

Bézier Reachable Polytopes: Efficient Certificates for Robust Motion Planning with Layered Architectures

Noel Csomay-Shanklin, and Aaron D. Ames

Abstract—Control architectures are often implemented in a layered fashion, combining independently designed blocks to achieve complex tasks. Providing guarantees for such hierarchical frameworks requires considering the capabilities and limitations of each layer and their interconnections at design time. To address this holistic design challenge, we introduce the notion of *Bézier Reachable Polytopes* – certificates of reachable points in the space of Bézier polynomial reference trajectories. This approach captures the set of trajectories that can be tracked by a low-level controller while satisfying state and input constraints, and leverages the geometric properties of Bézier polynomials to maintain an efficient polytopic representation. As a result, these certificates serve as a constructive tool for layered architectures, enabling long-horizon tasks to be reasoned about in a computationally tractable manner.

I. INTRODUCTION

Modern control systems overwhelmingly employ layered architectures, wherein independent blocks are combined to achieve complex behaviors [1]. Typically, each block is designed in isolation and their interconnections are established in an *ad-hoc* manner. While this separation enables tractable controller design, achieving joint feasibility between layers is non-trivial. To create safe and reliable autonomous systems, we need a cohesive theory that considers not only the individual behavior of each block, but also how their interaction effects overall performance and constraint satisfaction. In this work, we focus on advancing such a theory for layered architectures that include a trajectory generator (*planner*) and a feedback controller (*tracker*). Specifically, we leverage the geometric properties of Bézier polynomials to construct a certificate which enables the connection of such a planner-tracker setup with a high-level decision making layer while maintaining feasibility guarantees, as seen in Figure 1.

The planner-tracker paradigm is extremely common in robotic systems [2], [3], and has theoretical roots in hierarchically consistent control [4], approximate simulation relations [5], and bisimulation [6]. In such a framework, the planner ensures feasibility by adjusting the trajectories it generates based on a tracking certificate, i.e. a representation of what the tracking controller can reasonably accomplish. This concept of layers communicating through achievable performance metrics serves as the foundation for robust motion planning [7]. For linear systems, such

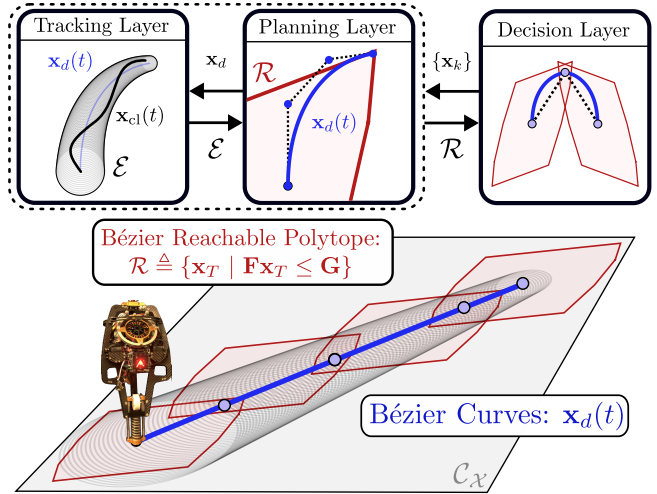


Fig. 1. A depiction of the layered architectures investigated in this work, where the reachable set of the combined planning and tracking layers can be represented via a linear inequality in the space of Bézier polynomials.

tracking certificates can be synthesized directly [8]. For nonlinear systems, generating tracking certificates is a more challenging task, and remains an active area of research. One option leverages Hamilton Jacobi reachability analysis to produce tracking upper bounds [9]. Alternatively, the linearization of the nonlinear system can be used to get approximate polytopic reachable sets [10]. Depending on the existing system structure, notions of Input to State Stability [11] can also be used to constructively produce tracking certificates for nonlinear control systems [12].

To extend the notion of guaranteed feasibility to a decision making layer, we require a certificate for the combined planner-tracker model, i.e. a representation of which states can be reached while satisfying state and input constraints. For discrete systems, this is often done by planning sequences of discrete actions based on motion primitives [13], [14]. Extending this to arbitrary continuous behaviors often requires solving two point boundary value problems [15], which can be computationally expensive. Importantly, however, the structure of the planning system can generally be imposed as a part of the design process. We leverage this design degree of freedom by enforcing the planner to generate Bézier polynomials; in doing so, we are able to ensure a computationally efficient reachable set representation.

This paper presents a theory for layered architectures that rely on Bézier curves, which have become increasingly popular for motion planning [12], [16], [17]. We take these ideas

Authors are with the Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA, Corresponding author: noelcs@caltech.edu.

This research is supported by Technology Innovation Institute (TII).

further by proving that if the planning layer is parameterized via Bézier polynomials, then the space of points which can be reached within a time interval is parameterizable via a polytope—the *Bézier Reachable Polytopes*. We show that these polytopes serve as performance certificates which enable holistic constraint satisfaction guarantees for layered architectures which utilize planning and tracking layers. We demonstrate the use of Bézier reachable polytopes in the context of completing long-horizon tasks on a simulated pendulum and experimentally on a physical 3D hopping robot with tight state and input constraints.

II. BACKGROUND

Consider the following nonlinear control system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \quad (1)$$

with state $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^N$, input $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}^M$, and whose dynamics $\mathbf{f} : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^N$ are assumed to be continuously differentiable in their arguments. The system (1) will be represented as the tuple $\Sigma = \{\mathcal{X}, \mathcal{U}, \mathbf{f}\}$. Due to the potential complexity of the dynamics \mathbf{f} , directly synthesizing control actions for challenging tasks may be intractable. To address this, control engineers often rely on *planning models*, which serve as template systems that enable desired behaviors to be constructed in a computationally tractable way. These models are defined as:

Definition 1. A system $\Sigma_d = \{\mathcal{X}_d, \mathcal{U}_d, \mathbf{f}_d\}$ is said to be a *planning model* for a system Σ if there exists a surjective mapping $\mathbf{\Pi} : \mathcal{X} \rightarrow \mathcal{X}_d$ and a right inverse $\mathbf{\Psi} : \mathcal{X}_d \rightarrow \mathcal{X}$ such that $\mathbf{\Pi} \circ \mathbf{\Psi} = \text{id}_{\mathcal{X}_d}$.

As the dimensionality of \mathcal{X}_d is typically much smaller than \mathcal{X} , there are many possible inverse mappings $\mathbf{\Psi}$, each of which induce an embedding of the reduced state space \mathcal{X}_d into the full state space \mathcal{X} . To link a full-order system with a planning model, we must define a feedback controller $\mathbf{k} : \mathcal{X} \times \mathcal{X}_d \times \mathcal{U}_d \rightarrow \mathcal{U}$ which aims to track the states of the planning model. This controller results in the following closed-loop system:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{k}(\mathbf{x}, \mathbf{x}_d, \mathbf{u}_d)) \triangleq \mathbf{f}_{\text{cl}}(\mathbf{x}, \mathbf{x}_d, \mathbf{u}_d), \quad (2)$$

which, given any initial condition $\mathbf{x}_0 \in \mathcal{X}$, has continuously differentiable solution $\mathbf{x}_{\text{cl}} : I \rightarrow \mathcal{X}$ over some interval $I \subset \mathbb{R}_{\geq 0}$ defined as:

$$\mathbf{x}_{\text{cl}}(t) \triangleq \mathbf{x}_0 + \int_0^t \mathbf{f}_{\text{cl}}(\mathbf{x}_{\text{cl}}(\tau), \mathbf{x}_d(\tau), \mathbf{u}_d(\tau)) d\tau.$$

A key desired property of this controller is its ability to maintain bounded tracking error:

Definition 2. Let Σ_d be a planning model for system Σ . Given a desired trajectory $\mathbf{x}_d(\cdot)$, a set-valued function $\mathcal{E} : \mathcal{U}_d \rightarrow \mathcal{P}(\mathcal{X})$ is a *tracking certificate* for the system Σ if:

$$\mathbf{x}_{\text{cl}}(t) \in \mathbf{\Psi}(\mathbf{x}_d(t)) \oplus \mathcal{E}(\mathbf{u}_d(t)),$$

where \oplus denotes the Minkowski sum.

Example 1. Let Σ represent the closed-loop system of a 3D hopping robot tracking a center of mass velocity command with whole-body model predictive control (MPC). In this scenario, the planning system Σ_d is that of a single integrator and the mapping $\mathbf{\Pi}$ projects the full state space of the hopper into the center of mass planar positions. The function \mathbf{k} and mapping $\mathbf{\Psi}$ define the process of MPC, which takes in desired velocity trajectories and produces joint-space trajectories which can be tracked with bounded error via PD control as a function of how much input the planning system applies.

Given a planning model Σ_d , we will be interested in characterizing the space of all desired trajectories for the system Σ which satisfy the following problem:

Problem 1. Consider a compact state constraint set $\mathcal{C}_{\mathcal{X}} \subset \mathcal{X}_d$ and compact input constraint set $\mathcal{C}_{\mathcal{U}} \subset \mathcal{U}$. Produce trajectories $\mathbf{x}_d(\cdot)$ which when tracked achieve the following:

- $\mathbf{\Pi}(\mathbf{x}_{\text{cl}}(t)) \in \mathcal{C}_{\mathcal{X}}$ for all $t \in I$,
- $\mathbf{k}(\mathbf{x}_{\text{cl}}(t), \mathbf{x}_d(t), \mathbf{u}_d(t)) \in \mathcal{C}_{\mathcal{U}}$ for all $t \in I$.

We will go about solving this problem by appropriately constraining the space of trajectories $\mathbf{x}_d(\cdot)$, wherein $\mathbf{x}_d(\cdot)$ will be a design parameter. Although planning models can have any system structure (and are useful as long as there exist an appropriate mapping $\mathbf{\Psi}$ and controller \mathbf{k}), in order to make constructive guarantees we make further assumptions about the planning dynamics. Specifically, consider a nonlinear planning model system with coordinates $\mathbf{q}_d \in \mathbb{R}^m$, state $\mathbf{x}_d = [\mathbf{q}_d^\top, \dot{\mathbf{q}}_d^\top, \dots, \mathbf{q}_d^{(\gamma-1)\top}]^\top \in \mathbb{R}^n$ for some $\gamma \in \mathbb{N}$, and control-affine dynamics of the form:

$$\dot{\mathbf{x}}_d = \begin{bmatrix} \mathbf{0} & \mathbf{I}_{n-m} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{x}_d + \begin{bmatrix} \mathbf{0} \\ \mathbf{f}_d(\mathbf{x}_d) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{g}_d(\mathbf{x}_d) \end{bmatrix} \mathbf{u}_d, \quad (3)$$

where \mathbf{I}_{n-m} is an identity matrix of size $n - m$, the $\mathbf{0}$ matrices are appropriately sized, $\mathbf{u}_d \in \mathbb{R}^m$ is the input, and the drift vector $\mathbf{f}_d : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and actuation matrix $\mathbf{g}_d : \mathbb{R}^n \rightarrow \mathbb{R}^{m \times m}$ are assumed to be locally Lipschitz continuous on \mathbb{R}^n . We define a dynamically feasible trajectory for such a system as:

Definition 3 (Dynamically Feasible Trajectory). Given a time interval $I \triangleq [0, T]$ for $T \in \mathbb{R}_{\geq 0}$, a piecewise continuously differentiable function $\mathbf{x}_d : I \rightarrow \mathbb{R}^n$ is a *dynamically feasible trajectory* for Σ_d if there is a piecewise continuous function $\mathbf{u}_d : I \rightarrow \mathbb{R}^m$ such that:

$$\dot{\mathbf{x}}_d(t) = \mathbf{f}_d(\mathbf{x}_d(t)) + \mathbf{g}_d(\mathbf{x}_d(t))\mathbf{u}_d(t), \quad (4)$$

for almost all $t \in I$.

In order to design dynamically feasible trajectories $\mathbf{x}_d(\cdot)$ for Σ_d , we must reason about integral curves of the planning model dynamics. To parameterize dynamically feasible trajectories of (3) via Bézier curves, we assume that the planning system Σ_d is fully actuated:

Assumption 1. We have that $\mathbf{f}_d(\mathbf{0}) = \mathbf{0}$ and the matrix $\mathbf{g}_d(\mathbf{x}_d)$ is invertible for all $\mathbf{x}_d \in \mathcal{X}_d$.

III. BÉZIER CURVES

A curve $\mathbf{b} : I \triangleq [0, T] \rightarrow \mathbb{R}^m$ for $T > 0$ is said to be a Bézier curve [18] of order $p \in \mathbb{N}$ if it is of the form:

$$\mathbf{b}(t) = \mathbf{p}\mathbf{z}(t),$$

where $\mathbf{z} : I \rightarrow \mathbb{R}^{p+1}$ is a Bernstein basis polynomial of degree p and $\mathbf{p} \in \mathbb{R}^{m \times p+1}$ are a collection of $p+1$ control points of dimension m . There exists a matrix $\mathbf{H} \in \mathbb{R}^{p+1 \times p+1}$ (as in [12]) which defines a linear relationship between control points of a curve \mathbf{b} and its derivative via:

$$\dot{\mathbf{b}}(t) = \mathbf{p}\mathbf{H}\mathbf{z}(t).$$

This enables us to define a state space curve $\mathbf{B} : I \rightarrow \mathbb{R}^n$:

$$\mathbf{B}(t) \triangleq \begin{bmatrix} \mathbf{b}(t) \\ \vdots \\ \mathbf{b}^{(\gamma-1)}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} \mathbf{p} \\ \vdots \\ \mathbf{p}\mathbf{H}^{\gamma-1} \end{bmatrix}}_{\triangleq \mathbf{P}} \mathbf{z}(t). \quad (5)$$

The columns of the matrix $\mathbf{P} \in \mathbb{R}^{n \times p+1}$, denoted as \mathbf{P}_j for $j = 0, \dots, p$, represent the collection of n dimensional control points of the Bézier curve \mathbf{B} in the state space. Furthermore, if we take $\mathbf{x}_d(\cdot) \equiv \mathbf{B}(\cdot)$ to represent a desired trajectory of Bézier curves, we observe that:

$$\dot{\mathbf{x}}_d = \begin{bmatrix} \mathbf{0} & \mathbf{I}_{n-m} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{x}_d + \begin{bmatrix} \mathbf{0} \\ \mathbf{f}_d(\mathbf{x}_d) \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{g}_d(\mathbf{x}_d) \end{bmatrix} \mathbf{u}_d,$$

for the continuous input signal:

$$\mathbf{u}_d = \mathbf{g}_d(\mathbf{x}_d)^{-1} \left(\mathbf{q}_d^{(\gamma)} - \mathbf{f}_d(\mathbf{x}_d) \right). \quad (6)$$

Therefore, any Bézier curve $\mathbf{B}(\cdot)$ constructed via (5) is a dynamically feasible trajectory for our planning model. As such, we can leverage Bézier curves towards the design of trajectories $\mathbf{x}_d(\cdot)$ satisfying Problem 1. Bézier curves enjoy a number of desirable properties:

Property 1 (Convex Hull [18]).

$$\mathbf{B}(t) \in \text{conv}(\{\mathbf{P}_j\}), \quad j = 0, \dots, p, \quad \forall t \in I.$$

Property 2 (Linear Bounding). For a vector $\mathbf{d} \in \mathbb{R}^k$ and a matrix $\mathbf{C} \in \mathbb{R}^{k \times n}$, we have:

$$\mathbf{C}\mathbf{P}_j \leq \mathbf{d}, \quad j = 0, \dots, p \implies \mathbf{C}\mathbf{B}(t) \leq \mathbf{d}, \quad \forall t \in I.$$

Proof. The convex hull property of Bézier curves implies that for any $t \in I$ and any row $\mathbf{c} \in \mathbb{R}^n$ of \mathbf{C} with corresponding value $d \in \mathbb{R}$ of \mathbf{d} , we may write:

$$\mathbf{c}\mathbf{B}(t) = \sum_{j=0}^p \lambda_j(t) \mathbf{c}\mathbf{P}_j$$

for some $\lambda_j(t) \geq 0$ and $\sum_{j=0}^p \lambda_j(t) = 1$. Therefore,

$$\mathbf{c}\mathbf{B}(t) \leq \sum_{j=0}^p \lambda_j(t) \max_j \mathbf{c}\mathbf{P}_j = \max_j \mathbf{c}\mathbf{P}_j \leq \mathbf{d},$$

as each \mathbf{P}_j term satisfies $\mathbf{c}\mathbf{P}_j \leq \mathbf{d}$ by assumption. \square




Fig. 2. A visual guide to the properties of Bézier curves.

We will specifically be interested in producing Bézier curves that connect initial conditions $\mathbf{x}_d(0) \in \mathcal{X}_d$ and terminal conditions $\mathbf{x}_d(T) \in \mathcal{X}_d$ in a fixed time T . Given such boundary conditions, a Bézier curve $\mathbf{B}(\cdot)$ which connects them must satisfy the following set of equality constraints:

$$\mathbf{b}^{(k)}(0) = \mathbf{p}\mathbf{H}^k \mathbf{z}(0) = \mathbf{q}_d^{(k)}(0), \quad k = 0, \dots, \gamma - 1, \quad (7)$$

$$\mathbf{b}^{(k)}(T) = \mathbf{p}\mathbf{H}^k \mathbf{z}(T) = \mathbf{q}_d^{(k)}(T), \quad k = 0, \dots, \gamma - 1. \quad (8)$$

These constraints lead to the following Property:

Property 3 (Boundary Values). Given a time $T > 0$, two points $\mathbf{x}_0, \mathbf{x}_T \in \mathbb{R}^n$, and order $p \geq 2\gamma - 1$, there exists a matrix $\mathbf{D} \in \mathbb{R}^{p+1 \times 2n}$ such that any curve $\mathbf{x}_d(\cdot)$ with control points satisfying:

$$\mathbf{p}\mathbf{D} = [\mathbf{x}_0^\top \quad \mathbf{x}_T^\top] \quad (9)$$

also satisfies $\mathbf{x}_d(0) = \mathbf{x}_0$ and $\mathbf{x}_d(T) = \mathbf{x}_T$.

Proof. We begin by noting that $\mathbf{z}(0) = [1 \quad \mathbf{0}_{1 \times p}]^\top$ and $\mathbf{z}(T) = [\mathbf{0}_{1 \times p} \quad 1]^\top$. Then, collecting the constraints in (7) and (8) yields:

$$\begin{aligned} \mathbf{p} [\mathbf{H}_0^0 \quad \mathbf{H}_0^1 \quad \dots \quad \mathbf{H}_0^{\gamma-1}] &= \mathbf{x}_0, \\ \mathbf{p} [\mathbf{H}_p^0 \quad \mathbf{H}_p^1 \quad \dots \quad \mathbf{H}_p^{\gamma-1}] &= \mathbf{x}_T. \end{aligned}$$

where \mathbf{H}_j^i denotes the j^{th} column of the matrix \mathbf{H} raised to the i^{th} power. It can be algebraically verified that \mathbf{H} has the form:

$$\mathbf{H}_0^i = \underbrace{[\star \dots \star]_{i+1}}_{i+1} \underbrace{[0 \dots 0]_{p-i}}_{p-i}, \quad \mathbf{H}_p^i = \underbrace{[0 \dots 0]_{p-i}}_{p-i} \underbrace{[\star \dots \star]_{i+1}}_{i+1},$$

with nonzero entries \star . Taking $\mathbf{D} \in \mathbb{R}^{p+1 \times 2n}$ as:

$$\mathbf{D} \triangleq [\mathbf{H}_0^0 \quad \mathbf{H}_0^1 \quad \dots \quad \mathbf{H}_0^{\gamma-1} \quad \mathbf{H}_p^0 \quad \mathbf{H}_p^1 \quad \dots \quad \mathbf{H}_p^{\gamma-1}],$$

in the case that $p \geq 2\gamma - 1$ the columns are linearly independent and thus the matrix \mathbf{D} has full column rank, implying that a solution \mathbf{p} exists (but is not unique unless $p = 2\gamma - 1$). \square

Remark 1. In the case that $p > 2\gamma - 1$, the constraint (9) is under-determined and can be resolved via a least squares solution, allowing for additional cost terms to be optimized.

Finally, we present one additional property which will be useful in increasing the resolution of Bézier curves and reduce the conservatism of their upper bounds. To do this, we introduce the notion of a refinement of the interval I as:

Definition 4. A k -refinement of an interval $[0, T]$ is a collection of times $\{T_i\}$ for $i = 0, \dots, k$ and associated intervals $\{[T_{i-1}, T_i]\}$ with $T_{i-1} < T_i$, $T_0 = 0$, and $T_k = T$.

From this, we can split a Bézier polynomial $\mathbf{B}(\cdot)$ into a sequence of B-splines:

Property 4 (Splitting [18]). Given the control points \mathbf{P} of a Bézier polynomial defined over the interval I and a k -refinement of I , there exists a collection of matrices $\{\mathbf{Q}_i\}$ for $i = 1, \dots, k$ such that $\mathbf{B}_Q(t) = \mathbf{P}\mathbf{Q}_i\mathbf{z}(t)$ satisfies $\mathbf{B}_Q(t) \triangleq \mathbf{B}(T_i + \frac{t}{T}(T_{i+1} - T_i))$ for all $t \in I$.

Finally, it will be useful to operate with the (column-wise) vectorized versions of \mathbf{p} and \mathbf{P} , defined as $\vec{\mathbf{p}} \triangleq \text{vec}(\mathbf{p}) \in \mathbb{R}^{m(p+1)}$ and $\vec{\mathbf{P}} \triangleq \text{vec}(\mathbf{P}) \in \mathbb{R}^{n(p+1)}$. With these new representations, we have the following equivalences:

$$\begin{aligned} \vec{\mathbf{P}} &= \vec{\mathbf{H}}\vec{\mathbf{p}} \\ \vec{\mathbf{D}}\vec{\mathbf{p}} &= \text{vec}([\mathbf{x}_0^\top \quad \mathbf{x}_T^\top]) \end{aligned}$$

with $\vec{\mathbf{H}}$ and $\vec{\mathbf{D}}$ the vectorized versions of \mathbf{H} and \mathbf{D} , respectively. With these tools, we will next discuss how to enforce state and input constraints on Bézier curves via linear constraints imposed on the control points.

IV. STATE AND INPUT CONSTRAINT SATISFACTION

To begin, we make the following assumption about the constraint sets \mathcal{C}_X and \mathcal{C}_U :

Assumption 2. The state constraint set is described by $\mathcal{C}_X = \{\mathbf{x}_d \in \mathcal{X}_d \mid \mathbf{C}\mathbf{x}_d \leq \mathbf{d}\}$ with $\mathbf{C} \in \mathbb{R}^{k \times n}$ and $\mathbf{d} \in \mathbb{R}^k$. Furthermore, we have that the input constraint set $\mathcal{C}_U \triangleq \{\mathbf{u} \in \mathbb{R}^m \mid \|\mathbf{u}\|_\infty \leq u_{\max}\}$ for $u_{\max} \in \mathbb{R}_{>0}$, i.e., we have a box input constraint.

The following constructions can also be performed with a positive diagonal weighting matrix $\mathbf{W} \in \mathbb{S}_{>0}^m$ to scale the box constraint on \mathbf{u} , such that $\|\mathbf{W}\mathbf{u}\|_\infty \leq u_{\max}$. Such constraints are extremely common in robotic systems. From this point on, $\|\cdot\|$ will represent the ∞ -norm unless otherwise stated. Given a tracking certificate set, we can define its upper bound $e: \mathcal{U}_d \rightarrow \mathbb{R}_{\geq 0}$ as:

$$e(\mathbf{u}_d) \triangleq \sup_{\mathbf{e} \in \mathcal{E}(\mathbf{u}_d)} \|\mathbf{e}\|. \quad (10)$$

If \mathcal{E} is described as the zero sublevel set of a function that is differentiable with respect to \mathbf{u}_d , then $e(\mathbf{u}_d)$ is locally Lipschitz with respect to \mathbf{u}_d . Along with this, we assume Lipschitz properties of $\mathbf{\Pi}$ and $\mathbf{\Psi}$:

Assumption 3. The functions $\mathbf{\Pi}$, $\mathbf{\Psi}$, and e are Lipschitz continuous over the domain \mathcal{C}_X with constants $L_{\mathbf{\Pi}}$, $L_{\mathbf{\Psi}}$ and L_e , respectively.

The remainder of the section will be devoted to proving the following statement:

Theorem 1. Let system Σ_d be a planning model for system Σ with tracking certificate \mathcal{E} . There exist matrices \mathbf{F} and \mathbf{G} such that any Bézier curve $\mathbf{B}: I \rightarrow \mathcal{X}_d$ with control points \mathbf{p} satisfying:

$$\mathbf{F}\vec{\mathbf{p}} \leq \mathbf{G},$$

when tracked results in the closed loop system satisfying $\mathbf{\Pi}(\mathbf{x}_{cl}) \in \mathcal{C}_X$ and $\mathbf{k}(\mathbf{x}_{cl}, \mathbf{x}_d, \mathbf{u}_d) \in \mathcal{C}_U$ for all $t \in I$.

Towards this goal, we first show that satisfying input constraints of the tracker can be reformulated as a linear constraint on state and input norms:

Lemma 1. Given a reference points $\bar{\mathbf{x}}_d \in \mathcal{X}_d$, enforcing the constraint:

$$[L_{\mathbf{k}}(1 + L_{\mathbf{\Psi}}) \quad L_{\mathbf{k}}(1 + L_e)] \begin{bmatrix} \|\mathbf{x}_d - \bar{\mathbf{x}}_d\| \\ \|\mathbf{u}_d\| \end{bmatrix} \leq u_{\max} - K(\bar{\mathbf{x}}_d),$$

with $K(\bar{\mathbf{x}}_d) \triangleq \|\mathbf{k}(\mathbf{\Psi}(\bar{\mathbf{x}}_d), \bar{\mathbf{x}}_d, \mathbf{0})\| + e(\mathbf{0})$ results in input constraints being satisfied, i.e. $\mathbf{k}(\mathbf{x}_{cl}, \mathbf{x}_d, \mathbf{u}_d) \in \mathcal{C}_U$.

Proof. Observe that the input \mathbf{k} can be bounded by:

$$\begin{aligned} \|\mathbf{k}(\mathbf{x}, \mathbf{x}_d, \mathbf{u}_d)\| &\leq L_{\mathbf{k}}(\|\mathbf{x} - \mathbf{\Psi}(\mathbf{x}_d)\| + \|\mathbf{\Psi}(\mathbf{x}_d) - \mathbf{\Psi}(\bar{\mathbf{x}}_d)\| \\ &\quad + \|\mathbf{u}_d\| + \|\mathbf{x}_d - \bar{\mathbf{x}}_d\|) + \|\mathbf{k}(\mathbf{\Psi}(\bar{\mathbf{x}}_d), \bar{\mathbf{x}}_d, \mathbf{0})\| \\ &\leq L_{\mathbf{k}}(1 + L_e)\|\mathbf{u}_d\| + L_{\mathbf{k}}(1 + L_{\mathbf{\Psi}})\|\mathbf{x}_d - \bar{\mathbf{x}}_d\| \\ &\quad + \|\mathbf{k}(\mathbf{\Psi}(\bar{\mathbf{x}}_d), \bar{\mathbf{x}}_d, \mathbf{0})\| + e(\mathbf{0}). \end{aligned}$$

Rearranging terms yields the desired result. \square

We will show in Lemma 3 that this constraint can be reformulated as a linear inequality constraint on the Bézier curve $\mathbf{x}_d(\cdot)$. Note that in order for this set to be nonempty, we must have that $u_{\max} - e(\mathbf{0}) > 0$. This requirement is one of feasibility of the tracking controller – if not, then the feedback controller applied over the tracking certificate \mathcal{E} is larger than the set \mathcal{U} meaning regardless of desired trajectory there could exist a perturbed state which would violate the input constraints. If this is the case, then the error tracking bound of the low-level controller needs to be improved before proceeding. In order to make a similar claim for state constraints, we present the following claim:

Lemma 2. Enforcing the constraint:

$$[\mathbf{C} \quad L_{\mathbf{\Pi}}L_e\mathbf{K}] \begin{bmatrix} \mathbf{x}_d \\ \|\mathbf{u}_d\| \end{bmatrix} \leq \mathbf{d} - L_{\mathbf{\Pi}}e(\mathbf{0})\mathbf{K} \quad (11)$$

with $\mathbf{K} \triangleq \sqrt{\text{diag}(\mathbf{C}\mathbf{C}^\top)}$ results in state constraints being satisfied, i.e. $\mathbf{\Pi}(\mathbf{x}_{cl}(t)) \in \mathcal{C}_X$.

Proof. Recall that applying the controller \mathbf{k} yields:

$$\mathbf{\Pi}(\mathbf{x}_{cl}) \in \Omega \triangleq \left\{ \mathbf{\Pi}(\mathbf{\Psi}(\mathbf{x}_d) + \mathbf{v}) \mid \mathbf{v} \in \mathcal{E}(\mathbf{u}_d) \right\},$$

which holds from Definition 2. From (10), we continue:

$$\begin{aligned} \Omega &\subset \left\{ \mathbf{\Pi}(\mathbf{\Psi}(\mathbf{x}_d) + \mathbf{v}) \mid \|\mathbf{v}\| \leq e(\mathbf{u}_d) \right\} \\ &\equiv \left\{ \mathbf{x}_d + \mathbf{y} \mid \mathbf{y} = \mathbf{\Pi}(\mathbf{\Psi}(\mathbf{x}_d) + \mathbf{v}) - \mathbf{x}_d, \|\mathbf{v}\| \leq e(\mathbf{u}_d) \right\} \\ &\subset \left\{ \mathbf{x}_d + \mathbf{y} \mid \|\mathbf{y}\| \leq \|\mathbf{\Pi}(\mathbf{\Psi}(\mathbf{x}_d) + \mathbf{v}) - \mathbf{x}_d\|, \|\mathbf{v}\| \leq e(\mathbf{u}_d) \right\}. \end{aligned}$$

Next, recalling that $\Pi \circ \Psi(\mathbf{x}_d) = \mathbf{x}_d$, we have:

$$\begin{aligned} \Omega &\subset \left\{ \mathbf{x}_d + \mathbf{y} \mid \|\mathbf{y}\| \leq L_{\Pi} \|\mathbf{v}\|, \|\mathbf{v}\| \leq e(\mathbf{u}_d) \right\} \\ &\subset \left\{ \mathbf{x}_d + \mathbf{y} \mid \|\mathbf{y}\| \leq L_{\Pi} e(\mathbf{u}_d) \right\} \end{aligned}$$

Therefore, defining $\rho \triangleq L_{\Pi} e(\mathbf{u}_d)$, we have:

$$\Pi(\mathbf{x}_{cl}) \in \mathbf{x}_d(t) \oplus B_{\rho}(0),$$

As such, if we can ensure $\mathbf{C}(\mathbf{x}_d + \mathbf{v}) \leq \mathbf{d}$ for all $\mathbf{v} \in B_{\rho}(0)$, we would have the desired result. Appealing to Lemma 4 in [12], we know that this is satisfied if:

$$\mathbf{C}\mathbf{x}_d \leq \mathbf{d} - L_{\Pi} e(\mathbf{u}_d) \sqrt{\text{diag}(\mathbf{C}\mathbf{C}^{\top})},$$

which is in turn satisfied if:

$$\mathbf{C}\mathbf{x}_d \leq \mathbf{d} - \left(L_{\Pi} L_e \|\mathbf{u}_d\| + L_{\Pi} e(\mathbf{0}) \right) \sqrt{\text{diag}(\mathbf{C}\mathbf{C}^{\top})},$$

which can be rearranged to achieve the desired result. \square

Now, we state the following Lemma, which will allow us to reformulate these state and input constraints via linear constraints on the Bézier curve:

Lemma 3. *Given a reference point $\bar{\mathbf{x}}_d \in \mathcal{X}_d$, a matrix $\mathbf{A} \in \mathbb{R}^{k \times n+2}$ and vector $\mathbf{b} \in \mathbb{R}^k$, there exists a matrix $\mathbf{L} \in \mathbb{R}^{4knm \times n}$ and a vector $\mathbf{h} \in \mathbb{R}^{4knm}$ such that:*

$$\mathbf{L} \begin{bmatrix} \mathbf{x}_d \\ \mathbf{q}_d^{(\gamma)} \end{bmatrix} \leq \mathbf{h} \implies \mathbf{A} \begin{bmatrix} \mathbf{x}_d \\ \|\mathbf{x}_d - \bar{\mathbf{x}}_d\| \\ \|\mathbf{u}_d\| \end{bmatrix} \leq \mathbf{b}.$$

Proof. We begin by bounding the term $\mathbf{u}_d(\cdot)$:

$$\|\mathbf{u}_d\| \leq \|\mathbf{g}_d(\mathbf{x}_d)^{-1}\| \|\mathbf{q}_d^{(\gamma)} - \mathbf{f}_d(\mathbf{x}_d)\|. \quad (12)$$

Taking $\bar{\mathbf{x}}_d \in \mathcal{X}_d$ to be a reference point in the planning state space, we can bound the first term by:

$$\|\mathbf{g}^{-1}(\mathbf{x}_d)\| \leq L_G \|\mathbf{x}_d - \bar{\mathbf{x}}_d\| + \|\mathbf{g}^{-1}(\bar{\mathbf{x}}_d)\|, \quad (13)$$

where L_G is a Lipschitz constant of \mathbf{g}^{-1} with respect to the ∞ -norm on $\mathcal{C}_{\mathcal{X}}$, which is well defined by the local Lipschitz continuity and nonzero assumptions on \mathbf{g} and the compactness of $\mathcal{C}_{\mathcal{X}}$. Similarly:

$$\|\mathbf{q}_d^{(\gamma)} - \mathbf{f}(\mathbf{x}_d)\| \leq L_f \|\mathbf{x}_d - \bar{\mathbf{x}}_d\| + \|\mathbf{q}_d^{(\gamma)} - \mathbf{f}(\bar{\mathbf{x}}_d)\|. \quad (14)$$

Now, let $\mathbf{a} \triangleq [\mathbf{a}_1 \ a_2 \ a_3]$ be a row of the constraint matrix \mathbf{A} with $\mathbf{a}_1 \in \mathbb{R}^n$ and $a_2, a_3 \in \mathbb{R}$ and $b \in \mathbb{R}$ the corresponding entry of the vector \mathbf{b} . Substituting (13) and (14) into (12), we can construct a quadratic form:

$$[a_2 \ a_3] \begin{bmatrix} \|\mathbf{x}_d - \bar{\mathbf{x}}_d\| \\ \|\mathbf{u}_d\| \end{bmatrix} \leq \boldsymbol{\sigma}^{\top} \mathbf{M} \boldsymbol{\sigma}_d + \mathbf{N}^{\top} \boldsymbol{\sigma}_d,$$

where $\boldsymbol{\sigma}_d \triangleq \left[\|\mathbf{x}_d - \bar{\mathbf{x}}_d\| \ \|\mathbf{q}_d^{(\gamma)} - \mathbf{f}(\bar{\mathbf{x}}_d)\| \right]^{\top}$ and:

$$\mathbf{M} = \frac{a_3}{2} \begin{bmatrix} 2L_G L_f & L_G \\ L_G & 0 \end{bmatrix}, \quad \mathbf{N} = \begin{bmatrix} a_3 L_f \|\mathbf{g}^{-1}(\bar{\mathbf{x}}_d)\| + a_2 \\ a_2 \|\mathbf{g}^{-1}(\bar{\mathbf{x}}_d)\| \end{bmatrix}.$$

Next, consider $\widehat{\mathbf{M}}$ as the projection of \mathbf{M} onto the positive semidefinite cone. With this, we can define the function

$h : \mathcal{X}_d \times \mathbb{R}^m \rightarrow \mathbb{R}$ as:

$$h(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) = \boldsymbol{\sigma}_d^{\top} \widehat{\mathbf{M}} \boldsymbol{\sigma}_d + \mathbf{N}^{\top} \boldsymbol{\sigma}_d + \mathbf{a}_1^{\top} \mathbf{x}_d.$$

Because \mathbf{M} is symmetric, we have that $\widehat{\mathbf{M}} \preceq \mathbf{M}$. As such, points in the set $\Omega \triangleq \{(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \mid h(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \leq b\}$ satisfy the desired inequality. Next, consider a function $\ell : \mathcal{X}_d \times \mathbb{R}^m \rightarrow \mathbb{R}$ of the form:

$$\ell(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) = \mathbf{c}^{\top} \boldsymbol{\sigma}_d + \mathbf{a}_1^{\top} \mathbf{x}_d,$$

for some vector $\mathbf{c} \in \mathbb{R}^2$, along with the following optimization program:

$$\begin{aligned} \delta^* &= \sup_{\delta \in \mathbb{R}} \delta \\ \text{s.t.} \quad &\ell(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \leq \delta \implies h(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \leq b \end{aligned}$$

In general, this set containment problem may be challenging to solve; however, given the specific problem structure this can be solved for in closed form (the details of which can be found in [19]). Then, we have that the set $\Lambda \triangleq \{(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \mid \ell(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \leq \delta^*\} \subset \Omega$; therefore points in Λ satisfy the desired constraints.

Finally, we will show that there exists a matrix $\mathbf{L}_i \in \mathbb{R}^{4nm \times n+m}$ and a vector $\mathbf{h}_i \in \mathbb{R}^{4nm}$ such that:

$$\mathbf{L}_i \begin{bmatrix} \mathbf{x}_d \\ \mathbf{q}_d^{(\gamma)} \end{bmatrix} \leq \mathbf{h}_i \implies \ell(\mathbf{x}_d, \mathbf{q}_d^{(\gamma)}) \leq \delta^*.$$

Based on the definition of $\boldsymbol{\sigma}_d$, the set Λ is given by:

$$\mathbf{c}^{\top} \begin{bmatrix} \max_i |\mathbf{x}_d - \bar{\mathbf{x}}_d|_i \\ \max_i |\mathbf{q}_d^{(\gamma)} - \mathbf{f}(\bar{\mathbf{x}})|_i \end{bmatrix} + \mathbf{a}_1^{\top} \mathbf{x}_d \leq \delta^*,$$

which, taking $\mathbf{c}^{\top} = [c_1, c_2]$, is equivalent to:

$$\underbrace{\begin{bmatrix} c_1 & c_1 & -c_1 & -c_1 \\ c_2 & -c_2 & c_2 & -c_2 \end{bmatrix}^{\top}}_{\triangleq \mathbf{F}^{\top}} \begin{bmatrix} (\mathbf{x}_d - \bar{\mathbf{x}})_i \\ (\mathbf{q}_d^{(\gamma)} - \mathbf{f}(\bar{\mathbf{x}}))_j \end{bmatrix} + \mathbf{a}_1^{\top} \mathbf{x}_d \leq \delta^*,$$

for all row pairs $i \leq n$ and $j \leq m$ and where $\delta^* \triangleq \delta^* \otimes \mathbf{1}$ with \otimes denoting the Kronecker product. Letting $\mathbf{L}_i \in \{0, 1\}^{4nm \times n+m}$ be matrices capturing the i, j permutations of the scaling matrix \mathbf{F}^{\top} above, we can reformulate this as:

$$[\mathbf{L}_1 \ \mathbf{L}_2] \begin{bmatrix} \mathbf{x}_d - \bar{\mathbf{x}} \\ \mathbf{q}_d^{(\gamma)} - \mathbf{f}(\bar{\mathbf{x}}) \end{bmatrix} + (\mathbf{a}_1^{\top} \otimes \mathbf{1}) \mathbf{x}_d \leq \delta^*,$$

which can be further rearranged as:

$$\underbrace{[\mathbf{L}_1 + \mathbf{a}_1^{\top} \otimes \mathbf{1} \ \mathbf{L}_2]}_{\triangleq \mathbf{L}_i} \begin{bmatrix} \mathbf{x}_d \\ \mathbf{q}_d^{(\gamma)} \end{bmatrix} \leq \delta^* + \underbrace{[\mathbf{L}_1 \ \mathbf{L}_2]}_{\triangleq \mathbf{h}_i} \begin{bmatrix} \bar{\mathbf{x}}_d \\ \mathbf{f}(\bar{\mathbf{x}}_d) \end{bmatrix}.$$

Repeating this process for each of the k rows of the constraint matrix \mathbf{A} yields the desired result. \square

The previous Lemma demonstrates that the inequalities on the desired trajectory $\mathbf{x}_d(\cdot)$ imposed by state and input constraints can be framed as affine constraints on the space of possible trajectories. As curves are infinite dimensional objects, traditional trajectory optimizers would generally

only approximately enforce these constraints. This is precisely where we see the usefulness of Bézier curves – we can exactly enforce these constraints on the continuous-time curve by reasoning about a discrete, low-dimensional collection of Bézier control points (as captured by Property 2). With this in mind, we are now equipped to prove the main statement of the section:

(*Proof of Theorem 1*). Enforcing the constraint in Lemma 1 will result in $\|\mathbf{k}(\mathbf{x}_{\text{cl}}, \mathbf{x}_d, \mathbf{u}_d)\|_\infty \leq u_{\text{max}}$. Furthermore, from Lemma 2, we know that enforcing (11) results in $\mathbf{\Pi}(\mathbf{x}_{\text{cl}}(t)) \in \mathcal{C}_X$. Combining these state and input constraints and leveraging Lemma 3 to produce matrices $\mathbf{L}_x, \mathbf{L}_u$ and vectors $\mathbf{h}_x, \mathbf{h}_u$ results in:

$$\begin{bmatrix} \mathbf{L}_u \\ \mathbf{L}_x \end{bmatrix} \begin{bmatrix} \mathbf{x}_d \\ \mathbf{q}_d^{(\gamma)} \end{bmatrix} \leq \begin{bmatrix} \mathbf{h}_u \\ \mathbf{h}_x \end{bmatrix}. \quad (16)$$

Based on Property 2, we know that if we enforce this constraint on the control points, it will be enforced for the continuous time curve. Therefore, instead we must enforce:

$$\begin{bmatrix} \mathbf{L}_u \\ \mathbf{L}_x \end{bmatrix} \begin{bmatrix} (\mathbf{P})_j \\ (\mathbf{pH}^\gamma)_j \end{bmatrix} \leq \begin{bmatrix} \mathbf{h}_u \\ \mathbf{h}_x \end{bmatrix}.$$

for $j = 0, \dots, p$. As this imposes linear constraints on the columns of \mathbf{p} , this can be vectorized and written as:

$$\mathbf{F}\bar{\mathbf{p}} \leq \mathbf{G},$$

where \mathbf{F} and \mathbf{G} are appropriate reformulations of (16) to account for the vectorization. Enforcing this constraint results in state and input constraint satisfaction as desired. \square

V. BÉZIER REACHABLE POLYTOPES

Given the constructions in Section IV, there exists an affine inequality that guarantees the existence of a Bézier polynomial which results in the closed-loop planner-tracker system satisfying state and input constraints. The matrix \mathbf{F} and vector \mathbf{G} represent an efficient oracle to check whether Bézier curves connecting initial and terminal points satisfy these constraints. Combining this affine constraint with Property 3 allows us to place constraints on the desired boundary conditions of the Bézier polynomial – that is, given an initial condition \mathbf{x}_0 , the set characterized by:

$$\mathcal{F}(\mathbf{x}_0) = \{\mathbf{x}_d \in \mathcal{X}_d \mid \mathbf{F}\bar{\mathbf{D}}^\dagger [\mathbf{x}_0^\top \quad \mathbf{x}_T^\top]^\top \leq \mathbf{G}\},$$

represents all terminal conditions for which there exists a feasible Bézier polynomial. As such, the set $\mathcal{F}(\mathbf{x}_0)$ can be thought of as the forward reachable set of the point \mathbf{x}_0 . Similarly, given a terminal condition \mathbf{x}_T , the backward reachable set is characterized by:

$$\mathcal{B}(\mathbf{x}_T) = \{\mathbf{x}_0 \in \mathcal{X}_d \mid \mathbf{F}\bar{\mathbf{D}}^\dagger [\mathbf{x}_0^\top \quad \mathbf{x}_T^\top]^\top \leq \mathbf{G}\}.$$

A depiction of the forward reachable set for a pendulum system and a variety of system parameters can be seen in Figure 3. As the error tracking tube \mathcal{E} varies in its dependence on \mathbf{u}_d , the reachable sets change shape to ensure that closed loop system still satisfies the desired constraints.

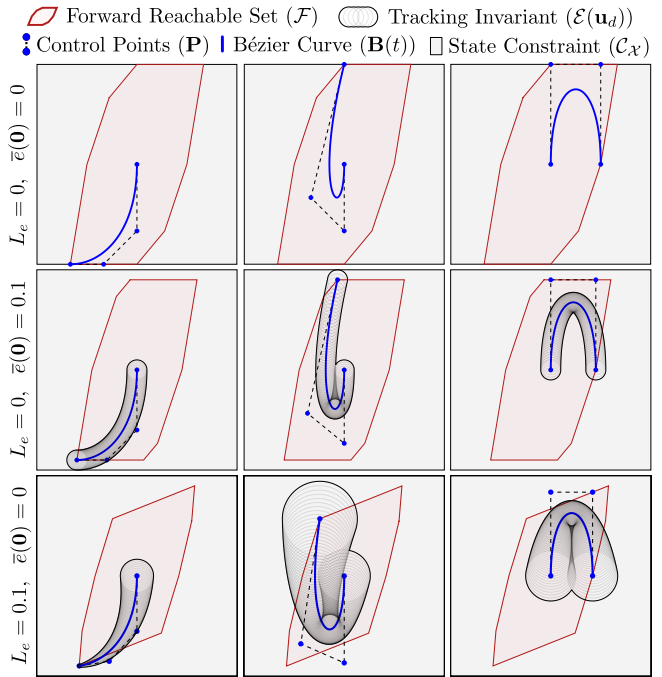


Fig. 3. A selection of Bézier curves and forward reachable sets. The top row depicts curves with exact tracking, the middle row with a fixed size tracking certificate, and the bottom row with a tracking certificate whose upper bound scales linearly with the planning input \mathbf{u}_d .

A. Reducing Conservatism

In the previous discussion, we used a reference point $\bar{\mathbf{x}}_d$ and bounded the deviation of a trajectory from this point. While this enables tractability, it creates conservatism in the bound as the same reference point was used over the entire trajectory $\mathbf{x}_d(\cdot)$. To resolve this conservatism, we would like to instead bound the trajectory with a collection of reference points $\{\bar{\mathbf{x}}_k\}$ spread out over the time interval $[0, T]$. Towards this goal, we leverage the notion of a k -refinement of the interval $[0, T]$ from Definition 4 as well as reference points $\{\bar{\mathbf{x}}_i\}$ for $i = 1, \dots, k$. With these, we can construct a piecewise constant reference trajectory $\bar{\mathbf{x}}(t) = \bar{\mathbf{x}}_i$ for $t \in [T_{i-1}, T_i]$ with $i = 1, \dots, k$. With this reference trajectory, we have the following:

Corollary 1. *Let system Σ_d be a planning model for a system Σ with tracking certificate \mathcal{E} , and consider a piecewise-constant trajectory $\bar{\mathbf{x}}(t)$ defined with respect to a k -refinement of the interval $[0, T]$. There exist matrices $\hat{\mathbf{F}}$ and $\hat{\mathbf{G}}$ such that any Bézier curve $\mathbf{B} : I \rightarrow \mathcal{X}_d$ with control points \mathbf{p} satisfying:*

$$\hat{\mathbf{F}}\bar{\mathbf{p}} \leq \hat{\mathbf{G}}, \quad (17)$$

when tracked results in the closed loop system satisfying $\mathbf{\Pi}(\mathbf{x}_{\text{cl}}(t)) \in \mathcal{C}_X$ and $\mathbf{k}(\mathbf{x}_{\text{cl}}(t), \mathbf{x}_d(t)) \in \mathcal{C}_U$ for all $t \in I$.

Proof. As refinement is linear in the control points, we can leverage the matrices from Theorem 1 and right multiply \mathbf{F} by $\hat{\mathbf{Q}}_i$, the vectorized version of the refinement matrix $\hat{\mathbf{Q}}_i$ for $i = 1, \dots, k$ to produce $\hat{\mathbf{F}}$. Taking $\hat{\mathbf{G}} = \hat{\mathbf{G}}$ yields the desired result. \square

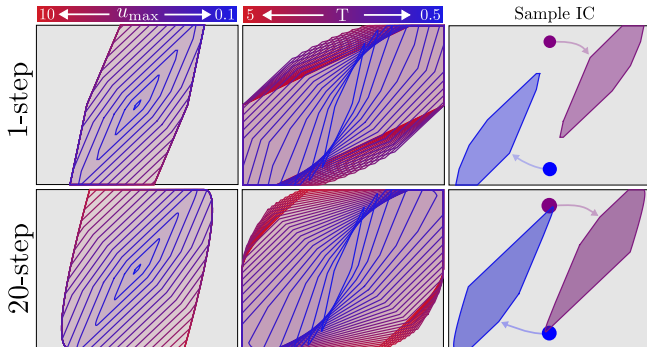


Fig. 4. A depiction of the forward reachable sets as a function of system parameters. (Top Row) 1-step reachable sets, as in Theorem 1. (Bottom Row) 20-step reachable sets, as in Corollary 1. (Left Column) Varying the input constraint \mathbf{u}_{\max} (Middle Column) Varying the time horizon T (Right Column) Varying the initial condition.

By enforcing the constraint in (17), we are able to ensure that the desired trajectory stays close to the piecewise constant reference trajectory, as opposed to a single reference point. This will reduce the conservatism of the bound, but requires increasing the number of constraints needed (and therefore faces of the polytope), demonstrating an obvious tradeoff. A depiction in the difference in resulting reachable sets can be seen in Figure 4. When a single point is used, the reachable set indicates the neighborhood around which that reference point can be feedback linearized, potentially requiring significant input over long time horizons. Instead, if we have a sequence of points, we can forward simulate the drift dynamics to produce reference trajectories, whereby the reachable set represents the neighborhood around the trajectory which we can converge to, thereby reducing conservatism. This notion is especially useful when using such reachable sets to represent an MPC layer, which often uses a sequence of reference points to linearize around.

VI. RESULTS

A. Simulation Results

We deploy the use of Bézier Reachable Polytopes towards the task of swinging up the pendulum. The duration of planning horizon needed to accomplish this task depends highly on how tight the input constraint for the system is. In this setup, the tracker was taken to be the feedback linearizing controller, and the planner produced trajectories on the pendulum dynamics. This planner-tracker was interfaced with a graph-search problem, which samples states uniformly from the state space and connects two vertices $\mathbf{v}_i, \mathbf{v}_j \in \mathcal{X}_d$ with an edge if the intersection of their forward and reachable sets were nonempty, i.e. $\mathcal{F}(\mathbf{v}_i) \cap \mathcal{B}(\mathbf{v}_j) \neq \emptyset$. This represents a graph of dynamically feasible Bézier curves, whereby a suitable Bézier curve between two boundary conditions can be found by solving a discrete graph search problem. As seen in Figure 5, when the low level input constraints are tight, the graph search has to produce a long sequence of points to achieve pendulum swingup. Instead, if the input constraints are loose, then a nearly direct swingup behavior can be achieved. In this

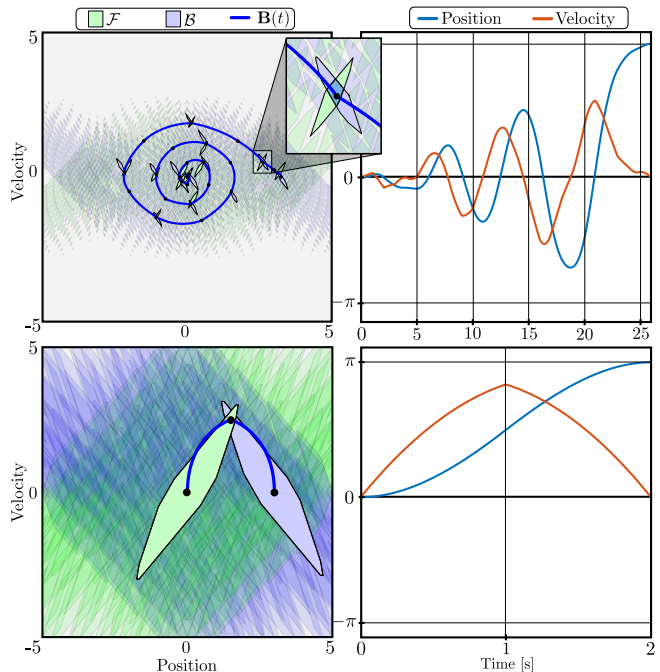


Fig. 5. The proposed method applied to the pendulum swingup problem. As the input bounds are tightened from 5 Nm (bottom) to 0.5 Nm (top), the resulting graph search trajectory increases in complexity and length.

way, we observe that the computational complexity of the decision making layer is imposed by the limitations of the underlying full order system. The code for this project is available at [19].

B. Hardware Results

We also deploy the Bézier Reachable Polytopes framework towards the control of a 3D hopping robot, ARCHER [20], as seen in Figure 6. Let $(\mathbf{p}, q) \in \mathbb{R}^3 \times \mathbb{S}^3$ denote the global position and quaternion of the robot, and $(\mathbf{v}, \boldsymbol{\omega}) \in \mathbb{R}^3 \times \mathfrak{s}^3$ the global linear velocity and body frame angular rates. The full state of the robot $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{20}$ contains these values, as well as foot and flywheel positions and velocities. Planning long-horizon tasks for this robot is extremely challenging due to the large number of passive degrees of freedom, tight input constraints, and hybrid dynamics. Separating the path planning problem into a layered architecture consisting of a tracking controller, a planner, and a decision layer enables this task to be split up, whereby behavior can be generated efficiently.

In this setup, we take the planning model to be a double integrator with state $\mathbf{x}_d \in \mathcal{X}_d \triangleq \mathbb{R}^4$ and input $\mathbf{u}_d \in \mathcal{U}_d \triangleq \mathbb{R}^2$. This planning model Σ_d can be corresponded with the hopping robot Σ by a projection map $\Pi : \mathcal{X} \rightarrow \mathbb{R}^4$ taken to be the restriction of the full order state to the center of mass x and y positions and velocities and an embedding Ψ , which is a Raibert-style controller that takes in desired center of mass state and input trajectories and produces desired orientation quaternions as:

$$q_d(\mathbf{x}, t) = \mathbf{K}_{\text{fb}}(\Pi(\mathbf{x}) - \mathbf{x}_d(t))$$

with desired angular rates $\boldsymbol{\omega}_d \equiv \mathbf{0}$. This desired quaternion

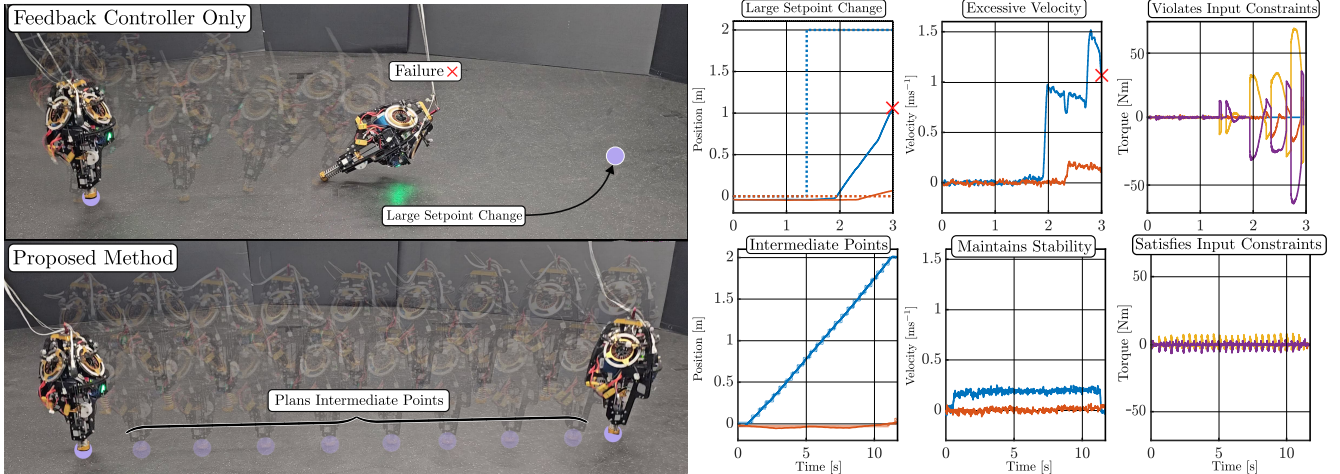


Fig. 6. Hardware results on the 3D hopping robot, ARCHER. When commanded to cross the room, a naive decision making layer provides a setpoint to the planner which is outside what is achievable by the low-level system, leading to system failure. Instead, if Bézier Reachable Polytopes are used, the decision layer provides a sequence of waypoints to the planner that results in completion of the objective while satisfying state and input constraints.

is then tracked by a low-level controller \mathbf{k} as:

$$\mathbf{k}(\mathbf{x}, q_d, \mathbf{u}_d) = -\mathbf{K}_p \text{Im}(q_d^{-1}q) - \mathbf{K}_d(\boldsymbol{\omega} - \boldsymbol{\omega}_d) + \mathbf{K}_f \mathbf{u}_d,$$

which runs at 1 kHz. As seen in Figure 6, if only the feedback layer is used, the system fails because the desired setpoint is outside the region of what can be accomplished by the tracking system. Instead, if the proposed method is used, the decision layer can autonomously produce a sequence of points which maintain stability and constraint satisfaction over the task.

VII. CONCLUSION

In this work, we introduced the concept of Bézier Reachable Polytopes, which provide a representation of the set of points that can be reached by planner-tracker control frameworks. By leveraging the properties of Bézier polynomials, we showed that this set can be efficiently represented via a polytopic constraint, enabling computationally tractable long-horizon planning to be achieved. Future work includes developing an abstract theory for such hierarchical control systems and their interconnections.

VIII. ACKNOWLEDGEMENTS

The authors would like to thank Andrew Taylor, Preston Culbertson, and Max Cohen for their many fruitful discussions and William Compton for his assistance both with theory with experiments.

REFERENCES

- [1] N. Matni, A. D. Ames, and J. C. Doyle, “A quantitative framework for layered multirate control: Toward a theory of control architecture,” *IEEE Control Systems Magazine*, vol. 44, no. 3, pp. 52–94, 2024.
- [2] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, “Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot,” *Autonomous robots*, vol. 40, pp. 429–455, 2016.
- [3] R. Grandia, F. Jenelten, S. Yang, F. Farshidian, and M. Hutter, “Perceptive locomotion through nonlinear model-predictive control,” *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 3402–3421, 2023.
- [4] G. Pappas, G. Lafferriere, and S. Sastry, “Hierarchically consistent control systems,” *IEEE Transactions on Automatic Control*, vol. 45, no. 6, pp. 1144–1160, 2000.
- [5] A. Girard and G. J. Pappas, “Hierarchical control system design using approximate simulation,” *Automatica*, vol. 45, no. 2, pp. 566–571, 2009.
- [6] A. van der Schaft, “Equivalence of dynamical systems by bisimulation,” *IEEE Transactions on Automatic Control*, vol. 49, no. 12, pp. 2160–2172, 2004.
- [7] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017.
- [8] D. Q. Mayne, M. M. Seron, and S. V. Raković, “Robust model predictive control of constrained linear systems with bounded disturbances,” *Automatica*, vol. 41, no. 2, pp. 219–224, 2005.
- [9] M. Chen, S. L. Herbert, H. Hu, Y. Pu, J. F. Fisac, S. Bansal, S. Han, and C. J. Tomlin, “Fastrack: a modular framework for real-time motion planning and guaranteed safe tracking,” *IEEE Transactions on Automatic Control*, vol. 66, no. 12, pp. 5861–5876, 2021.
- [10] A. Wu, S. Sadraddini, and R. Tedrake, “R3t: Rapidly-exploring random reachable set tree for optimal kinodynamic planning of nonlinear hybrid systems,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4245–4251.
- [11] E. D. Sontag, “Input to state stability: Basic concepts and results,” in *Nonlinear and Optimal Control Theory*. Springer, 2008, pp. 163–220.
- [12] N. Csomay-Shanklin, A. J. Taylor, U. Rosolia, and A. D. Ames, “Multi-rate planning and control of uncertain nonlinear systems: Model predictive control and control lyapunov functions,” *arXiv:2204.00152*, 2022.
- [13] B. Donald, P. Xavier, J. Canny, and J. Reif, “Kinodynamic motion planning,” *Journal of the ACM*, vol. 40, no. 5, pp. 1048–1066, Nov. 1993.
- [14] S. M. LaValle and J. J. Kuffner, “Randomized Kinodynamic Planning,” *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, May 2001.
- [15] D. J. Webb and J. van den Berg, “Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics,” in *2013 IEEE International Conference on Robotics and Automation*. Karlsruhe, Germany: IEEE, May 2013, pp. 5054–5061.
- [16] T. Marcucci, P. Nobel, R. Tedrake, and S. Boyd, “Fast Path Planning Through Large Collections of Safe Boxes,” May 2023, arXiv:2305.01072 [cs, eess].
- [17] M. E. Flores Contreras, “Real-time trajectory generation for constrained nonlinear dynamical systems using non-uniform rational b-spline basis functions,” Ph.D. dissertation, California Institute of Technology, 2008.
- [18] M. Kamermans, “A primer on bézier curves,” (*online book*), 2020. [Online]. Available: <https://pomax.github.io/bezierinfo/>
- [19] “Code,” 2024. [Online]. Available: <https://github.com/noelc-s/BezierTubes>
- [20] E. R. Ambrose, “Creating ARCHER: A 3D Hopping Robot with Flywheels for Attitude Control,” Ph.D. dissertation, California Institute of Technology, 2022.