointing to vertex  $v_i^n$  can be satisfied, meaning that  $v_i^n$  is indeed reachable. In case 2, based on the BTPG execution policy, the other edge in the bi-pair could be selected, in which case the current edge would be discarded, making that path disappear. Both cases contradict the assumption that  $v_i^n$  is unreachable.  $\square$

Next, we determine whether the conflict vertex  $v_j^m$  of a bi-pair edge within the cycle is unreachable. If it is, this bi-pair edge will not be selected, resulting in a non-deadlock cycle.

**Theorem 1.** *A cycle in BTPG is a non-deadlock cycle if and only if it contains a bi-pair edge  $(v_i^m, v_j^n)$  such that  $v_{i-1}^m$  is unreachable.*

*Proof.* We first prove by contradiction that if the cycle satisfies the given conditions, it is a non-deadlock cycle. Assume that this cycle could lead to a deadlock. According to the BTPG execution policy, since  $v_{i-1}^m$  is unreachable, the bi-pair edge could not be selected, meaning that the cycle would never be encountered during execution, leading to a contradiction. Thus, the cycle must be a non-deadlock cycle.

We then prove the other direction also by contradiction: Assume there is a non-deadlock cycle that does not satisfy the given conditions, meaning that all edges in this cycle can be executed (i.e., all bi-pair edges in the cycle can be selected to execute). Agents may encounter this cycle during execution, leading to a deadlock.  $\square$

Based on Theorem 1, we can find new types of non-deadlock cycles that BTPG-o does not cover. For example, in Figure 2 when evaluating if edge  $(v_{d+1}^3, v_{f-1}^4)$  can be reversed to  $(v_f^4, v_d^3)$ , BTPG-o would classify the cycle  $v_d^3 \rightarrow \dots \rightarrow v_e^3 \rightarrow v_c^2 \rightarrow v_a^1 \rightarrow v_f^4 \rightarrow v_d^3$  as a deadlock cycle based on Property 1. However, since  $v_d^3$  is an unreachable vertex and there is a path from  $v_d^3$  to  $v_{c-1}^2$  only through type-1 and normal type-2 edges,  $v_{c-1}^2$  is unreachable, and this cycle is a non-deadlock cycle as the bi-pair edge originating from  $v_c^2$  cannot be chosen during execution.

### 3.1 Reachable deadlock cycle

So far, we have discussed non-deadlock cycles and how they can be identified. This implicitly defines deadlock cycles (i.e., those cycles which are not non-deadlock cycles). However, we can define an even stricter subset of deadlock cycles that are faster to detect.

In a deadlock situation, no agent can proceed—each gets stuck at the vertex just before entering the deadlock cycle. We refer to such cycles as reachable deadlock cycles.

---

**Algorithm 1:** BTPG-max: Underline highlights the main difference from BTPG-o. Note this version omits an outer while loop discussed in Section 4.3.

---

**Input:** TPG  $G = (V, E_1 \cup E_2)$   
**Output:** BTPG  $\mathcal{G}$

```

1  $E_{pair} \leftarrow \emptyset;$  // set of bi-pair edges
2  $E_{group} \leftarrow Grouping(E_2);$ 
3 for  $g$  in  $E_{group}$  (or until  $TimeOut$ ) do
4    $\tilde{g} \leftarrow \{(v_{j+1}^m, v_i^n) | (v_{i+1}^n, v_j^m) \in g\};$ 
5    $E_{pair} \leftarrow E_{pair} \cup \{g, \tilde{g}\};$ 
6   for  $e = (v_{j+1}^m, v_i^n)$  in  $\tilde{g}$  do
7      $\mathcal{G} \leftarrow (V, E_{group} \cup E_{pair});$ 
8      $VertexStatus[v] \leftarrow reachable | \forall v \in V;$ 
9      $UpdateVertexStatus(VertexSta$ 
10     $tus, \mathcal{G}, v_i^n);$ 
11    if  $DeadlockCycleDetection(\mathcal{G}, v_i^n,$ 
12     $v_{j+1}^m, \{e\}, VertexStatus)$  then
13       $E_{pair} \leftarrow E_{pair} \setminus \{g, \tilde{g}\};$ 
14      break;
15 return  $\mathcal{G} = (V, E_{group} \cup E_{pair} \cup E_1);$ 

```

---

**Definition 4** (Reachable deadlock cycle in BTPG). A reachable deadlock cycle in BTPG is a cycle such that for every (normal or bi-pair) type-2 edge  $(v_i^m, v_j^n)$ ,  $v_{j-1}^n$  is where robot  $n$  gets stuck during execution.

This means that when a deadlock occurs, all the agents involved in the deadlock will stop at the vertices just one before the deadlock. So in Definition 4, if any vertex like  $v_{j-1}^n$  is marked as unreachable, then the cycle is not a reachable deadlock cycle because stuck agents cannot be at those locations when deadlock happens (based on Definition 3). With this stricter property of a deadlock cycle, we can speed up the cycle detection process. Every possible deadlock has a corresponding reachable deadlock cycle, so we can design an algorithm that identifies only reachable deadlock cycles to determine whether a type-2 edge can be reversed.

Back to the previous example, when we check the reversed edge  $(v_f^4, v_d^3)$  and traverse the path  $v_d^3 \rightarrow \dots \rightarrow v_e^3 \rightarrow v_c^2$  during cycle detection, we can skip the last edge  $(v_e^3, v_c^2)$ . Because if  $v_d^3$  is in the deadlock cycle, there is a path from  $v_d^3$  to  $v_{c-1}^2$  through only the type-1 and normal type-2 edges, then  $v_{c-1}^2$  is unreachable.  $v_{c-1}^2$  is unreachable means that  $v_c^2$  cannot be in a reachable deadlock cycle, so we can skip the edge  $(v_e^3, v_c^2)$ . This operation will not affect the number of bi-pairs found but can accelerate the algorithm.

## 4 BTPG-max

We propose BTPG-max, which, like BTPG-o, checks type-2 edges one by one and converts a type-2 edge into a bi-pair if adding the other edge in the bi-pair does not introduce deadlocks. The main difference between BTPG-o and BTPG-max lies in cycle detection—BTPG-max considers only reachable deadlock cycles using Theorem 1 and Def-

---

**Algorithm 2:** DeadlockCycleDetection: Underline highlights the main difference from BTPG-o.

---

**Input:** (1) BTPG  $\mathcal{G} = (V, E_{group} \cup E_{pair} \cup E_1)$ , (2) current vertex for expansion  $v_i^n$ , (3) origin vertex  $v_o$ , (4) set(s) of edges  $E_{vis}$  along the current DFS branch, (5) the set of status of vertices.

**Output:** true or false

```

1 if  $v_i^n = v_o$  then
2   if  $E_{vis} \subseteq E_2$  and  $|E_{vis}| > 2$  then
3     return false ▷ rotation cycle
4   else return true;
5 for  $v_j^m$  in  $\{v_j^m \in V \mid (v_i^n, v_j^m) \in E_1 \cup E_2\}$  do
6    $e \leftarrow (v_i^n, v_j^m)$ ;
7    $v_e^n \leftarrow \text{getEarliestOutNode}(e.group)$ ;
8   if  $e \in E_2$  and
      $\text{VertexStatus}[v_{j-1}^m] \neq \text{'reachable'}$ 
     then continue;
9   if  $e \in E_{pair}$  then
10    if  $\text{VertexStatus}[v_{e-1}^n] \neq \text{'reachable'}$ 
        then continue;
11     $\text{UpdateVertexStatus}(\text{VertexSta}$ 
         $\text{tus}, \mathcal{G}, v_j^m)$ ;
12    if  $\text{DeadlockCycleDetection}(\mathcal{G}, v_j^m, v_o,$ 
         $E_{vis} \cup \{e\}, \text{VertexStatus})$  then return
        true;
13 return false;
```

---

inition 4.

The main idea of cycle detection is to perform a depth-first search rooted at the edge we are checking to find a reachable deadlock cycle. To ensure that the cycles we identify meet the conditions of a reachable deadlock cycle, we propose requirements for traversing type-2 edges (normal or bi-pair) during cycle detection below. During cycle detection, we can omit those edges that do not satisfy the following requirements as those edges are guaranteed not in the current reachable deadlock cycle.

#### 4.1 Requirements of traversing a type-2 edge

**Normal type-2 edges** According to Definition 4, if type-2 edge  $(v_i^n, v_j^m)$  is in a reachable deadlock cycle, then agent  $m$  must be in  $v_{j-1}^m$  when this deadlock occurs. Therefore, vertex  $v_{j-1}^m$  must be reachable for agent  $m$ .

**Bi-pair type-2 edges** According to Theorem 1, if edge  $(v_i^n, v_j^m)$  is in a bi-pair, then in addition to vertex  $v_{j-1}^m$  being reachable, vertex  $v_{i-1}^n$  must also be reachable.

#### 4.2 Effects of traversing a type-2 edge

After traversing a type-2 edge  $(v_i^n, v_j^m)$  during cycle detection, meaning that we suppose that this edge is in a reachable deadlock cycle, it will lead to the following effect: All locations of  $v_k^m$  with  $k < j - 1$  must have been visited. According to Definition 4, this means that the agent cannot get stuck

at  $v_k^m$ . We therefore mark these vertex as “visited” to denote that the agents in deadlock cannot be at these vertices.

#### 4.3 BTPG-max

BTPG-max follows the general approach of BTPG construction: From a valid TPG, check each type-2 edge (Line 6, Algorithm 1) to see if it forms a reachable deadlock cycle (Line 10, Algorithm 1). If not, it is converted into a bi-pair. Since the construction approach being the same as BTPG-o, where each type-2 edge is checked one by one, BTPG-max inherits the anytime property of BTPG-o. This means that the algorithm can be stopped at any time, and the BTPG always remains valid.

In BTPG-max, `VertexStatus` records the status (reachable, unreachable, visited) of each vertex in the BTPG and is used to check the requirements when traversing a type-2 edge. All vertices are initially marked as reachable when checking an edge (Line 8, Algorithm 1). Then `DeadlockCycleDetection` is invoked (Algorithm 1, Line 10) which checks for reachable deadlock cycles. `DeadlockCycleDetection` Algorithm 2 is a depth-first search which calls `UpdateVertexStatus` to update the status of each vertex (by modifying `VertexStatus`) via a depth-first search based on the effects of traversing a type-2 edge (Algorithm 2, Line 11), i.e., marks vertices as unreachable or visited according to Section 4.2.

An important detail is that `UpdateVertexStatus` can only pass through type-2 edges that have been checked not to be bi-pairs. In Figure 2, if edge  $(v_d^3, v_b^2)$  is later converted into a bi-pair, even if this type-2 edge cannot be satisfied, according to the BTPG execution policy, agent 2 can choose the other direction and pass through node  $v_b^2$  to reach  $v_{c-1}^2$ . In this case, the cycle  $v_d^3 \rightarrow \dots \rightarrow v_e^3 \rightarrow v_c^2 \rightarrow v_a^1 \rightarrow v_f^4 \rightarrow v_d^3$  becomes a deadlock cycle, reversing the previous result that allowed edge  $(v_{d+1}^3, v_{f-1}^4)$  to be reversed.<sup>1</sup>

Similar to BTPG-o, once a new bi-pair is found, it may allow previously non-convertible edges to be converted into bi-pairs because `UpdateVertexStatus` can propagate unreachability through more type-2 edges during cycle detection. Therefore, to find the locally maximal number of bi-pairs, a while loop is needed to repeatedly check if normal type-2 edges can be converted until no new bi-pairs are found. However, as mentioned above, transforming a normal type-2 edge into a bi-pair can also affect edges previously determined to be bi-pairs. Thus, in the second and subsequent iterations of this while loop, the condition for converting a normal type-2 edge into a bi-pair must include that it should not affect edges that were previously determined to be bi-pairs. However, in practical use, if this locally maximal property is not critical, one can implement a simpler version without the while loop. Empirically, we observed that the number of bi-pairs identified in subsequent iterations is usually very small and has a negligible impact on the results. We demonstrate this finding in Section 5. For clarity and simplicity, we have omitted the while loop in Algorithm 1.

We note that BTPG-max is “locally maximal” but does not necessarily find a global maximal/optimal solution as the

<sup>1</sup>Our appendix walks through BTPG-max on this example.

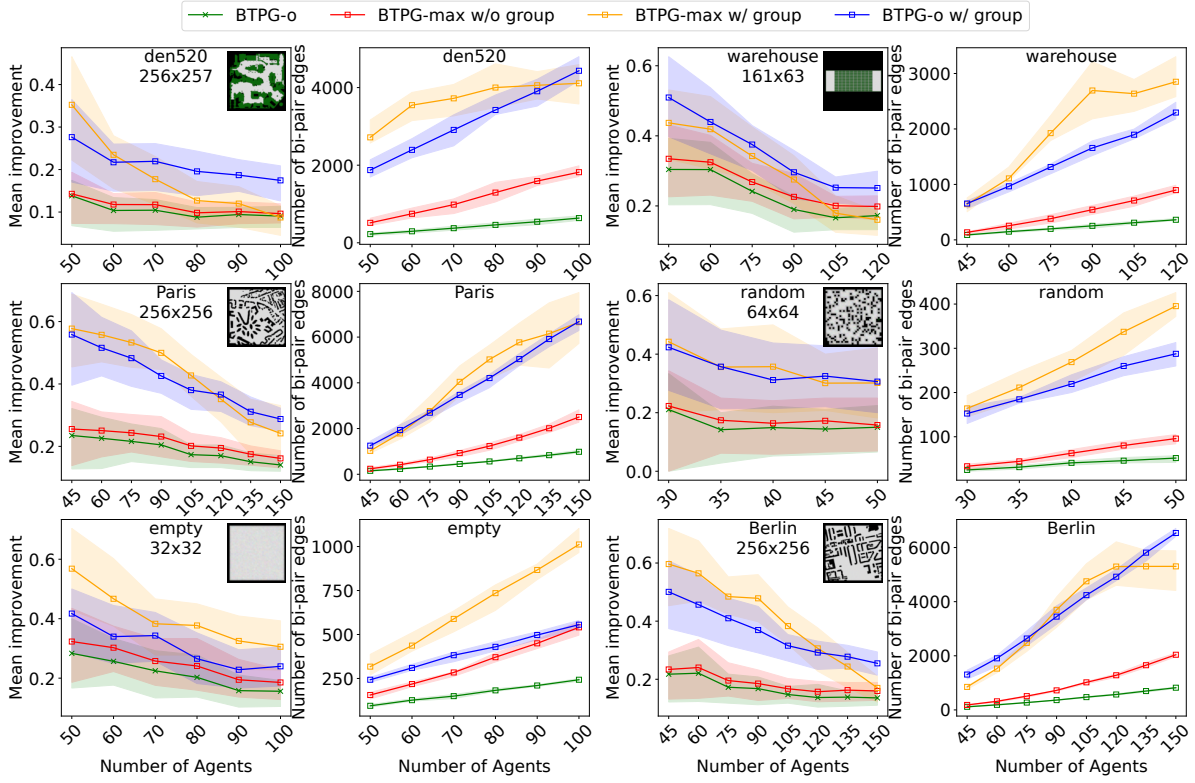


Figure 3: Comparison of BTPG-o, BTPG-max w/o grouping, and BTPG-max w/ groups.

order of converting type-2 edges into bidirectional pairs affects the result of converting future type-2 edges.

#### 4.4 BTPG-max with grouping

BTPG-o identifies two common groups of type-2 edges where each individual edge in the group cannot be converted into a bi-pair. The first group corresponds to agents following each other across consecutive locations (Figure 4 top row, agent 2 is following agent 1 across  $BCD$ ), while the second group corresponds to agents crossing consecutive locations in opposite directions (Figure 4 bottom row, agent 1 crosses  $BCD$  then agent 2 crosses  $DCB$ ). Although individual edges cannot be reversed, all edges in the group can be reversed simultaneously (Berndt et al. 2023). We incorporate this finding into our approach.

To determine if an edge group can be reversed, we need to reverse all edges in the group simultaneously and then verify that each one can be reversed. If any edge in the group cannot be reversed, the whole group cannot be reversed.

In cycle detection, passing through an edge in one bi-pair edge group means that all edges in the group must be consistent with the direction it passes through. Therefore, we need to ensure that the location that allows the agent to choose this direction is reachable, i.e., the start of each group needs to be reachable. In Figure 4, this corresponds to nodes B and D, respectively. More specifically, if the edges  $(v_i^n, v_j^m), (v_{i+1}^n, v_{j\pm 1}^m), \dots$  are in a bi-pair edge group and during cycle detection we want to pass through any one of them, then we need to check if the vertex  $v_{i-1}^n$  is reachable. `getEarliestOutNode` (Algorithm 2, Line 7) han-

dles this logic and returns the vertex with the earliest type-2 edge in the group.

#### 4.5 Time Complexity

Algorithm 2 is a recursive Depth-First Search (DFS) which visits each edge at most once unless an edge/vertex is marked as visited or unreachable (these vertices/edges are skipped according to Line 8 and Line 10). Vertices are marked visited or unreachable via `UpdateVertexStatus`. Thus, each edge is traversed once either through a recursive call in Line 12, or through a recursive call in `UpdateVertexStatus` (Line 11). Thus the total time complexity of Algorithm 2 is  $O(|E|)$ .

Now moving to Algorithm 1. Algorithm 1 calls `DeadlockCycleDetection` at most  $|E_2|$  times (where  $E_2$  is the set of type-2 edges). Finally, Algorithm 1 is invoked at most  $|E_2|$  times in the outer `while` loop (not written in Algorithm 1 but described in Section 4.3). Therefore, the total worst-case time complexity of the BTPG-max algorithm is:  $O(|E_2|^2 \cdot |E|)$ . We emphasize this is a worst-case time complexity and that in practice it can converge faster. Also, as mentioned earlier, BTPG-max is an anytime algorithm which allows it to be stopped early while still returning a valid BTPG.

## 5 Empirical Evaluation

We use the optimal MAPF solver CBSH2-RTC (Li et al. 2021) to generate the MAPF solution for each MAPF instance. We then convert each MAPF solution into a TPG. We run our proposed method BTPG-max with and without

	den520		warehouse		Paris		random		empty		Berlin	
	O	max	O	max	O	max	O	max	O	max	O	max
Mean improvement	<b>9.1%</b>	8.9%	<b>18.3%</b>	16.0%	15.5%	<b>24.5%</b>	14.5%	<b>32.7%</b>	19.8%	<b>40.0%</b>	15.5%	<b>17.2%</b>
Median improvement	<b>8.1%</b>	7.0%	<b>18.3%</b>	15.2%	15.8%	<b>24.4%</b>	11.8%	<b>28.6%</b>	18.3%	<b>29.3%</b>	15.3%	<b>15.9%</b>
Max improvement	21.6%	<b>27.9%</b>	31.2%	<b>38.3%</b>	26.3%	<b>44.2%</b>	63.6%	<b>76.9%</b>	42.8%	<b>60.0%</b>	24.6%	<b>40.3%</b>
Min improvement	<b>2.2%</b>	0.9%	<b>8.8%</b>	5.1%	<b>6.8%</b>	4.9%	0.0%	<b>5.8%</b>	5.0%	<b>6.3%</b>	<b>7.0%</b>	1.0%
# Type-2 edges	25,766		15,228		24,506		1,153		3,044		26,025	
# Bi-Pairs found	1,041	<b>4,110</b>	532	<b>2,851</b>	1,679	<b>6,649</b>	69	<b>395</b>	360	<b>1,011</b>	1,389	<b>5,303</b>
# Used Bi-Pairs	<b>68</b>	62	<b>56</b>	55	120	<b>169</b>	5.6	<b>14</b>	38	<b>51</b>	110	<b>128</b>

Table 1: Statistics of BTPG-o and BTPG-max w/ grouping. The number of agents for the six maps selected for the statistics are the largest ones, respectively. o: BTPG-optimized; max: BTPG-max w/ grouping. All data in the bottom block are averages of 10 scenarios for each map. Used Bi-Pairs: Bi-Pairs that are used (reversed) in the simulation.




Figure 4: Two cases of grouping: before and after reversing.

grouping and compare them to the baseline BTPG-o. All BTPG methods have a 10-minute time cutoff. Finally, we simulate the execution policies of TPG and BTPG, where 10% of the agents have a 30% chance of being delayed by 5 timesteps at each non-delayed timestep, same as the settings used in the BTPG-o paper (Su, Veerapaneni, and Li 2024) for a fair comparison.

We run a total of 3,900 simulations across six different benchmark maps (Figure 3). Each map features 5 to 8 different agent counts, with 10 random instances per agent count. To simulate delays, each instance was run with 10 random seeds. The maximum number of agents per map is the highest number that CBSH2-RTC can solve within a 2-minute time limit. All experiments were run on a PC with a 3.40 GHz Intel i7-14700K CPU and 64 GB RAM.

**Improvement** We adopt the improvement calculation from Su, Veerapaneni, and Li (2024),  $\frac{T_{TPG} - T_{BTPG}}{T_{TPG} - T_{Ideal}}$ .  $T_{TPG}$  and  $T_{BTPG}$  represent the mean execution times for TPG and BTPG, respectively, corresponding to the average number of timesteps agents take to reach their goal locations.  $T_{Ideal}$  is a lower bound of the mean execution time, calculated as the sum of execution timesteps specified in the original MAPF solution and the total delay timesteps of the delayed agents, divided by the number of agents.

Figure 3 plots the results of BTPG-o and BTPG-max with and without grouping. Comparing BTPG-o without grouping (green) and BTPG-max without grouping (red), we see that BTPG-max consistently finds more bi-pairs and has a higher mean improvement than BTPG-o. Comparing BTPG-o without and with grouping (green vs blue), and BTPG-max without and with grouping (red vs yellow), we see the advantage of grouping as it significantly improves both the number of bi-pairs found and the overall improvement.

Interestingly, we see that the number of bi-pairs found

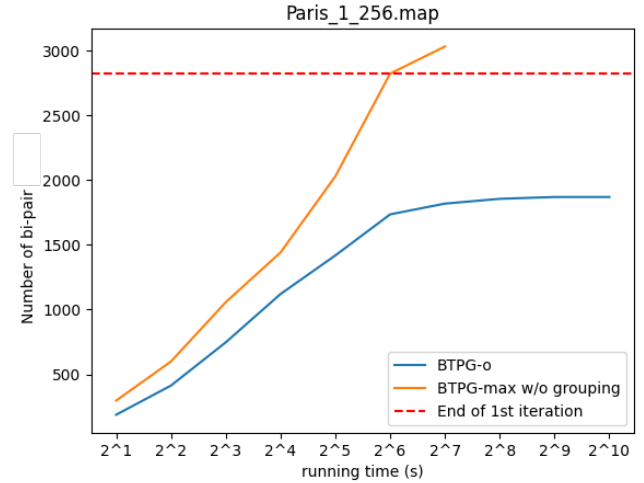


Figure 5: Anytime behavior of BTPG-o/max w/o grouping in the map Paris\_1\_256 with 150 agents.

and the mean improvement are not necessarily related. For example, in the warehouse map BTPG-max with grouping finds more bi-pairs than BTPG-o with grouping, but BTPG-o with grouping has a larger mean improvement. Table 1, which compares BTPG-o and BTPG-max with grouping, shows a similar story. BTPG-max consistently finds 3-5x more bi-pairs on all 6 maps but only has better median improvements on 4 of the 6 maps. We believe this discrepancy is due to the order of type-2 edges that BTPG-o picks. BTPG-o checks type-2 edges in the order they are encountered during execution, while BTPG-max converts more earlier edges into bi-pairs. During execution though, the locations further back on the paths are more likely to be affected by delays. Therefore, in the larger maps where BTPG-max times out (e.g. den520 and warehouse map (Table 1)), even though the number of bi-pairs is larger, the number of used bi-pairs is smaller, so the improvement is smaller.

**Anytime Performance** Figure 5 shows the anytime performance of BTPG-o vs BTPG-max without grouping on Paris\_1\_256 with 150 agents. The x-axis is the BTPG creation cut-off time in seconds while the y-axis is the number of bi-pair found. We see that BTPG-max has strictly superior

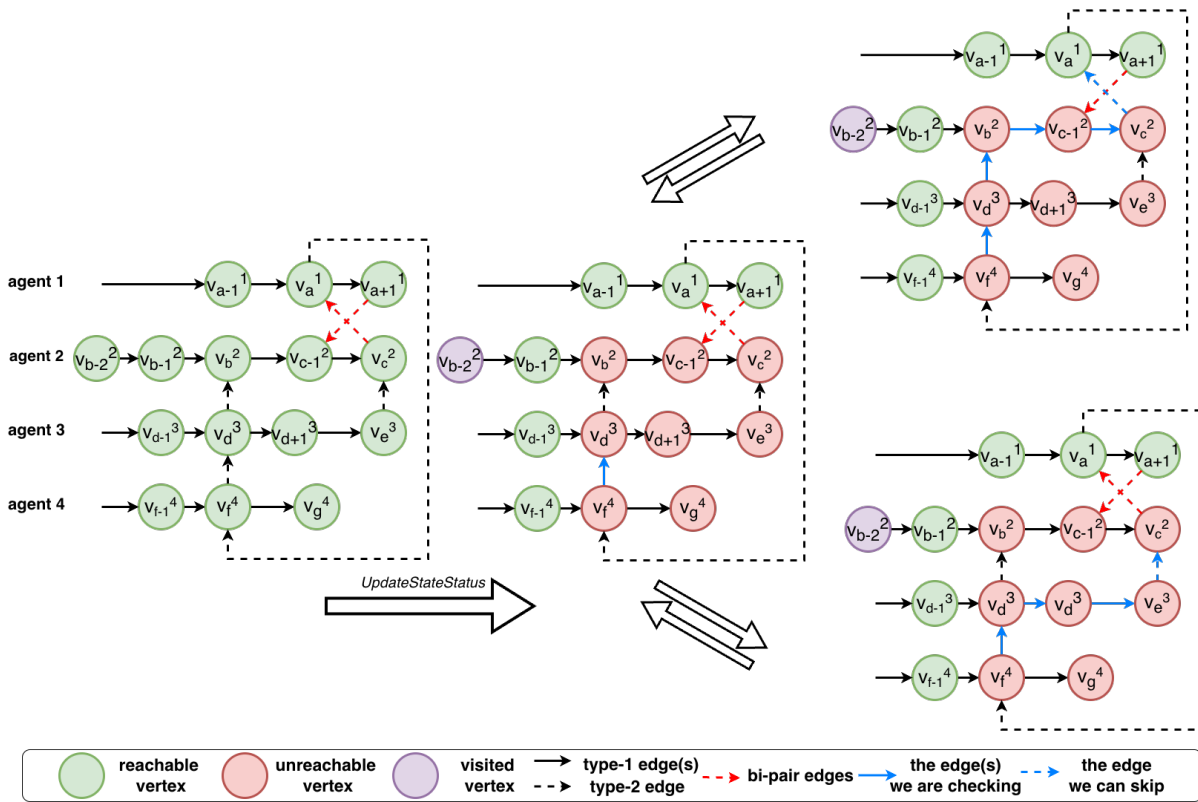


Figure 6: BTPG-max example.

anytime performance than BPTG-o; BTPG-max finds more bi-pairs faster and saturates at a higher final value.

The red horizontal dashed line denotes the end of the 1st iteration (after we have gone through all the type-2 edges once) and the start of the next iterations. We see that the vast majority of edges in BTPG-max are found during this first iteration with very few found afterwards.

Overall our results show that BTPG-max w/o grouping is strictly better (in bi-pairs found, mean improvement, anytime performance) than BTPG-o w/o grouping. Also, we find that grouping has a significant impact on performance and augments both BTPG-o and BTPG-max significantly.

## 6 Conclusion

Our work introduces BTPG-max, a method for identifying a locally maximal number of bi-pairs in BTPGs. Our key theoretical contribution is defining all possible non-deadlock cycles in BTPGs, provided in Theorem 1, which lays the foundation for distinguishing deadlock cycles in BTPG construction. Building on this, BTPG-max checks all possible deadlock cycles to construct locally maximal BTPGs to maximize execution flexibility. Empirically, BTPG-max consistently discovers more bi-pairs than BTPG-o and, combined with the grouping strategy, yields significantly better execution performance.

There are several possible directions for future work, such as making BTPG-max an online method invoked during execution or improving its anytime performance. Since check-

ing type-2 edges is semi-independent, a parallel version of BTPG-max is also promising. Finally, more sophisticated grouping strategies or incorporating prior knowledge of delays into the BTPG framework could yield further gains.

## A Run-through of BTPG-max

Figure 6 demonstrates the BTPG-max process. In Algorithm 2, `VertexStatus` is used to track the status of each vertex in the BTPG. At first, every vertex in the BTPG is reachable (Line 8, Algorithm 1). If we want to check if the reversed edge  $(v_f^4, v_d^3)$  is valid, `VertexStatus` should be updated (Line 9, Algorithm 1). If edge  $(v_f^4, v_d^3)$  is in the deadlock cycle, then vertices  $\{v_x^3 | x \geq d\}$  should be unreachable, and vertices  $\{v_x^3 | x < d - 1\}$  should be set to “visited” based on the effect of traversing a type-2 edge. Then, since  $v_d^3$  cannot be reached and edge  $(v_d^3, v_b^2)$  is a normal type-2 edge and cannot be satisfied, the vertices  $\{v_x^2 | x \geq b\}$  are also unreachable. `UpdateVertexStatus` uses simple DFS to update the conditions of each vertex based on the effects of traversing a type-2 edge.

For the top right case, after passing through vertices  $v_d^3$ ,  $v_b^2$ , and  $v_{c-1}^2$ , because vertex  $v_{c-1}^2$  is unreachable and the edge  $(v_c^2, v_{c-1}^2)$  is an edge in a bi-pair, BTPG-max will skip this edge (Line 10, Algorithm 2) based on the requirements of traversing a type-2 edge. Similarly, for the below right one, when traversing the vertices  $v_d^3$  and  $v_e^3$ , since the node  $v_{c-1}^2$  is unreachable, the edge  $(v_e^3, v_c^2)$  will be skipped (Line 8, Algorithm 2).

## Acknowledgments

This work is in part supported by the National Science Foundation (NSF) under grant numbers 2328671 and 2441629, as well as a gift from Amazon.

## References

- Atzmon, D.; Stern, R.; Felner, A.; Wagner, G.; Bartak, R.; and Zhou, N.-F. 2018. Robust Multi-Agent Path Finding. In *Proceedings of the International Symposium on Combinatorial Search*, volume 9, 2–9.
- Berndt, A.; Van Duijkeren, N.; Palmieri, L.; Kleiner, A.; and Keviczky, T. 2023. Receding horizon re-ordering of multi-agent execution schedules. *IEEE Transactions on Robotics*, 40: 1356–1372.
- Chen, Z.; Harabor, D. D.; Li, J.; and Stuckey, P. J. 2021. Symmetry Breaking for k-Robust Multi-Agent Path Finding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12267–12274.
- Coskun, A.; O’Kane, J.; and Valtorta, M. 2021. Deadlock-Free Online Plan Repair in Multi-Robot Coordination with Disturbances. In *Proceedings of the International FLAIRS Conference*, volume 34.
- Feng, Y.; Paul, A.; Chen, Z.; and Li, J. 2024. A real-time rescheduling algorithm for multi-robot plan execution. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, 201–209.
- Hoening, W.; Kumar, T. K.; 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*, volume 26, 477–485.
- Jiang, H.; Lin, M.; and Li, J. 2025. Speedup Techniques for Switchable Temporal Plan Graph Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 23212–23221.
- Kottinger, J.; Geft, T.; Almagor, S.; Salzman, O.; and Lahijanian, M. 2024. Introducing Delays in Multi Agent Path Finding. In *Proceedings of the International Symposium on Combinatorial Search*, volume 17, 37–45.
- Li, J.; Harabor, D.; Stuckey, P. J.; Ma, H.; Gange, G.; and Koenig, S. 2021. Pairwise symmetry reasoning for multi-agent path finding search. *Artificial Intelligence*, 301: 103574.
- Liu, Y.; Tang, X.; Cai, W.; and Li, J. 2024. Multi-agent path execution with uncertainty. In *Proceedings of the International Symposium on Combinatorial Search*, volume 17, 64–72.
- 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*, volume 31, 3605–3612.
- 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*, volume 38, 17559–17566.
- Wagner, A.; Veerapaneni, R.; and Likhachev, M. 2022. Minimizing coordination in multi-agent path finding with dynamic execution. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 18, 61–69.
- Wu, Y.; Veerapaneni, R.; Li, J.; and Likhachev, M. 2024. From Space-Time to Space-Order: Directly Planning a Temporal Planning Graph by Redefining CBS. arXiv:2404.15137.