

# Multi-Agent Path Finding Modulo Theory with Continuous Movements and the Sum of Costs Objective

Pavel Surynek<sup>[0000-0001-7200-0542]</sup>

Faculty of Information Technology  
Czech Technical University in Prague  
Thákurova 9, 160 00 Praha 6, Czechia  
pavel.surynek@fit.cvut.cz

**Abstract.** Multi-agent path finding with continuous movements and time (denoted MAPF<sup>R</sup>) is addressed. The task is to navigate agents that move smoothly between predefined positions to their individual goals so that they do not collide. Recently a novel solving approach for obtaining makespan optimal solutions called SMT-CBS<sup>R</sup> based on *satisfiability modulo theories* (SMT) has been introduced. We extend the approach further towards the sum-of-costs objective which is a more challenging case in the yes/no SMT environment due to more complex calculation of the objective. The new algorithm combines collision resolution known from conflict-based search (CBS) with previous generation of incomplete SAT encodings on top of a novel scheme for selecting decision variables in a potentially uncountable search space. We experimentally compare SMT-CBS<sup>R</sup> and previous CCBS (continuous conflict-based search) algorithm for MAPF<sup>R</sup>.

**Keywords:** path finding, multiple agents, robotic agents, logic reasoning, satisfiability modulo theory, sum-of-costs optimality

## 1 Introduction

In *multi-agent path finding* (MAPF) [15, 27, 24, 30, 37, 26, 25, 6] the task is to navigate agents from given starting positions to given individual goals. The problem takes place in an undirected graph  $G = (V, E)$  where agents from set  $A = \{a_1, a_2, \dots, a_k\}$  are placed in vertices with at most one agent per vertex. The navigation task can be then expressed formally as transforming an initial configuration of agents  $\alpha_0 : A \rightarrow V$  to a goal configuration  $\alpha_+ : A \rightarrow V$  using instantaneous movements across edges assuming no collision occurs.

To reflect various aspects of real-life applications, variants of MAPF have been introduced such as those considering *kinematic constraints* [9], *large agents* [17], *generalized costs* of actions [36], or *deadlines* [19] - see [18, 28] for more variants. Particularly in this work we are dealing with an extension of MAPF introduced only recently [1, 33] that considers continuous movements and time (MAPF<sup>R</sup>). Agents move smoothly along predefined curves interconnecting predefined positions placed arbitrarily in some continuous space. It is natural in MAPF<sup>R</sup> to assume geometric agents of various shapes that occupy certain volume in the space - circles in the 2D space, polygons, spheres in

the 3D space etc. In contrast to MAPF, where the collision is defined as the simultaneous occupation of a vertex or an edge by two agents, collisions are defined as any spatial overlap of agents' bodies in  $\text{MAPF}^{\mathcal{R}}$ .

The motivation behind introducing  $\text{MAPF}^{\mathcal{R}}$  is the need to construct more realistic paths in many applications such as controlling fleets of robots or aerial drones [7, 10] where continuous reasoning is closer to the reality than the standard MAPF.

The contribution of this paper consists in generalizing the previous makespan optimal approach for  $\text{MAPF}^{\mathcal{R}}$  [31, 33] that uses satisfiability modulo theory (SMT) reasoning [5, 20] for the sum-of-costs objective. The SMT paradigm constructs decision procedures for various complex logic theories by decomposing the decision problem into the propositional part having arbitrary Boolean structure and the complex theory part that is restricted on the conjunctive fragment. Our SMT-based algorithm called  $\text{SMT-CBS}^{\mathcal{R}}$  combines the Conflict-based Search (CBS) algorithm [8, 25] with previous algorithms for solving the standard MAPF using incomplete encodings [32] and continuous reasoning.

## 1.1 Previous Work

Using reductions of planning problems to propositional satisfiability has been coined in the SATPlan algorithm and its variants [11–14]. Here we are trying to apply a similar idea in the context of  $\text{MAPF}^{\mathcal{R}}$ . So far  $\text{MAPF}^{\mathcal{R}}$  has been solved by a modified version of CBS that tries to solve MAPF lazily by adding collision avoidance constraints on demand. The adaptation of CBS for  $\text{MAPF}^{\mathcal{R}}$  consists in implementing continuous collision detection while the high-level framework of the algorithm remains the same as demonstrated in the CCBS algorithm [1] (*continuous conflict-based search*).

We follow the idea of CBS too but instead of searching the tree of possible collision eliminations at the high-level we encode the requirement of having collision free paths as a propositional formula [4] and leave it to the SAT solver as done in [34]. We construct the formula *lazily* by adding collision elimination refinements following [32] where the lazy construction of incomplete encodings has been suggested for the standard MAPF within the algorithm called SMT-CBS. SMT-CBS works with propositional variables indexed by *agent*  $a$ , *vertex*  $v$ , and *time step*  $t$  with the meaning that if the variable is *TRUE*  $a$  in  $v$  at time step  $t$ . In  $\text{MAPF}^{\mathcal{R}}$  we however face major technical difficulty that we do not know necessary decision (propositional) variables in advance and due to continuous time we cannot enumerate them all. Hence we need to select from a potentially uncountable space those variables that are sufficient for finding the solution.

The previous application of SMT in  $\text{MAPF}^{\mathcal{R}}$  [33] focused on the makespan optimal solutions where the shortest duration of the plan is required. The **sum-of-costs** is another important objective used in the context of MAPF [26, 36]. Calculated as the summation over all agents of times they spend moving before arriving to the goal. Due to its more complex calculation, the sum-of-costs objective is more challenging to be integrated in the SMT-based solving framework.

## 1.2 MAPF with Continuous Movements and Time

We use the definition of MAPF with continuous movements and time denoted  $\text{MAPF}^{\mathcal{R}}$  from [1].  $\text{MAPF}^{\mathcal{R}}$  shares components with the standard MAPF: undirected graph  $G = (V, E)$ , set of agents  $A = \{a_1, a_2, \dots, a_k\}$ , and the initial and goal configuration of agents:  $\alpha_0 : A \rightarrow V$  and  $\alpha_+ : A \rightarrow V$ . A simple 2D variant of  $\text{MAPF}^{\mathcal{R}}$  is as follows:

**Definition 1.** ( $\text{MAPF}^{\mathcal{R}}$ ) *Multi-agent path finding with continuous time and space is a 5-tuple  $\Sigma^{\mathcal{R}} = (G = (V, E), A, \alpha_0, \alpha_+, \rho)$  where  $G, A, \alpha_0, \alpha_+$  are from the standard MAPF and  $\rho$  determines continuous extensions:*

- $\rho.x(v), \rho.y(v)$  for  $v \in V$  represent the position of vertex  $v$  in the 2D plane
- $\rho.speed(a)$  for  $a \in A$  determines constant speed of agent  $a$
- $\rho.radius(a)$  for  $a \in A$  determines the radius of agent  $a$ ; we assume that agents are circular discs with omni-directional ability of movements

For simplicity we assume circular agents with constant speed and instant acceleration. The major difference from the standard MAPF where agents move instantly between vertices (disappears in the source and appears in the target instantly) is that smooth continuous movement between a pair of vertices (positions) along the straight line interconnecting them takes place in  $\text{MAPF}^{\mathcal{R}}$ . Hence we need to be aware of the presence of agents at some point in the 2D plane at any time.

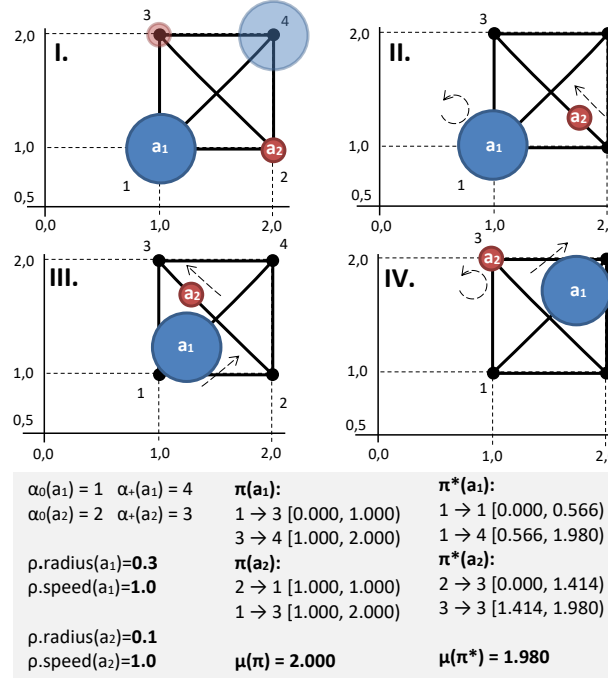
Collisions may occur between agents in  $\text{MAPF}^{\mathcal{R}}$  due to their volume; that is, they collide whenever their bodies **overlap**. In contrast to MAPF, collisions in  $\text{MAPF}^{\mathcal{R}}$  may occur not only in a single vertex or edge being shared by colliding agents but also on pairs of edges (lines interconnecting vertices) that are too close to each other and simultaneously traversed by large agents.

A solution to given  $\text{MAPF}^{\mathcal{R}}$   $\Sigma^{\mathcal{R}}$  is a collection of temporal plans for individual agents  $\pi = [\pi(a_1), \pi(a_2), \dots, \pi(a_k)]$  that are **mutually collision-free**. A temporal plan for agent  $a \in A$  is a sequence  $\pi(a) = [((\alpha_0(a), \alpha_1(a)), [t_0(a), t_1(a)]); ((\alpha_1(a), \alpha_2(a)), [t_1(a), t_2(a)]); \dots; ((\alpha_{m(a)-1}(a), \alpha_{m(a)}(a)), [t_{m(a)-1}(a), t_{m(a)}(a)])]$  where  $m(a)$  is the length of individual temporal plan and each pair  $(\alpha_i(a), \alpha_{i+1}(a)), [t_i(a), t_{i+1}(a)]$  corresponds to traversal event between a pair of vertices  $\alpha_i(a)$  and  $\alpha_{i+1}(a)$  starting at time  $t_i(a)$  and finished at  $t_{i+1}(a)$ .

It holds that  $t_i(a) < t_{i+1}(a)$  for  $i = 0, 1, \dots, m(a) - 1$ . Moreover consecutive events in the individual temporal plan must correspond to edge traversals or waiting actions, that is:  $\{\alpha_i(a), \alpha_{i+1}(a)\} \in E$  or  $\alpha_i(a) = \alpha_{i+1}(a)$ ; and times must reflect the speed of agents for non-wait actions.

The duration of individual temporal plan  $\pi(a)$  is called an *individual makespan*; denoted  $\mu(\pi(a)) = t_{m(a)}$ . The overall *makespan* of  $\pi$  is defined as  $\max_{i=1}^k \{\mu(\pi(a_i))\}$ . The individual makespan is sometimes called an *individual cost*. A *sum-of-cost* for given temporal plan  $\pi(a)$  is defined as  $\sum_{i=1}^k \mu(\pi(a_i))$ . An example of  $\text{MAPF}^{\mathcal{R}}$  and makespan/sum-of-costs optimal solution is shown in Figure 1.

Through straightforward reduction of MAPF to  $\text{MAPF}^{\mathcal{R}}$  it can be observed that finding a makespan or sum-of-costs optimal solution with continuous time is an NP-hard problem [22, 38].



**Fig. 1.** An example of MAPF<sup>R</sup> instance with two agents. A feasible makespan sub-optimal solution  $\pi$  (makespan  $\mu(\pi) = 2.0$ ) and makespan optimal solution  $\pi^*$  (makespan  $\mu(\pi^*) = 1.980$ ) are shown.

## 2 Solving MAPF with Continuous Time

Let us recall CCBS [1], a variant of CBS [25] modified for MAPF<sup>R</sup>. The idea of CBS algorithms is to resolve conflicts lazily.

### 2.1 Conflict-based Search

The high-level of CCBS searches a *constraint tree* (CT) using a priority queue ordered according to the sum-of-costs in the breadth first manner. CT is a binary tree where each node  $N$  contains a set of collision avoidance constraints  $N.\text{cons}$  - a set of triples  $(a_i, (u, v), [t_0, t_+))$  forbidding agent  $a_i$  to start smooth traversal of edge  $\{u, v\}$  (line) at any time between  $[t_0, t_+)$ , a solution  $N.\pi$  - a set of  $k$  individual temporal plans, and the sum-of-costs  $N.\xi$  of  $N.\pi$ .

The low-level in CCBS associated with node  $N$  searches for individual temporal plan with respect to set of constraints  $N.\text{cons}$ . For given agent  $a_i$ , this is the standard single source shortest path search from  $\alpha_0(a_i)$  to  $\alpha_+(a_i)$  that at time  $t$  cannot start to traverse any  $\{(u, v) \in E \mid (a_i, (u, v), [t_0, t_+)) \in N.\text{cons} \wedge t \in [t_0, t_+))\}$ . Various intelligent single source shortest path algorithms such as SIPP [21] can be used here.

CCBS stores nodes of CT into priority queue OPEN sorted according to the ascending makespan. At each step CBS takes node  $N$  with the lowest makespan from OPEN

---

**Algorithm 1:** CCBS algorithm for solving MAPF<sup>R</sup> for the sum-of-costs objective.

---

```

1 CBSR( $\Sigma^{\mathcal{R}} = (G = (V, E), A, \alpha_0, \alpha_+, \rho)$ )
2    $R.cons \leftarrow \emptyset$ 
3    $R.\pi \leftarrow \{\text{shortest temporal plan from } \alpha_0(a_i) \text{ to } \alpha_+(a_i) \mid i = 1, 2, \dots, k\}$ 
4    $R.\xi \leftarrow \sum_{i=1}^k \mu(N.\pi(a_i))$ 
5   OPEN  $\leftarrow \emptyset$ 
6   insert  $R$  into OPEN
7   while OPEN  $\neq \emptyset$  do
8      $N \leftarrow \min_{\xi}(\text{OPEN})$ 
9     remove- $\text{Min}_{\xi}(\text{OPEN})$ 
10    collisions  $\leftarrow$  validate-Plans( $N.\pi$ )
11    if collisions =  $\emptyset$  then
12      return  $N.\pi$ 
13    let  $(m_i \times m_j) \in$  collisions where  $m_i = (a_i, (u_i, v_i), [t_i^0, t_i^+])$  and
14       $m_j = (a_j, (u_j, v_j), [t_j^0, t_j^+])$ 
15       $([\tau_i^0, \tau_i^+]; [\tau_j^0, \tau_j^+]) \leftarrow$  resolve-Collision( $m_i, m_j$ )
16      for each  $m \in \{(m_i, [\tau_i^0, \tau_i^+]), (m_j, [\tau_j^0, \tau_j^+])\}$  do
17        let  $((a, (u, v), [t_0, t_+]), [\tau_0, \tau_+]) = m$ 
18         $N'.cons \leftarrow N.cons \cup \{(a, (u, v), [\tau_0, \tau_+])\}$ 
19         $N'.\pi \leftarrow N.\pi$ 
20        update( $a, N'.\pi, N'.cons$ )
21         $N'.\xi \leftarrow \sum_{i=1}^k \mu(N'.\pi(a_i))$ 
22        insert  $N'$  into OPEN

```

---

and checks if  $N.\pi$  represents non-colliding temporal plans. If there is no collision, the algorithm returns valid solution  $N.\pi$ . Otherwise the search branches by creating a new pair of nodes in CT - successors of  $N$ . Assume that a collision occurred between  $a_i$  traversing  $(u_i, v_i)$  during  $[t_i^0, t_i^+)$  and  $a_j$  traversing  $(u_j, v_j)$  during  $[t_j^0, t_j^+)$ . This collision can be avoided if either agent  $a_i$  or agent  $a_j$  waits after the other agent passes. We can calculate for  $a_i$  so called maximum *unsafe interval*  $[\tau_i^0, \tau_i^+)$  such that whenever  $a_i$  starts to traverse  $(u_i, v_i)$  at some time  $t \in [\tau_i^0, \tau_i^+)$  it ends up colliding with  $a_j$  assuming  $a_j$  did not try to avoid the collision. Hence  $a_i$  should wait until  $\tau_i^+$  to tightly avoid the collision with  $a_j$ . Similarly we can calculate maximum unsafe interval for  $a_j$ :  $[\tau_j^0, \tau_j^+)$ . These two options correspond to new successor nodes of  $N$ :  $N_1$  and  $N_2$  that inherit set of constraints from  $N$  as follows:  $N_1.cons = N.cons \cup \{(a_i, (u_i, v_i), [\tau_i^0, \tau_i^+])\}$  and  $N_2.cons = N.cons \cup \{(a_j, (u_j, v_j), [\tau_j^0, \tau_j^+])\}$ .  $N_1.\pi$  and  $N_2.\pi$  inherits plans from  $N.\pi$  except those for agents  $a_i$  and  $a_j$  respectively that are recalculated with respect to the constraints. After this  $N_1$  and  $N_2$  are inserted into OPEN.

## 2.2 A Satisfiability Modulo Theory Approach

A recent algorithm called SMT-CBS<sup>R</sup> [33] rephrases CCBS as problem solving in *satisfiability modulo theories* (SMT) [5, 35]. The basic use of SMT divides the satisfiability problem in some complex theory  $T$  into a propositional part that keeps the Boolean structure of the problem and a simplified procedure  $DECIDE_T$  that decides fragment of  $T$  restricted on *conjunctive formulae*. A general  $T$ -formula  $\Gamma$  being decided for satisfiability is transformed to a *propositional skeleton* by replacing its atoms with propositional variables. The standard SAT solver then decides what variables should be assigned *TRUE* in order to satisfy the skeleton - these variables tells what atoms hold in  $\Gamma$ .  $DECIDE_T$  then checks if the conjunction of atoms assigned *TRUE* is valid with respect to axioms of  $T$ . If so then satisfying assignment is returned. Otherwise a conflict from  $DECIDE_T$  (often called a *lemma*) is reported back to the SAT solver and the skeleton is extended with new constraints resolving the conflict. More generally not only new constraints are added to resolve the conflict but also new atoms can be added to  $\Gamma$ .

$T$  will be represented by a theory with axioms describing movement rules of MAPF<sup>R</sup>; a theory we will denote  $T_{MAPF^R}$ .  $DECIDE_{MAPF^R}$  can be naturally represented by the plan validation procedure from CCBS (validate-Plans).

## 2.3 RDD: Real Decision Diagram

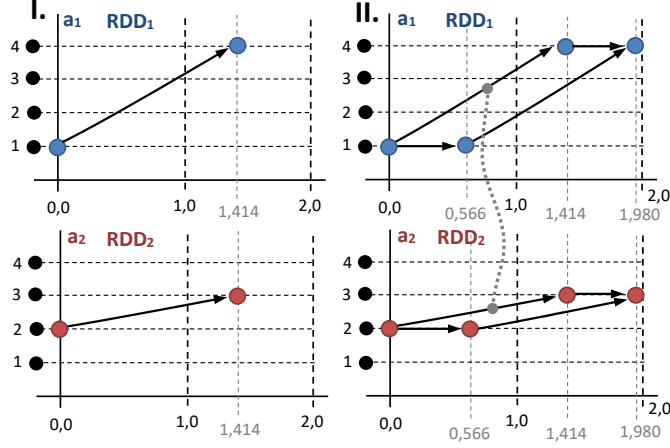
The key question in the propositional logic-based approach is what will be the decision variables. In the standard MAPF, time expansion of  $G$  for every time step can be done resulting in a multi-value decision diagram (MDD) [34] representing possible positions of agents at any time step. Since MAPF<sup>R</sup> is no longer discrete we cannot afford to use a decision variable for every time moment. We show how to restrict the decision variables on finitely many important moments only without compromising soundness nor optimality of the approach.

Analogously to MDD, we introduce *real decision diagram* (RDD).  $RDD_i$  defines for agent  $a_i$  its space-time positions and possible movements. Formally,  $RDD_i$  is a directed graph  $(X^i, E^i)$  where  $X_i$  consists of pairs  $(u, t)$  with  $u \in V$  and  $t \in \mathbb{R}_0^+$  is time and  $E_i$  consists of directed edges of the form  $((u, t_u); (v, t_v))$ . Edge  $((u, t_u); (v, t_v))$  correspond to agent's movement from  $u$  to  $v$  started at  $t_u$  and finished at  $t_v$ . Waiting in  $u$  is possible by introducing edge  $((u, t_u); (v, t'_u))$ . Pair  $(\alpha_0(a_i), 0) \in X_i$  indicates start and  $(\alpha_+(a_i), t)$  for some  $t$  corresponds to reaching the goal position.

RDDs for individual agents are constructed with respect to collision avoidance constraints. If there is no collision avoidance constraint then  $RDD_i$  simply corresponds to a shortest temporal plan for agent  $a_i$ . But if a collision avoidance constraint is present, say  $(a_i, (u, v), [\tau_0, \tau_+))$ , and we are considering movement starting in  $u$  at  $t$  that interferes with the constraint, then we need to generate a node into  $RDD_i$  that allows agent to wait until the unsafe interval passes by, that is node  $(u, \tau^+)$  and edge  $((u, \tau^+); (u, \tau^+))$  are added.

Similarly for wait constraints  $(a_i, (u, u), [\tau_0, \tau_+))$  that forbid waiting in  $u$  during  $[\tau_0, \tau_+)$ . In such a case, we need to anticipate the constraint before entering  $u$ , that is we can wait until  $\tau_+ - t_x$  in the source vertex before entering  $u$  where  $t_x$  is the time needed to traverse the edge towards  $u$ .

The process of building RDDs is formalized is described in details in [33]. An example of RDDs is shown in Figure 2.



**Fig. 2.** Real decision diagrams (RDDs) for agents  $a_1$  and  $a_2$  from MAPF<sup>R</sup> from Figure 1. Decisions corresponding to shortest paths for agents  $a_1$  and  $a_2$  moving diagonally towards their goals are shown:  $a_1 : 1 \rightarrow 4$ ,  $a_2 : 2 \rightarrow 3$  (left). This however results in a collision whose resolution is either waiting for agent  $a_1$  in vertex 1 from 0.000 until 0.566 or waiting for agent  $a_2$  in vertex 2 from 0.000 until 0.566; reflected in the next RDDs (right). Mutex is depicted using dotted line connecting arcs form RDD<sub>1</sub> and RDD<sub>2</sub>.

## 2.4 SAT Encoding from RDD

We introduce a decision variable for each node and edge  $[RDD_1, \dots, RDD_k]$ ;  $RDD_i = (X^i, E^i)$ : we have variable  $X_u^t(a_i)$  for each  $(u, t) \in X^i$  and  $\mathcal{E}_{u,v}^{t_u, t_v}(a_i)$  for each directed edge  $((u, t_u); (v, t_v)) \in E^i$ . The meaning of variables is that  $X_u^t(a_i)$  is *TRUE* if and only if agent  $a_i$  appears in  $u$  at time  $t$  and similarly for edges:  $\mathcal{E}_{u,v}^{t_u, t_v}(a_i)$  is *TRUE* if and only if  $a_i$  moves from  $u$  to  $v$  starting at time  $t_u$  and finishing at  $t_v$ .

MAPF<sup>R</sup> rules are encoded on top of these variables so that eventually we want to obtain formula  $\mathcal{F}(\mu)$  that encodes existence of a solution of makespan  $\mu$  to given MAPF<sup>R</sup>. We need to encode that agents do not skip but move along edges, do not disappear or appear from nowhere etc. We show below constraints stating that if agent  $a_i$  appears in vertex  $u$  at time step  $t_u$  then it has to leave through exactly one edge connected to  $u$  (constraint (2) although Pseudo-Boolean can be encoded using purely propositional means):

$$X_u^{t_u}(a_i) \Rightarrow \bigvee_{(v, t_v) \mid ((u, t_u), (v, t_v)) \in E^i} \mathcal{E}_{u,v}^{t_u, t_v}(a_i), \quad (1)$$

$$\sum_{(v, t_v) \mid ((u, t_u), (v, t_v)) \in E^i} \mathcal{E}_{u,v}^{t_u, t_v}(a_i) \leq 1 \quad (2)$$

$$\mathcal{E}_{u,v}^{t_u, t_v}(a_i) \Rightarrow \mathcal{X}_v^{t_v}(a_i) \quad (3)$$

Analogously to (2) we have constraint allowing a vertex to accept at most one agent through incoming edges; plus we need to enforce agents starting in  $\alpha_0$  and finishing in  $\alpha_+$ . Let us summarize soundness of the encoding in the following proposition (proof omitted).

**Proposition 1.** *Any satisfying assignment of  $\mathcal{F}(\mu)$  correspond to valid individual temporal plans for  $\Sigma^{\mathcal{R}}$  whose makespans are at most  $\mu$ .*

We a-priori do not add constraints for eliminating collisions; these are added lazily after assignment/solution validation. Hence,  $\mathcal{F}(\mu)$  constitutes an *incomplete model* for  $\Sigma^{\mathcal{R}}$ :  $\Sigma^{\mathcal{R}}$  is solvable within makespan  $\mu$  then  $\mathcal{F}(\mu)$  is satisfiable. The opposite implication does not hold since satisfying assignment of  $\mathcal{F}(\mu)$  may lead to a collision.

From the perspective of SMT, the propositional level does not understand geometric properties of agents so cannot know what simultaneous variable assignments are invalid. This information is only available at the level of theory  $T = \text{MAPF}^{\mathcal{R}}$  through  $\text{DECIDE}_{\text{MAPF}^{\mathcal{R}}}$ . We also leave the bounding of the sum-of-costs at the level of  $\text{DECIDE}_{\text{MAPF}^{\mathcal{R}}}$ .

## 2.5 Lazy Encoding of Mutex Refinements and Sum-of-Costs Bounds

The SMT-based algorithm itself is divided into two procedures: SMT-CBS <sup>$\mathcal{R}$</sup>  representing the main loop (Algorithm 2) and SMT-CBS-Fixed <sup>$\mathcal{R}$</sup>  solving the input MAPF <sup>$\mathcal{R}$</sup>  for a fixed maximum makespan  $\mu$  and sum-of-costs  $\xi$  (Algorithm 3).

Procedures *encode-Basic* and *augment-Basic* in Algorithm 3 build formula  $\mathcal{F}(\mu)$  according to given RDDs and the set of collected collision avoidance constraints. New collisions are resolved **lazily** by adding *mutexes* (disjunctive constraints). A collision is avoided in the same way as in CCBS; that is, one of the colliding agent waits. Collision eliminations are tried until a valid solution is obtained or until a failure for current  $\mu$  and  $\xi$  which means to try bigger makespan and sum-of-costs.

For resolving a collision we need to: **(1)** eliminate simultaneous execution of colliding movements and **(2)** augment the formula to enable avoidance (waiting). Assume a collision between agents  $a_i$  traversing  $(u_i, v_i)$  during  $[t_i^0, t_i^+)$  and  $a_j$  traversing  $(u_j, v_j)$  during  $[t_j^0, t_j^+)$  which corresponds to variables  $\mathcal{E}_{u_i, v_i}^{t_i^0, t_i^+}(a_i)$  and  $\mathcal{E}_{u_j, v_j}^{t_j^0, t_j^+}(a_j)$ . The collision can be eliminated by adding the following **mutex** (disjunction) to the formula:  $\neg \mathcal{E}_{u_i, v_i}^{t_i^0, t_i^+}(a_i) \vee \neg \mathcal{E}_{u_j, v_j}^{t_j^0, t_j^+}(a_j)$ . Satisfying assignments of the next  $\mathcal{F}(\mu)$  can no longer lead to this collision. Next, the formula is augmented according to new RDDs that reflect the collision - decision variables and respective constraints are added.

After resolving all collisions we check whether the sum-of-costs bound is satisfied by plan  $\pi$ . This can be done easily by checking if  $\mathcal{X}_u^{t_u}(a_i)$  variables across all agents together yield higher cost than  $\xi$  or not. If cost bound  $\xi$  is exceeded then corresponding nogood is recorded and added to  $\mathcal{F}(\mu)$  and the algorithm continues by searching for a new satisfying assignment to  $\mathcal{F}(\mu)$ . The nogood says that  $\mathcal{X}_u^{t_u}(a_i)$  variables that jointly exceed  $\xi$  cannot be simultaneously set to *TRUE*.



---

**Algorithm 2:** High-level of SMT-CBS<sup>R</sup> for the sum-of-costs objective.

---

```

1 SMT-CBSR( $\Sigma^R = (G = (V, E), A, \alpha_0, \alpha_+, \rho)$ )
2    $constraints \leftarrow \emptyset$ 
3    $\pi \leftarrow \{\pi^*(a_i) \text{ a shortest temporal plan from } \alpha_0(a_i) \text{ to } \alpha_+(a_i) \mid i = 1, 2, \dots, k\}$ 
4    $\mu \leftarrow \max_{i=1}^k \mu(\pi(a_i)); \xi \leftarrow \sum_{i=1}^k \mu(\pi(a_i))$ 
5   while TRUE do
6      $(\pi, constraints, \mu_{next}, \xi_{next}) \leftarrow \text{SMT-CBS-Fixed}^R(\Sigma^R, constraints, \mu, \xi)$ 
7     if  $\pi \neq \text{UNSAT}$  then
8       return  $\pi$ 
9      $\mu \leftarrow \mu_{next}; \xi \leftarrow \xi_{next}$ 

```

---

The set of pairs of collision avoidance constraints is propagated across entire execution of the algorithm. Constraints originating from a single collision are grouped in pairs so that it is possible to introduce mutexes for colliding movements discovered in previous steps.

Algorithm 2 shows the main loop of SMT-CBS<sup>R</sup>. The algorithm checks if there is a solution for  $\Sigma^R$  of makespan  $\mu$  and sum-of-costs  $\xi$ . It starts at the lower bound for  $\mu$  and  $\xi$  obtained as the duration of the longest from shortest individual temporal plans ignoring other agents and the sum of these lengths respectively.

Then  $\mu$  and  $\xi$  are iteratively increased in the main loop following the style of SAT-Plan [14]. The algorithm relies on the fact that the solvability of MAPF<sup>R</sup> w.r.t. cumulative objective like the sum-of-costs or makespan behaves as a non decreasing function. Hence trying increasing makespan and sum-of-costs eventually leads to finding the optimum provided we do not skip any relevant value.

We need to ensure important property in the makespan/sum-of-costs increasing scheme: any solution of sum-of-costs  $\xi$  has the makespan of at most  $\mu$ . The next sum-of-costs to try is obtained by taking the current sum-of-costs plus the smallest duration of the continuing movement (lines 17-27 of Algorithm 3).

The following proposition is a direct consequence of soundness of CCBS and soundness of the encoding (Proposition 1) and soundness of the makespan/sum-of-costs increasing scheme (proof omitted).

**Proposition 2.** *The SMT-CBS<sup>R</sup> algorithm returns sum-of-costs optimal solution for any solvable MAPF<sup>R</sup> instance  $\Sigma^R$ .*

### 3 Experimental Evaluation

We implemented SMT-CBS<sup>R</sup> in C++ to evaluate its performance and compared it with CCBS<sup>1</sup>. SMT-CBS<sup>R</sup> was implemented on top of Glucose 4 SAT solver [2] which

<sup>1</sup> To enable reproducibility of presented results we will provide complete source code of our solvers on the author's website.

**Algorithm 3:** Low-level of SMT-CBS<sup>R</sup>


---

```

1 SMT-CBS-FixedR( $\Sigma^R, cons, \mu, \xi$ )
2   RDD  $\leftarrow$  build-RDDs( $\Sigma^R, cons, \mu$ );  $\mathcal{F}(\mu) \leftarrow$  encode-Basic(RDD,  $\Sigma^R, cons, \mu$ )
3   while TRUE do
4     assignment  $\leftarrow$  consult-SAT-Solver( $\mathcal{F}(\mu)$ )
5     if assignment  $\neq$  UNSAT then
6        $\pi \leftarrow$  extract-Solution(assignment)
7       collisions  $\leftarrow$  validate-Plans( $\pi$ )
8       if collisions =  $\emptyset$  then
9         while TRUE do
10          nogoods  $\leftarrow$  validate-Cost( $\pi, \xi$ )
11          if nogoods =  $\emptyset$  then
12            return ( $\pi, \emptyset, UNDEF, UNDEF$ )
13           $\mathcal{F}(\mu) \leftarrow \mathcal{F}(\mu) \cup$  nogoods
14          assignment  $\leftarrow$  consult-SAT-Solver( $\mathcal{F}(\mu)$ )
15          if assignment = UNSAT then
16            ( $\mu_{next}, \xi_{next}$ )  $\leftarrow$  calc-Next-Bounds( $\mu, \xi, cons, RDD$ )
17            return (UNSAT, cons,  $\mu_{next}, \xi_{next}$ )
18           $\pi \leftarrow$  extract-Solution(assignment)
19        else
20          for each ( $m_i \times m_j$ )  $\in$  collisions where  $m_i = (a_i, (u_i, v_i), [t_i^0, t_i^+])$  and
21             $m_j = (a_j, (u_j, v_j), [t_j^0, t_j^+])$  do
22               $\mathcal{F}(\mu) \leftarrow \mathcal{F}(\mu) \wedge (\neg \mathcal{E}_{u_i, v_i}^{t_i^0, t_i^+}(a_i) \vee \neg \mathcal{E}_{u_j, v_j}^{t_j^0, t_j^+}(a_j))$ 
23              ( $[\tau_i^0, \tau_i^+]; [\tau_j^0, \tau_j^+]$ )  $\leftarrow$  resolve-Collision( $m_i, m_j$ )
24              cons  $\leftarrow cons \cup \{[(a_i, (u_i, v_i), [\tau_i^0, \tau_i^+]); (a_j, (u_j, v_j), [\tau_j^0, \tau_j^+])]\}$ 
25          RDD  $\leftarrow$  build-RDDs( $\Sigma^R, cons, \mu$ );  $\mathcal{F}(\mu) \leftarrow$  augment-Basic(RDD,  $\Sigma^R, cons$ )
26        else
27          ( $\mu_{next}, \xi_{next}$ )  $\leftarrow$  calc-Next-Bounds( $\mu, \xi, cons, RDD$ )
28          return (UNSAT, cons,  $\mu_{next}, \xi_{next}$ )

```

---

ranks among the best SAT solvers according to recent SAT solver competitions [3]. The solver is consulted in the incremental mode if the formula is extended with new clauses. In case of CCBS, we used the existing C++ implementation [1].

### 3.1 Benchmarks and Setup

SMT-CBS<sup>R</sup> and CCBS were tested on benchmarks from the movinai.com collection [29]. We tested algorithms on three categories of benchmarks:

- (i) **small** empty grids (presented representative benchmark `empty-16-16`),
- (ii) **medium** sized grids with regular obstacles (presented `maze-32-32-4`),
- (iii) **large** game maps (presented `ost003d`).

In each benchmark, we interconnected cells using the  $2^K$ -neighborhood [23] for  $K = 3, 4, 5$  - the same style of generating benchmarks as used in [1] ( $K = 2$  corresponds to MAPF hence omitted). Instances consisting of  $k$  agents were generated by taking first  $k$  agents from random scenario files accompanying each benchmark on movinai.com. Having 25 scenarios for each benchmarks this yields to 25 instances per number of agents.

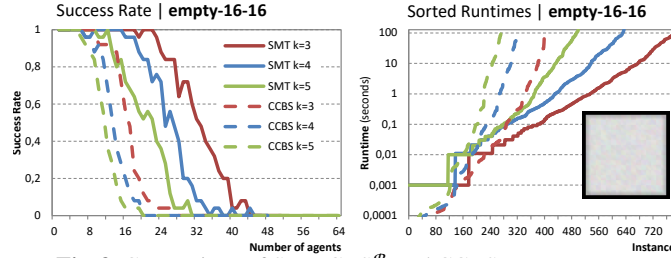


Fig. 3. Comparison of SMT-CBS $^{\mathcal{R}}$  and CCBS on empty-16-16.

Part of the results obtained in our experimentation is presented in this section<sup>2</sup>. For each presented benchmark we show *success rate* as a function of the number of agents. That is, we calculate the ratio out of 25 instances per number of agents where the tested algorithm finished under the timeout of 120 seconds. In addition to this, we also show concrete runtimes sorted in the ascending order. Results for one selected representative benchmark from each category are shown in Figures 3, 4, and 5.

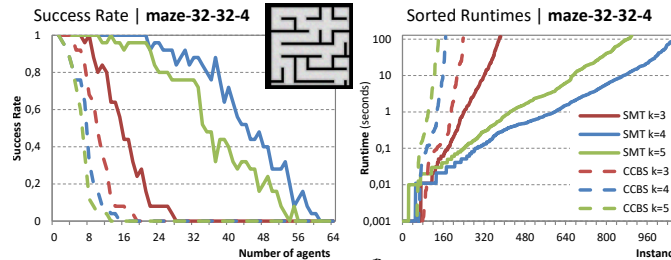


Fig. 4. Comparison of SMT-CBS $^{\mathcal{R}}$  and CCBS on maze-32-32-4.

The observable trend is that the difficulty of the problem increases with increasing size of the  $K$ -neighborhood with notable exception of maze-32-32-4 for  $K = 4$  and  $K = 5$  which turned out to be easier than  $K = 3$  for SMT-CBS $^{\mathcal{R}}$ .

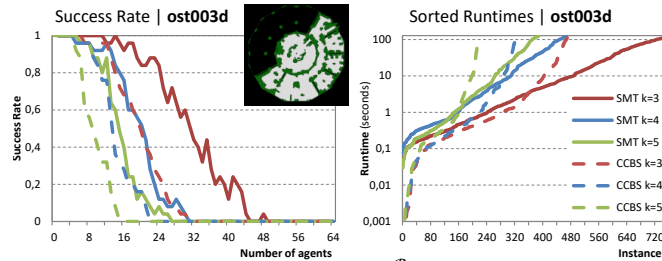
Throughout all benchmarks SMT-CBS $^{\mathcal{R}}$  tends to outperform CCBS. The dominance of SMT-CBS $^{\mathcal{R}}$  is most visible in medium sized benchmarks. CCBS is, on the other hand, faster in instances containing few agents. The gap between SMT-CBS $^{\mathcal{R}}$  and CCBS is smallest in large maps where SMT-CBS $^{\mathcal{R}}$  struggles with relatively big

<sup>2</sup> All experiments were run on a system with Ryzen 7 3.0 GHz, 16 GB RAM, under Ubuntu Linux 18.

overhead caused by the big size of the map (the encoding is proportionally big). Here SMT-CBS<sup>R</sup> wins only in hard cases.

## 4 Discussion and Conclusion

We extended the approach based on *satisfiability modulo theories* (SMT) for solving MAPF<sup>R</sup> from the makespan objective towards the sum-of-costs objective. Our approach builds on the idea of treating constraints lazily as suggested in the CBS algorithm but instead of branching the search after encountering a conflict we refine the propositional model with the conflict elimination disjunctive constraint as it has been done in previous application of SMT in the standard MAPF. Bounding the sum-of-costs is done in similar lazy way through introducing nogoods incrementally. If it is detected that a conflict free solution exceeds given cost bound then decisions that jointly induce cost greater than given bound are forbidden via a nogood (that is, at least one of these decisions must not be taken).



**Fig. 5.** Comparison of SMT-CBS<sup>R</sup> and CCBS on ost003d.

We compared SMT-CBS<sup>R</sup> with CCBS [1], currently the only alternative algorithm for MAPF<sup>R</sup> that modifies the standard CBS algorithm, on a number of benchmarks. The outcome of our comparison is that SMT-CBS<sup>R</sup> performs well against CCBS. The best results SMT-CBS<sup>R</sup> are observable on medium sized benchmarks with regular obstacles. We attribute the better runtime results of SMT-CBS<sup>R</sup> to more efficient handling of disjunctive conflicts in the underlying SAT solver through *propagation*, *clause learning*, and other mechanisms. On the other hand SMT-CBS<sup>R</sup> is less efficient on large instances with few agents.

We plan to extend the RDD generation scheme to directional agents where we need to add the third dimension in addition to space (vertices) and time: *direction* (angle). The work on MAPF<sup>R</sup> could be further developed into multi-robot motion planning in continuous configuration spaces [16].

## Acknowledgement

This research has been supported by GAČR - the Czech Science Foundation, grant registration number 19-17966S. We would like to thank anonymous reviewers for their valuable comments.

## References

1. Andreychuk, A., Yakovlev, K.S., Atzmon, D., Stern, R.: Multi-agent pathfinding with continuous time. In: Proceedings of IJCAI 2019. pp. 39–45 (2019)
2. Audemard, G., Simon, L.: Predicting learnt clauses quality in modern SAT solvers. In: IJCAI. pp. 399–404 (2009)
3. Balyo, T., Heule, M.J.H., Jarvisalo, M.: SAT competition 2016: Recent developments. In: AAI 2017. pp. 5061–5063 (2017)
4. Biere, A., Heule, M., van Maaren, H., Walsh, T.: Handbook of Satisfiability. IOS Press (2009)
5. Bofill, M., Palahí, M., Suy, J., Villaret, M.: Solving constraint satisfaction problems with SAT modulo theories. *Constraints* **17**(3), 273–303 (2012)
6. Botea, A., Surynek, P.: Multi-agent path finding on strongly biconnected digraphs. In: Proceedings of AAI 2016. pp. 2024–2030 (2015)
7. Cáp, M., Novák, P., Vokřínek, J., Pechoucek, M.: Multi-agent RRT: sampling-based cooperative pathfinding. In: Proceedings of AAMAS 2013. pp. 1263–1264 (2013)
8. Felner, A., Li, J., Boyarski, E., Ma, H., L.Cohen, Kumar, T.K.S., Koenig, S.: Adding heuristics to conflict-based search for multi-agent path finding. In: Proceedings of ICAPS 2018. pp. 83–87 (2018)
9. Höning, W., Kumar, T.K.S., Cohen, L., Ma, H., Xu, H., Ayanian, N., Koenig, S.: Summary: Multi-agent path finding with kinematic constraints. In: Proceedings of IJCAI 2017. pp. 4869–4873 (2017)
10. Janovsky, P., Cáp, M., Vokřínek, J.: Finding coordinated paths for multiple holonomic agents in 2-d polygonal environment. In: Proceedings of AAMAS 2014. pp. 1117–1124 (2014)
11. Kautz, H.A.: Deconstructing planning as satisfiability. In: Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, 2006. pp. 1524–1526. AAAI Press (2006)
12. Kautz, H.A., Selman, B.: Planning as satisfiability. In: Proceedings ECAI 1992. pp. 359–363 (1992)
13. Kautz, H.A., Selman, B.: Pushing the envelope: Planning, propositional logic and stochastic search. In: Proceedings of AAAI 1996. pp. 1194–1201 (1996)
14. Kautz, H.A., Selman, B.: Unifying sat-based and graph-based planning. In: Proceedings of IJCAI 1999. pp. 318–325 (1999)
15. Kornhauser, D., Miller, G.L., Spirakis, P.G.: Coordinating pebble motion on graphs, the diameter of permutation groups, and applications. In: FOCS, 1984. pp. 241–250 (1984)
16. LaValle, S.M.: Planning algorithms. Cambridge University Press (2006)
17. Li, J., Surynek, P., Felner, A., Ma, H., Koenig, S.: Multi-agent path finding for large agents. In: Proceedings of AAAI 2019. AAAI Press (2019)
18. Ma, H., Koenig, S., Ayanian, N., Cohen, L., Höning, W., Kumar, T.K.S., Uras, T., Xu, H., Tovey, C.A., Sharon, G.: Overview: Generalizations of multi-agent path finding to real-world scenarios. CoRR **abs/1702.05515** (2017), <http://arxiv.org/abs/1702.05515>
19. Ma, H., Wagner, G., Felner, A., Li, J., Kumar, T.K.S., Koenig, S.: Multi-agent path finding with deadlines. In: Proceedings of IJCAI 2018. pp. 417–423 (2018)
20. Nieuwenhuis, R.: SAT modulo theories: Getting the best of SAT and global constraint filtering. In: Proceedings of CP 2010. pp. 1–2 (2010)
21. Phillips, M., Likhachev, M.: SIPP: safe interval path planning for dynamic environments. In: Proceedings of ICRA 2011. pp. 5628–5635 (2011)
22. Ratner, D., Warmuth, M.K.: Nxn puzzle and related relocation problem. *J. Symb. Comput.* **10**(2), 111–138 (1990)
23. Rivera, N., Hernández, C., Baier, J.A.: Grid pathfinding on the  $2k$  neighborhoods. In: Proceedings of AAAI 2017. pp. 891–897 (2017)

24. Ryan, M.R.K.: Exploiting subgraph structure in multi-robot path planning. *J. Artif. Intell. Res. (JAIR)* **31**, 497–542 (2008)
25. Sharon, G., Stern, R., Felner, A., Sturtevant, N.: Conflict-based search for optimal multi-agent pathfinding. *Artif. Intell.* **219**, 40–66 (2015)
26. Sharon, G., Stern, R., Goldenberg, M., Felner, A.: The increasing cost tree search for optimal multi-agent pathfinding. *Artif. Intell.* **195**, 470–495 (2013)
27. Silver, D.: Cooperative pathfinding. In: *AIIDE*. pp. 117–122 (2005)
28. Stern, R.: Multi-agent path finding - an overview. In: Osipov, G.S., Panov, A.I., Yakovlev, K.S. (eds.) *Artificial Intelligence - 5th RAAI Summer School. Lecture Notes in Computer Science*, vol. 11866, pp. 96–115. Springer (2019)
29. Sturtevant, N.R.: Benchmarks for grid-based pathfinding. *Computational Intelligence and AI in Games* **4**(2), 144–148 (2012)
30. Surynek, P.: A novel approach to path planning for multiple robots in bi-connected graphs. In: *ICRA 2009*. pp. 3613–3619 (2009)
31. Surynek, P.: Multi-agent path finding with continuous time and geometric agents viewed through satisfiability modulo theories (SMT). In: Surynek, P., Yeoh, W. (eds.) *Proceedings of the Twelfth International Symposium on Combinatorial Search, SOCS 2019*. pp. 200–201. AAAI Press (2019)
32. Surynek, P.: Unifying search-based and compilation-based approaches to multi-agent path finding through satisfiability modulo theories. In: *Proceedings of IJCAI 2019*. pp. 1177–1183 (2019)
33. Surynek, P.: On satisfiability modulo theories in continuous multi-agent path finding: Compilation-based and search-based approaches compared. In: Rocha, A.P., Steels, L., van den Herik, H.J. (eds.) *Proceedings of the 12th International Conference on Agents and Artificial Intelligence, ICAART 2020, Volume 2, 2020*. pp. 182–193. SCITEPRESS (2020)
34. Surynek, P., Felner, A., Stern, R., Boyarski, E.: Efficient SAT approach to multi-agent path finding under the sum of costs objective. In: *ECAI*. pp. 810–818 (2016)
35. Tinelli, C.: Foundations of satisfiability modulo theories. In: *Proceedings of WoLLIC 2010*. p. 58 (2010)
36. Walker, T.T., Sturtevant, N.R., Felner, A.: Extended increasing cost tree search for non-unit cost domains. In: *Proceedings of IJCAI 2018*. pp. 534–540 (2018)
37. Wang, K., Botea, A.: MAPP: a scalable multi-agent path planning algorithm with tractability and completeness guarantees. *JAIR* **42**, 55–90 (2011)
38. Yu, J., LaValle, S.M.: Optimal multi-robot path planning on graphs: Structure and computational complexity. *CoRR* **abs/1507.03289** (2015)