

MGCBS: An Optimal and Efficient Algorithm for Solving Multi-Goal Multi-Agent Path Finding Problem

Mingkai Tang¹, Yuanhang Li¹, Hongji Liu¹, Yingbing Chen¹, Ming Liu² and Lujia Wang²

¹Hong Kong University of Science and Technology

²Hong Kong University of Science and Technology (Guangzhou)

{mtangag, yliog, hliucq, ychengz}@connect.ust.hk, {eelium, eewanglj}@hkust-gz.edu.cn

Abstract

With the expansion of the scale of robotics applications, the multi-goal multi-agent pathfinding (MG-MAPF) problem began to gain widespread attention. This problem requires each agent to visit pre-assigned multiple goal points at least once without conflict. Some previous methods have been proposed to solve the MG-MAPF problem based on Decoupling the goal Vertex visiting order search and the Single-agent pathfinding (DVS). However, this paper demonstrates that the methods based on DVS cannot always obtain the optimal solution. To obtain the optimal result, we propose the Multi-Goal Conflict-Based Search (MGCBS), which is based on Decoupling the goal Safe interval visiting order search and the Single-agent pathfinding (DSS). Additionally, we present the Time-Interval-Space Forest (TIS Forest) to enhance the efficiency of MGCBS by maintaining the shortest paths from any start point at any start time step to each safe interval at the goal points. The experiment demonstrates that our method can consistently obtain optimal results and execute up to 7 times faster than the state-of-the-art method in our evaluation.

1 Introduction

With the development of the robotic industry, the multi-agent system has attracted more and more attention [Salzman and Stern, 2020; Stern *et al.*, 2019; Tjiharjadi *et al.*, 2022]. One of the critical problems to be solved is the multi-agent path finding (MAPF) problem. The MAPF problem requires planning a conflict-free path for each agent from its starting point to its goal point. MAPF is involved in many practical application scenarios in the real world, such as aircraft towing vehicles [Morris *et al.*, 2016], video games [Ma *et al.*, 2017b] and traffic management [Choudhury *et al.*, 2022; Dresner and Stone, 2008].

In the MAPF problem, each agent can be assigned only one goal point. This setting does not meet the needs of some large-scale robot applications. For example, in an automated warehouse scenario, each robot may need to deliver multiple goods in one trip. In this case, the robot needs to be provided with a collision-free path with multiple goal points. This

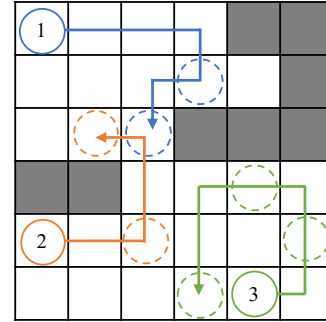


Figure 1: An example of the MG-MAPF with three agents. The grey cells represent the impassable areas occupied by obstacles. The solid circles indicate the start point of the agent, and the number marked inside is the agent ID. The dotted circles represent the goal points of the agent in the corresponding color. The colored arrow marks the path that can visit all goal points from the start point of the agent.

problem can be modeled as a multi-goal multi-agent pathfinding (MG-MAPF) problem. The solver of MG-MAPF needs to calculate a collision-free path for each agent so that the agent can visit each of its goals at least once in an arbitrary visiting order. Figure 1 shows an example of MG-MAPF.

Solving the MG-MAPF problem optimally, even in degenerate scenarios, can be time-consuming. The open loop traveling salesman problem [Applegate *et al.*, 2011], which is widely recognized as an NP-hard problem, can be reduced to a subset of the MG-MAPF that only considers one agent. Furthermore, the classical MAPF problem, which is proved to be NP-hard [Yu and LaValle, 2013], can also be reduced to a subset of the MG-MAPF, which considers only one goal for each agent. Therefore, it can be concluded that the problem of optimally solving the MG-MAPF is NP-hard.

Some methods have been proposed to solve the MG-MAPF problem by Decoupling the goal Vertex visiting order search and the Single-agent pathfinding (DVS), such as the Hamiltonian Conflict-based Search [Surynek, 2021], which is the current state-of-the-art (SOTA) method for the MG-MAPF problem. We named the MG-MAPF solver that is based on DVS as the DVS method. To find the shortest path to visit all goals for a single agent under constraints, DVS methods search a goal vertex visiting order iteratively. At each iteration, an unvisited goal is enumerated and tried to append to

the end of the order. The corresponding path is constructed by concatenating the shortest paths under constraints between two neighbor goals on the goal vertex visiting order.

However, this paper provides evidence through the case study and experiment that DVS methods cannot always obtain the optimal result for the MG-MAPF problem. We introduce a new approach, the Multi-Goal Conflict-Based Search (MGCBS), to solve the MG-MAPF problem optimally and efficiently. In contrast to the DVS method, MGCBS is based on Decoupling the goal Safe interval visit order search and the Single-agent pathfinding (DSS). The safe interval (SI) refers to a safe configuration with a maximal time period, where the ‘maximal’ means that if this time period were to be extended by a time step in any direction, the collision would occur [Phillips and Likhachev, 2011]. The goal safe interval (GSI) refers to the safe interval at the goal vertex. Formally, let $([t_0^{SI}, t_1^{SI}], s^{SI})$ be the SI at vertex s for the maximal time interval $[t_0^{SI}, t_1^{SI}]$. $([t_0^{SI}, t_1^{SI}], s^{SI})$ is an GSI iff s^{SI} is a goal vertex. We name the MG-MAPF solver that is based on DSS as the DSS method. The DSS method searches a GSI visiting order so that at least one GSI is visited at each goal. The corresponding path is obtained by concatenating the shortest path to visit each GSI following the order. In addition, we propose a data structure, the Time-Interval-Space Forest (TIS Forest), to reduce redundant calculations of multiple queries in the low-level solver of MGCBS by maintaining the shortest paths and their length from any start vertex at any start time to each GSI. Overall, the main contributions of this paper are as follows.

1. We present that the DVS methods cannot always obtain the optimal solution by case study and experiment.
2. We introduce a two-level approach, MGCBS, to solve MG-MAPF, achieving high computational efficiency in obtaining optimal solutions.
3. We present the TIS Forest, which maintains the shortest paths to each GSI from any start vertex at any start time, minimizing redundant calculations of multiple queries.
4. We provide the theoretical proof of the optimality and completeness of MGCBS.
5. We conducted a comprehensive experimental evaluation and compared our proposed method with the current SOTA method. Compared with the SOTA method, our method can consistently obtain the optimal solution while achieving a maximum speedup ratio of up to 7 in our evaluation.

2 Related Work

Some variants of the MAPF problem, which consider more than one goal point for an agent, have recently been studied. The Multi-Agent Pickup-and-Delivery (MAPD) problem and its variants, which require the agent to pick up the object in one location and deliver it to another location, were studied in [Čáp *et al.*, 2015; Ma *et al.*, 2017a; Xu *et al.*, 2022]. In [Zhang *et al.*, 2022], the Multi-Agent Path Finding with Precedence Constraints (MAPF-PC) problem was proposed, where the visiting order of goal points

needs to satisfy some precedence constraints. The Multi-Agent Simultaneous Multi-Goal Sequencing and Path Finding (MSMP) problem, which requires assigning goals to each agent before pathfinding, was solved in [Ren *et al.*, 2021; Ren *et al.*, 2022]. However, the settings of the above problems differ from the MG-MAPF problem, making their methods not directly usable in the MG-MAPF problem.

The MG-MAPF problem was firstly discussed in [Surynek, 2021], and two solutions were proposed: Hamiltonian Conflict-based Search (HCBS) and Satisfiability Modulo Theories Conflict-based Search (SMT-HCBS). The HCBS, which can be categorized as a DVS method, typically runs faster than SMT-HCBS due to the leverage of its heuristic function. However, it has room for improvement in optimality and efficiency. For optimality, our method improves upon HCBS by using DSS. For efficiency, our method uses the TIS Forest to reduce redundant calculations of the multiple queries of the shortest path for each agent.

3 Problem Definition

The MG-MAPF problem is defined as follows. A set of agents $A = \{a_1, a_2, \dots, a_k\}$ can move on an undirected graph $G(V, E)$ where each edge is of unit length. Let n_i be the number of goals of a_i . The task of agent a_i can be described as $\Omega_a = \{s_i, g_i\}$ where s_i is the start vertex and $g_i = \{g_i^1, g_i^2, \dots, g_i^{n_i}\}$ denotes the goal vertices of the agent. At each time step, agents can choose to wait at the current vertex or move to an adjacent vertex. The agent will stay at one of its goal vertices after completing all movements without incurring any additional cost. The solution to the MG-MAPF problem is a collection of agents’ collision-free paths, where the agent can start from its start vertex and visit all its goal vertices at least once with arbitrary visiting order. A collision occurs when two agents are located in the same vertex (vertex conflict) or moving along the same edge (edge conflict) at the same time step. The cost of an agent’s path is the total time steps used to visit all goals. We use the summation of the cost (SOC) as the objective of the problem, meaning that the cost of the solution is the summation of the cost of each agent’s path.

4 Case Study

This section will provide an example of the MG-MAPF problem to illustrate that the DVS method is not optimal.

Theorem 1. *The optimality of the methods based on decoupling the goal vertex visiting order search and single-agent pathfinding cannot be guaranteed in the MG-MAPF problem.*

Proof. An example that the DVS method cannot find the optimal solution is shown in Figure 2. In this example, agent 2 only has one possible goal vertex visiting order, while agent 1 has two possible situations. If agent 1 chooses to go to g_1^1 before going to g_1^2 , the lower bound of the total cost is $207 + 7 = 214$. If agent 1 chooses to travel to g_1^2 before g_1^1 , the lower bound of the total cost will be $111 + 7 = 118$. Now we consider to adopt the second choice. In this case, an edge conflict between vertex D to vertex E occurs at the time step 6, which is in the path of agent 1 from g_1^2 to g_1^1 . Based on

the definition of DVS, the path from g_1^2 to g_1^1 may be adjusted to avoid the conflict. For example, in HCBS, constraints are created on the path of g_1^2 to g_1^1 of agent 1. However, the path of agent 1 from the start vertex s_1 to the goal vertex g_1^2 will remain unchanged because the time step of conflict surpasses the time step of reaching g_1^2 . One of the minimal-cost solutions generated by the DVS method is that agent 1 goes to g_1^2 before g_1^1 , and agent 2 remains at vertex C for 8 time steps to give paths to agent 1 before proceeding directly to g_1^1 . The total cost is $111 + 15 = 126$. However, there exists an alternative solution with a lower total cost that deviates from the shortest path for agent 1 starting from s_1 towards g_1^2 . Agent 1 remains at vertex s_1 for 2 time steps. Subsequently, it proceeds to vertex g_1^2 and then travels to vertex g_1^1 . In this solution, the total cost is $113 + 7 = 120$, which is better than the minimal cost solution given by the DVS method among all goal vertex visiting orders. \square

The reason why the DVS method is not optimal is because it incorrectly assumes that the path reaching the goal vertex earlier is always better than the path reaching it later. Under the given goal vertex visiting order, it only generates one path to reach each goal at the earliest possible time step. However, in some cases, the optimal path might not reach a subset of goal vertices at the earliest possible time step. The DVS method misses these paths.

To obtain the optimal path, we propose to search based on DSS, which only assumes that the path reaching the GSI earlier is always better than the path reaching it later. In contrast to the DVS method, DSS can obtain multiple paths under a given goal vertex visiting order when there are multiple GSIs at some goal vertices. Considering the same case above, when the path between g_1^2 to g_1^1 is adjusted and causes conflicts at g_1^2 at a time step larger than the earliest time step to reach g_1^2 , more than one GSI will appear in g_1^2 , providing the potential to search the non-earliest path to g_1^1 from s_1 .

It is observed that the DVS method is more likely to obtain a non-optimal result in crowded scenarios. On the one hand, in scenarios with few agents and obstacles, the agent can temporarily move to the neighbor vertex to give paths to another agent and return to the original vertex once another agent has passed. In this way, the DVS method can generate a path from an earlier SI to a later SI without missing any potential optimal paths. On the other hand, in crowded scenarios, the agent may not be able to return to the original vertex quickly. It makes the agent infeasible for some time steps in the later SI, making the DVS method miss some potential optimal paths, which can be obtained by the DSS method.

5 Methodology

We propose an optimal and efficient two-level method, the MGCBS, for the MG-MAPF problem based on DSS. To improve the efficiency of the search process, we introduce the TIS Forest data structure to reduce redundant calculations of multiple queries. A TIS Forest corresponds to an agent and is composed of several Time-Interval-Space Trees (TIS Tree), which corresponds to a specific GSI of the agent. The TIS Tree can be used to get the shortest path and its length to the corresponding GSI from any start vertex at any start time step.

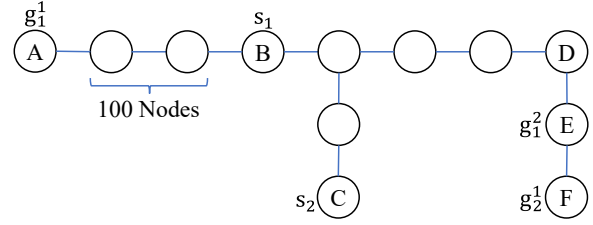


Figure 2: An example of the MG-MAPF with two agents is represented in a planar graph. The agent 1 starts at vertex B and has two goals located at vertex A and vertex E . The agent 2 starts at the vertex C and only has one goal at vertex F .

5.1 MGCBS

The MGCBS is a two-level solver that can effectively calculate the optimal path for each agent to visit all goals without conflict. The high-level solver employs a constraint tree to manage conflicts between agents, while the low-level solver uses an A*-based solver to find the best GSI visiting order.

High-level Solver

The high-level solver of the MGCBS is an extension of the high-level solver of the conflict-based search (CBS) [Sharon *et al.*, 2015]. It builds a constraint tree to solve conflicts between different agents. The constraints consist of vertex constraints, which prevent an agent from occupying a vertex at a specific time step, and edge constraints, which prohibit an agent from traversing an edge at a given time step. Each constraint tree node saves the TIS Forests, the constraint set, each agent’s path, and the value of SOC. The implementation of the TIS Forest will be discussed in subsection 5.2.

In the high-level solver, a distance table D is built to store the distance from each vertex to each goal vertex for each agent, which will be used in the low-level solver. It is observed that there are no vertex constraints at the root node of the constraint tree, causing only one GSI with a whole time interval at each goal vertex. Therefore, D can be built by querying the TIS Forest at the root node of the constraint tree.

The differences between the high-level solver of basic CBS and MGCBS are the operations on the TIS Forest and the distance table. At the beginning of the high-level solver of MGCBS, the TIS Forests in the root node are built and used to construct the distance table. Then, the TIS Forests, the distance table, and the agents’ task information are fed to the low-level solver to calculate the path that visits all goal vertices without considering other agents. When a new constraint tree node is generated, its TIS Forests are copied from its parent node, and one of them is reconstructed by the new constraint set. The new TIS Forest, the distance table, and the agent’s task information are put into the low-level solver to compute the shortest path under the constraints.

Algorithm 1 shows the pseudocode of the high-level solver. In lines 1 \sim 9, the root node is constructed and put into the open set. The TIS Forests are built in line 4, and the distance table is constructed in line 5. The initial path of each agent is calculated in line 6. In lines 11 \sim 16, the minimum cost node is found and checked whether it contains a conflict. If no, the final solution is found. Otherwise, the constraints are

Algorithm 1 $MGCBS_{high}$

Input: agents A , graph G

- 1: $R \leftarrow$ new node
- 2: $R.cons \leftarrow \emptyset$
- 3: **for** each agent a in A **do**
- 4: $R.tisf[a] \leftarrow BuildTISForest(G, a, \emptyset)$
- 5: $D[a_i] \leftarrow BuildDistanceTable(R.tisf[a])$
- 6: $R.paths[a] \leftarrow$
 $MGCBS_{low}(R.tisf[a], D[a], a, \emptyset)$ // Algorithm 2
- 7: **end for**
- 8: $R.cost \leftarrow$ calculate the SOC of $R.paths$.
- 9: $OPEN \leftarrow \{R\}$
- 10: **while** $OPEN \neq \emptyset$ **do**
- 11: $N \leftarrow$ minimum cost node from $OPEN$.
- 12: $OPEN \leftarrow OPEN \setminus \{N\}$
- 13: $F \leftarrow$ the earliest collision in N
- 14: **if** F is *None* **then**
- 15: **return** $N.paths$
- 16: **end if**
- 17: $C \leftarrow$ build constraints from F
- 18: **for** constraint $c = (a, t, v)/(a, t, v_i, v_j)$ in C **do**
- 19: $P \leftarrow$ new node
- 20: $P.tisf, P.paths \leftarrow N.tisf, N.paths$
- 21: $P.cons \leftarrow N.cons \cup c$
- 22: $P.tisf[a] = BuildTISForest(G, a, P.cons)$
- 23: $P.paths[a] \leftarrow MGCBS_{low}(P.tisf[a],$
 $D[a], a, P.cons)$ // Algorithm 2
- 24: **if** $P.paths[a]$ is not *NULL* **then**
- 25: $P.cost \leftarrow$ calculate the SOC of $P.paths$
- 26: $OPEN \leftarrow OPEN \cup \{P\}$
- 27: **end if**
- 28: **end for**
- 29: **end while**
- 30: **return** *NULL*

built according to the earliest conflict. In lines 19 ~ 27, new nodes are constructed for each constraint, and the corresponding constraint is added to the constraint set. The TIS Trees that are related to the constrained agent are reconstructed according to the new constraint set in line 22. The low-level solver is called to build the path of a single agent in line 23.

Low-level Solver

The low-level solver computes the shortest path for a single agent under constraints, ensuring that each goal is visited at least once based on DSS. The low-level solver comprises two stages: the GSI visiting order search stage and the path-generating stage.

The GSI visiting order stage search uses an A* solver for the best GSI visiting order. The state in the search can be represented by (L, p) , where L is the visited goal set and p is the SI where the agent is currently located. The cost of the state (L, p) is the length of the path that has visited all the goals in L and is currently located in p . Let the g be the cost value which is the same as the current time step, the h be the heuristic value, and the f be the evaluation value. The cost of the minimum spanning tree (MST) of the currently located vertex and all unvisited goal vertices is used as the h . The distance table D can be utilized to construct the MST. During the search, the state (L, p) can transfer to $(L \cup \{q.v\}, q)$, where the q is a GSI whose vertex is not in L , and the $q.v$

Algorithm 2 $MGCBS_{low}$

Input: TIS Forest $tisf$, distance table D , agent a , constraint set $cons$

- 1: $R \leftarrow$ new node
- 2: $R.L \leftarrow \emptyset$
- 3: $R.p \leftarrow GetEarliestSI(a.start, cons)$
- 4: $R.g \leftarrow 0$
- 5: $R.h \leftarrow Cost(GetMST(D, R.L, (R.p).v))$
- 6: $R.f \leftarrow R.g + R.h$
- 7: $OPEN, CLOSED \leftarrow \{R\}, \emptyset$
- 8: **while** $OPEN \neq \emptyset$ **do**
- 9: $N \leftarrow$ minimum f node from $OPEN$.
- 10: $OPEN \leftarrow OPEN \setminus \{N\}$
- 11: $CLOSED \leftarrow CLOSED \cup \{N\}$
- 12: **if** $N.L$ contains all goals of a **then**
- 13: $order \leftarrow GenerateGoalVisitedOrder(N)$
- 14: $path \leftarrow GeneratePath(tisf, order)$
- 15: **return** $path$
- 16: **end if**
- 17: **for** q in $GetAllPotentialGSI(a, cons, N.L)$ **do**
- 18: $M \leftarrow GetNode(OPEN, CLOSED, N.L \cup \{q.v\}, q)$
- 19: **if** $M \notin CLOSED$ **then**
- 20: $tist \leftarrow GetTISTree(q, tisf)$
- 21: $len \leftarrow GetPathLength(tist, (N.p).v, N.g)$
- 22: **if** $M \in OPEN$ **then**
- 23: $M.g \leftarrow \min(M.g, N.g + len)$
- 24: $M.f \leftarrow M.g + M.h$
- 25: **else**
- 26: $M.g \leftarrow N.g + len$
- 27: $M.h \leftarrow Cost(GetMST(D, M.L, (M.p).v))$
- 28: $M.f \leftarrow M.g + M.h$
- 29: $OPEN \leftarrow OPEN \cup \{M\}$
- 30: **end if**
- 31: **end if**
- 32: **end for**
- 33: **end while**
- 34: **return** *NULL*

refers to the vertex corresponding to the q . It means that the agent moves from the $p.v$ at the start time step g to the GSI q , whose vertex is an unvisited goal vertex $q.v$. The transfer cost is the shortest path length from the $p.v$ at time step g to the GIS q under the constraint set. It can be directly obtained from the TIS Forest. Specially, if $q.v$ is the final unvisited goal, only the latest GSI at $q.v$ can be chosen as the next GSI for avoiding conflicts after the agent finishes at the final goal.

The path-generating stage is executed when the minimum-cost state that visits all goals is found. We backtrack the final state to the initial state to get the GSI visiting order and use the TIS Forest to build the path based on the order.

Algorithm 2 shows the pseudocode of the low-level solver. In lines 1 ~ 7, the initial node is initialized and put into the open set, while the $R.p$ is the earliest SI at the start vertex of the agent. The heuristic value of the node is calculated by the cost of the MST in line 5. At each iteration, the node with the minimum f value is popped from the open set (lines 9 ~ 11). If the current node N has already visited all goals, extract the GSI visiting order and then build the path (lines 12 ~ 16). Otherwise, all GSI at unvisited goals is enumerated with the exception of the last unvisited goal, for which only the latest GSI is considered (lines 17). Let q be the current enumerated

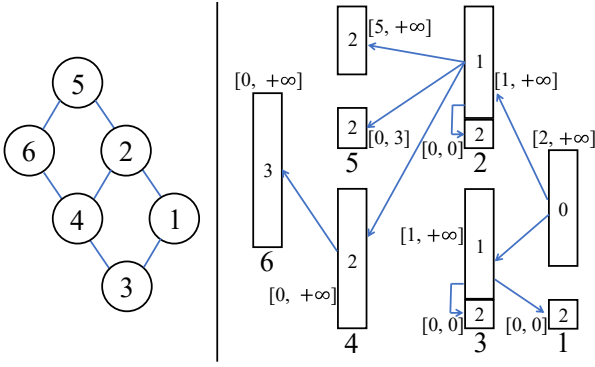


Figure 3: The left subfigure shows the graph on which the TIS Tree was built. The number in the circle is the vertex ID. The right subfigure shows the TIS Tree whose seed GSI is $[2, +\infty]$ at vertex 1. The constraint set includes two vertex constraints, including vertex 1 at time step 1, and vertex 5 at time step 4. The block represents the node in the tree, where the number in the block is the cost of the node. The number under the block is the vertex ID. The bracketed number represents the time interval of the TIS state of the node. The arrow represents the edge of the TIS Tree pointing from the parent node to the child node. The self-point arrow at vertices 2 and 3 indicates waiting for one time step.

GSI. A node M is constructed where $M.L = N.L \cup \{q.v\}$ and $M.p = q$ (line 18). The TIS Tree corresponding to q is filtered from the TIS Forest and used to obtain the path length from $(N.p).v$ at the time step $N.g$ to q (lines 20 ~ 21). M is updated if it can be improved by N (lines 22 ~ 30).

5.2 Time-Interval-Space Forest

The shortest path length to the GSI is frequently queried in the low-level solver. If the MGCBS employs the A* algorithm to obtain the path length, it will become time-consuming. Observing that the shortest path length to a specific GSI may be queried multiple times during the search, we consider utilizing a data structure to reduce the redundant computation. It is not trivial because the start time of each query is unknown before the search and can only be obtained after the path to the previous GSI is generated. Considering that the constraints are related to the time step, the optimal path of different start time steps might be diverse, making it hard to reuse the result of the previous search.

We propose TIS Forest to reduce redundant calculations. Each TIS Forest corresponds to an agent. It consists of several TIS Trees, each corresponding to a GSI of the agent. The TIS Tree maintains the shortest path and its length from any start vertex at any start time to the GSI. We name that GSI as the seed GSI of the TIS Tree.

Let (t, s) be a time-space state (TS state) representing that the agent is located at the vertex s at the time step t . Let $([t_0, t_1], s)$ be a time-interval-space state (TIS state), representing a collection of TS states at the same vertex, i.e., $([t_0, t_1], s) = \{(t, s) | t \in [t_0, t_1]\}$. A TIS state is safe if and only if it does not contain any TS states under vertex constraints. It should be noted that the SI is a special type of TIS state, while the TIS state does not always need to be maximal.

Each node in the TIS Tree represents a safe TIS state, containing the TS states that take the same vertex sequence as the shortest path to the seed GSI. Therefore, different TS states in a TIS state have the same shortest path length to the seed GSI. We define the cost of the node as the shortest path length of the TS states in the TIS state to the seed GSI. There might be more than one node at a vertex when some vertex constraints exist at the vertex.

The TIS Tree is constructed by a Dijkstra-based algorithm [Dijkstra, 1959]. Initially, we find all maximal TIS states (i.e., SI) at each vertex and create a node for each of them. We refer to the node corresponding to the seed GSI as the seed node. If the node is the seed node, its cost is set to 0; otherwise, its cost is set to $+\infty$. At the beginning, all nodes are put into an unvisited set U . At each iteration of the search, the node with the minimum cost in U pops out and is used to improve the path of its neighbor through the reverse edge. Specifically, let $([t_0, t_1], v)$ be the TIS state of the current minimum cost node n and $cost(n)$ be the cost of n . Let v' be a vertex that can take one action to transfer to v , i.e., $v' \in Nbr(v)$ and $Nbr(v) = \{v' | (v', v) \in E\} \cup \{v\}$. Now we consider how to improve the path of the nodes at v' by $([t_0, t_1], v)$. We construct a TIS state set denoted by B at vertex v' , such that each TS state in the TIS state is safe and can transfer to a TS state in $([t_0, t_1], v)$ at a time step. For example, $([t_0, t_1], v) = ([2, 8], v)$ and here is a vertex constraint at vertex v' at time step 3 and an edge constraint from vertex v' to vertex v at time step 6, the B will be $\{([1, 2], v'), ([4, 5], v'), ([7, 7], v')\}$. We enumerate all nodes at v' and all TIS states in B . Let $([t'_0, t'_1], v')$ be the TIS state of the current enumerated node n' and $([t_0^B, t_1^B], v')$ be the TIS state in B . If $cost(n') > cost(n) + 1$ and $[t'_0, t'_1]$ can be fully covered by $[t_0^B, t_1^B]$, the n' can be improved by setting n as its parent and $cost(n') = cost(n) + 1$. If $cost(n') > cost(n) + 1$ and only a part of $[t'_0, t'_1]$ is covered by $[t_0^B, t_1^B]$, n' will be divided into several new nodes according to the coverage, and only the new node whose time interval is fully covered by $[t_0^B, t_1^B]$ can be improved by n . When a node needs to be divided, the U is updated by deleting the original node and adding the new nodes. The search will stop when the U is empty. Figure 3 shows an example of the TIS Tree.

After the construction, the TIS Tree can be used to query the shortest path from any start vertex at any start time to the seed GSI by the following steps. Firstly, we construct a TS state according to the start time and vertex. Secondly, we find out the node whose TIS state contains the TS state. Finally, we backtrack this node to the seed node and construct the shortest path through the vertices of passed nodes. Furthermore, we can get the shortest path length directly using the cost of the node found in the second step without building the path explicitly. When we need to obtain the shortest path or its length to a GSI by a TIS Forest, we can extract the TIS Tree that corresponds to the GSI and use it for the query.

The concept of TIS Forest may seem similar to the SIPP algorithm [Phillips and Likhachev, 2011] as both the TIS Forest and the SIPP algorithm search based on time intervals. However, their underlying principles differ. The SIPP algorithm utilizes a forward search, with each SI containing a single

dominant time step, rendering all other time steps unimportant for the search. This property allows the SIPP algorithm to reduce the total number of nodes during the search. In contrast, the TIS Forest employs a backward search. On the one hand, we cannot use a mechanism similar to the SIPP algorithm to construct the TIS Forest because the precise goal-reaching time step is unknown beforehand, and the dominant time step doesn't exist. On the other hand, the TIS Forest is built by the principle that the shortest path to the seed GSI from all TS states within a TIS state have the same vertex sequence, and all TS states within a TIS state can be expanded simultaneously along a reverse edge.

6 Theoretical Analysis

In this section, we will prove the optimality and completeness of MGCBS.

Lemma 1. *For two TS states in the same SI, the minimum completion time of the path to visit a set of goals at least once from the earlier TS state will not exceed the time from the later TS state.*

Proof. We prove it by contradiction. Let (t_0, s) and (t_1, s) be two TS states in the same SI where $t_1 > t_0$. Define $C(t, s)$ as the minimum completion time of the path to visit a set of goals at least once from a TS state (t, s) . Suppose, towards a contradiction, that $C(t_1, s) < C(t_0, s)$. Consider the optimal path that achieves the completion time $C(t_0, s)$. This path can be modified to stay at vertex s until time t_1 and then follow the same sequence of vertices as the optimal path starting from (t_1, s) . The modified path would visit all goals no later than the optimal path from (t_1, s) , yielding a completion time that is at most $C(t_1, s)$, in direct contradiction to the supposition that $C(t_1, s) < C(t_0, s)$. \square

Lemma 2. *The path maintained in the TIS Tree is the optimal path to the seed GSI of the TIS tree under given constraints.*

Proof. The TIS Tree is constructed based on a backward version of the Dijkstra algorithm [Dijkstra, 1959]. The optimality of the TIS Tree can be guaranteed by the optimality of the Dijkstra algorithm. \square

Lemma 3. *The function `GeneratePath` in Algorithm 2 can obtain the optimal path following a given GSI visiting order under given constraints.*

Proof. The resulting path of `GeneratePath` is constructed by iteratively concatenating the subpath to the next GSI by the TIS Tree, and the length of the subpath is shortest according to the Lemma 2. The current Lemma can be proved by induction. The search starts from a TS state with time step 0. Assume that using this construction method can obtain the path with the minimum completion time following the first k GSI visiting order. According to Lemma 1, concatenating the shortest subpath to the $k + 1$ GSI can obtain the path with the minimum completion time following the first $k + 1$ GSI visiting order. Therefore, by the principle of induction, the final resulting path has the minimum completion time following the whole GSI visiting order. \square

Theorem 2. *MGCBS is an optimal solver of the MG-MAPF problem.*

Proof. In the low-level solver, the heuristic function is admissible and satisfies the consistency assumption according to the property of MST. The optimality of the low-level solver can be guaranteed by the Lemma 3 and the optimality of A* [Hart *et al.*, 1968]. According to the optimality of CBS and the low-level solver, the optimality of MGCBS can be guaranteed. \square

Theorem 3. *MGCBS is a complete solver of the MG-MAPF problem.*

Proof. The completeness can be proven by following steps. Firstly, based on the TIS forest design, if there is a feasible path r from one start/goal vertex u at time step t_u to reach one goal vertex v at time step t_v , there will not be a vertex constraint on vertex v at time step t_v . A TIS Tree must exist, whose seed GSI includes time step t_v at vertex v . This node can be iteratively expanded following the reverse direction of r to reach the node whose TIS state includes time step t_u at vertex u . Therefore, the TIS Tree is complete. Secondly, if there is a feasible solution for a multi-goal single-agent pathfinding problem, their corresponding GSI visiting order could be searched in the low-level solver, and the path could be constructed according to the completeness of A* and TIS Tree. Thirdly, MGCBS is complete according to the completeness of CBS and the low-level solver. \square

7 Experiment

We verify our proposed algorithm's effectiveness and optimality on the 4-neighbor grid maps. We randomly sample each agent's start and goal points in the grid map, ensuring that the start points of the different agents are distinct. The test computer is equipped with an I9-7900X CPU with 3.3 GHz and 32GB RAM. The code is publicly available at <https://github.com/tangmingkai/MGCBS>.

We use the following four algorithms for the experiments.

- **HCBS (A1):** A three-level algorithm [Surynek, 2021]. The highest level solver is a CBS algorithm, and the middle level solver is an A* algorithm that searches a goal vertex visiting order. The lowest level solver is a classic A* for single-agent pathfinding. It is the current SOTA method based on DVS.
- **MGCBS without TIS Forest (A2):** MGCBS with a modified low-level solver that uses the A* algorithm to obtain the shortest path and its length to a GSI.
- **MGCBS (A3):** Our propose method.
- **CBS + A* (A4):** An optimal algorithm that is the CBS with a coupled low-level search using an A* algorithm to search for the shortest path visiting all the goals.

7.1 Experiment for Efficiency

From the MAPF benchmark [Sturtevant, 2012], three grid maps from small to large are selected, namely 'maze-32-32-4', 'lak303d', and 'orz900d'. They are marked as M1, M2, and M3, as shown in Figure 4. We generate 100 instances for

Map	k	A1	A2	A3
M1	2	98%	98%	100%
	4	69%	69%	86%
	6	22%	22%	50%
	8	4%	4%	9%
M2	2	100%	100%	98%
	4	93%	93%	91%
	6	66%	66%	74%
	8	20%	20%	50%
M3	2	0%	0%	100%
	4	0%	0%	73%
	6	0%	0%	8%
	8	0%	0%	0%

 Table 1: Table of the success rate. k is the number of agents.

map	k	A1	A2	A3
M1	2	4.37(-)	4.52(0.97)	0.63(6.94)
	4	27.39(-)	27.84(0.98)	12.16(2.25)
	6	53.44(-)	53.69(1.00)	37.47(1.43)
	8	58.75(-)	58.75(1.00)	56.41(1.04)
M2	2	8.87(-)	8.87(1.00)	3.79(2.34)
	4	23.73(-)	23.79(1.00)	12.23(1.94)
	6	44.40(-)	44.40(1.00)	24.60(1.80)
	8	57.18(-)	57.23(1.00)	42.26(1.35)
M3	2	60.00(-)	60.00(1.00)	17.73(3.38)
	4	60.00(-)	60.00(1.00)	42.96(1.40)
	6	60.00(-)	60.00(1.00)	58.73(1.02)

 Table 2: Table of the average running time and the speedup ratio. The first number in the cell shows the average running time (sec), and the number in parentheses indicates the speedup ratio relative to A1. k is the number of agents. The result of $k = 8$ on M3 is omitted because there are no successful instances.

each number of agents in the range of $\{2, 4, 6, 8\}$. We fix each agent’s goal number to 12. We employ the average running time and the success rate as evaluation metrics for comparing the performance of A1, A2, and A3. A test case is unsuccessful if its running time exceeds 60 seconds, in which case the running time will be directly set to 60 seconds.

Table 1 and Table 2 show the success rate, the average running time, and the speedup ratio to A1 of the other two on M1, M2, and M3. In most instances, A3 outperforms A1 and A2 regarding average running time and success rate. Specifically, the speedup ratio of A3 to A1 is nearly 7 when the number of agents is two on M1. When the number of agents is smaller, the acceleration is more pronounced because when the number of agents is large, unsuccessful instances will smooth the speedup ratio. On M3, A1 and A2 cannot solve any instances, while A3 solves all instances successfully when the number of agents is 2. In small numbers of instances on M2 with few agents, the speedup by using TIS Forest to find the shortest path cannot overlap the construction overhead, making the success rate lower than A1 and A2. However, regarding the average running time, A3 performs the best on the same map and the number of agents. In some instances, A2 runs slightly slower than A1 because A2, which searches the GSI visiting order in the low-level solver, has a higher computation com-

Map	Algorithm	SN	DN	MRE	ARE
M4	A1	841	29	17.39%	0.19%
	A3	852	0	0.00%	0.00%
M5	A1	896	4	7.14%	0.03%
	A3	896	0	0.00%	0.00%

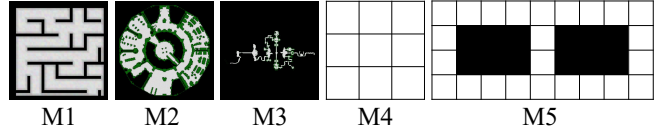
 Table 3: Table for evaluating the solution quality of A1 and A3. SN is the number of instances both the target algorithm and A4 can successfully run. DN is the number of instances where the solution cost produced by the target algorithm differs from those produced by A4. MRE and ARE are maximal and average relative error, respectively.


Figure 4: The grid maps used in the experiment.

plexity than A1, which only searches the goal vertex visiting order.

7.2 Experiment for Optimality

We use two self-defined grid maps, M4 and M5, in Figure 4, to evaluate the optimality of our proposed algorithm. We generate 100 instances for the number of agents and goals ranging from 2 to 4, with a total of 900 instances for each grid map. We utilize relative error to evaluate the solution quality with A4 and the target algorithm (A1 and A3), only considering instances where both A4 and the target algorithm are successful. Specifically, we count the instance number where their costs differ and compute the maximum and average values of the relative error.

Table 3 displays the number of instances in which the cost differs between the target algorithms and A4 across all successfully solved instances, as well as the maximum and average relative error. In some instances, A1 cannot obtain the optimal path with a maximal relative error exceeding 17%, while A3 can obtain optimal results among all instances.

8 Conclusion

This work used the case study and experiment to demonstrate that the method based on decoupling the goal vertex visiting order search and the single-agent pathfinding is not optimal for the multi-goal multi-agent pathfinding problem. Hence, we proposed a two-level optimal and efficient solver, MGCBS, decoupling the goal safe interval visiting order search and the single-agent pathfinding. To obtain the shortest path and its length to each goal safe interval efficiently, we proposed the Time-Interval-Space Forest to maintain the shortest path from any start vertex at any start time step to the goal safe interval. Experiments have shown that MGCBS can consistently obtain the optimal result and significantly outperform the SOTA decoupled method regarding running speed.

Acknowledgments

This work was supported by Guangdong Basic and Applied Basic Research Foundation (No. 2021B1515120032), and Guangzhou-HKUST(GZ) Joint Funding Program (No. 2024A03J0618).

References

- [Applegate *et al.*, 2011] David L Applegate, Robert E Bixby, Vašek Chvátal, and William J Cook. The traveling salesman problem. In *The Traveling Salesman Problem*. Princeton university press, 2011.
- [Čáp *et al.*, 2015] Michal Čáp, Peter Novák, Alexander Kleiner, and Martin Selecký. Prioritized planning algorithms for trajectory coordination of multiple mobile robots. *IEEE transactions on automation science and engineering*, 12(3):835–849, 2015.
- [Choudhury *et al.*, 2022] Shushman Choudhury, Kiril Solovey, Mykel Kochenderfer, and Marco Pavone. Coordinated multi-agent pathfinding for drones and trucks over road networks. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 272–280, 2022.
- [Dijkstra, 1959] Edsger W Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [Dresner and Stone, 2008] Kurt Dresner and Peter Stone. A multiagent approach to autonomous intersection management. *Journal of artificial intelligence research*, 31:591–656, 2008.
- [Hart *et al.*, 1968] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968.
- [Ma *et al.*, 2017a] Hang Ma, Jiaoyang Li, TK Satish Kumar, and Sven Koenig. Lifelong multi-agent path finding for online pickup and delivery tasks. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 837–845, 2017.
- [Ma *et al.*, 2017b] Hang Ma, Jingxing Yang, Liron Cohen, TK Satish Kumar, and Sven Koenig. Feasibility study: Moving non-homogeneous teams in congested video game environments. In *Thirteenth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2017.
- [Morris *et al.*, 2016] Robert Morris, Corina S Pasareanu, Kasper Luckow, Waqar Malik, Hang Ma, TK Satish Kumar, and Sven Koenig. Planning, scheduling and monitoring for airport surface operations. In *Workshops at the Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [Phillips and Likhachev, 2011] Mike Phillips and Maxim Likhachev. Sipp: Safe interval path planning for dynamic environments. In *2011 IEEE international conference on robotics and automation*, pages 5628–5635. IEEE, 2011.
- [Ren *et al.*, 2021] Zhongqiang Ren, Sivakumar Rathinam, and Howie Choset. Ms*: A new exact algorithm for multi-agent simultaneous multi-goal sequencing and path finding. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11560–11565. IEEE, 2021.
- [Ren *et al.*, 2022] Zhongqiang Ren, Sivakumar Rathinam, and Howie Choset. Conflict-based steiner search for multi-agent combinatorial path finding. *Proceedings of Robotics: Science and Systems, New York City, NY, USA*, 2022.
- [Salzman and Stern, 2020] Oren Salzman and Roni Stern. Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1711–1715, 2020.
- [Sharon *et al.*, 2015] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence*, 219:40–66, 2015.
- [Stern *et al.*, 2019] Roni Stern, Nathan R Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne T Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Satish Kumar, et al. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Twelfth Annual Symposium on Combinatorial Search*, 2019.
- [Sturtevant, 2012] Nathan R Sturtevant. Benchmarks for grid-based pathfinding. *IEEE Transactions on Computational Intelligence and AI in Games*, 4(2):144–148, 2012.
- [Surynek, 2021] Pavel Surynek. Multi-goal multi-agent path finding via decoupled and integrated goal vertex ordering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12409–12417, 2021.
- [Tjiharjadi *et al.*, 2022] Semuil Tjiharjadi, Szalinsyah Razali, and Hamzah Asyrani Sulaiman. A systematic literature review of multi-agent pathfinding for maze research. *Journal of Advances in Information Technology Vol*, 13(4), 2022.
- [Xu *et al.*, 2022] Qinghong Xu, Jiaoyang Li, Sven Koenig, and Hang Ma. Multi-goal multi-agent pickup and delivery. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9964–9971. IEEE, 2022.
- [Yu and LaValle, 2013] Jingjin Yu and Steven M LaValle. Structure and intractability of optimal multi-robot path planning on graphs. In *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [Zhang *et al.*, 2022] Han Zhang, Jingkai Chen, Jiaoyang Li, Brian C Williams, and Sven Koenig. Multi-agent path finding for precedence-constrained goal sequences. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 1464–1472, 2022.