

UC Berkeley

UC Berkeley Previously Published Works

Title

Scalable Lazy SMT-Based Motion Planning

Permalink

<https://escholarship.org/uc/item/8z44m3zx>

ISBN

978-1-5090-1837-6

Authors

Shoukry, Yasser

Nuzzo, Pierluigi

Saha, Indranil

et al.

Publication Date

2016-12-01

DOI

10.1109/cdc.2016.7799298

Peer reviewed

Scalable Motion Planning Using Lazy SMT-Based Solving

Yasser Shoukry^{†§}

Pierluigi Nuzzo[†]

Indranil Saha^{**}

Alberto L. Sangiovanni-Vincentelli[†]

Sanjit A. Seshia[†]

George J. Pappas^{*}

Paulo Tabuada[§]

Abstract—We present a scalable robot motion planning algorithm for reach-avoid problems. We assume a discrete-time, linear model of the robot dynamics and a workspace described by a set of obstacles and a target region, where both the obstacles and the region are polyhedra. Our goal is to construct a trajectory, and the associated control strategy, that steers the robot from its initial point to the target while avoiding obstacles. Differently from previous approaches, based on the discretization of the continuous state space or uniform discretization of the workspace, our approach, inspired by the lazy satisfiability modulo theory paradigm, decomposes the planning problem into smaller subproblems, which can be efficiently solved using specialized solvers. At each iteration, we use a coarse, obstacle-based discretization of the workspace to obtain candidate high-level, discrete plans that solve a set of Boolean constraints, while completely abstracting the low-level continuous dynamics. The feasibility of the proposed plans is then checked via a convex program, under constraints on both the system dynamics and the control inputs, and new candidate plans are generated until a feasible one is found. To achieve scalability, we show how to generate succinct explanations for the infeasibility of a discrete plan by exploiting a relaxation of the convex program that allows detecting the earliest possible occurrence of an infeasible transition between workspace regions. Simulation results show that our algorithm favorably compares with state-of-the-art techniques and scales well for complex systems, including robot dynamics with up to 50 continuous states.

I. INTRODUCTION

Algorithmic control synthesis from formal specifications captured by logic formalisms, such as Linear Temporal Logic (LTL) [1], holds considerable promise for providing correct-by-construction controllers for a rich set of tasks [2]–[8] and safety-critical applications in robotics (e.g., in navigation, manipulation, and surgery) and autonomous systems (e.g., unmanned aircraft and self-driving cars). However, the complexity of today’s systems poses a set of unprecedented challenges to synthesis techniques.

A major difficulty stems from the need to reason about the tight integration of discrete abstractions (as in *task planning*) with continuous motions (*motion planning*) [9]. This integration can become daunting for complex, high-dimensional systems, since a vast discrete/continuous space must be searched while accounting for complex geometries, motion dynamics, collision avoidance, and temporal goals. In complex systems, effective discrete planning techniques may produce solutions

that are not realizable due to constraints imposed by dynamics; on the other hand, effective methods for generating collision-free and dynamically-feasible trajectories may end up with violating the constraints imposed by the task planner.

In this paper, we address these challenges by focusing on an essential motion planning problem, which is embedded in almost all robotics applications, i.e., the reach-avoid problem. Given the robot dynamics, described as a discrete-time linear system, an initial state, a description of the workspace, and a target region, we aim at planning a collision-free and dynamically-feasible motion trajectory that steers the robot from its initial point to the target region. However, our approach can be directly extended to motion planning from generic LTL specifications by leveraging the bounded model checking encoding technique for LTL model checking by Biere et al. [10] to encode the discrete planning problem.

A growing body of work has focused, over the years, on the synthesis of reactive controllers to perform high-level tasks. A first category of techniques in the context of motion planning utilizes a discrete abstraction of the system, often obtained by partitioning the continuous state space into polytopes, and an automata theoretic approach to synthesize the controller [2]–[6]. However, these approaches are subject to the *curse of dimensionality* and become usually impractical for systems with more than five continuous states [11]. Moreover, some of these approaches assume the availability of low-level feedback controllers that are capable of generating feasible trajectories that are compatible with each automaton action, which may not always be the case for complex robotic systems. A second category of approaches attempts at synthesizing the high-level planner together with the associated low-level controller, by either leveraging mixed integer linear programming (MILP) encodings of task specifications [12], [13] or sampling-based methods [14], [15]. MILP-based planners can leverage the empirical performance of state-of-the-art solvers to solve for both the discrete and continuous constraints at the same time; however, they still tend to be impractical when the problem size grows. On the other hand, sampling-based techniques tend to perform poorly on the obstacle avoidance problem in the presence of narrow passages [16], and do not have, in general, control over the number of hops of the generated trajectory.

In this paper, we propose a scalable solution for the integration of task planning and robot motion planning. Differently from previous approaches, based on the discretization of the continuous state space or uniform discretization of the workspace, our algorithm, inspired by the *lazy Satisfiability Modulo Theory (SMT)* paradigm [17], aims at decomposing the planning problem into smaller subproblems, involving only Boolean or only convex constraints, which can be efficiently solved using specialized solvers.

Affiliations: [†]Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, CA; [§]Department of Electrical Engineering, University of California, Los Angeles, CA; ^{*}Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA; ^{**}Department of Computer Science and Engineering, Indian Institute of Technology, Kanpur, India.

This work was partially sponsored by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA and by the NSF project ExCAPE: Expeditions in Computer Augmented Program Engineering.

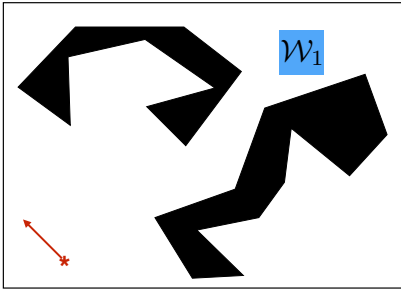


Fig. 1. Pictorial representation of the workspace, obstacles, and target for a reach-avoid problem. The initial state of the robot, including both position and angle, is represented by the red star and arrow.

Our methodology differs from classical approaches to reach-avoid problems [18]–[20], e.g., based on the solution of a Hamilton-Jacobi-Isaacs equation. Rather than formulating a complete, general optimization problem, which may be computationally challenging, we focus on solving a special case accurately and efficiently. We then aim at leveraging this result as a building block to solve more general problems, e.g., by supporting complex LTL specifications, through abstraction and refinement techniques. In this respect, our methods are also inspired by the counterexample-based control approaches [21]. Simulation results show that our algorithm favorably compares with other state-of-the-art techniques. For brevity, we omit here the proofs of the main technical results and report them in an extended version of the paper [22].

II. PROBLEM FORMULATION

We consider a robot that moves in a workspace $\mathcal{W} \subset \mathbb{R}^w$ where w can be 2 or 3, corresponding, respectively, to a 2-dimensional or 3-dimensional workspace. We use $\|a\|$ to denote the infinity norm of a and formulate the reach-avoid motion planning problem as follows.

A. Robot Model

We assume the robot dynamics is described by a discrete-time, input-constrained, linear system of the form:

$$x_{t+1} = Ax_t + Bu_t, \quad (\text{II.1})$$

$$x_0 = \bar{x}, \quad \|u_t\| \leq \bar{u} \quad \forall t \in \mathbb{N} \quad (\text{II.2})$$

where $x_t \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state of the robot at time $t \in \mathbb{N}$, $u_t \in \mathcal{U} \subseteq \mathbb{R}^m$ is the robot input, \bar{x} is the robot initial state and \bar{u} is the input constraint. The matrices A and B represent the robot dynamics and have appropriate dimensions. For a robot with nonlinear dynamics that is either differentially flat or feedback linearizable, the state space model (II.1) corresponds to its feedback linearized dynamics.

B. Workspace

We assume the robot must avoid a set of *obstacles* $\mathcal{O} = \{\mathcal{O}_1, \dots, \mathcal{O}_o\}$, with $\mathcal{O}_i \subset \mathbb{R}^w$, and represent the *workspace* as $\mathcal{W} = \mathcal{W}_0 \cup \mathcal{W}_G$, where \mathcal{W}_G is a *region of interest* (the target region or *Goal*) and \mathcal{W}_0 is the *free space*, characterized by the absence of both obstacles and target. As pictorially represented in Fig. 1, both the target region and the obstacles are assumed to be polygons.

To better describe the interplay between discrete planner and continuous planner in our algorithm, it is also useful to uniquely associate to the target region and the free space defined above an atomic proposition in the set $\Pi = \{\pi_0, \pi_G\}$. We then denote by $h_{\mathcal{W} \rightarrow \Pi} : \mathcal{W} \rightarrow \Pi$ the map from each point $w \in \mathcal{W}$ to the atomic proposition $\pi_i \in \Pi$ that evaluates to one (true) at w . Moreover, a subset of the robot state variables, describing its position (coordinates), is also used to describe \mathcal{W} . Therefore, we denote as $h_{\mathcal{X} \rightarrow \mathcal{W}} : \mathcal{X} \rightarrow \mathcal{W}$ the natural projection of the state x onto the workspace \mathcal{W} , and by $h_{\mathcal{X} \rightarrow \Pi}$ the map from the robot state to the set of atomic propositions, obtained after projecting the state onto the workspace, i.e., $h_{\mathcal{X} \rightarrow \Pi}(x) = h_{\mathcal{W} \rightarrow \Pi}(h_{\mathcal{X} \rightarrow \mathcal{W}}(x))$.

Finally, given the set $\overline{\mathcal{W}} = \{\mathcal{W}_0, \mathcal{W}_G\}$, we introduce an *adjacency function* $Adj : \overline{\mathcal{W}} \times \overline{\mathcal{W}} \rightarrow \mathbb{B}$ over the pairs of non-overlapping elements in $\overline{\mathcal{W}}$ such that $Adj(\mathcal{W}_i, \mathcal{W}_j) = 1$ if \mathcal{W}_i and \mathcal{W}_j are adjacent and 0 otherwise¹. Because of the one-to-one correspondence between elements in $\overline{\mathcal{W}}$ and atomic propositions in Π , we also write $Adj(\pi_i, \pi_j) = 1$ if π_i and π_j are associated with adjacent regions in $\overline{\mathcal{W}}$ and 0 otherwise. Moreover, for all i , $Adj(\pi_i, \pi_i) = 1$ holds.

C. Problem Definition

Definition 2.1 (Input Problem Instance): An input problem instance is defined as the tuple $\mathcal{P} = \langle \mathcal{W}, \Pi, Adj, (A, B), \bar{x}, \bar{u} \rangle$, where:

- \mathcal{W} is the workspace,
- Π is the set of atomic propositions corresponding to the target and the free space,
- Adj is the adjacency function defining the connectivity of the different regions in the workspace,
- (A, B) is the robot dynamics,
- \bar{x} is the initial state of the robot,
- \bar{u} is the bound on the robot inputs.

Definition 2.2 (Trajectory): A *trajectory* of a robot for an input problem instance $\mathcal{P} = \langle \mathcal{W}, \Pi, Adj, (A, B), \bar{x}, \bar{u} \rangle$ is defined as a pair of finite sequences (x, ρ) where $x = x_0 x_1 x_2 \dots x_{L+1}$, with $x_i \in \mathcal{X}$, is a sequence of states and $\rho = \rho_0 \rho_1 \rho_2 \dots \rho_{L+1}$, with $\rho_i \in \Pi$, is a sequence of propositions associated with the workspace regions, and such that $h_{\mathcal{X} \rightarrow \Pi}(x_i) = \rho_i$ for any $0 \leq i \leq L + 1$. Because of the one-to-one correspondence between workspace regions and atomic propositions, we also use $\rho^{\mathcal{W}}$ to denote the sequence of regions associated with ρ , and call ρ (or $\rho^{\mathcal{W}}$) the *region trajectory*.

Definition 2.3 (Valid trajectory): For an input problem instance $\mathcal{P} = \langle \mathcal{W}, \Pi, Adj, (A, B), \bar{x}, \bar{u} \rangle$, a trajectory (x, ρ) is called a *valid trajectory*, if the following holds:

- **Initial state constraint:** $x_0 = \bar{x}$,
- **Dynamics and input constraints:** there exists u_i such that $x_{i+1} = Ax_i + Bu_i$ and $\|u_i\| \leq \bar{u}$,
- **Workspace constraints:** $Adj(\rho_i, \rho_{i+1}) = 1 \quad \forall i : 0 \leq i \leq L + 1$,
- **Final state constraints:** $\rho_{L+1} = \pi_G$.

¹Two polyhedra in \mathbb{R}^w are adjacent if they share a facet of dimension $w - 1$.

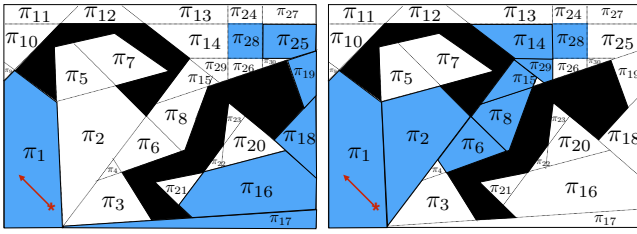


Fig. 2. Coarse discretization of the free space for the workspace in Fig. 1, for the same configuration of obstacles and target region. On the left side, a candidate path to satisfy the reach-avoid specification with $\pi_G = \pi_{28}$ is highlighted in blue. On the right side, an alternative path to the same goal is proposed.

We now formally define the motion planning problem that we solve in this paper.

Problem 2.4 (Motion Planning Problem): Given an input problem instance $\mathcal{P} = \langle \mathcal{W}, \Pi, Adj, (A, B), \bar{x}, \bar{u} \rangle$, synthesize a valid trajectory for the robot.

III. SMT-BASED SOLUTION

The problem of synthesizing a robot motion plan under constraints on the continuous dynamics is traditionally solved using expensive discretizations of the state space, which typically lead to state explosion as the number of continuous states and the number of obstacles increases. Our strategy aims, instead, at creating coarser abstractions of both the state space and the workspace, thus effectively decoupling the problem of *generating an obstacle-free path* from the one of *checking its physical realizability*. By leveraging the *lazy satisfiability modulo theory* (SMT) paradigm [17], we then partition the planning problem into two smaller subproblems involving reasoning, respectively, on sets of discrete and continuous variables from the original problem. These subproblems can be efficiently solved using specialized techniques, following a similar approach as in the CALCS [23] and IMHOTEP-SMT solvers [24]–[26].

As illustrated in Algorithm 1 and Fig. 2, we start by computing a multi-resolution discretization of the free space \mathcal{W}_0 based on the obstacles and the target. Unlike grid-based methods, where the workspace is discretized using a grid (or mesh) of (small) uniform resolution, the coarse abstraction used by our method avoids state explosion. Our decomposition procedure is similar to the ones previously proposed in the literature for triangular [27] or polygonal [18] representations. After the discretization step, we generate a new set Π^* of atomic propositions, representing the new abstraction, and the corresponding adjacency function, denoted by Adj^* .

Based on the above discretization, our solution follows an iterative approach combining a SAT solver (DIS-PLAN) and a theory solver (CON-PLAN). At each iteration, we start by generating a candidate high-level path ρ that satisfies the set of constraints of the reach-avoid problem, encoded using a Boolean formula ξ . Since this path is only defined over the set of Boolean propositions Π^* , its computation will ignore the robot dynamics and the input constraints. Because of the coarse abstraction of the workspace, the number of atomic propositions in Π^* is only determined by the obstacle configuration in the workspace and does not depend on the

Algorithm 1 SMT-BASED MOTION PLANNER

- 1: $(\Pi^*, \mathcal{W}^*, Adj^*) := \text{WKSP.ABSTRACT}(\Pi, \mathcal{W}, Adj)$;
 - 2: Initialize the horizon: $L := 1$;
 - 3: **while** Trajectory is not found **do**
 - 4: $(\text{DSTATUS}, \rho) := \text{DIS-PLAN}(\Pi^*, Adj^*, \xi)$;
 - 5: **if** DSTATUS == UNSAT **then**
 - 6: Increase horizon: $L := L + 1$;
 - 7: **else**
 - 8: $(\text{CSTATUS}, x, u) := \text{CON-PLAN.CHECK}(\bar{x}, \bar{u}, \rho)$;
 - 9: **if** CSTATUS == Infeasible **then**
 - 10: $\phi_{\text{ce}} := \text{CON-PLAN.COUNTEREXAMPLE}(\bar{x}, \bar{u}, \rho)$;
 - 11: $\xi := \xi \wedge \phi_{\text{ce}}$;
 - 12: **return** (ρ, x, u) ;
-

state dimension of the continuous dynamics. This step can then be performed efficiently using off-the-shelf SAT solvers.

We can then check the *feasibility* of the generated path ρ with respect to the system dynamics (A, B) , the control inputs \bar{u} , and the robot initial state \bar{x} , by casting it as a *convex* optimization problem. In fact, while generating a path that satisfies both the robot dynamics and the constraints imposed by the reach-avoid problem may be, in general, non-convex, our SMT-based architecture is able to reason about this complex problem by decomposing it into a sequence of smaller subproblems, each combining a Boolean satisfiability problem, defined only over Π^* , and a convex optimization problem, defined only over the real-valued variables x (state) and u (input). If both the Boolean and the real-valued constraints are satisfied, we return a valid trajectory consisting of the proposed plan and the corresponding state and control input trajectories. Otherwise, the proposed sequence ρ is marked as infeasible and new candidate plans are generated, such as the ones in Fig. 2, until a feasible one is found.

In this iterative scheme, learning “succinct explanations” that can capture the root causes for the infeasibility of a plan, and rule out the largest possible number of invalid plans per iteration, is instrumental to achieve fast convergence and scalability. To do so, we exploit convex programming to check the feasibility of a plan in terms of a conjunction of continuous constraints while minimizing a certain cost. We then use a relaxation of the continuous trajectory feasibility problem with slack variables to detect the earliest possible occurrence of an infeasible transition between two workspace regions and suggest the generation of plans that can avoid such a transition.

In what follows, we provide details on both the discrete and continuous plan generation mechanisms, including the implementation of DIS-PLAN and CON-PLAN, as well as on the generation of succinct infeasibility proofs.

IV. GENERATION OF THE HIGH-LEVEL DISCRETE PLAN

As represented in Fig. 1, a classic reach-avoid specification defines an initial point, a *Goal* (target) region ($\mathcal{W}_G = \mathcal{W}_1$), and a set of obstacles to avoid. Given an input problem instance $\mathcal{P} = \langle \mathcal{W}, \Pi, Adj, (A, B), \bar{x}, \bar{u} \rangle$ for this problem, DIS-PLAN performs a multi-resolution discretization of the free space and generates a formula that represents any valid trajectory

(x, ρ) of the robot. The decision variables for the formula are given by the sequence of atomic propositions associated with the regions to be occupied by the robot. Given the new set Π^* of atomic propositions associated to the workspace regions after discretization of the free space, and the corresponding adjacency function Adj^* , we represent the region trajectory ρ as $\rho = (\rho_0 \rho_1 \dots \rho_{L+1})$, where $\rho_0 = \bar{\rho} = h_{\mathcal{X} \rightarrow \Pi^*}(\bar{x})$ is the atomic proposition associated with the initial state of the system (π_1 in Fig. 2), and ρ_{L+1} is the *Goal* region (π_{28} in Fig. 2). For instance, for the scenario in Fig. 2, we obtain $\Pi^* = \{\pi_1, \dots, \pi_{30}\}$.

The set of constraints for the workspace can be captured by the following formula:

$$\xi \equiv (\rho_0 = \bar{\rho}) \wedge (\rho_{L+1} = Goal) \wedge \bigwedge_{t=1}^{L+1} \rho_t \in \mathcal{N}(\rho_{t-1}),$$

where $\mathcal{N}(\rho_i) = \{\pi_j \in \Pi^* \mid Adj^*(\pi_i, \pi_j) = 1\}$ denotes the set of regions that are adjacent (neighbors) to ρ_i . The above formula enforces that the trajectory starts with the initial region, associated with $\bar{\rho}$, and proceeds to the *Goal* region while only visiting regions that are adjacent. We observe that obstacle avoidance is implicitly encoded by the fact that Π^* and Adj^* are defined only over the regions of interest and the free space. To support generic LTL specifications, DIS-PLAN can use the bounded model checking encoding technique for LTL model checking [10] to generate high-level plans that satisfy the specifications. The Boolean formula ξ can then be solved using a SAT solver to generate a model ρ .

V. GENERATION OF THE CONTINUOUS TRAJECTORY

Given a region trajectory ρ specifying $(L + 1)$ polyhedra to be visited by the robot, the continuous planner CON-PLAN simultaneously performs two tasks: (1) it checks whether ρ is feasible given the constraints on the robot dynamics (CON-PLAN.CHECK); if ρ is not feasible, it generates counterexample formulas describing the “minimal” set of inconsistent constraints (CON-PLAN.COUNTEREXAMPLE). To detail these tasks, we first consider the following definition.

Definition 5.1: Let B' be a matrix chosen such that the map $[B \ B']$ is surjective. Let \bar{s} be defined as the least upper bound on the value of the slack variable $s = s^u + s^v$ such that the following constraints:

$$\begin{aligned} x &= Ax' + Bu + B'v \\ h_{\mathcal{X} \rightarrow \mathcal{W}}(x) &\in \rho^{\mathcal{W}} & h_{\mathcal{X} \rightarrow \mathcal{W}}(x') &\in \rho^{\mathcal{W}'} \\ \|u\| &\leq \bar{u} + s^u & \|v\| &\leq s^v \\ 0 &\leq s^u & 0 &\leq s^v \end{aligned}$$

are feasible for any two adjacent regions $\mathcal{W}, \mathcal{W}'$ and for any two states $x \in \mathcal{W}$ and $x' \in \mathcal{W}'$.

The value of \bar{s} can be easily pre-computed offline for a given workspace and obstacle configuration, and a given discretization associated with the set Π^* of atomic propositions. Then, for a constant tolerance $\epsilon \in \mathbb{R}_{>0}$ and the same choice of B' we define the following problem:

Problem 5.2:

$$\begin{aligned} \min_{\substack{u_0, \dots, u_L \in \mathbb{R}^m \\ v_0, \dots, v_L \in \mathbb{R}^m \\ s_0^u, \dots, s_L^u \in \mathbb{R} \\ s_0^v, \dots, s_L^v \in \mathbb{R} \\ x_1, \dots, x_{L+1} \in \mathbb{R}^n}} & \sum_{i=0}^L s_i^u + s_i^v \\ \text{subject to} & \\ x_0 &= \bar{x}, \\ h_{\mathcal{X} \rightarrow \mathcal{W}}(x_i) &\in \rho_i^{\mathcal{W}}, & i &= 1, \dots, L + 1 \\ x_{i+1} &= Ax_i + Bu_i + B'v_i & i &= 0, \dots, L \\ \|u_i\| &\leq \bar{u} + s_i^u, & i &= 0, \dots, L \\ \|v_i\| &\leq s_i^v, & i &= 0, \dots, L \\ 0 &\leq s_i^u, \quad 0 \leq s_i^v & i &= 0, \dots, L \\ \frac{\bar{s}}{\epsilon} \left(\sum_{k=0}^{i-1} s_k^u + s_k^v \right) &\leq s_i^u + s_i^v & i &= 1, \dots, L \end{aligned}$$

Because the region $\rho_i^{\mathcal{W}}$, for all i , is a polyhedron, Problem 5.2 is a linear program that can be efficiently solved. If the condition $\sum_{i=0}^L (s_i^u + s_i^v) \leq \epsilon$ is satisfied, then the high-level plan is feasible and the state and input trajectories generated by Problem 5.2 are valid trajectories. If, instead, $\sum_{i=0}^L (s_i^u + s_i^v) \leq \epsilon$ is not satisfied, then we obtain a *counterexample*, i.e., an infeasible trajectory that can be used to augment the original SAT encoding with additional Boolean constraints that forbid the current assignment. However, as anticipated in Sec. III, we aim, instead, at generating succinct infeasibility proofs that can precisely detect the origin of inconsistency and rule out a broader class of assignments to Boolean variables. It is possible to generate a more compact formula by detecting the earliest possible occurrence of an infeasible transition between two regions in ρ . The following theorem guarantees that this infeasibility proof can be retrieved from the information in the slack variables s_i^u, s_i^v in Problem 5.2.

Theorem 5.3: Let $\bar{s} \in \mathbb{R}_{>0}$, as in Definition 5.1, and $\epsilon \in \mathbb{R}_{>0}$ satisfy $\bar{s} \geq \epsilon$. Define the function $\text{ZEROPREFIX}_\epsilon : \mathbb{R}_{\geq 0}^{L+1} \rightarrow \mathbb{N}$ as:

$$\text{ZEROPREFIX}_\epsilon(s_0, s_1, \dots, s_L) = \min k \text{ s.t. } \sum_{i=0}^k s_i > \epsilon.$$

Then, an optimal solution of Problem 5.2 is also an optimal solution of the following optimization problem:

$$\begin{aligned} \max_{\substack{u_0, \dots, u_L \in \mathbb{R}^m \\ v_0, \dots, v_L \in \mathbb{R}^m \\ s_0^u, \dots, s_L^u \in \mathbb{R} \\ s_0^v, \dots, s_L^v \in \mathbb{R} \\ x_1, \dots, x_{L+1} \in \mathbb{R}^n}} & \text{ZEROPREFIX}_\epsilon((s_0^u + s_0^v), \dots, (s_L^u + s_L^v)) \\ \text{subject to} & \\ x_0 &= \bar{x}, \\ h_{\mathcal{X} \rightarrow \mathcal{W}}(x_i) &\in \rho_i^{\mathcal{W}}, & i &= 1, \dots, L + 1 \\ x_{i+1} &= Ax_i + Bu_i + B'v_i & i &= 0, \dots, L \\ \|u_i\| &\leq \bar{u} + s_i^u, & i &= 0, \dots, L \end{aligned}$$

Algorithm 2 ($\text{CSTATUS}, x, u, \phi_{\text{ce}} = \text{CON-PLAN}(\rho)$)

```

1: Solve Problem 5.2;
2:  $\forall i : s_i^* = s_i^u + s_i^v$ ;
3: if  $\sum_{i=0}^L s_i^* = 0$  then
4:    $\text{CSTATUS} = \text{feasible}$ ;
5:   return ( $\text{CSTATUS}, x^*, u^*, 1$ )
6: else
7:    $\text{CSTATUS} = \text{infeasible}$ ;
8:   Let  $k^* := \text{ZEROPREFIX}_\epsilon(s_0^* \dots s_{L+1}^*)$ ;
9:    $\phi_{\text{ce}} := \bigvee_{i=0}^{k^*+1} \neg \rho_i$ ;
10:  return ( $\text{CSTATUS}, x^*, u^*, \phi_{\text{ce}}$ );

```

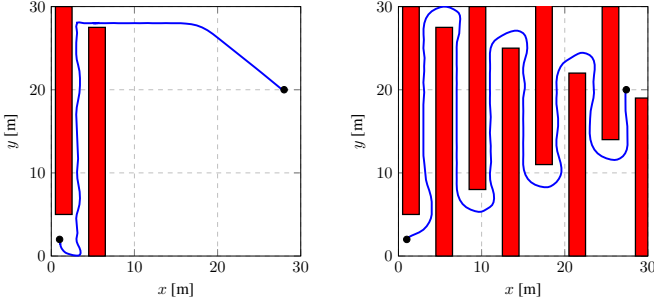


Fig. 3. Trajectories generated by the SMT-based motion planner for a maze-like workspace with different numbers of passages.

$$\begin{aligned} \|v_i\| &\leq s_i^v, & i = 0, \dots, L \\ 0 &\leq s_i^u, \quad 0 \leq s_i^v & i = 0, \dots, L \end{aligned}$$

Intuitively, for small ϵ , $\text{ZEROPREFIX}_\epsilon$ returns the number of zero elements at the beginning of a sequence $s = s_0, \dots, s_L$, i.e., the length of its “zero prefix.” Using this function, we can then look for sequences of slack variables that maximize the number of initial elements set to zero and introduce nonzero elements only when necessary. Generating a compact formula then amounts to finding the earliest occurrence of a nonzero slack, i.e., the earliest occurrence of an infeasible transition between two regions. The counterexample takes the following form:

$$\phi_{\text{ce}} := \bigvee_{i=0}^{k^*+1} \neg \rho_i, \quad (\text{V.1})$$

where k^* is equal to the ZEROPREFIX of the slack variables sequence generated by solving Problem 5.2. This result is summarized in Algorithm 2. By leveraging Theorem 5.3, we can state the following guarantees of Algorithm 1.

Theorem 5.4: The SMT-based motion planning Algorithm 1 generates a valid trajectory (x, ρ) for the reach-avoid motion planning problem $\mathcal{P} = \langle \mathcal{W}, \Pi, \text{Adj}, (A, B), \bar{x}, \bar{u} \rangle$.

VI. RESULTS

We developed our theory solver in MATLAB and interfaced it with the SAT solver SAT4J [28], to generate the discrete plans, and CPLEX, to solve the LPs and generate the counterexamples. All the experiments were executed on an Intel Core i7 3.4-GHz processor with 8 GB of memory.

TABLE I

COMPARISON OF THE RUN TIME PERFORMANCE OF ALGORITHM 1 WITH RESPECT TO THE RRT ALGORITHM [29] AND THE LTL OPT TOOLBOX [13] FOR A MAZE-LIKE WORKSPACE (E.G., SEE FIG. 3).

Number of passages	SMT-Based Motion Planner [s]			RRT [s]	LTL OPT [s]
	Discrete abstraction	DIS-PLAN	CON-PLAN		
1	1.9975	0.1360	0.2542	10.4381	> 7200
2	7.1461	1.1290	0.9294	122.3017	time out
3	19.3267	3.6495	1.0053	423.6957	time out
4	43.0985	4.0913	1.9204	1002.4193	time out

A. Case Study 1: Dubin’s Vehicle

We demonstrate the effectiveness of our motion planning algorithm by applying it to a reach-avoid problem for a Dubin’s vehicle (also known as differential drive robot). The kinematics of this robot can be transformed into a linear chain of integrators using dynamic feedback linearization. A discrete-time linear model is then computed from the feedback linearized model. As shown in Table I, we consider a $30\text{m} \times 30\text{m}$ maze-like workspace with increasing number of passages, ranging from one to four. Since automata theoretic approaches to control synthesis are known to be subject to state explosion, we directly compare the performance of our SMT-based motion planner against the rapidly exploring random tree (RRT) algorithm with dynamics and input constraints [29] and the LTL OPT toolbox [13], which implements a one-shot MILP encoding of the specification. LTL OPT is configured to use CPLEX as in our tool. We run each experiment 10 times and report the average execution time of each of the three algorithms. Fig. 3 shows some of the generated trajectories.

As shown in Table I, Algorithm 1 scales better than other techniques. We observe that most of the execution time is spent in the generation of the workspace discretization. Solving multiple instances of SAT and LPs is efficiently performed thanks to the proposed solver architecture and its underlying abstractions. On our benchmarks, Algorithm 1 is at least an order of magnitude faster than the RRT method, which is known to significantly slow down in the presence of narrow passages. Specifically, in the maze with one passage, RRT explored 58397 samples to find a feasible trajectory, i.e., two orders of magnitude more samples than the length of the final trajectory. On the contrary, each passage is compactly represented just as a polyhedron in our setup.

The execution time of LTL OPT exceeded the time-out threshold (4 hours) in all the cases except for the 1-passage maze. The degradation in performance may be due to the very large number of variables (several thousands) of the resulting MILP, the number of variables depending on the length of the trajectory. Thanks to the separation between real-valued and Boolean variable reasoning, Algorithm 1 requires, instead, solving a set of very efficient LPs.

B. Case Study 2: Scalability Results

Curse of dimensionality is known to be a major concern when applying formal methods to control synthesis and robotic applications. In this case study, we assess the effectiveness of the algorithms introduced in Sec. III-V in terms of scalability.

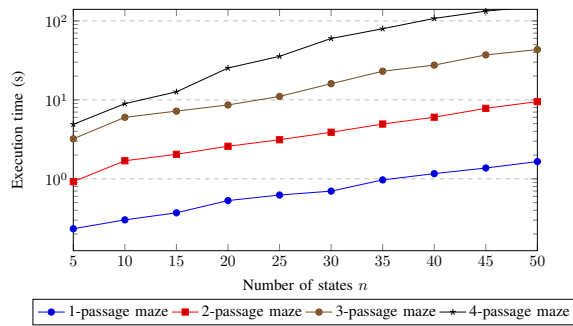


Fig. 4. Execution time of the motion planner for different maze-like workspace configurations as a function of the number of continuous states.

We consider again the maze-like workspace with increasing number of passages as in Sec. VI-A, and we investigate the execution time as we increase the number of states n . For each test case, we randomly generate the matrices A and B and average the computation time over 10 runs of the same experiment. Albeit not offering a statistical significant sample, our results are representative of the several simulations performed while testing our solver.

Fig. 4 shows, on a logarithmic scale, the execution time of the proposed motion planner as n increases. In all tests, we report the cumulative time due to both the discrete and continuous planners. The time to compute the discrete abstractions is equal to the one reported in Tab. I. Thanks to the proposed architecture, the dimension of the continuous state space only affects the number of variables of the linear programs in Sec. V. Even for systems with up to 50 state variables, the execution time ranges from 1.64 s (maze with one passage) to 152 s (maze with 4 passages), showing the potential of our approach to be deployed on complex robotic systems.

VII. CONCLUSIONS

We presented a scalable algorithm for reach-avoid robot motion planning under the assumption of discrete-time, linear dynamics and workspaces described by unions of polyhedra. The proposed algorithm scales well for systems with up to 50 continuous states. Future work includes extending the proposed techniques to motion planning from generic LTL specifications, multi-robot motion planning, and planning in the presence of uncertainties in the dynamics and bounded disturbances.

REFERENCES

- [1] A. Pnueli, "The temporal logic of programs," in *FOCS*, 1977, pp. 46–57.
- [2] P. Tabuada and G. J. Pappas, "Linear time logic control of discrete-time linear systems," *IEEE Trans. Automatic Control*, vol. 51, no. 12, pp. 1862–1877, 2006.
- [3] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Trans. Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [4] G. E. Fainekos, A. Girard, H. Kress-Gazit, and G. J. Pappas, "Temporal logic motion planning for dynamic robots," *Automatica*, vol. 45, no. 2, pp. 343–352, 2009.
- [5] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [6] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning," *IEEE Trans. Automatic Control*, vol. 57, no. 11, pp. 2817–2830, 2012.

- [7] P. Nuzzo, H. Xu, N. Ozay, J. Finn, A. Sangiovanni-Vincentelli, R. Murray, A. Donze, and S. Seshia, "A contract-based methodology for aircraft electric power system design," *IEEE Access*, vol. 2, pp. 1–25, 2014.
- [8] I. Saha, R. Ramaithitima, V. Kumar, G. J. Pappas, and S. A. Seshia, "Automated composition of motion primitives for multi-robot systems from safe LTL specifications," in *Int. Conf. Intelligent Robots and Systems*, 2014, pp. 1525–1532.
- [9] E. Plaku and S. Karaman, "Motion planning with temporal-logic specifications: Progress and challenges," *AI Communications*, no. Preprint, pp. 1–12.
- [10] A. Biere, K. Heljanko, T. Junttila, T. Iatvala, and V. Schuppan, "Linear encoding of bounded LTL model checking," *Logical Methods in Computer Science*, vol. 2, no. 5:5, pp. 1–64, 2006.
- [11] M. Rungger, M. Mazo Jr, and P. Tabuada, "Specification-guided controller synthesis for linear systems and safe linear-time temporal logic," in *Proc. Int. Conf. Hybrid Systems: Computation and Control*. ACM, 2013, pp. 333–342.
- [12] S. Karaman and E. Frazzoli, "Linear temporal logic vehicle routing with applications to multi-UAV mission planning," *Int. J. Robust and Nonlinear Control*, vol. 21, no. 12, pp. 1372–1395, 2011.
- [13] E. M. Wolff, U. Topcu, and R. M. Murray, "Optimization-based trajectory generation with linear temporal logic specifications," in *IEEE Int. Conf. Robotics and Automation*. IEEE, 2014, pp. 5319–5325.
- [14] S. Karaman and E. Frazzoli, "Sampling-based motion planning with deterministic μ -calculus specifications," in *Proc. IEEE Conf. Decision and Control*. IEEE, 2009, pp. 2222–2229.
- [15] A. Bhatia, L. E. Kavraki, and M. Y. Vardi, "Sampling-based motion planning with temporal goals," in *IEEE Int. Conf. Robotics and Automation*. IEEE, 2010, pp. 2689–2696.
- [16] L. Zhang and D. Manocha, "An efficient retraction-based rrt planner," in *IEEE Int. Conf. Robotics and Automation*. IEEE, 2008, pp. 3743–3750.
- [17] C. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli, *Satisfiability Modulo Theories, Chapter in Handbook of Satisfiability*. IOS Press, 2009.
- [18] J. Ding, E. Li, H. Huang, and C. J. Tomlin, "Reachability-based synthesis of feedback policies for motion planning under bounded disturbances," in *Int. Conf. Robotics and Automation*. IEEE, 2011, pp. 2160–2165.
- [19] K. Margellos and J. Lygeros, "Hamilton-jacobi formulation for reach-avoid differential games," *IEEE Trans. Automatic Control*, vol. 56, no. 8, pp. 1849–1861, 2011.
- [20] Z. Zhou, R. Takei, H. Huang, and C. J. Tomlin, "A general, open-loop formulation for reach-avoid games," in *Proc. Conf. Decision and Control*. IEEE, 2012, pp. 6501–6506.
- [21] T. A. Henzinger, R. Jhala, and R. Majumdar, *Proc. Int. Colloquium on Automata, Languages and Programming*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2003, ch. Counterexample-Guided Control, pp. 886–902.
- [22] Y. Shoukry, I. Nuzzo, A. Saha, A. L. Sangiovanni-Vincentelli, S. A. Seshia, G. J. Pappas, and P. Tabuada, "Scalable Motion Planning From LTL Specifications Using Satisfiability Modulo Convex Optimization," *ArXiv e-prints*, 2016.
- [23] P. Nuzzo, A. Puggelli, S. A. Seshia, and A. Sangiovanni-Vincentelli, "CaICS: SMT solving for non-linear convex constraints," in *Int. Conf. Formal Methods in Computer-Aided Design*, Oct 2010, pp. 71–79.
- [24] Y. Shoukry, P. Nuzzo, A. Puggelli, A. L. Sangiovanni-Vincentelli, S. A. Seshia, and P. Tabuada, "Secure State Estimation Under Sensor Attacks: A Satisfiability Modulo Theory Approach," *ArXiv e-prints*, Dec. 2014, [online] <http://adsabs.harvard.edu/abs/2014arXiv1412.4324S>.
- [25] Y. Shoukry, P. Nuzzo, A. Puggelli, A. L. Sangiovanni-Vincentelli, S. A. Seshia, M. Srivastava, and P. Tabuada, "IMHOTEP-SMT: A Satisfiability Modulo Theory Solver for Secure State Estimation," in *Proc. Int. Workshop on Satisfiability Modulo Theories*, July 2015.
- [26] Y. Shoukry, P. Nuzzo, N. Bezzo, A. L. Sangiovanni-Vincentelli, S. A. Seshia, and P. Tabuada, "Secure state reconstruction in differentially flat systems under sensor attacks using satisfiability modulo theory solving," in *Conf. Decision and Control*, 2015, pp. 3804–3809.
- [27] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, and G. J. Pappas, "Symbolic planning and control of robot motion [grand challenges of robotics]," *IEEE Robotics & Automation Magazine*, vol. 14, no. 1, pp. 61–70, 2007.
- [28] D. L. Berre and A. Parrain, "The Sat4j library, release 2.2," *Journal on Satisfiability, Boolean Modeling and Computation*, vol. 7, pp. 59–64, 2010.
- [29] S. M. LaValle and J. J. Kuffner, "Randomized kinodynamic planning," *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.