

Delayed Proofs

For easier reference, we copy the execution procedure and the general algorithm framework below.

Procedure 1: Execute $(\mathcal{G} = (\mathcal{V}, \mathcal{E}))$
 Lines highlighted in blue are activated to compute the cost of a TPG, and can be omitted for the mere purpose of execution.

```

1 Define a counter cost;
2 Function INITEXEC( $\mathcal{G}$ )
3   cost  $\leftarrow$  0;
4   Mark all vertices in  $\mathcal{V}_0 = \{v_0^i : i \in \mathcal{A}\}$  as
   satisfied;
5   Mark all remaining  $v \in (\mathcal{V} \setminus \mathcal{V}_0)$  as unsatisfied;
6 Function STEPEXEC( $\mathcal{G}, i$ )
7   if  $\forall k : v_k^i$  satisfied then
8     return NULL;
9   cost  $\leftarrow$  cost + 1;
10   $v \leftarrow v_k^i : v_k^i$  unsatisfied and  $\forall k' < k, v_{k'}^i$ 
   satisfied;
11  forall  $(u, v) \in \mathcal{E}$  do
12    if  $u$  unsatisfied then
13      return NULL
14  return v
15 Function EXEC( $\mathcal{G}$ )
16  INITEXEC( $\mathcal{G}$ );
17  while there exists unsatisfied vertex in  $\mathcal{V}$  do
18    Define a set  $\mathcal{S} \leftarrow \emptyset$ ;
19    forall agent  $i \in \mathcal{A}$  do
20      Add STEPEXEC( $\mathcal{G}, i$ ) into  $\mathcal{S}$ ;
21    forall  $v \in \mathcal{S}$  do
22      if  $v \neq \text{NULL}$  then
23        Mark  $v$  as satisfied;
24  return cost;
  
```

Proofs for Section Temporal Plan Graph

Lemma 1. Executing \mathcal{G} encounters a deadlock if and only if there exists a cycle in \mathcal{G} .

Proof. Assuming that a deadlock is encountered in the t^{th} iteration of the while-loop (on line 17 of Procedure 1). Let \mathcal{V}' denote that set of all vertices that are unsatisfied, which is non-empty by the definition of deadlock. Assume towards contradiction that \mathcal{G} is acyclic, then there exists a topological ordering (i.e. a linear ordering of its vertices such that for every directed edge (u, v) , u comes before v in the ordering) of \mathcal{V}' . In this case, \mathcal{S} must contain the first vertex v_{first} in the ordering of \mathcal{V}' , because it satisfies both conditions in STEP_{EXEC}:

1. $v_{\text{first}} = \arg \min_k (v_k^i : \text{agent } i \in \mathcal{A}, v_k^i \text{ is unsatisfied})$, and
2. For all $(u, v_{\text{first}}) \in \mathcal{E}$, u is satisfied.

Algorithm 2: Replanning

HEURISTIC, TERMINATE, CYCLEDETECTION, and BRANCH are modules that will be specified later. \mathcal{X} denotes some auxiliary information accompanying a TPG, whose format is defined by the set of modules.

```

Input: TPG  $\mathcal{G}_{\text{root}} = (\mathcal{V}, \mathcal{E}_1, (\mathcal{S}_{\mathcal{E}_2}, \mathcal{N}_{\mathcal{E}_2}))$ 
Output: TPG  $\mathcal{G}_{\text{result}}$ 
1 Initialize an empty priority queue  $\mathcal{Q}$ ;
2  $h_{\text{root}} \leftarrow \text{HEURISTIC}(\mathcal{G}_{\text{root}}, \mathcal{X}_{\text{init}})$ ;
3  $\mathcal{Q}.push((\mathcal{G}_{\text{root}}, \mathcal{X}_{\text{init}}), 0, h_{\text{root}})$ ;
4 while  $\mathcal{Q}$  is not empty do
5    $((\mathcal{G}, \mathcal{X}), g, h) \leftarrow \mathcal{Q}.pop()$ ;
6    $(g', \mathcal{X}', (v_k^i, v_s^j)) \leftarrow \text{BRANCH}(\mathcal{G}, \mathcal{X})$ ;
7   if TERMINATE( $\mathcal{G}, \mathcal{X}'$ ) then
8     fix all edges in  $\mathcal{S}_{\mathcal{E}_2}$  of  $\mathcal{G}$ ;
9     return  $\mathcal{G}$ ;
10   $\mathcal{G}_f \leftarrow \text{fix}(\mathcal{G}', (v_k^i, v_s^j))$ ;
11  if not CYCLEDETECTION( $\mathcal{G}_f, (v_k^i, v_s^j)$ ) then
12     $h_f \leftarrow \text{HEURISTIC}(\mathcal{G}_f, \mathcal{X}')$ ;
13     $\mathcal{Q}.push((\mathcal{G}_f, \mathcal{X}'), g + g', h_f)$ ;
14   $\mathcal{G}_r \leftarrow \text{reverse}(\mathcal{G}', (v_k^i, v_s^j))$ ;
15  if not CYCLEDETECTION( $\mathcal{G}_r, (v_{s+1}^j, v_k^i)$ ) then
16     $h_r \leftarrow \text{HEURISTIC}(\mathcal{G}_r, \mathcal{X}')$ ;
17     $\mathcal{Q}.push(\mathcal{G}_r, \mathcal{X}'), g + g', h_r)$ ;
18 throw exception "No solution found";
  
```

This contradicts the deadlock condition that $\mathcal{S} = \{\text{NULL}\}$. \square

Proposition 2 (Collision-Free). Let \mathcal{G} be a TPG constructed from a MAPF solution. Assuming \mathcal{G} is executed as in Procedure 1 and an agent i is moved to its k^{th} location l_k^i at timestep t iff vertex v_k^i is satisfied in the t^{th} iteration of the while-loop on line 17, any two agents i, j never collide.

This proposition is similar to lemma 4 in (Hönig et al. 2016) with slightly different terms. We include a proof for completeness.

Proof. Assume towards contradiction that i and j collide because they are at the same location at the same timestep t , i.e. after the t^{th} iteration, $v_k^i = \arg \max_k (v_k^i : v_k^i \text{ satisfied})$, $v_s^j = \arg \max_s (v_s^j : v_s^j \text{ satisfied})$, and $l_k^i = l_s^j$. Either edge (v_{k+1}^i, v_s^j) or (v_{s+1}^j, v_k^i) should be in \mathcal{E} . But $(v_{k+1}^i, v_s^j) \notin \mathcal{E}$ as otherwise v_s^j cannot be satisfied since v_{k+1}^i is unsatisfied; similarly $(v_{s+1}^j, v_k^i) \notin \mathcal{E}$ as otherwise v_k^i cannot be satisfied since v_{s+1}^j is unsatisfied. This shows a contradiction.

If they collide because i leaves a location at a timestep t , and j enters the same location at timestep t , then $l_{k-1}^i = l_s^j$ and vertices v_k^i and v_s^j are satisfied exactly in the t^{th} iteration of the loop. However, this is impossible since either (v_k^i, v_s^j) or (v_{s+1}^j, v_{k-1}^i) is in \mathcal{E} , but the out-going vertex in neither of these edges are satisfied before the t^{th} iteration. \square

Corollary 3. Let \mathcal{G} be a collision-free TPG. If we replace an arbitrary Type 2 edge (v_{s+1}^j, v_k^i) in it with (v_{k+1}^i, v_s^j) , the TPG remains to be collision-free.

Proof. This is observed from the fact that the proof of the above proposition argues the non-existence of two pairs of edges $(v_{k+1}^i, v_s^j) - (v_{s+1}^j, v_k^i)$ and $(v_k^i, v_{k+1}^j) - (v_{s+1}^j, v_{k-1}^i)$. Since in both pairs, the two edges are equal upon the replacement, the proof remains exactly the same after we perform an arbitrary replacement. \square

Proposition 4. Let \mathcal{G} be a TPG constructed from a MAPF solution \mathcal{P} , the cost of \mathcal{G} is no greater than the sum of travel time for agents following \mathcal{P} .

We rely on the following lemma to prove this proposition.

Lemma A. If an arbitrary agent i is planned to move to location l_k^i at a time $t = t_k^i$ in \mathcal{P} , then vertex v_k^i is either satisfied or can be satisfied in the t^{th} iteration of the while-loop.

Proof (of Lemma A). We induct on t . When $t = 0$, this holds by the functionality of $\text{INIT}_{\text{EXEC}}$. For $t > 0$, we consider the non-trivial case that the index k of v_k^i is non-zero. For the Type 1 edge of v_k^i , since $t_{k-1}^i \leq t_k^i - 1 = t - 1$, v_{k-1}^i is satisfied by an inductive hypothesis. Let v_s^j be an arbitrary Type 2 in-neighbor of v_k^i , by construction of Type 2 edges, $t_{s-1}^j < t_k^i$, so $t_{s-1}^j \leq t - 1$. This shows that all in-neighbors of v_k^i must be satisfied after the $(t - 1)^{\text{th}}$ iteration, thus v_k^i can be satisfied in the t^{th} iteration if it has not been satisfied yet. \square

Proof (of Proposition 4). Lemma A shows that if \mathcal{P} plans an agent i to enter its goal location at time t_{zi}^i , then all vertices of i are satisfied after the $(t_{zi}^i)^{\text{th}}$ iteration, i.e. agent i contributes to cost by at most t_{zi}^i units. Therefore $\text{cost} \leq \sum_{i \in \mathcal{A}} t_{zi}^i$, which is the sum of travel times of all agents in \mathcal{P} . \square

Corollary 5 (Deadlock-Free). If a TPG \mathcal{G} is constructed from a MAPF solution \mathcal{P} , then it is deadlock-free.

Proof. If \mathcal{G} contains a deadlock, then its execution would enter the while-loop for infinitely many iterations, and in each iteration, cost strictly increases. Thus $\text{cost} = \infty$. Yet the sum of travel time of \mathcal{P} is always finite, contradicting Proposition 16. \square

Proofs for Section Switchable TPG

Theorem 6. Let \mathcal{G} be a switchable TPG constructed as in Construction 1, there always exists a finite-cost, collision-free standard TPG that can be produced from \mathcal{G} .

Proof. One naive solution $\mathcal{G}_{\text{naive}}$ is produced by fixing all switchable edges in \mathcal{G} . $\mathcal{G}_{\text{naive}}$ has a finite cost (i.e. is deadlock-free) because by Corollary 5, an initial TPG \mathcal{G}_0 constructed from a MAPF solution is deadlock-free. And by Lemma 1, a TPG has a deadlock iff it contains a cycle, so it suffices to argue that step 2 and 3 in Construction 1 does not introduce a new cycle. This holds since step 2 has no

effect once we fix all switchable edges. Step 3 behaves as expanding a pre-existing edge (v_{k-1}^i, v_k^i) into a line of connecting edges. If a cycle exists in $\mathcal{G}_{\text{naive}}$, it either involves edges in this line or not. In the latter case, this cycle would exist exactly in \mathcal{G}_0 , which is impossible. In the former case, the entire line has to be contained in this cycle, in which case (v_{k-1}^i, v_k^i) along with the remaining component of this cycle would form a cycle in \mathcal{G}_0 , which is impossible.

Executing $\mathcal{G}_{\text{naive}}$ is collision-free because the exact same proof of Proposition 2 shows that two agents cannot collide if none of them is the delayed agent i or if the most-recently satisfied vertex of i is not a dummy vertex. So we may assume without loss of generality that agent i and j collide when i has already entered location l_{k-1}^i . However, such a collision is impossible since any vertex v_s^j for $j \neq i$ corresponding to the same location cannot be satisfied before v_k^i is satisfied. \square

Proofs for Section Algorithm

Lemma 7. Let $\mathcal{G}_{\text{switch}}$ be a switchable TPG and \mathcal{G} be an arbitrary standard TPG produced from $\mathcal{G}_{\text{switch}}$. The partial cost of $\mathcal{G}_{\text{switch}}$ is no greater than the cost of \mathcal{G} .

Proof. Let \mathcal{G}_{red} be the reduced standard TPG of $\mathcal{G}_{\text{switch}}$ that contains only its non-switchable edges. Consider running Procedure 1 on \mathcal{G} and \mathcal{G}_{red} , respectively. Since an edge appears in \mathcal{G}_{red} must also appear in \mathcal{G} , we can inductive 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 \mathcal{G}_{red} as well. Therefore the total timestep to satisfy all vertices in \mathcal{G}_{red} cannot exceed that in \mathcal{G} . \square

Proofs for Section Graph-based Modules

Theorem 11. Given a TPG, compute the longest path from vertex v_0^i to vertex v_{zi}^i for each $i \in \mathcal{A}$. Taking the sum of lengths of all such longest paths, this equals the cost of this TPG.

Proof. We again refer back to Procedure 1. Fix a longest path from vertex v_0^i to vertex v_{zi}^i . We prove by induction that the distance from v_0^i to a vertex v_s^j on this longest path is equal to the number of iterations required in the while-loop (line 17) in Procedure 1 to satisfy v_k^i . In the base case, the distance from v_0^i to itself is indeed 0. In the inductive step, assuming v_s^j is satisfied in the $t - 1^{\text{th}}$ iteration, and the longest path from v_0^i to v_s^j is $t - 1$. Then the next vertex $v_{s'}^{j'}$ on the longest path is satisfied in the t^{th} iteration, because:

- #iterations $\geq t$ since v_s^j is a in-neighbor of $v_{s'}^{j'}$ which needs to be satisfied before $v_{s'}^{j'}$.
- #iterations $\leq t$ since otherwise there must be another in-neighbor of $v_{s'}^{j'}$ that is not yet satisfied in the $t - 1^{\text{th}}$ iteration, which is going to compose a longer path than the one we look at.

Therefore the cost computed by Procedure 1 which equals the sum of iterations for all agents to reach their goal vertex is exactly the sum of lengths of longest paths \square