

Offline Time-Independent Multi-Agent Path Planning

Keisuke Okmura*, François Bonnet, Yasumasa Tamura and Xavier Défago

Tokyo Institute of Technology

{okumura.k, bonnet.f, tamura.y, defago.x}@coord.c.titech.ac.jp

Abstract

This paper studies a novel planning problem for multiple agents that cannot share holding resources, named *OTIMAPP (Offline Time-Independent Multi-Agent Path Planning)*. Given a graph and a set of start-goal pairs, the problem consists in assigning a path to each agent such that every agent eventually reaches their goal without blocking each other, regardless of how the agents are being scheduled at runtime. The motivation stems from the nature of distributed environments that agents take actions fully asynchronous and have no knowledge about those exact timings of other actors. We present solution conditions, computational complexity, solvers, and robotic applications.

1 Introduction

The eventual goal of collective path planning for multiple agents is to make each agent in a shared workspace be on their respective goal status. This problem becomes non-trivial when agents cannot pass through each other, i.e., each agent occupies some resources in the space while the others are blocked to access these resources at that time. We see such situations in fleet operations of warehouses [Wurman *et al.*, 2008], intersection management for self-driving cars [Dresner and Stone, 2008], multi-robot 3D printing systems [Zhang *et al.*, 2018], packet-switched networks with limited buffer spaces [Tel, 2000], and lock operations of transactions on distributed databases [Knapp, 1987], to name just a few.

In such multi-agent systems, each agent inherently takes and finishes actions (or moves) *at their own timings independently and unpredictably from other actors*, regardless of centralized or decentralized controls. This is due to the nature of *distributed environments* such as message delay or clock shift/drift, as well as uncaptured individual differences between agents like frictions of physical robots. Nevertheless, the cutting-edge research on pathfinding for multiple agents, known as Multi-Agent Path Finding (MAPF) [Stern *et al.*, 2019] that aims at finding a set of collision-free paths on graphs, heavily rely on timing assumptions. Typical MAPF



Figure 1: Example of OTIMAPP. A graph is depicted with black lines. Two agents (i, j) and their paths are colored. *left*: Both agents stop progression permanently due to mutual exclusion (i.e., no collision) if i moved two steps before j moves. *right*: As long as each agent follows a respective path, both agents eventually reach their last vertex; these paths constitute an OTIMAPP solution.

assumes that agents take actions just at the same time. Not to mention, such “timed” schedules contradict the nature of distributed environments. Even worse, on-time execution of offline planning is too optimistic with more agents.

One counter approach to the timing uncertainties is runtime supports by online monitoring, re-planning, and intervention, e.g., [Van Den Berg *et al.*, 2011; Ma *et al.*, 2017; Atzmon *et al.*, 2020b; Okumura *et al.*, 2021]. This approach however requires runtime effort and additional infrastructures (e.g., steady network and monitoring systems) to manage agents’ status in real-time. Moreover, how to realize such schemes in large systems is not trivial at all.

Instead, this paper studies a novel planning problem in which agents spontaneously take actions without any timing assumptions. The problem requests a set of paths (i.e., solution) ensuring that all agents eventually reach their destinations without blocking each other permanently. To see this, consider the situation in Fig. 1(left). This plan runs a risk of execution failure; if the agent j gets delayed for any reason while the agent i moves two steps to the right, then each agent blocks each other and neither agent can progress on its respective path. In contrast, in Fig. 1(right), regardless of how the two agents are scheduled, both agents eventually reach their destinations unless they permanently stop the progression. We call the corresponding problem *Offline Time-Independent Multi-Agent Path Planning (OTIMAPP)*.

The contribution of this paper is to establish the foundation of OTIMAPP for both theory and practice. Specifically, the topics are categorized into two:

- *We formalize and analyze OTIMAPP.* Section 3 identifies a necessary and sufficient condition for a *solution*, i.e., a set of paths that makes all agents reach their goals without timing assumptions. This is based on characterization of *deadlocks*.

*Contact Author

Section 4 conducts a series of complexity analyses and reveals that (1) finding a solution is NP-hard on directed graphs, (2) finding a solution is NP-hard on undirected graphs when solutions are restricted to simple paths, and (3) verifying a solution is co-NP-complete.

► *We present algorithms to solve OTIMAPP and demonstrate the utility of OTIMAPP via robotic applications.* Section 5 presents two approaches to derive solutions: prioritized planning (PP) and deadlock-based search (DBS). Both algorithms are respectively derivative from basic MAPF algorithms [Erdmann and Lozano-Perez, 1987; Sharon *et al.*, 2015] and rely on a newly developed procedure to detect deadlocks within a set of paths. Section 6 shows that either PP or DBS can compute large OTIMAPP instances to some extent. Furthermore, we show that solutions keep robots’ moves efficient in an adverse environment for timing assumptions compared to existing approaches with runtime supports [Ma *et al.*, 2017; Okumura *et al.*, 2021]. Moreover, we demonstrate that solutions are executable with physical robots in both a centralized style and a decentralized style with only local interactions, without cumbersome procedures of online interventions.

In the remainder, all omitted proofs including sketches are available in the appendix. The appendix, code, and movie are available on <https://kei18.github.io/otimapp>. Related work will be discussed at the end.

2 Problem Definition

An *OTIMAPP instance* is given by a graph $G = (V, E)$, a set of agents $A = \{1, 2, \dots, N\}$, an injective initial state function $s : A \mapsto V$, and an injective goal state function $g : A \mapsto V$. An OTIMAPP instance on digraphs is similar to the undirected case.

An *execution schedule* is an infinite sequence of agents. An *OTIMAPP execution* is defined by an OTIMAPP instance, an execution schedule \mathcal{E} , and a set of paths $\{\pi_1, \dots, \pi_N\}$ as follows. The agents are *activated* in turn according to \mathcal{E} . Upon activation and until reaching the end of its path π_i , an agent i takes a single step along π_i if the vertex is vacant or stays at its current location otherwise. After reaching the end of the path, the agent only stays. \mathcal{E} is called *fair* when every agent appears infinitely-many times in \mathcal{E} .

An *OTIMAPP problem* is to decide whether there is a set of paths $\{\pi_1, \dots, \pi_N\}$ such that (1) each path for an agent i begins from $s(i)$ and ends at $g(i)$, (2) for any fair execution schedule, all agents reach the end of their paths (i.e., goals) in a finite number of activations. A *solution* is a set of paths satisfying these two.

Other Notations Let s_i and g_i denote $s(i)$ and $g(i)$, respectively. A location for an agent i is associated with a *progress index* $clock_i \in \{1, \dots, |\pi_i|\}$ and represented as $\pi_i[clock_i]$, where $\pi_i[j]$ is the j -th vertex in π_i . Every progress index starts at one and is incremented each time the agent moves a step along its path. The progress index is non-decreasing and no longer increases after reaching the end of the path. We use $S[-1]$ to denote the last element of the sequence S .



Figure 2: Examples of unreachable potential deadlocks. *left*: cyclic; $((i, j, k), (3, 1, 2))$. *right*: terminal; $(i, j, 2)$.

Rationale and remarks Any solution must deal with all timing uncertainties because execution schedules are unknown when offline planning. We assume that agents are activated sequentially and that each activation is atomic. However, there is no loss of generality as long as an agent can atomically reserve its destination before each move. Indeed, several robots acted simultaneously in our demos. Throughout the paper, we assume that each path π_i starts from s_i and ends at g_i to focus on analyses related to schedules.

3 Solution Analysis

Given a set of paths, our first question is to determine whether it is a solution. This section derives a necessary and sufficient condition for solutions. For this purpose, we introduce four types of *deadlocks*, categorized as; *cyclic* or *terminal*; *potential* or *reachable*. Informally, a cyclic deadlock is a situation where agent i wants to move to the current vertex of j , who wants to move to ... of i . A terminal deadlock is a situation where agent i reaches its destination and blocks the progress of another agent j . A potential deadlock is called reachable when there exists an execution schedule leading to the deadlock.

Definition 3.1 (potential cyclic deadlock). *Given an OTIMAPP instance and a set of paths $\{\pi_1, \dots, \pi_N\}$, a potential cyclic deadlock is a pair of tuples $((i, j, k, \dots, l), (t_i, t_j, t_k, \dots, t_l))$ such that $\pi_i[t_i + 1] = \pi_j[t_j] \wedge \pi_j[t_j + 1] = \pi_k[t_k] \wedge \dots \wedge \pi_l[t_l + 1] = \pi_i[t_i]$. The elements of the first tuple are without duplicates.*

Definition 3.2 (potential terminal deadlock). *Given an OTIMAPP instance and a set of paths $\{\pi_1, \dots, \pi_N\}$, a potential terminal deadlock is a tuple (i, j, t_j) such that $\pi_i[-1] = \pi_j[t_j]$ and $i \neq j$.*

Definition 3.3 (reachable cyclic deadlock). *A potential cyclic deadlock $((i, j, \dots, l), (t_i, t_j, \dots, t_l))$ is reachable when there is an execution schedule leading to a situation where $clock_i = t_i \wedge clock_j = t_j \wedge \dots \wedge clock_l = t_l$. This deadlock is called a reachable cyclic deadlock.*

Definition 3.4 (reachable terminal deadlock). *A potential terminal deadlock (i, j, t_j) is reachable when there is an execution schedule leading to a situation where $clock_i = |\pi_i| \wedge clock_j = t_j - 1$. This deadlock is called a reachable terminal deadlock.*

We refer to both reachable (or potential) cyclic/terminal deadlocks by reachable (resp. potential) deadlocks and simply use “deadlock” whenever the context is obvious. At least one execution schedule is required to verify whether a potential deadlock is reachable. For instance, in Fig. 1 (left),

a schedule (i, i, \dots) is evidence. A potential deadlock is not always reachable as illustrated in Fig. 2.

Theorem 3.5 (necessary and sufficient condition). *Given an OTIMAPP instance, a set of path $\{\pi_1, \dots, \pi_N\}$ is a solution if and only if there are (1) no reachable terminal deadlocks and (2) no reachable cyclic deadlocks.*

Proof sketch. Verifying that they are necessary is straightforward. To see that they are sufficient, consider a potential function $\phi := \sum_{i \in A} (|\pi_i| - \text{clock}_i)$ defined over a configuration $\{\text{clock}_1, \dots, \text{clock}_N\}$. Observe that ϕ is non-increasing and $\phi = 0$ means that all agents have reached their goals. Furthermore, when $\phi > 0$, ϕ is guaranteed to decrease if each agent is activated at least once. \square

4 Computational Complexity

This section studies the complexity of OTIMAPP. In particular, we address two questions: the difficulty to find solutions (Sec. 4.1) and the difficulty to verify solutions (Sec. 4.2). Our main results are that both problems are computationally intractable; the former is NP-hard and the latter is co-NP-complete. Both proofs are based on reductions from the 3-SAT problem, deciding satisfiability for a formula in conjunctive normal form with three literals in each clause.

4.1 Finding Solutions

We distinguish directed graphs and undirected graphs to analyze the complexity. The following proof is partially inspired by the NP-hardness of MAPF on digraphs [Nebel, 2020].

Theorem 4.1 (complexity on digraphs). *OTIMAPP on directed graphs is NP-hard.*

Proof. The proof is a reduction from the 3-SAT problem. Figure 3 is an example of the reduction from a formula $(x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee x_3)$.

A. Construction of an OTIMAPP instance. We introduce two gadgets, called *variable decider* and *clause constrainer*. The OTIMAPP instance contains one variable decider for each variable and one clause constrainer for each clause.

The variable decider for a variable x_i assigns x_i to true or false. This gadget contains one agent χ_i with two paths to reach its goal: *left* or *right*. Taking a left path corresponds to assigning x_i to false, and vice versa. For the j -th clause C^j in the formula, when its k -th literal is either x_i or $\neg x_i$, we further add one agent c_k^j to the gadget. Its start and goal are positioned on the *right* side from χ_i when the literal is a negation; otherwise, on the *left* side. When several such agents are positioned on one side, let them connect (see the gadget for x_2). c_k^j has two alternate paths to reach its goal: a path within the variable decider or a path via a clause constrainer. The former is available only when χ_i takes a path of the opposite direction to avoid a reachable cyclic deadlock.

The clause constrainer for a clause C^j connects the start and the goal of c_k^j . The gadget contains a triangle. Each literal c_k^j enters this triangle from a distinct vertex and exits from another vertex. As a result, this gadget prevents three literals in C^j from being false simultaneously; if not so, three agents enter the gadget and there is a reachable cyclic deadlock.

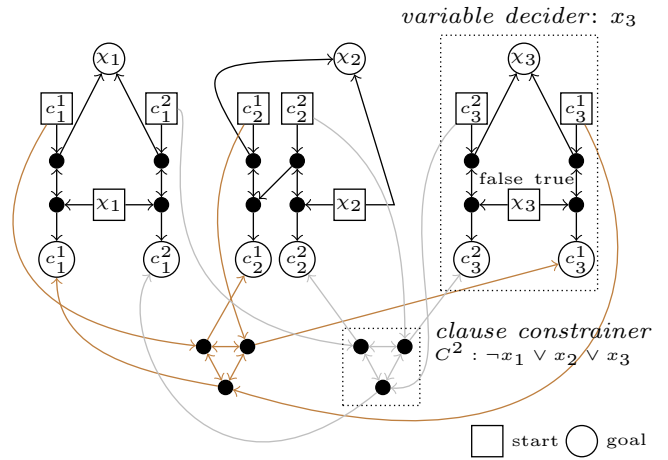


Figure 3: An OTIMAPP instance reduced from the 3-SAT formula $(x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee x_3)$.




Figure 4: A *oneway constrainer*. This gadget transforms an undirected edge (u, v) to a directed one. Any agent is allowed to move only the way from u to v when limiting solutions to simple paths.

The number of agents, vertices, and edges are all polynomial with respect to the size of the formula.

B. The formula is satisfiable if OTIMAPP has a solution: the use of one clause constrainer by three agents leads to a reachable cyclic deadlock. Thus, at least one literal for each clause becomes true in any OTIMAPP solution.

C. OTIMAPP has a solution if the formula is satisfiable: If satisfiable, let χ_i take a path that follows the assignment. Let c_k^j take a path within the variable decider when χ_i takes the opposite direction; otherwise, use the clause constrainer. Since three agents never enter one clause constrainer due to satisfiability, those paths constitute a solution. \square

For undirected graphs, we limit solutions to those containing only simple paths.¹

Theorem 4.2 (complexity on undirected graphs). *For OTIMAPP on undirected graphs, it is NP-hard to find a solution with simple paths.*

Proof sketch. We add a new gadget called *oneway constrainer*, which transforms an undirected edge to a virtually directed one, to the proof of the NP-hardness on digraphs (Thm. 4.1). We derive the claim by replacing all directed edges, except for bidirectional edges, with this gadget. Figure 4 illustrates it, including two new agents: z_1 and z_2 . In

¹We recently proved that it is NP-hard for the general case of undirected graphs. The formal proof will appear soon.

this gadget, any agents outside of the gadget are allowed to move only in the direction from u to v . \square

4.2 Verification

The co-NP completeness of the verification relies on the following lemma, stating that finding cyclic deadlocks is computationally intractable. Its entire proof is delivered in the Appendix.

Lemma 4.3 (complexity of detecting cyclic deadlocks). *Determining whether a set of paths contains either reachable or potential cyclic deadlocks is NP-complete.*

We then derive the complexity result since a solution has no reachable deadlocks.

Theorem 4.4 (complexity of verification). *Verifying a solution of OTIMAPP is co-NP-complete.*

Proof. Thm. 3.5 states that a solution has no reachable terminal/cyclic deadlocks. Verifying no terminal deadlocks is in co-NP; a terminal deadlock is verified in polynomial time with an execution schedule. Verifying no potential deadlocks is co-NP-complete according to Lemma 4.3. \square

5 Solvers

We now focus on how to solve OTIMAPP. In practice, it is difficult to use the necessary and sufficient condition (Thm. 3.5) because we have to find corresponding schedules. This motivates to build a relaxed sufficient condition.

Theorem 5.1 (relaxed condition). *Given an OTIMAPP instance, a set of path $\{\pi_1, \dots, \pi_N\}$ is a solution when there are (1) no use of other goals, i.e., $g_j \notin \pi_i$ for all $i \neq j$ except for $s_i = g_j$, and (2) no potential cyclic deadlocks.*

It is straightforward to see that the above conditions are respectively sufficient for the two conditions in Thm. 3.5. Given a set of paths, “no use of other goals” is easy to check while “no potential cyclic deadlocks” is intractable to compute (Lemma 4.3). Nevertheless, *detecting potential cyclic deadlock is the heart of solving OTIMAPP*. Thus, we first explain how to detect potential cyclic deadlocks. After that, two algorithms to solve OTIMAPP are presented.

5.1 Detection of Potential Deadlocks

Due to the space limit, we only describe the intuition behind the algorithm. The details are in the Appendix (Alg. 3). We first introduce a *fragment*, a candidate of potential cyclic deadlocks.

Definition 5.2 (fragment). *Given a set of paths $\{\pi_1, \dots, \pi_N\}$, a fragment is a tuple $((i, j, k, \dots, l), (t_i, t_j, t_k, \dots, t_l))$ such that $\pi_i[t_i + 1] = \pi_j[t_j] \wedge \pi_j[t_j + 1] = \pi_k[t_k] \wedge \dots = \pi_l[t_l]$. The elements of the first tuple are without duplicates.*

We say that a fragment *starts* from a vertex u when $\pi_i[t_i] = u$ and a fragment *ends* at a vertex v when $\pi_l[t_l + 1] = v$. A fragment that ends at its start (i.e., $\pi_l[t_l + 1] = \pi_i[t_i]$) is a potential cyclic deadlock.

Using fragments, we construct an algorithm to detect a potential cyclic deadlock in a set of paths if it exists. This is based on induction on π_i . The induction hypothesis for i is that there are no potential cyclic deadlocks for

induction	key	new fragments
$\{\pi_1\}$	u	$[(1), (1), (u, v)]$
	v	$[(1), (2), (v, w)]$
$\{\pi_1, \pi_2\}$	u	$[(1, 2), (1, 1), (u, v, x)]$
	v	$[(2), (1), (v, x)]$
	x	$[(2), (2), (x, y)]$
$\{\pi_1, \pi_2, \pi_3\}$	u	$[(1, 2, 3), (1, 1, 2), (u, v, x, u)]$
	v	$[(2, 3), (1, 2), (v, x, u)]$
	x	$[(3), (2), (x, u)], [(3, 1), (2, 1), (x, u, v)]$
	z	$[(3), (1), (z, x)], [(3, 2), (1, 2), (z, x, y)]$

Table 1: Example of detecting potential cyclic deadlocks. We describe the update of Θ_s for $\pi_1 = (u, v, w)$, $\pi_2 = (v, x, y)$, $\pi_3 = (z, x, u)$. The table uses [(agents), (progress indexes), (path)] as a notation of fragment, where “path” is a corresponding sequence of vertices of the fragment. The algorithm halts with a blue-colored fragment, a detected potential cyclic deadlock.

$\{\pi_1, \dots, \pi_{i-1}\}$ and all fragments for them are identified. All new fragments about π_i are categorized into three groups: (1) a fragment only with π_i , (2) a fragment that extends existing fragments, and (3) a fragment that connects existing two fragments. In either case, if a newly created fragment ends at its start, this is a deadlock.

The algorithm realizes this procedure by managing two tables that store fragments: Θ_s and Θ_t . Both tables take one vertex as a key. One entry in Θ_s stores all fragments starting from the vertex. One entry in Θ_t stores all fragments ending at the vertex. Table 1 presents an example to detect deadlocks.

5.2 Prioritized Planning (PP)

Prioritized planning [Erdmann and Lozano-Perez, 1987; Silver, 2005] is neither complete nor optimal, but it is computationally cheap hence a popular approach to MAPF. It plans paths sequentially while avoiding collisions with previously planned paths. Instead of inter-agent collisions, solvers for OTIMAPP have to care about potential cyclic deadlocks.

Algorithm 1 is prioritized planning for OTIMAPP, named *PP*. When planning a single-agent path, *PP* avoids using (1) goals of other agents and (2) edges causing potential cyclic deadlocks [Line 3]. The latter is detected by storing all fragments created by previously computed paths. For this purpose, *PP* uses the adaptive version of Alg. 3 [Line 5] in the Appendix. A path satisfying the constraints can be found by ordinary pathfinding algorithms. If not, *PP* returns **FAILURE**. The correctness of *PP* is derived from Thm. 5.1.

PP is simple but incomplete. In particular, the planning order of agents is crucial; an instance may be solved or may not be solved as illustrated in Fig. 5. One resolution is repeating *PP* with random priorities until the problem is solved; let call this *PP*⁺. However, finding good orders can be challenging because there are $|A|!$ patterns. This motivates us to develop a search-based solver, described in the next.

5.3 Deadlock-based Search (DBS)

We present *deadlock-based search (DBS)* to solve OTIMAPP, based on a popular search-based MAPF solver called *conflict-based search (CBS)* [Sharon et al., 2015]. CBS uses a two-

Algorithm 1 PP: Prioritized Planning

Input: an OTIMAPP instance

Output: a solution $\{\pi_1, \dots, \pi_N\}$ or **FAILURE**

```

1:  $\Theta_s, \Theta_t \leftarrow \emptyset$ 
2: for  $i = 1 \dots |A|$  do
3:    $\pi_i \leftarrow$  a path for agent  $i$  while avoiding the use of
     ·  $g_j, \forall j \neq i$ , except for  $s_i$ 
     ·  $(u, v) \in E$  s.t.  $\exists \theta \in \Theta_t[u]$  and  $\theta$  starts from  $v$ 
        $\triangleright$  avoiding cyclic deadlocks for  $\pi_j, j < i$ 
4:   if  $\pi_i$  is not found then return FAILURE
5:   update  $\Theta_s$  and  $\Theta_t$  with  $\pi_i$  using Algorithm 3
6: end for
7: return  $\{\pi_1, \dots, \pi_N\}$ 
    
```




Figure 5: Example that the planning order affects the solvability. When i plans prior to j , PP results in success with solid lines. PP fails if j plans first and takes the dotted line.

level search. The high-level search manages collisions between agents. When a collision occurs between two agents at some time and location, two possible resolutions are depending on which agent gets to use the location at that time. Following this observation, CBS constructs a binary tree where each node includes constraints prohibiting to use space-time pairs for certain agents. In the low-level search, agents find a single path constrained by the corresponding high-level node.

Instead of collisions, DBS considers potential cyclic deadlocks. When detecting a deadlock in a set of paths, a resolution is that one of the agents in the deadlock avoids using the edge. Thus, the constraints identify which agents prohibit using which edges in which orientation.

Algorithm 2 describes the high-level search of DBS. Each node in the high-level search contains *constraints*, a list of tuples consisting of one agent and two vertices (representing “from vertex” and “to vertex”), and *paths* as a solution candidate. The root node does not have any constraints [Line 1]. Its paths are computed following “no use of other goals” of Thm. 5.1 [Line 2]. Then, the node is inserted into a priority queue *OPEN* [Line 3]. In the main loop [Line 4–13], DBS repeats; (1) Picking up one node [Line 5]. (2) Checking a deadlock and creating constraints [Line 6]. (3) Returning a solution if the paths contain no deadlocks [Line 7]. (4) If not, creating successors and inserting them to *OPEN* [Line 8–12]. DBS returns **FAILURE** when *OPEN* becomes empty [Line 14]. We complement several details below.

► *Line 5:* *OPEN* is a priority queue and needs the order of nodes. DBS works in any order but good orders reduce the search effort. As effective heuristics, we use the descending order of the number of deadlocks with two agents, which is

Algorithm 2 DBS: Deadlock-based Search

Input: an OTIMAPP instance

Output: a solution $\{\pi_1, \dots, \pi_N\}$ or **FAILURE**

```

1:  $R.constraints \leftarrow \emptyset$ 
2:  $R.paths \leftarrow$  find paths with “no use of other goals”
3: insert  $R$  to  $OPEN$   $\triangleright OPEN$ : priority queue
4: while  $OPEN \neq \emptyset$  do
5:    $N \leftarrow OPEN.pop()$ 
6:    $C \leftarrow$  get constraints of  $N$  using Algorithm 3
7:   if  $C = \emptyset$  then return  $N.paths$ 
8:   for  $(i, u, v) \in C$  do
9:      $N' \leftarrow \{constraints : N.constraints + (i, u, v),$ 
        $paths : N.paths\}$ 
10:    update  $\pi_i$  in  $N'.paths$  to follow  $N'.constraints$ 
11:    if  $\pi_i$  is found then insert  $N'$  to  $OPEN$ 
12:  end for
13: end while
14: return FAILURE
    
```

computed within a reasonable time.

► *Line 6:* Let $((i, j, k, \dots, l), (t_i, t_j, t_k, \dots, t_l))$ be a returned deadlock by Alg. 3. Then, create constraints $(i, \pi_i[t_i], \pi_i[t_i + 1]), (j, \pi_j[t_j], \pi_j[t_j + 1]), \dots, (l, \pi_l[t_l], \pi_l[t_l + 1])$.

► *Line 10:* forces one path π_i in the node to follow the new constraints. This low-level search must follow “no use of other goals,” furthermore, all edges in the constraints for i . If not found, DBS discards the corresponding successor.

Theorem 5.3 (DBS). *DBS returns a solution when solutions satisfying Thm. 5.1 exist; otherwise returns FAILURE.*

Example We describe an example of DBS using Fig. 5. Assume that the initial path of i is the solid blue line and the path for j is the dashed red line [Line 2]. This node is inserted into *OPEN* [Line 3] and is expanded immediately [Line 5]. There is one potential cyclic deadlock in the paths then two constraints are created: either i or j avoids using the shared edge [Line 10]. Two child nodes are generated, however, the node that changes j ’s path is invalid because there is no such path without the use of the goal of j . Another one is valid; j takes the solid red line. Therefore, one node is added to *OPEN* from the root node. In the next iteration, this newly added node is expanded. There are no potential cyclic deadlocks in this node; DBS returns its paths as a solution.

Optimization Although this paper focuses on a feasibility problem, DBS can adapt to optimization problems. As objective functions, total path length and maximum path length in a solution can be defined. Those optimization problems are solved optimally by DBS when it prioritizes high-level search nodes with smaller scores, as commonly done in CBS. Note that metrics that assess time aspects such as total traveling time used in MAPF studies are significantly affected by execution schedules; the adaptation is not trivial.

6 Evaluation

This section empirically demonstrates that OTIMAPP solutions are computable to some extent (Sec. 6.1) and they




Figure 6: Stress test on 4-connected grids. The success rate is based on 25 identical instances. DBS* includes detected instances that are unsolvable for DBS before timeout, which is not possible for PP(+). We also present accumulated runtime with a fixed number of agents over 100 instances, and runtime profiling (median) of each solver over success instances for both solvers.

are useful in adverse environments about timings (Sec. 6.2) through the simulation experiments. We also present OTIMAPP execution with robots (Sec. 6.3). The simulator was coded in C++ and the experiments were run on a desktop PC with Intel Core i9 2.8 GHz CPU and 64 GB RAM.

6.1 Stress Test

Setup Each solver was tested with a timeout of 5 min on four-connected undirected grids picked up from [Stern *et al.*, 2019], as a graph G . We also tested random graphs, shown in the Appendix. All instances were generated by setting a start s_i and a goal g_i randomly while ensuring that s_i and g_i have at least one path without the use of other goals; otherwise, it violates “no use of other goals” of Thm. 5.1. Note that unsolvable instances might still be included.

Result Fig. 6 presents the results. The main findings are: (1) Both solvers can solve instances with tens of agents in various maps within a reasonable time. (2) PP often fails due to priority orders (e.g., Fig. 5) while PP+ and DBS can overcome such limitations to some extent. (3) A bottleneck of each solver is the procedure of detecting potential cyclic deadlocks, an NP-hard problem (Lemma. 4.3). This also leads to similar success rates of PP+ and DBS.

6.2 Delay Tolerance

We next show that OTIMAPP solutions (if found) are useful in a simulated environment with stochastic delays of agents’ moves built on conventional MAPF, called MAPF-DP (with Delay Probabilities) [Ma *et al.*, 2017]. Given a graph and start-goal pairs for each agent, the aim of MAPF is to move agents to their goals without collisions. Collisions occur when two agents occupy the same vertex or traverse the

$ A = 35$	$\bar{p} = 0.2$	$\bar{p} = 0.5$	$\bar{p} = 0.8$
MCPs+ECBS	1015 (1004,1026)	1422 (1404,1440)	2551 (2507,2596)
Causal-PIBT	986 (976,995)	1238 (1225,1250)	1841 (1816,1866)
OTIMAPP	941 (931,951)	1178 (1165,1190)	1730 (1707,1752)
$\bar{p} = 0.5$	$ A = 20$	$ A = 40$	$ A = 60$
MCPs+ECBS	724 (711,736)	1698 (1678,1716)	2938 (2911,2964)
Causal-PIBT	662 (653,671)	1466 (1453,1479)	2425 (2405,2444)
OTIMAPP	639 (631,648)	1395 (1383,1408)	2328 (2311,2345)

Table 2: Total traveling time on MAPF-DP. All settings used *random-32-32-10*. For each setting, we first picked up 10 instances that OTIMAPP solutions were found by PP+. For each instance and approach, we then performed 50 trials while changing the random seed. Thus, the scores are means on 500 executions, accompanied with 95% confidence intervals. *upper*: Results of changing \bar{p} while fixing $|A|$. *lower*: Results of changing $|A|$ while fixing \bar{p} . Note that the probability that someone delays increases with more agents.

same edge simultaneously. Time is discrete. All agents synchronously take actions, i.e., either move to an adjacent vertex or stay at the current location. MAPF-DP emulates the imperfect execution of MAPF by introducing the possibility p_i of unsuccessful moves for agent i (remaining there).

Setup The delay probabilities p_i were chosen uniformly at random from $[0, \bar{p}]$, where \bar{p} is the upper bound of p_i . The higher \bar{p} means that agents delay often, and vice versa. The metric is the total traveling time of agents; smaller values mean less wasting time at runtime. We tested the following two as baselines: (1) MCPs [Ma *et al.*, 2017] force agents to preserve order relations of visiting one vertex in an offline MAPF plan at runtime. The plan was obtained by ECBS [Barer *et al.*, 2014]. (2) Causal-PIBT [Okumura *et al.*, 2021] is online time-independent planning, that is, each agent repeats one-step planning and action adaptively to surrounding current situations. The other details are in the Appendix.

Result Table 2 shows that the execution of OTIMAPP solutions outperforms the alternatives. This is because: (1) Unlike MCPs, OTIMAPP solutions are free from temporal dependencies of offline plans that one agent delays are possibly fatal. (2) Unlike Causal-PIBT, agents follow long-term plans and avoid possible congested locations.

Discussion Although finding OTIMAPP solutions is challenging, Table 2 motivates us to compute them. Meanwhile, the other approaches can solve larger instances with more agents (e.g., $|A| = 200$) and with much smaller planning time than solving OTIMAPP. Moreover, there are situations where OTIMAPP has no solutions while the others can find feasible plans because OTIMAPP assumes no intervention at runtime. One future direction pursues to fill these gaps.

6.3 Robot Demonstrations

We present two OTIMAPP execution styles: (1) a *centralized* control using the *toio* robots (<https://toio.io>) and (2) a *decentralized* one with only local interactions using a multi-robot platform [Kameyama *et al.*, 2021]. A solution was obtained by DBS. Figure 7 is snapshots. A video is available online.



Figure 7: An OTIMAPP execution with 10 robots in an 8×8 grid. Colored arrows represent an OTIMAPP solution.

In both cases, robots move without any synchronization procedures but are ensured to eventually reach their goals thanks to the nature of OTIMAPP. Moreover, for the latter, any actor has no methods to know the entire configuration, which cannot be addressed by conventional execution strategies.

7 Related Work

A *deadlock* [Coffman *et al.*, 1971] is a widely recognized phenomenon not limited to robotics; a system state that several components claim resources that others hold, then block each other permanently. Strategies to cope with deadlocks are categorized into prevention, detection/recovery, and avoidance [Silberschatz *et al.*, 2006; Fanti and Zhou, 2004]; *OTIMAPP is prevention*. A non-deadlock state that is “inevitable” to reach deadlocks is called *unsafe* [Silberschatz *et al.*, 2006]. Meanwhile, reachable deadlocks of OTIMAPP correspond to states that are “possible” to reach deadlocks. The notion of potential terminal deadlock is related to well-formed instances of MAPF [Čáp *et al.*, 2015], that is, for each start-goal pair, a path exists that traverses no other starts and goals. The notion of reachable cyclic deadlock is mentioned as nonlive states/sets for deadlock management in automated manufacturing systems [Fanti and Zhou, 2004] or in a multi-robot scheduling problem [Mannucci *et al.*, 2021].

The *multi-agent pathfinding (MAPF)* problem [Stern *et al.*, 2019] aims at finding a set of collision-free paths on a graph. Many studies on MAPF consider timing uncertainties because they are inevitable in multi-agent scenarios. However, current methods largely rely on additional assumptions on the travel speed of agents or assume delays to follow some probability distributions [Wagner and Choset, 2017; Mansouri *et al.*, 2019; Peltzer *et al.*, 2020; Atzmon *et al.*, 2020a]. Failing to represent the inherent uncertainty in the domain means the system behavior can be unpredictable. Alternative approaches are online intervention during execution, e.g., forcing agents to preserve temporal dependencies of offline planning via communication [Ma *et al.*, 2017; Hömig *et al.*, 2019; Atzmon *et al.*, 2020b]. Another direction is online time-independent planning [Okumura *et al.*, 2021] that incrementally moves agents based on current situations. OTIMAPP shares the concept of time independence but aims at offline planning without or less runtime effort.

In graph theory, the (vertex) disjoint path problem and its variants [Robertson and Seymour, 1985] are partly related to ours in the sense that a set of disjoint paths clearly satisfies the solution condition of OTIMAPP, but the reverse does not.

8 Conclusion

This paper studied a novel path planning problem called OTIMAPP, motivated by the nature of distributed environments (i.e., timing uncertainties) that multi-agent systems must address. We focused on robotic applications in evaluation but believe that OTIMAPP can be leveraged to other resource allocation problems with mutual exclusion, e.g. distributed databases, which is our future direction.

Acknowledgments

We thank the anonymous reviewers for their many insightful comments and suggestions. This work was partly supported by JSPS KAKENHI Grant Numbers 20J23011, 21K11748, and 21H03423. Keisuke Okumura thanks the support of the Yoshida Scholarship Foundation.

References

- [Atzmon *et al.*, 2020a] Dor Atzmon, Roni Stern, Ariel Felner, Nathan R Sturtevant, and Sven Koenig. Probabilistic robust multi-agent path finding. In *Proceedings of International Conference on Automated Planning and Scheduling (ICAPS)*, 2020.
- [Atzmon *et al.*, 2020b] Dor Atzmon, Roni Stern, Ariel Felner, Glenn Wagner, Roman Barták, and Neng-Fa Zhou. Robust multi-agent path finding and executing. *Journal of Artificial Intelligence Research (JAIR)*, 2020.
- [Barer *et al.*, 2014] Max Barer, Guni Sharon, Roni Stern, and Ariel Felner. Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Proceedings of Annual Symposium on Combinatorial Search (SOCS)*, 2014.
- [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 (T-ASE)*, 2015.
- [Coffman *et al.*, 1971] Edward G Coffman, Melanie Elphick, and Arie Shoshani. System deadlocks. *ACM Computing Surveys (CSUR)*, 1971.
- [Dresner and Stone, 2008] Kurt Dresner and Peter Stone. A multiagent approach to autonomous intersection management. *Journal of Artificial Intelligence Research (JAIR)*, 2008.
- [Erdmann and Lozano-Perez, 1987] Michael Erdmann and Tomas Lozano-Perez. On multiple moving objects. *Algorithmica*, 1987.
- [Fanti and Zhou, 2004] Maria Pia Fanti and MengChu Zhou. Deadlock control methods in automated manufacturing systems. *IEEE Transactions on systems, man, and cybernetics-part A: systems and humans*, 2004.
- [Hömig *et al.*, 2019] Wolfgang Hömig, Scott Kiesel, Andrew Tinka, Joseph W Durham, and Nora Ayanian. Persistent and robust execution of mapf schedules in warehouses. *IEEE Robotics and Automation Letters (RA-L)*, 2019.

- [Kameyama *et al.*, 2021] Shota Kameyama, Keisuke Okumura, Yasumasa Tamura, and Xavier Défago. Active modular environment for robot navigation. In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [Knapp, 1987] Edgar Knapp. Deadlock detection in distributed databases. *ACM Computing Surveys (CSUR)*, 1987.
- [Ma *et al.*, 2017] Hang Ma, TK Satish Kumar, and Sven Koenig. Multi-agent path finding with delay probabilities. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 2017.
- [Mannucci *et al.*, 2021] Anna Mannucci, Lucia Pallottino, and Federico Pecora. On provably safe and live multirobot coordination with online goal posting. *IEEE Transactions on Robotics (T-RO)*, 2021.
- [Mansouri *et al.*, 2019] Masoumeh Mansouri, Bruno Lacerda, Nick Hawes, and Federico Pecora. Multi-robot planning under uncertain travel times and safety constraints. In *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 2019.
- [Nebel, 2020] Bernhard Nebel. On the computational complexity of multi-agent pathfinding on directed graphs. In *Proceedings of International Conference on Automated Planning and Scheduling (ICAPS)*, 2020.
- [Okumura *et al.*, 2021] Keisuke Okumura, Yasumasa Tamura, and Xavier Défago. Time-independent planning for multiple moving agents. In *Proceedings of AAAI Conference on Artificial Intelligence (AAAI)*, 2021.
- [Peltzer *et al.*, 2020] Oriana Peltzer, Kyle Brown, Mac Schwager, Mykel J Kochenderfer, and Martin Sehr. Stt-cbs: A conflict-based search algorithm for multi-agent path finding with stochastic travel times. 2020.
- [Robertson and Seymour, 1985] Neil Robertson and Paul D Seymour. Disjoint paths—a survey. *SIAM Journal on Algebraic Discrete Methods*, 1985.
- [Sharon *et al.*, 2015] Guni Sharon, Roni Stern, Ariel Felner, and Nathan R Sturtevant. Conflict-based search for optimal multi-agent pathfinding. *Artificial Intelligence (AIJ)*, 2015.
- [Silberschatz *et al.*, 2006] Abraham Silberschatz, Peter B Galvin, and Greg Gagne. *Operating system concepts*. John Wiley & Sons, 2006.
- [Silver, 2005] David Silver. Cooperative pathfinding. *AIIDE*, 2005.
- [Stern *et al.*, 2019] Roni Stern, Nathan Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Kumar, et al. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of Annual Symposium on Combinatorial Search (SOCS)*, 2019.
- [Tel, 2000] Gerard Tel. Deadlock-free packet switching. In *Introduction to distributed algorithms*, chapter 5. 2000.
- [Van Den Berg *et al.*, 2011] Jur Van Den Berg, Stephen J Guy, Ming Lin, and Dinesh Manocha. Reciprocal n-body collision avoidance. In *Robotics Research*. 2011.
- [Wagner and Choset, 2017] Glenn Wagner and Howie Choset. Path planning for multiple agents under uncertainty. In *Proceedings of International Conference on Automated Planning and Scheduling (ICAPS)*, 2017.
- [Wurman *et al.*, 2008] Peter R Wurman, Raffaello D’Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI magazine*, 2008.
- [Zhang *et al.*, 2018] Xu Zhang, Mingyang Li, Jian Hui Lim, Yiwei Weng, Yi Wei Daniel Tay, Hung Pham, and Quang-Cuong Pham. Large-scale 3d printing by a team of mobile robots. *Automation in Construction*, 2018.