


Multi-Agent Path Finding with Continuous Time Using SAT Modulo Linear Real Arithmetic

Tomáš Kolárik^{1,2}^a, Stefan Ratschan¹^b and Pavel Surynek²^c

¹*Institute of Computer Science of the Czech Academy of Sciences, Czech Republic*

²*Faculty of Information Technology, Czech Technical University in Prague, Czech Republic*

Keywords: Multi-Agent Path Finding, Satisfiability Modulo Theories.

Abstract: This paper introduces a new approach to solving a continuous-time version of the multi-agent path finding problem. The algorithm translates the problem into an extension of the classical Boolean satisfiability problem, satisfiability modulo theories (SMT), that can be solved by off-the-shelf solvers. This enables the exploitation of conflict generalization techniques that such solvers can handle. Computational experiments show that the new approach scales better with respect to the available computation time than state-of-the-art approaches and is usually able to avoid their exponential behavior on a class of benchmark problems modeling a typical bottleneck situation.

1 INTRODUCTION

Multi-agent path finding (MAPF) (Silver, 2005; Ryan, 2008; Luna and Bekris, 2011) is the problem of navigating agents from their start positions to given individual goal positions in a shared environment so that agents do not collide with each other. The standard discrete variant of the MAPF problem is modeled using an undirected graph in which k agents move instantaneously between its vertices. The space occupancy by agents is modeled by the requirement that at most one agent reside per vertex and via movement rules that forbid conflicting moves that traverse the same edge in opposite directions.


Standard discrete MAPF however lacks expressiveness for various real life problems where continuous time and space play an important role such as robotics applications and/or traffic optimization (Felner et al., 2017; Ma, 2022).


This drawback of standard MAPF has been mitigated by introducing various generalizations such as MAPF with continuous time (MAPF_R) (Andreychuk et al., 2022). This allows more accurate modeling of the target application problem without introducing denser and larger discretizations. Especially in applications, where agents correspond to robots, it is


important to consider graph edges that interconnect vertices corresponding to more distant positions. It is unrealistic to consider unit time for such edges as done in the standard MAPF, hence general duration of actions must be adopted. The action duration often corresponds to the length of edges which implies fully continuous reasoning over the time domain.

In this paper, we show how to solve the MAPF_R problem by directly translating it to a satisfiability modulo theories (SMT) (Barrett and Tinelli, 2018) problem. SMT extends the classical Boolean satisfiability problem to formulas whose atomic subformulas may not only be formed by propositional variables, but also by first-order predicate language formulas whose syntax is restricted to certain predicate and function symbols and whose semantics by interpreting those predicate and function symbols according to a given first-order theory. This allows the exploitation of decision procedures for such theories by so-called SMT solvers. Typical theories include the linear theory of integers, the theory of bit-vectors, and the theory of free function symbols. In this paper, we will use the theory of quantifier free linear real arithmetic (QF_LRA), which will allow us to reason about time in MAPF modeled in a continuous manner.

State-of-the-art approaches for MAPF_R such as Continuous-time Conflict-based Search (CCBS) (Andreychuk et al., 2022), a generalization of Conflict-based Search (CBS) (Sharon et al., 2015) that represents one of the most popular algorithms for MAPF,

^a <https://orcid.org/0000-0002-7207-5197>

^b <https://orcid.org/0000-0003-1710-1513>

^c <https://orcid.org/0000-0001-7200-0542>

search for optimal plans. However, in real-world applications, where the formalized MAPF problem results from an approximation of the original application problem, an overly strong emphasis on optimality is often pointless. Moreover, it may result in non-robust plans that are difficult to realize in practice (Atzmon et al., 2020). Hence we aim for a sub-optimal method whose level of optimality can be adapted according to the needs of the given application domain.

Unlike methods based on CCBS that approaches the optimum from below by iterating through plans that still contain collisions, our method approaches the optimum from above, iterating through collision-free plans. This has the advantage that—after finding its first plan—our method can be interrupted at any time, still producing a collision-free, and hence feasible plan. This anytime behavior is highly desirable in practice (Li et al., 2021).

Another advantage over existing methods is the fact that the objective function is a simple expression handed over to the underlying SMT solver. This allows any objective function that the SMT solver is able to handle without the need for any algorithmic changes.

We did experiments comparing our method with the state-of-the-art approaches CCBS and SMT-CCBS (Andreychuk et al., 2022) on benchmark problems with various numbers of agents. The results show that our method is more sensitive to time-outs than the existing approaches, typically being able to solve more instances than existing approaches for high time-outs and less for lower time-outs. Future improvement of computer efficiency will consequently make the method even more competitive.

Moreover, for one class of benchmark problems—modeling a bottleneck situation where all agents have to queue for passing a single node, the new method usually avoids the exponential behavior of CCBS and SMT-CCBS whose run-times explode from a certain number of agents on. Such bottleneck situations frequently occur for certain types of application problems (e.g., in traffic problems or navigation of characters in computer games through tunnels and the like).

Further Related Work: Existing methods for generalized variants of MAPF with continuous time include variants of Increasing Cost Tree Search (ICTS) (Walker et al., 2018) where durations of individual actions can be non-unit. The difference from our generalization is that agents do not have an opportunity to wait an arbitrary amount of time but wait times are predefined via discretization. Similar discretization has been introduced in the Conflict-based Search al-

gorithm (Cohen et al., 2019). Since discretization in case of ICTS as well as CBS brings inaccuracies of representation of the time, it is hard to define optimality. Moreover, a more accurate discretization often increases the number of actions, which can lead to an excessively large search space.

Our method for MAPF_R comes from the stream of compilation-based methods for MAPF, where the MAPF instance is compiled to an instance in a different formalism for which an off-the-shelf efficient solver exists. Solvers based on formalisms such as Boolean satisfiability (SAT) (Surynek, 2019), Constraint Programming (CP) (Ryan, 2010), or Mixed-integer Linear Programming (MILP) (Lam et al., 2022) exist. The advantage of these solvers is that any progress in the solver for the target formalism can be immediately reflected in the MAPF solver that it is based on.

The earlier MAPF method related to SMT (the SMT-CBS algorithm) (Surynek, 2019) separates the rules of MAPF into two logic theories, one theory for conflicts between agents and one theory for the rest of the MAPF rules. The two theories are used to resolve conflicts between agents lazily similarly as it is done by the CBS algorithm.

The application of SAT and SMT solvers to planning problem different from MAPF is not new (Rintanen, 2021; Leofante, 2023; Cashmore et al., 2020), usually in the context of temporal and numerical planning—extensions of the classical planning problem with numerical variables. SMT solvers have been used for specific planning problems with multiple agents (Kolárik and Ratschan, 2023), employing however a synchronous model that identifies each step of the unrolled planning problem with a fixed time period.

Acknowledgments. The work of Tomáš Kolárik was supported by the project 22-31346S of the Czech Science Foundation GA ČR and by CTU project SGS20/211/OHK3/3T/18. The work of Stefan Ratschan was supported by the project 21-09458S of the Czech Science Foundation GA ČR and institutional support RVO:67985807. The work of Pavel Surynek was supported by the project 22-31346S of the Czech Science Foundation GA ČR.

2 PROBLEM DEFINITION

We follow the definition of multi-agent path finding with continuous time (MAPF_R) from (Andreychuk et al., 2022).

We define a MAPF_R problem by the tuple $(G, M, A, s, g, \text{coord})$, where $G = (V, E)$ is a directed graph with V modeling important positions in the environment and E modeling possible transitions between the positions, M is a metric space that models the continuous environment, $A = \{a_1, a_2, \dots, a_k\}$ is a set of agents, functions $s : A \rightarrow V$ and $g : A \rightarrow V$ define start and goal vertices for the agents, and $\text{coord} : V \rightarrow M$ assigns each vertex a coordinate in metric space M .

The edges E define a set of possible move actions, where each $e = (u, v) \in E$ is assigned a duration $e_D \in \mathbb{R}_{>0}$ and a motion function $e_M : [0, e_D] \rightarrow M$ where $e_M(0) = \text{coord}(u)$ and $e_M(e_D) = \text{coord}(v)$. In addition to this, there is infinite set of wait actions associated with each vertex $v \in V$ such that an agent can wait in v any amount of time. The motion function for a wait action is constant and equals to $\text{coord}(v)$ throughout the duration of the action.

Collisions between agents are defined via a collision-detection predicate $\text{ISCOLLISION} \subseteq A \times A \times M \times M$ such that $\text{ISCOLLISION}(a_i, a_j, m_i, m_j)$ if and only if the bodies of agents a_i and a_j overlap at coordinates m_i and m_j . For this purpose, we assume that the bodies are open sets and overlapping is understood to be strict. Hence agents are permitted to touch if they are assumed to have a closed boundary which is not defined as a collision.

The algorithm described in this paper is abstract in the sense that it does not explicitly restrict the class of motion actions. Instead it assumes that it is possible to do collision detection and avoidance, as described in Section 5. This is possible, for example, if the agents and motion functions are described by polynomials, due to the fact that the theory of real closed fields allows quantifier elimination. Note that this allows the modeling of non-constant agent speed and of movements along non-linear curves. Still, in our implementation, for reasons of efficiency and ease of implementation, the motion functions are required to be linear.

Given a sequence of actions $\pi = (e_1, e_2, \dots, e_n)$, we generalize the duration and motion functions from individual actions to overall π which we denote by π_D and π_M , respectively. Let $\pi[:n'] = (e_1, e_2, \dots, e_{n'})$ denote the prefix of the sequence of actions, then $\pi_D = \sum_{i=1}^n e_{iD}$ and analogously $\pi[:n']_D = \sum_{i=1}^{n'} e_{iD}$. The motion function π_M needs to take into account the relative time of individual motion functions e_{iM} , that is: $\pi_M(t) = e_{1M}(t)$ for $t \leq e_{1D}$, ..., $\pi_M(t) = e_{n'M}(t - \pi[:n'-1]_D)$ for $\pi[:n'-1]_D \leq t \leq \pi[:n']_D$, ..., $\pi_M(t) = e_{nM}(e_{nD})$ for $\pi_D < t$. The last case means that the agent stops after executing the sequence of actions and stays at the coordinate of the goal vertex.

Definition 1. *There is a collision between sequences of actions π_i and π_j if and only if $\exists t \in [0, \max\{\pi_{iD}, \pi_{jD}\}]$ such that $\text{ISCOLLISION}(a_i, a_j, \pi_{iM}(t), \pi_{jM}(t))$.*

Definition 2. *A pre-plan of a given MAPF_R problem $(G, M, A, s, g, \text{coord})$ is a collection of sequences of actions $\pi_1, \pi_2, \dots, \pi_k$ s.t. for every $i \in \{1, \dots, k\}$, $\pi_i(0) = s(a_i)$ and $\pi_i(\pi_{iD}) = g(a_i)$. A plan for given MAPF_R problem P is a pre-plan of P whose sequences are pair-wise collision free.*

We define several types of cost functions that we denote by $\text{cost}(\Pi)$, for a given plan Π . For example, we will work with sum-of-costs (in this case, $\text{cost}(\Pi) = \sum_{i=1}^k \text{cost}(\pi_{iD})$), or makespan.

For a MAPF_R problem P , we denote by $\text{opt}(P)$ its optimal plan and by $\text{opt}_{\text{pre}}(P)$ its optimal pre-plan. Clearly $\text{cost}(\text{opt}_{\text{pre}}(P)) \leq \text{cost}(\text{opt}(P))$, but $\text{opt}_{\text{pre}}(P)$ is much easier to compute than $\text{opt}(P)$ since it directly follows from the plans of the individual agents.

For our approach, the following two observations are essential:

- Multiple subsequent wait actions can always be merged into a single one without changing the overall motion.
- It is always possible to insert a wait action of zero length between two subsequent move actions—again without changing the overall motion.

Due to this, we can restrict the search space to plans for which each wait action is immediately followed by a move action and vice versa. Our SMT encoding will then be able to encode wait and move actions in pairs, which motivates counting the number of steps of plans by just counting move actions. Hence, for a sequence of actions π we denote by $|\pi|$ the number of move actions in the sequence, and for a plan Π , we call $|\Pi| := \max_{i=1}^k |\pi_i|$ the number of steps of plan Π .

3 ALGORITHM

Our goal is to encode the planning problem in an SMT theory that is rich enough to model time and to represent conflict generalization constraints. Since SMT solvers only encode a fixed number of steps, we have to use a notion of optimality that takes this into account. Hence the first optimization criterion is the number of steps, and the second criterion cost, which we optimize up to a given $\delta > 0$:

Definition 3. *A plan Π satisfying a MAPF_R problem P is minstep δ -optimal iff*

- $|\Pi| = \min\{|\Pi'| \mid \Pi' \text{ is a plan of } P\}$, and
- $\text{cost}(\Pi) \leq (1 + \delta) \inf\{\text{cost}(\Pi') \mid |\Pi'| = |\Pi|\}$.

The result is Algorithm 1. It searches from below for a plan of minimal number of steps, and then minimizes cost for the given number of steps using iterative bisection. For this, it uses a function *findplan* that searches for a plan with a fixed number of states whose cost is between some minimal and maximal cost and that we will present in more details below in Algorithm 2.

Algorithm 1: Main algorithm MAPF-LRA.

MAPF-LRA(P, δ) $\rightarrow p_{opt}$

Input:

- a MAPF_R problem $P = (G, M, A, s, g, coord)$
- $\delta \in \mathbb{R}_{>0}$

Output:

- p_{opt} : a minstep δ -optimal plan for P

$h \leftarrow |opt_{pre}(P)|$

$t_{min} \leftarrow cost(opt_{pre}(P))$

$C \leftarrow \emptyset$

$(p, C) \leftarrow findplan(P, h, t_{min}, \infty, C)$

while $p = \perp$ **do**

$h \leftarrow h + 1$

$(p, C) \leftarrow findplan(P, h, t_{min}, \infty, C)$

$p_{opt} \leftarrow p$

while $cost(p_{opt}) > (1 + \delta)t_{min}$ **do**

let $\hat{t} \in (t_{min}, cost(p_{opt}))$

$(p, C) \leftarrow findplan(P, h, t_{min}, \hat{t}, C)$

if $p = \perp$ **then** $t_{min} \leftarrow \hat{t}$ **else** $p_{opt} \leftarrow p$

return p_{opt}

When using a SAT solver to implement the function *findplan*, it would be an overkill to encode the whole planning problem at once, since we would have to encode the avoidance of a huge number of potential collisions. Instead, we will encode this information on demand, initially looking for a pre-plan, and adding information on collision avoidance only based on collisions that have already occurred.

However, whenever a collision occurs, we do not only avoid the given collision, but also collisions that are in some sense similar. We will call this a *generalization* of a collision which we will also formalize in Section 5.

So denote by $\varphi_{P, h, \underline{t}, \bar{t}}$ an SMT formula encoding the existence of a pre-plan Π of MAPF_R-problem P with number of steps h and cost in $[\underline{t}, \bar{t}]$, that is, $|\Pi| = h$ and $cost(\Pi) \in [\underline{t}, \bar{t}]$ (see Section 4 for details). We will use an SMT solver to solve those formulas and assume that for any formula φ encoding a planning problem, $sat(\varphi)$ either returns the pre-plan satisfying φ or \perp if φ is not satisfiable.

The result is Algorithm 2 below.

Algorithm 2: Function *findplan* that searches for a plan with a bounded cost.

findplan(h, t_{min}, t_{max}, C) $\rightarrow (p, C')$

Input:

- $h \in \mathbb{N}_0$
- $t_{min} \in \mathbb{R}_{\geq 0}$
- $t_{max} \in \mathbb{R}_{\geq 0} \cup \{\infty\}$
- C : a set of formulas that every plan must satisfy

Output:

- p : either a plan Π with $|\Pi| = h$ and $cost(\pi) \in [t_{min}, t_{max}]$, or \perp , if such a plan does not exist
- C' : a set of formulas that every plan must satisfy

$p \leftarrow sat(\varphi_{P, h, t_{min}, t_{max}} \wedge \bigwedge_{\varphi_c \in C} \varphi_c)$

while $\neg[p = \perp \vee p \text{ is collision-free}]$ **do**

let $\varphi_{c_{new}}$ represent the generalization of collisions in p

$C \leftarrow C \cup \{\neg\varphi_{c_{new}}\}$

$p \leftarrow sat(\varphi_{P, h, t_{min}, t_{max}} \wedge \bigwedge_{\varphi_c \in C} \varphi_c)$

return (p, C)

Note that if $p \neq \perp$, the pre-plan p may have several collisions. The algorithm leaves it open for which of those collisions to add collision avoidance information into the formula $\varphi_{c_{new}}$. The algorithm leaves it open, as well, how much to generalize a found collision occurring at a certain point in time. In our approach, we use a specific choice here that we will describe in Section 5.

Theorem 1. *The main algorithm is correct, and if \hat{t} is chosen as $(1 - c)t_{min} + ct_{max}$, for some fixed $c \in (0, 1)$, then it also terminates.*

Proof. Since $|opt_{pre}(P)|$ is a lower bound on the number of steps of any plan of P , $h \leq \min\{|\Pi| \mid \Pi \text{ is a plan of } P\}$ at the beginning of the first while loop. After termination of the first while loop, $h = \min\{|\Pi| \mid \Pi \text{ is a plan of } P\}$. Moreover, the second while loop does not change h , and hence the result of the algorithm certainly satisfies the first condition of Definition 3. Throughout the first loop, t_{min} is a lower bound on all collision free plans, and throughout the second loop, it is a lower bound on all collision free plans that take h steps, and p_{opt} contains a h -step plan. Hence, after termination of the second loop also the second condition of Definition 3 holds.

Finally, if $c \in (0, 1)$, $cost(p_{opt}) - t_{min}$ goes to zero as the second-while loop iterates. Hence the termination condition of this loop must eventually be satisfied. \square

4 SMT ENCODING

In this section, we present an encoding of the planning problem from Section 2 as an SMT formula in the quantifier-free theory of linear real arithmetic QF_LRA. Here we concentrate on the formula $\Phi_{P,h,t,\bar{t}}$ that models time constraints of the agents and their paths in graph G but does *not* model collisions of the agents and metric space M . We leave the SMT encoding of collision avoidance to the next section.

Variables. As usual in planning applications of SAT solvers (Rintanen, 2021), we unroll the planning problem in a similar way as in Bounded Model Checking (Biere, 2021), where each step $0, 1, \dots, h$ corresponds to one wait and one move action. As a consequence, unrolling over h steps corresponds to search for a pre-plan Π with $|\Pi| = h$.

Note that h corresponds to the maximum of the steps of all agents, so an agent that already reached the goal in step $j < h$ may remain in the same state in steps j, \dots, h , without any further actions.

Each agent is modeled using a separate set of Boolean and real-valued variables. For each agent $a \in A$ and discrete step j , we define variables $V_a^{[j]}$, $T_a^{[j]}$, $w_a^{[j]}$ and $m_a^{[j]}$: We model the vertex position of the agent by $V_a^{[j]}$ which is a Boolean encoding of a vertex $v \in V$ using $O(|V|)$ or $O(\log(|V|))$ Boolean variables. We will use the notation $V_a^{[j]} = v$ to denote a constraint that expresses that an agent occupies vertex $v \in V$ at the beginning of discrete step j . We will also use $V_a^{[j]} \neq v$ as an abbreviation for $\neg(V_a^{[j]} = v)$.

Next, we model time constraints of the agent using real variables $T_a^{[j]}$, $w_a^{[j]}$, and $m_a^{[j]}$. The variables $T_a^{[j]}$ model the absolute time when the agent occupies a vertex that corresponds to $V_a^{[j]}$, before it takes further actions within discrete step j (or later). The variables $w_a^{[j]}$ model the duration of wait actions and the variables $m_a^{[j]}$ the duration of move actions.

Finally, we use an auxiliary real variable λ that, for a pre-plan Π , corresponds to $cost(\Pi)$. The objective is to minimize this variable. There may be arbitrary linear constraints on the variable, allowing specification of rich cost functions.

Constraints. We define (1) initial and goal conditions on the agents, (2) constraints that ensure that the agents follow paths through the graph G , and (3) time constraints that correspond to occurrences of the agents at vertices of the graph. For that, we only use the variables defined above.

The initial and goal conditions ensure that each agent a visits the start and goal vertex at the beginning and end of the plan, respectively: $V_a^{[0]} = s(a) \wedge V_a^{[h]} = g(a)$.

To ensure that the agents follow paths through the graph G , we use for each agent a and $j < h$ a constraint ensuring that the pair of vertex positions $V_a^{[j]}$ and $V_a^{[j+1]}$ corresponds to an edge of the graph G . However, this is not necessarily the case for an already finished agent, that is, if $V_a^{[j]} = g(a)$, then also $V_a^{[j+1]} = g(a)$ is allowed.

The time constraints ensure that the initial value of time of all agents is zero: $T_a^{[0]} = 0$. For $j > 0$, they assume that during each discrete step, an agent may first wait and then it moves, resulting in the constraint $T_a^{[j]} = T_a^{[j-1]} + w_a^{[j-1]} + m_a^{[j-1]}$. For the waiting times, we require $w_a^{[j]} \geq 0$. In addition, we ensure that at least one agent starts to move at the beginning of a plan without waiting, asserting $\bigvee_{a \in A} w_a^{[0]} = 0$. For the moving times, if $j < h$ we ensure that $m_a^{[j]}$ corresponds to the duration of the edge between $V_a^{[j]}$ and $V_a^{[j+1]}$. In addition, $m_a^{[j]} = 0$ if $j = h$ or $V_a^{[j]} = g(a) \wedge V_a^{[j+1]} = g(a)$.

Note that agents are modeled asynchronously, meaning that for a pair of agents $a, b \in A$, $T_a^{[j]}$ and $T_b^{[j]}$ corresponds *not* necessarily to the same moment in time. This implies that comparing times and the corresponding positions of agents, in order to check whether there are collisions, cannot be done in a straightforward way, and in the worst case, variables corresponding to all discrete steps must be examined.

We present two variants of cost functions: sum of costs, defined as $\lambda = \sum_{a \in A} T_a^{[h]}$, and makespan, defined as $\lambda = \max_{a \in A} T_a^{[h]}$. To ensure that the formula $\Phi_{P,h,t,\bar{t}}$ satisfies the bounds of the cost function, we simply require $\lambda \geq t \wedge \lambda \leq \bar{t}$. An example of an alternative cost function is $\lambda = \sum_{a \in A} \sum_{j=0}^{h-1} (2m_a^{[j]} + w_a^{[j]})$ which prefers minimizing moving times over waiting times and can therefore result in more power-optimal plans.

Building the formula $\Phi_{P,h,t,\bar{t}}$ from scratch after each increase of the number of steps h would be inefficient. Hence we build the formula incrementally. However, some parts of the formula (e.g., the cost functions or constraints such as $V_a^{[h]} = g(a)$), explicitly depend on h , and hence need to be updated when h is increased. Here we use the feature of modern SMT solvers, that allow the user to cancel constraints asserted after a previously specified milestone, and to reuse the rest.

5 COLLISION DETECTION AND AVOIDANCE

Whenever the algorithm *findplan* from Section 3 computes a pre-plan that still contains a collision, it represents the generalization of collisions by a formula $\varphi_{c_{new}}$ whose negation it then adds to the formula passed to the SMT solver. We will now discuss how to first detect collisions and how to then construct the formula $\varphi_{c_{new}}$ generalizing detected collisions. Here, we will assume precise arithmetic, deferring the discussion of implementation in finite computer arithmetic to Section 6.

Collision Detection. Assume that two agents a and b follow their motion functions $\alpha_{Ma} : [0, \alpha_{Da}] \rightarrow M$ and $\alpha_{Mb} : [0, \alpha_{Db}] \rightarrow M$ with durations α_{Da} and α_{Db} , respectively, corresponding to either a move or a wait action. Assume that the agents start the motions at points in time $\hat{\tau}_a$ and $\hat{\tau}_b$, respectively. To determine whether there is a collision, we will use the abstract predicate ISCOLLISION introduced in Section 2. Based on this, we can check for a collision of two agents that follow motion functions starting at certain times:

Definition 4. For two motion functions α_{Ma} and α_{Mb} with respective starting times $\hat{\tau}_a$ and $\hat{\tau}_b$, $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b)$ iff

$$\exists t \in [\hat{\tau}_a, \hat{\tau}_a + \alpha_{Da}] \cap [\hat{\tau}_b, \hat{\tau}_b + \alpha_{Db}] \cdot \text{ISCOLLISION}(a, b, \alpha_{Ma}(t - \hat{\tau}_a), \alpha_{Mb}(t - \hat{\tau}_b)).$$

We will now discuss the construction of the formula $\varphi_{c_{new}}$ that generalizes collisions of pre-plans found in Algorithm 2. A found pre-plan may result in several such collisions. We start with generalizing one of them and consider two cases: The case of a collision between two moving agents, and the case of a collision between a waiting and a moving agent. We can ignore the case when both agents are waiting: Such a conflict either should have been avoided already in the previous discrete steps, or the agents must in the case of a conflict overlap right at the beginning of a pre-plan, resulting in a trivially infeasible plan.

Collisions While Moving. In this case, one of the two agents has to wait until the conflict vanishes. We are interested in waiting the minimal time and hence define $\text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b) :=$

$$\inf\{\tau_a \mid \hat{\tau}_a < \tau_a, \neg \text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \tau_a, \hat{\tau}_b)\}.$$

Note that $\hat{\tau}_a < \text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b) \leq \hat{\tau}_b + \alpha_{Db}$. Here, the lower bound is a consequence of the assumption that agents are open sets which makes collisions happen in the interior of those sets. The upper bound follows from the fact that $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_b + \alpha_{Db}, \hat{\tau}_b)$ is always false.

Assume that we detected a conflict between two move actions starting at times $\hat{\tau}_a$ and $\hat{\tau}_b$ and hence $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b)$. We know that any value of τ_a with $\hat{\tau}_a \leq \tau_a < \text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b)$ also leads to a conflict. In addition—letting the second agent wait—any value of τ_b with $\hat{\tau}_b \leq \tau_b < \text{SAFE}_{b,a}(\alpha_{Mb}, \alpha_{Ma}, \hat{\tau}_b, \hat{\tau}_a)$ also leads to a conflict.

However, we know even more. For seeing this, observe that $\text{INCONFLICT}_{a,b}$ is invariant wrt. translation along the time-axes, that is, for every $\Delta \in \mathbb{R}$, $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b)$ iff $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a + \Delta, \hat{\tau}_b + \Delta)$ which can be seen by simply translating the witness t from Definition 4 by the same value Δ . Due to this, the same conflict happens for all pairs (τ_a, τ_b) with the same relative distance as the relative distance of $(\hat{\tau}_a, \hat{\tau}_b)$. Hence we know that both

$$\hat{\tau}_a - \hat{\tau}_b \leq \tau_a - \tau_b < \text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b) - \hat{\tau}_b$$

and

$$\hat{\tau}_b - \hat{\tau}_a \leq \tau_b - \tau_a < \text{SAFE}_{b,a}(\alpha_{Mb}, \alpha_{Ma}, \hat{\tau}_b, \hat{\tau}_a) - \hat{\tau}_a$$

lead to a conflict.

Multiplying the second inequality by -1 , we get

$$\hat{\tau}_a - \text{SAFE}_{b,a}(\alpha_{Mb}, \alpha_{Ma}, \hat{\tau}_b, \hat{\tau}_a) < \tau_a - \tau_b \leq \hat{\tau}_a - \hat{\tau}_b$$

and combining the result with the first inequality, we get

$$\begin{aligned} \hat{\tau}_a - \text{SAFE}_{b,a}(\alpha_{Mb}, \alpha_{Ma}, \hat{\tau}_b, \hat{\tau}_a) &< \tau_a - \tau_b \\ \wedge \tau_a - \tau_b &< \text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b) - \hat{\tau}_b. \end{aligned}$$

For applying this to the variables of the SMT encoding described in Section 4, we denote this formula by $\varphi^{mm}(a, b, \alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b, \tau_a, \tau_b)$, replacing the arguments by the corresponding terms from the SMT encoding. More concretely, we observe that the start of a move action of agent a at step j_a is modeled by the term $T_a^{[j_a]} + w_a^{[j_a]}$ and the start of a move action of agent b at step j_b by the term $T_b^{[j_b]} + w_b^{[j_b]}$.

Now assume that we detected a conflict of two agents a and b moving along respective edges (u_a, v_a) and (u_b, v_b) , starting in discrete steps j_a and j_b and times $\hat{T}_a^{[j_a]} + \hat{w}_a^{[j_a]}$ and $\hat{T}_b^{[j_b]} + \hat{w}_b^{[j_b]}$ (the hats denoting the values assigned to the respective variables). In this case, $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{T}_a^{[j_a]} + \hat{w}_a^{[j_a]},$

$\hat{T}_b^{[j_b]} + \hat{w}_b^{[j_b]}$), and the formula $\varphi_{c_{new}}$ has the form

$$\begin{aligned} V_a^{[j_a]} &= u_a \wedge V_a^{[j_a+1]} = v_a \\ \wedge V_b^{[j_b]} &= u_b \wedge V_b^{[j_b+1]} = v_b \\ \wedge \varphi^{mm}(a, b, (u_a, v_a)_M, (u_b, v_b)_M, \\ &\quad \hat{T}_a^{[j_a]} + \hat{w}_a^{[j_a]}, \hat{T}_b^{[j_b]} + \hat{w}_b^{[j_b]}, \\ &\quad T_a^{[j_a]} + w_a^{[j_a]}, T_b^{[j_b]} + w_b^{[j_b]}). \end{aligned}$$


This means that there are 6 possibilities how to resolve such a conflict (changing one of the four vertices of edges along which the two move actions took place or changing one of the two starting times of the move actions).

Now we also discuss conflicts where a waiting agent participates.

Collisions While Waiting. We also have to ensure that no collisions happen while an agent a is waiting. In principle, the motion function α_{M_a} can also be constant, and hence one might be tempted to just specialize the formula for two moving agents to this case. However, unlike move actions, wait actions do not have fixed durations, but their duration is a consequence of the timing of the previous and following move action. We take this into account, generalizing the given conflict over arbitrarily long wait actions.

So assume an agent a waiting at a point $x_a \in M$ and an agent b following motion function α_{M_b} starting from time $\hat{\tau}_b$. Assume that a collision happens at a certain point in time \hat{t} . So we have $\hat{\tau}_b \leq \hat{t} \leq \hat{\tau}_b + \alpha_{D_b} \wedge \text{ISCOLLISION}(a, b, x_a, \alpha_{M_b}(\hat{t} - \hat{\tau}_b))$.

Let $\hat{\tau}_a$ be the end of the move action of the waiting agent a before this waiting period and let τ'_a be the starting time of the move action of the waiting agent after this waiting period.



So we compute the beginning of the collision

$$\underline{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b} := \inf\{\hat{\tau}_b \leq t \leq \hat{t} \mid \text{ISCOLLISION}(a, b, x_a, \alpha_{M_b}(t - \hat{\tau}_b))\}$$

and its end $\bar{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b} :=$

$$\sup\{\hat{t} \leq t \leq \hat{\tau}_b + \alpha_{D_b} \mid \text{ISCOLLISION}(a, b, x_a, \alpha_{M_b}(t - \hat{\tau}_b))\}.$$

So for any wait action of agent a starting at τ_a and ending at τ'_a and any move action of agent b starting at τ_b , the collision happens if the upper

bound $\bar{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b}$ is after the end of the previous action and the lower bound $\underline{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b}$ is before the beginning of the next action. The result is

$$\begin{aligned} \tau_a - \tau_b &< \bar{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b} - \hat{\tau}_b \\ &\wedge \underline{t}_{a,b,x_a,\alpha_{M_b},\hat{\tau}_b} - \hat{\tau}_b < \tau'_a - \tau_b, \end{aligned}$$

which we will denote by $\varphi^{wm}(a, b, x_a, \alpha_{M_b}, \hat{\tau}_b, \tau_a, \tau'_a, \tau_b)$.

Now we again apply this to the variables of the SMT encoding described in Section 4, replacing the arguments of the formula $\varphi^{wm}(a, b, x_a, \alpha_{M_b}, \hat{\tau}_b, \tau_a, \tau'_a, \tau_b)$ by their corresponding terms from the SMT encoding. So when we detect a conflict between an agent a that waits at vertex u_a at time step j_a and an agent b moving along an edge (u_b, v_b) at time step j_b , starting at $\hat{T}_b^{[j_b]} + \hat{w}_b^{[j_b]}$, the formula $\varphi_{c_{new}}$ has the form

$$\begin{aligned} V_a^{[j_a]} &= u_a \\ \wedge V_b^{[j_b]} &= u_b \wedge V_b^{[j_b+1]} = v_b \\ \wedge \varphi^{wm}(a, b, \text{coord}(u_a), (u_b, v_b)_M, \\ &\quad \hat{T}_b^{[j_b]} + \hat{w}_b^{[j_b]}, \\ &\quad T_a^{[j_a]}, T_a^{[j_a]} + w_a^{[j_a]}, T_b^{[j_b]} + w_b^{[j_b]}). \end{aligned}$$

We ended up with a conflict clause that actually does not depend on a previous move action. In the case where the wait action does not have a next move action, the conflict clause can be modified in a straightforward way.

Further Generalization. To fully exploit the computational effort that is necessary to find pre-plans, we generate generalizations for all conflicts a found pre-plan contains. Hence we check for conflicts between all pairs of agents, discrete steps and corresponding move and wait actions.

We further generalize the conflicts such that we also check for other conflicts of the given pair of agents when taking a move action from the same source to a different target vertex.

We did not find it useful to furthermore generalize the conflicts to further pairs of agents and/or discrete steps.

6 IMPLEMENTATION

Collision Detection and Avoidance. We implemented the predicates and functions introduced in Section 5 as follows:

- We assume that each agent a is abstracted into an open disk with a fixed radius $r_a \in \mathbb{R}^{>0}$. Hence

for agents a and b , $\text{ISCOLLISION}(a, b, c_a, c_b)$ iff $\|c_a - c_b\| < r_a + r_b$, where $c_a, c_b \in M$ are respective centers of the disks of the agents.

- We assume that the agents move with constant velocity following straight lines of the edges. As a result, $\text{INCONFLICT}_{a,b}$ corresponds to checking whether a quadratic inequality has a solution in the intersection of the time intervals from Definition 4.
- Exploiting the observation that $\text{SAFE}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \hat{\tau}_a, \hat{\tau}_b) \in (\hat{\tau}_a, \hat{\tau}_b + \alpha_{Db}]$ we compute the resulting value by binary search of the switching point τ_a in the interval for which $\neg \text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \tau_a, \hat{\tau}_b)$ and $\text{INCONFLICT}_{a,b}(\alpha_{Ma}, \alpha_{Mb}, \tau_a - \varepsilon, \hat{\tau}_b)$, for a small enough $\varepsilon > 0$.

Floating-point Numbers. Our simulations of collisions of agents are based on floating-point computation whereas SMT solvers treat linear real arithmetic precisely, using rational numbers for all computation. There are two main issues here:

- Conversion of a floating point number to a rational number may result in huge integer values for the numerator and denominator, although the intended value is very close to a simple rational or even integer number.
- Conversion of rational numbers to floating point numbers, and the following computation in floating point representation may incur approximation errors (e.g., due to round-off or discretization). For example, this may lead to the situation where—in the case of a collision between two moving agents—the added conflict clause does not require the waiting agent to wait long enough to completely avoid the same collision. Hence a very close collision may re-appear, and the same situation may repeat itself several times.

We overcome these deficiencies with an overapproximation of the conflict intervals along with simplification of the resulting rational numbers using simple continued fractions and best rational approximation: in the case that a floating-point value x is respectively a lower or a higher bound of a conflicting interval, the result corresponds to best rational approximation from $(x - \varepsilon, x]$ or $[x, x + \varepsilon)$, respectively, for an ε that is large enough. This not only avoids the re-appearance of the same conflict, but also maps floating-point values that are close to each other to the same rational numbers, avoiding the appearance of tiny differences between rational numbers that tend to clog the SMT solver.

Heuristics. The formula passed to the SMT solver often represents a highly underconstrained problem, spanning a huge solution space. Due to this, it is essential that the SMT solver chooses a solution in a goal oriented way in order to maximize the chances of hitting upon a δ -optimal plan, or at least to concentrate search on the most promising part of the solution space which also concentrates the addition of conflict avoidance clauses to this part. For this we prefer transitions to vertices that lie on shorter paths to the goal over transitions to vertices that lie on longer paths. This can be easily precomputed using Dijkstra’s algorithm for all vertices with a fixed start and goal.

Nonetheless, using such heuristics does not evade the problem of encoding all the transitions into the formula, which floods the SMT solver with a lot of constraints that are not essential at arriving at the desired plan. Also, conversion of the resulting formula to CNF might be expensive.

Tools. We implemented the resulting algorithm on top of MathSAT5 (Cimatti et al., 2013) SMT solver. We incrementally build the formula described in Section 4 using API. However, since we do not even require the SMT solver itself to handle optimization, it is easy to replace the API calls to another SMT solver that handles QF_LRA. We also implemented a visualization tool of MAPF_R problems. Our tools are available online^{1,2} and are open-source.

7 COMPUTATIONAL EXPERIMENTS

We compare run-times of our implementation from Section 6 denoted as SMT-LRA and state-of-the-art tools CCBS and SMT-CCBS, both presented in (Andreychuk et al., 2022), which define the MAPF_R problem in a similar way. These tools search for optimal plans, which is more difficult than searching for sub-optimal plans, such as minstep δ -optimal plans in our case. However, as discussed in the introduction, not only that the price for getting optimal plans may be too high, but the resulting plans may also be not an ideal fit in practice, due to possible requirements on flexible dispatchability of the plans, and due to the fact that the dynamics of the agents may not be modeled accurately. Based on these observations, we consider the comparisons to be practically reasonable.

There are also differences concerning the function that is being optimized which is sum of costs in the

¹https://gitlab.com/Tomaqa/mapf_r

²https://github.com/Tomaqa/mapf_r-visualizer

case of CCBS and makespan in the case of SMT-CCBS. We support both of these cost functions in the form of a parameter. While there are certainly instances where the choice of the cost function qualitatively matters, (Andreychuk et al., 2022) showed that both the tools yield similar respective costs of the resulting plans within their benchmarks. Hence, comparing such tools with different objectives also makes sense.

In the following experiments, we will use a similar setup to (Andreychuk et al., 2022), that is, similar to both of the presented state-of-the-art tools.

We will start by the description of the benchmarks. Finally, we will present and discuss computational results of the performed experiments.

7.1 Description of Benchmarks

A benchmark is specified by a graph and a set of agents, each defined by a radius and a starting and goal vertex. The following benchmarks use the same radius for all agents, and hence we will not discuss radii any more.

We did experiments with three classes of problems: `empty`, `roadmap` and `bottleneck`. Benchmarks `empty` and `roadmap` are adopted from (Andreychuk et al., 2022) and correspond to MAPF maps from the Moving AI repository. Our `bottleneck` benchmark is an additional simple experiment which identifies a weakness of the state-of-the-art approaches. More detailed description of the benchmarks follows below.

However, we exclude benchmark `roadmap` here. We include all benchmarks along with further figures, plots and tables in the extended version of the paper (<https://arxiv.org/abs/2312.08051>).

We did not include benchmarks with large graphs (i.e. with high number of vertices or edges), because our current algorithm encodes the whole graph into the formula, as discussed within heuristics in Section 6.

Empty Room. This benchmark is based on a graph that represents an empty square room with 16×16 vertices—the result of grid approximations of MAPF maps from the Moving AI repository (Andreychuk et al., 2022).

Interconnection of the vertices depend on a neighborhood parameter n , which defines that each vertex has exactly 2^n neighbors (with the exception of boundary vertices). For example, $n = 2$ corresponds to square grid, $n = 3$ extends the square grid of diagonal edges, etc. Using such a graph, it may be necessary to include a high amount of useless edges in

order to cover a suitable number of realistic movements of the agents. On the other hand, it might be possible to exploit the fact that such graphs are highly symmetric.

The resulting benchmark `empty` represents a model with no physical obstacles. Still, with an increasing number of agents k , the number of possible collisions of the agents grows significantly, because most of the shortest paths lead via central regions of the graph.

Bottleneck. Benchmark `bottleneck` models the problem of steering k agents from k initial vertices through a single transfer vertex to k goal vertices. Hence the transfer vertex represents a bottleneck every agent has to pass through. We place the initial and goal vertices (i.e., altogether $2k$ vertices) on a circle whose center is formed by the transfer vertex.

The task here boils down to just choosing a certain order of the agents. Resolving such a benchmark problem can still result in an exponential complexity in the number of agents k —if just various permutations of the agents are tried, without a thorough exclusion of the conflicting time intervals of particular agents.

7.2 Experimental Setup

In the case of benchmark `empty`, we observe whether particular experiments finish within a given timeout. The set of experiments contains instances where the number of agents ranges from 1 to 64 (none of the tools managed to finish with more agents within the selected timeouts). For each number of agents $k \in \{1, \dots, 64\}$, each start and goal vertex of each agent is pre-generated in 25 random variants. Here, when generating the variants for agent $k + 1$, all the previous k agents are reused and only the positions of the new one are generated randomly. The result is $64 \times 25 = 1600$ instances for each neighborhood n .

The subject of our interest is how the evaluated algorithms scale with time, so we ran all the experiments with different timeouts ranging from 30 seconds up to 16 minutes (with exponential growth) and observed how many instances finished in time. We will show the results in the form of box plots.

Our tool SMT-LRA is in addition parametrized by a cost function—either makespan or sum of costs—and by a sub-optimal coefficient $\delta \in \{1, \frac{1}{2}, \frac{1}{4}\}$. In the plots, the parameters are denoted in the form (C, δ) , where C is either `M` (makespan) or `S` (sum of costs). In all experiments, the higher δ was, the more instances were solved. Hence, to make the box plots more compact, we merged all the variants of δ related to the

same cost function such that the boxes of the variants with lower δ overlay the boxes of the variants with higher δ . Also, higher values of δ correspond to lighter colors. For example, boxes $\delta = \frac{1}{4}$ overlay boxes $\delta = \frac{1}{2}$, but the magnitudes of both boxes are still visible since the number of solved instances is always lower in the case of $\delta = \frac{1}{4}$ than in the case of $\delta = \frac{1}{2}$. As a result, for each timeout in the box plot, our tool always takes two columns, each consisting of three (overlaid) boxes.

In the case of benchmark `bottleneck`, we only focus on runtime of the evaluated tools for some numbers of agents ranging from 2 to 30. In the case of our implementation, we again include all the variants of parameters mentioned above. We used timeout 30 min to set some upper boundary on the runtimes of the tools.

We executed all the benchmarks on a machine with Intel(R) Xeon(R) Gold 6254 CPU @ 3.10GHz, with 180 cores and 1TB memory. To unify the runtime environment, we reused and adapted the scripts from the previous experiments (Andreychuk et al., 2022) which are a part of the SMT-CCBS tool. These scripts do not exploit all the available resources of the machine, though. Still, none of the evaluated tools use parallel computation—the cores are only used to run multiple benchmarks concurrently.

7.3 Results

Empty Room. Recall that benchmark `empty` is parametrized by its neighborhood n which means that vertices have approximately 2^n neighbors. We did experiments with $n \in \{2, 3, 4, 5\}$, all of which are shown in box plots in Figure 1.

We firstly focus on the comparison of the selected cost functions in the case of our tool. We see that usually the cases that optimize sum of costs perform better than the cases optimizing makespan. Observe, though, that the results are similar in the case of $\delta = \frac{1}{4}$ which correspond to the boxes at the base. We explain these as that in the case of sum of costs there are more possibilities how to reduce the cost of the plan than in the case of makespan where the cost usually depends on just one agent, regarding the symmetry of the graph. We assume that at the same time this is the reason why, in the case of makespan, there are lower increases of the number of solved instances with growing δ compared to the variant which optimizes sum of costs. Also notice that in the cases of neighborhood $n = 2$ and especially $n = 5$, there are quite low performance growths with increasing δ . However, (Andreychuk et al., 2022) showed that these corner cases of n are actually the least useful bench-

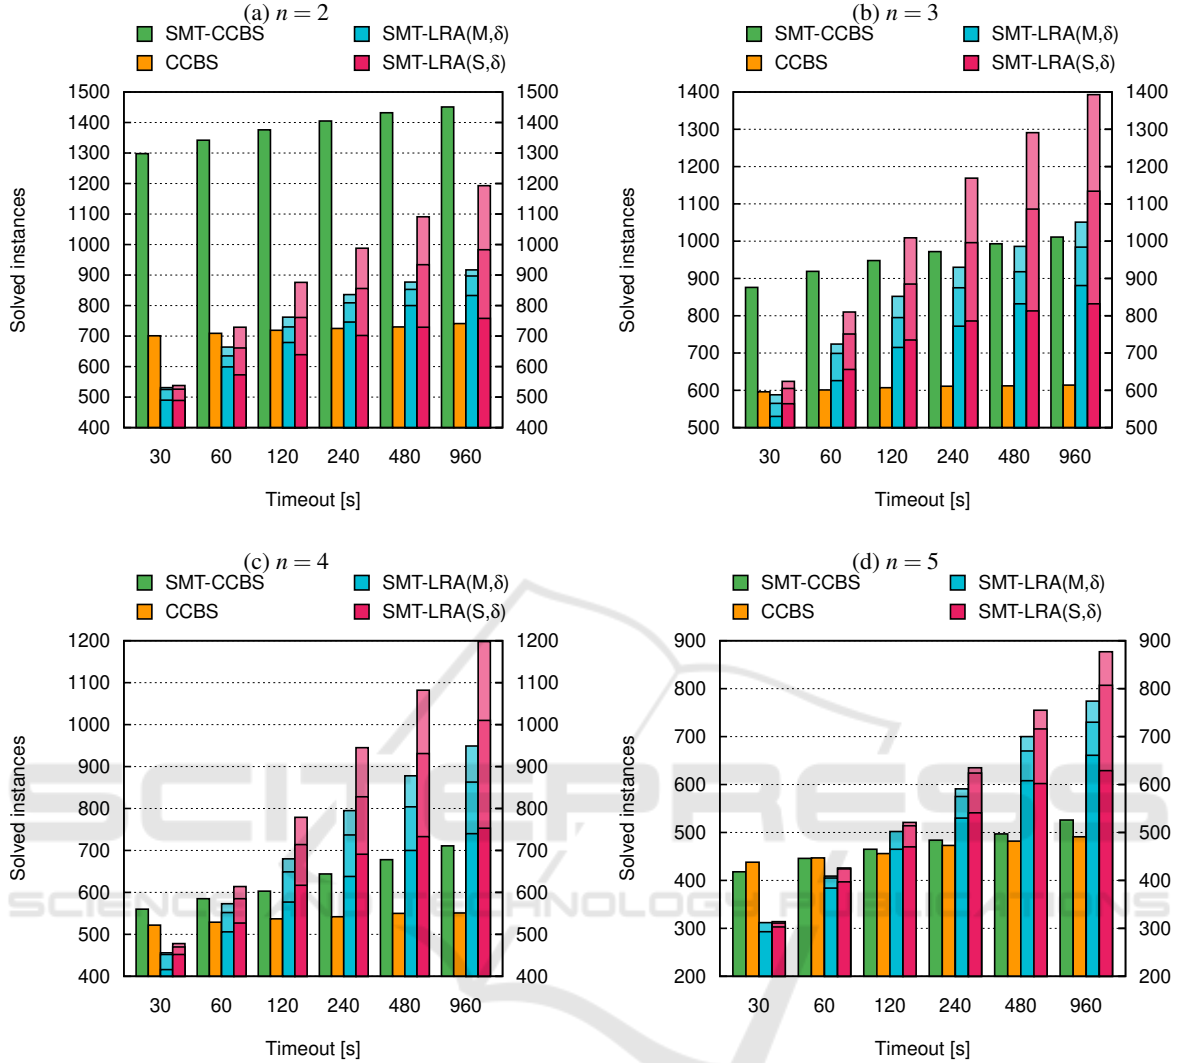
marks: benchmarks with $n = 3$ offer much faster plans than in the case of $n = 2$, and $n = 5$ on the other hand provide only very low improvement over the case of $n = 4$. All in all, our approach scales well with the growing timeouts, in every case of neighborhood n , cost function and parameter δ .

Now we also focus on the state-of-the-art tools, where we will actually confirm the observations made in (Andreychuk et al., 2022). These tools are consistent in the sense that the lower parameter n , the faster is their algorithm—because there are less possible paths to the goals. In our case, the observation holds as well, but with one exception in the case of $n = 2$ vs. $n = 3$, where the runtime of the experiments with the lowest neighborhood is higher. The reason is that the graphs with higher n allow that the shortest paths to the goals take less edges—which in our case becomes more important than the number of possible choices, because our current algorithm is sensitive to the number of edges in the graph which we all encode into the formula.

The state-of-the-art tools usually perform better than our tool when the timeout is less or equal to 1 min. SMT-CCBS performs especially well in the case of $n = 2$ because it maps a lot of time points to the same values since many of them are integer values. We however consider this case to be the least useful benchmark referring the earlier discussion and in addition since the square grids are not too realistic models and can also be handled by standard MAPF approaches (which are currently much faster than MAPF_R approaches). Although SMT-CCBS scales better with time than CCBS, the highest growth of the number of solved instances still occurs in the case of our approach, even in the cases of $\delta = \frac{1}{4}$ which correspond to the boxes of our algorithm at the base. In the cases of $n \geq 3$ and $\delta = 1$, which correspond to the upmost boxes of our algorithm, we outperform the state-of-the-art tools if the timeout is high.

Bottleneck. We summarize the runtimes of benchmark `bottleneck` of particular tools in Table 1. In the case of SMT-CCBS, we excluded the built-in verification of the solutions which here seemed to be very time-consuming. In our case we merge all the corresponding cases of parameter δ into single column since the runtimes were almost the same regardless the parameter. We further merge both cost functions `M` and `S` into one column s.t. the respective runtimes are separated by a pipe character.

It is clear that runtimes of both state-of-the-art solvers exhibit exponential relationship with the number of agents k , while our algorithm is much less sensitive. For example, CCBS is fastest until $k = 6$ but

Figure 1: Comparison of the number of solved instances of benchmark `empty` with given n wrt. a given timeout.

after that point our SMT-LRA dominates the runtime. The reason is that we resolve the conflicts of agents using the learning mechanism of generalized conflict clauses where the timing constraints efficiently exclude inappropriate orderings of the agents, making the benchmark fairly easy for our approach—which is consistent with the observation that such a problem is indeed trivial, as discussed in the description of benchmarks. For example, the problem is easily solvable using an ad-hoc approach. Nevertheless, such bottlenecks may appear as a part of more complex problems where a sophisticated algorithm instead of an ad-hoc should be used.

8 CONCLUSION

We have demonstrated how to solve the continuous-time MAPF problem (MAPF_R) by direct translation to SAT modulo linear real arithmetic. While the approach insists only on sub-optimality up to a certain factor, it shows several advantages over state-of-the-art algorithms, especially better scaling wrt. an increasing time budget and its ability to efficiently handle bottleneck situations. Our approach also allows for easy change of the objective function to another. The downside is a certain basic translation effort, especially for problems depending on large graphs.

In future work we will explore a lazy approach to translation that only generates the information neces-

Table 1: Comparison of the runtimes in seconds of benchmark bottleneck with a given number of agents.

k	SMT-CCBS	CCBS	SMT-LRA(M S)	
2	0.00	0.00	0.01	0.01
3	0.47	0.00	0.02	0.02
4	> 1800	0.00	0.04	0.02
5	?	0.01	0.03	0.03
6	?	0.05	0.07	0.05
7	?	0.45	0.08	0.06
8	?	3.52	0.15	0.08
9	?	43.99	0.15	0.11
10	?	720.27	0.22	0.14
11	?	> 1800	0.26	0.17
15	?	?	0.46	0.42
20	?	?	0.79	0.97
30	?	?	4.88	5.63

sary for solving the current problem instance. This will be especially relevant in practical applications where similar problems have to be solved repeatedly. Moreover, we will generalize the method to problems with non-linear motion functions, allowing both non-linear geometry of the involved curves and the modeling of non-linear dynamical phenomena such as acceleration of agents. The method will also benefit from the fact that the efficiency of SMT solvers is currently improving with each year.

REFERENCES

- Andreychuk, A., Yakovlev, K. S., Surynek, P., Atzmon, D., and Stern, R. (2022). Multi-agent pathfinding with continuous time. *Artif. Intell.*, 305:103662.
- Atzmon, D., Stern, R., Felner, A., Sturtevant, N. R., and Koenig, S. (2020). Probabilistic robust multi-agent path finding. In *Proc. of the Thirtieth International Conference on Automated Planning and Scheduling, 2020*, pages 29–37. AAAI Press.
- Barrett, C. and Tinelli, C. (2018). Satisfiability modulo theories. In Clarke, E. M., Henzinger, T. A., Veith, H., and Bloem, R., editors, *Handbook of Model Checking*. Springer International Publishing.
- Biere, A. (2021). Bounded model checking. In (Biere et al., 2021).
- Biere, A., Heule, M., van Maaren, H., and Walsh, T., editors (2021). *Handbook of Satisfiability*. IOS Press, 2nd edition.
- Cashmore, M., Magazzeni, D., and Zehtabi, P. (2020). Planning for hybrid systems via satisfiability modulo theories. *Journal of Artificial Intelligence Research*, 67:235–283.
- Cimatti, A., Griggio, A., Schaafsma, B., and Sebastiani, R. (2013). The MathSAT5 SMT solver. In Piterman, N. and Smolka, S., editors, *Proc. of TACAS*, volume 7795 of LNCS. Springer.
- Cohen, L., Uras, T., Kumar, T. K. S., and Koenig, S. (2019). Optimal and bounded-suboptimal multi-agent motion planning. In *Proc. of the Twelfth International Symposium on Combinatorial Search, SOCS 2019*, pages 44–51. AAAI Press.
- Felner, A., Stern, R., Shimony, S. E., Boyarski, E., Goldenberg, M., Sharon, G., Sturtevant, N. R., Wagner, G., and Surynek, P. (2017). Search-based optimal solvers for the multi-agent pathfinding problem: Summary and challenges. In *Proc. of the Tenth International Symposium on Combinatorial Search, SOCS 2017*, pages 29–37. AAAI Press.
- Kolárik, T. and Ratschan, S. (2023). Railway scheduling using boolean satisfiability modulo simulations. In Chechik, M., Katoen, J.-P., and Leucker, M., editors, *Formal Methods*, number 14000 in LNCS, pages 56–73, Cham. Springer International Publishing.
- Lam, E., Bodic, P. L., Harabor, D., and Stuckey, P. J. (2022). Branch-and-cut-and-price for multi-agent path finding. *Comput. Oper. Res.*, 144:105809.
- Leofante, F. (2023). OMTPlan: a tool for optimal planning modulo theories. *Journal on Satisfiability, Boolean Modeling and Computation*, 14(1):17–23.
- Li, J., Chen, Z., Harabor, D., Stuckey, P. J., and Koenig, S. (2021). Anytime multi-agent path finding via large neighborhood search. In *Proc. of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI 2021*, pages 4127–4135. ijcai.org.
- Luna, R. and Bekris, K. E. (2011). Push and swap: Fast cooperative path-finding with completeness guarantees. In *IJCAI 2011, Proc. of the 22nd International Joint Conference on Artificial Intelligence, 2011*, pages 294–300. IJCAI/AAAI.
- Ma, H. (2022). Intelligent planning for large-scale multi-agent systems. *AI Mag.*, 43(4):376–382.
- Rintanen, J. (2021). Planning and SAT. In (Biere et al., 2021).
- Ryan, M. (2010). Constraint-based multi-robot path planning. In *IEEE International Conference on Robotics and Automation, ICRA 2010*, pages 922–928. IEEE.
- Ryan, M. R. K. (2008). Exploiting subgraph structure in multi-robot path planning. *J. Artif. Intell. Res.*, 31:497–542.
- Sharon, G., Stern, R., Felner, A., and Sturtevant, N. R. (2015). Conflict-based search for optimal multi-agent pathfinding. *Artif. Intell.*, 219:40–66.
- Silver, D. (2005). Cooperative pathfinding. In *Proc. of the First Artificial Intelligence and Interactive Digital Entertainment Conference, 2005*, pages 117–122. AAAI Press.
- Surynek, P. (2019). Unifying search-based and compilation-based approaches to multi-agent path finding through satisfiability modulo theories. In *Proc. of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, pages 1177–1183. ijcai.org.
- Walker, T. T., Sturtevant, N. R., and Felner, A. (2018). Extended increasing cost tree search for non-unit cost domains. In *Proc. of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*, pages 534–540. ijcai.org.