

Dynamically Feasible Path Planning in Cluttered Environments via Reachable Bézier Polytopes

Noel Csomay-Shanklin, William D. Compton, Aaron D. Ames

Abstract—The deployment of robotic systems in real world environments requires the ability to quickly produce paths through cluttered, non-convex spaces. These planned trajectories must be both kinematically feasible (i.e., collision free) and dynamically feasible (i.e., satisfy the underlying system dynamics), necessitating a consideration of both the free space and the dynamics of the robot in the path planning phase. In this work, we explore the application of *reachable Bézier polytopes* as an efficient tool for generating trajectories satisfying both kinematic and dynamic requirements. Furthermore, we demonstrate that by offloading specific computation tasks to the GPU, such an algorithm can meet tight real time requirements. We propose a layered control architecture that efficiently produces collision free and dynamically feasible paths for nonlinear control systems, and demonstrate the framework on the tasks of 3D hopping in a cluttered environment.

I. INTRODUCTION

The development of control strategies which enable robots to operate in cluttered environments has been a focal challenge in control, dating back to the early days of mobile robots [1]. While it is theoretically possible to accomplish this task in a single shot via policy-design [2]–[5], this poses significant challenges including reasoning about high dimensional state and action spaces, nonlinear dynamics, and having a lack of safety and feasibility guarantees. As a result, it is often more practical to deploy hierarchical control architectures, whereby high level path planning algorithms are paired with low-level tracking controllers. Typically, kinematically feasible (i.e., collision free) paths are passed down the control stack, and dynamic feasibility (i.e., satisfaction of the underlying system dynamics) is an assumed property of the low-level controller. This top-down approach, however, provides minimal guarantees of state and input constraint satisfaction, as violations of this assumption at the low level may cause significant deviations from the planned high-level trajectory.

Constructing a pipeline which is able to guarantee feasibility necessitates reasoning about both the dynamics and free space information during the planning phase; algorithms which plan for both kinematic and dynamic feasibility are termed *kinodynamic planners* [6]. Kinodynamic planners can broadly be placed into two categories: sampling-based and optimization-based methods. In the sampling-based paradigm, one main way of producing kinodynamically feasible paths is by sampling policies from a discrete collection of predefined primitives [7], [8]. This method

The authors are with the Control and Dynamical Systems Department at the California Institute of Technology. Email: {noelcs@caltech.edu}. This research is supported by Technology Innovation Institute (TII).




Fig. 1: The proposed framework performing path planning around segmented obstacles. A coarse dynamically feasible path is solved for on a graph with Bézier polynomial edges, and is further refined with MPC.

suffers from the fact that the predefined policies induce a bias and may not make meaningful progress towards the goal. Another approach is to randomly sample new points and connect them to the nearest node on the graph if a two point boundary value problem is feasible [9], which can be viewed as an extension to the classic RRT [10] method of kinematic path planning. This method requires solving as many such problems as there are nodes in the graph, which can be expensive to compute for high dimensional nonlinear systems.

Sampling-based methods face two major challenges – path suboptimality due to the probability of sampling the true optimal path being zero, and bias, either through the discretized policy design, or Voronoi bias introduced by using a Euclidean metric [11]. To address suboptimality, various methods have proposed refinements to the path which increase the optimality [9]. To address the notion of bias, the notion of reachability has been introduced into sampling-based planners [12]. In particular, authors have developed reachability-guided RRT variants [13]–[15] that avoid solving boundary value problems during graph construction. For nonlinear systems, however, this top-down approach provides guarantees of feasibility which are often approximate and rely on using local linearizations of the nonlinearities; this requires the system to remain sufficiently close to the nominal trajectory during execution.

In contrast to building a graph of sampled states, optimization-based approaches instead try to directly generate optimal paths to the goal state. These methods often

require the solution of large mixed integer programs to generate optimal paths to the goal, either via a keep-out philosophy [16], or a keep-in philosophy via convex decomposition [17]–[19]. To address the benefits and drawbacks sampling-based and optimization techniques, various approaches have been proposed to combine these methods [20]–[22]. Despite these advancements, the computational burden associated with these approaches hinders their deployment in obstacle-rich environments in real time. As a means to mitigate these computational limitations, layered approaches to control synthesis have proven effective [23].

In this work, we leverage a layered control architecture for producing kinodynamically feasible trajectories for nonlinear systems. Motivated by [24], we employ Bézier basis polynomials for the task of solving paths through nonconvex spaces. In contrast to this and other works, significant effort is put towards guaranteeing dynamic feasibility of the generated paths in order to produce a guaranteed feasible pipeline. Specifically, we leverage the notion of *Bézier reachable polytopes*, i.e., polytopic reachable sets in the space of polynomial Bézier reference trajectories, as a primitive for feasible trajectory tracking. Importantly, as the dynamics are abstracted via the reachable set of the mid level controller, the planner does not need to include differential constraints. This work builds on [25] by incorporating a high level path planner in coordination with a mid and low level controller. As seen in Figure 1, we demonstrate the effectiveness of this framework by planning paths in cluttered environments on the high-dimensional 3D hopping robot, ARCHER [26].

II. BACKGROUND

A. Reduced Order Models and Problem Statement

We begin with the following nonlinear control system:

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

where $\mathbf{x} \in \mathbb{R}^N$ is the state, $\mathbf{u} \in \mathbb{R}^M$ is the input, and the vector $\mathbf{f} : \mathbb{R}^N \times \mathbb{R}^M \rightarrow \mathbb{R}^N$ is assumed to be continuously differentiable in each argument. Directly synthesizing control actions that achieve complex tasks such as path planning may be challenging, in part due to the nonlinearities of the dynamics. To alleviate this complexity, roboticists often leverage reduced-order models, which serve as template systems that enable desired behaviors to be constructed in a computationally tractable way. To this end, consider the planning dynamics:

$$\dot{\mathbf{x}}_d = \begin{bmatrix} \mathbf{0} & \mathbf{I}_{n-m} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{x}_d + \begin{bmatrix} \mathbf{0} \\ \mathbf{I} \end{bmatrix} \mathbf{u}_d, \quad (2)$$

where $\mathbf{x}_d \in \mathbb{R}^n$ is the reduced-order state and $\mathbf{u}_d \in \mathbb{R}^m$ is the reduced-order input. Note that the the subsequent discussion, the planning dynamics in 2 can be extended to the case of full-state feedback linearizable systems (as in [25]), but for the sake of simplicity we restrict our attention to integrators.

Let $I = [0 \ T] \subset \mathbb{R}_{\geq 0}$ be a time interval, $\mathbf{u}_d(\cdot)$ an input trajectory defined over I , and $\mathbf{x}_d : I \rightarrow \mathbb{R}^n$ be an integral curve of the reduced-order model (2). We can correspond a reduced order system with a full order system through a

surjective mapping $\mathbf{\Pi} : \mathbb{R}^N \rightarrow \mathbb{R}^n$, and feedback controller $\mathbf{k} : \mathbb{R}^N \times \mathbb{R}^n \rightarrow \mathbb{R}^M$, which takes in reduced order model trajectories and produces control actions for the full order system. A desired property of this controller is its ability to maintain bounded tracking error:

Definition 1. The set $\mathcal{E} : I \rightarrow \mathcal{X}$ is a *tracking invariant* for a desired trajectory $\mathbf{x}_d : I \rightarrow \mathcal{X}_d$ for the system (1) if:

$$\mathbf{\Pi}(\mathbf{x}(t)) \in \mathbf{x}_d(t) \oplus \mathcal{E}(t),$$

where \oplus denotes the Minkowski sum.

Many existing works discuss ways to produce controllers with tracking invariant sets \mathcal{E} [27], [28]; here, we assume such a tracking controller is given. For a desired trajectory of the reduced-order model $\mathbf{x}_d(\cdot)$, a suitable tracking controller \mathbf{k} , and an initial condition $\mathbf{x}(0) \in \mathbb{R}^N$, we can define the closed-loop system as:

$$\mathbf{x}_{\text{cl}}(t) = \mathbf{x}(0) + \int_0^t \mathbf{f}(\mathbf{x}(\tau), \mathbf{k}(\mathbf{x}(\tau), \mathbf{x}_d(\tau))) d\tau. \quad (3)$$

We can now introduce the main problem statement:

Problem 1. Given a reduced order model (2), with state constraint $\mathcal{X}_d \subset \mathbb{R}^n$, a full order model (1) with input bounds $\mathcal{U} \subset \mathbb{R}^M$, a worst case tracking error bound $\bar{\mathcal{E}} \supset \mathcal{E}(t) \forall t$, and an initial and goal location for the full order model $\mathbf{x}(0)$ and \mathbf{x}_G , find a trajectory $\mathbf{x}_d(t)$ such that:

$$\lim_{t \rightarrow \infty} \mathbf{\Pi}(\mathbf{x}_{\text{cl}}(t)) \in \mathbf{\Pi}(\mathbf{x}_G) \oplus \bar{\mathcal{E}}, \quad (4a)$$

$$\mathbf{\Pi}(\mathbf{x}_{\text{cl}}(t)) \in \mathcal{X}_d \quad \forall t, \quad (4b)$$

$$\mathbf{k}(\mathbf{x}_{\text{cl}}(t), \mathbf{x}_d(t)) \in \mathcal{U} \quad \forall t. \quad (4c)$$

Simply put, we are searching for a reduced order model trajectory such that the closed-loop full-order model will navigate to a goal position while respecting state and input constraints. To solve this problem, we will use Bézier curves to parameterize the space of trajectories passed from the planning level to the feedback level.

B. Bézier Curves and Graphs

A curve $\mathbf{b} : I \rightarrow \mathbb{R}^m$ is said to be a Bézier curve [29] 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 $\mathbf{H} \in \mathbb{R}^{p+1 \times p+1}$ 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 $\mathbf{B} : I \rightarrow \mathbb{R}^n$:

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




Fig. 2: The path planning framework presented in Algorithm 1. From left to right: a) A Bézier graph is constructed, b) it is cut based on the present obstacles, c) a path is solved, and d) the path is refined with MPC.

where $\gamma > 0$ is the relative degree of an output $\mathbf{b}(t)$ for the system (2) and is defined such that $m\gamma = n$. 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 $\mathbf{B}(\cdot)$ in the state space.

Given a Bézier polynomial control input $\mathbf{u}_d(\cdot)$, the integral curve $\mathbf{x}_d(\cdot)$ can be exactly described by the Bézier polynomial (5) of degree γ larger than that of $\mathbf{u}_d(\cdot)$. Furthermore, given two boundary conditions, there exist unique Bézier curves (of minimal degree $p = 2\gamma - 1$) connecting them:

Lemma 1 ([30]). *Given two points $\mathbf{x}_0, \mathbf{x}_T \in \mathbb{R}^n$, there exists a unique matrix \mathbf{D} such that $\mathbf{P} = \mathbf{D} [\mathbf{x}_0 \quad \mathbf{x}_T]$ are the control points of a Bézier curve $\mathbf{x}_d(\cdot)$ satisfying $\mathbf{x}_d(0) = \mathbf{x}_0$ and $\mathbf{x}_d(T) = \mathbf{x}_T$.*

This statement allows us to plan dynamically feasible curves simply by reasoning about discrete points in state space. This concept is paired with a key property of Bézier curves – namely, the continuous time curve is bounded by the convex hull of the control points:

Property 1 (Convex Hull [29]). A Bézier curve \mathbf{x}_d is contained in the convex hull of its control points:

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

Property 2 (Path Length [24]). The path length of a Bézier curve $\mathbf{x}_d(\cdot)$ is upper bounded by the norm distance of its control points:

$$\int_0^T \|\dot{\mathbf{x}}_d(\tau)\|_2 d\tau \leq \sum_{j=0}^{p-1} \|\mathbf{P}_{j+1} - \mathbf{P}_j\|_2$$

Property 3 (Subdivision [29]). A Bézier curve \mathbf{x}_d defined on the time interval $I = [0 \ T]$ can be subdivided into arbitrarily many Bézier curves given a subdivision of I .

Besides the standard aforementioned properties, the reachable set of Bézier polynomials which when tracked satisfy state and input constraints of the tracking system is efficiently representable:

Theorem 1 ([30]). *Given a convex state constraint set $\mathcal{X}_d \subset \mathbb{R}^n$, input constraint set $\mathcal{U} \subset \mathbb{R}^M$ and error tracking bound $\bar{\epsilon}$, there exist matrices \mathbf{F} and \mathbf{G} such that any two points $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n$ satisfying:*

$$\mathbf{F} [\mathbf{x}_1^\top \quad \mathbf{x}_2^\top]^\top \leq \mathbf{G}$$

implies the existence of a Bézier curve \mathbf{x}_d with $\mathbf{x}_d(0) = \mathbf{x}_1$, $\mathbf{x}_d(T) = \mathbf{x}_2$, such that, when tracked, the closed loop system satisfies $\Pi(\mathbf{x}_{cl}(t)) \in \mathcal{X}_d$ and $\mathbf{k}(\mathbf{x}_{cl}(t), \mathbf{x}_d(t)) \in \mathcal{U}$.

The matrices \mathbf{F} and \mathbf{G} can be thought of as reachable set oracles, as they determine if there exists dynamically feasible curves between two points. This notion of Bézier reachable sets will be used to solve Problem 1 by first efficiently generating a graph of dynamically feasible Bézier curves.

A graph \mathcal{G} is described by a tuple (V, E) with vertices V connected by edges E . For the purpose of our analysis, the vertices will represent points in the reduced-order state space \mathbb{R}^n . A directed edge between two vertices $e = (v_1, v_2)$ for $v_1, v_2 \in V$ will represent the existence of a dynamically feasible trajectory connecting v_1 to v_2 .

III. KINODYNAMIC BÉZIER GRAPHS

The goal of this section will be to produce desired trajectories $\mathbf{x}_d(\cdot)$ which satisfy Problem 1 when obstacles are present, i.e. the free space is non-convex. The approach will be to produce a graph of dynamically feasible Bézier curves leveraging the *Bézier reachable polytopes* introduced in Theorem 1, cut the edges that intersect obstacles, and perform a graph search that will serve as a feasible warm start for a trajectory optimization program. The paradigm is outlined in Algorithm (1) and shown in Figure 2, and starts with building a graph of Bézier curves.

A. Build Graph

The compact state space of interest \mathcal{X}_d represents the free space of the problem setup, and will serve as a seed which we can prune in the presence of obstacles. For the sake of limiting algorithmic complexity, we make the following assumption about the description of the free space:

Assumption 1. The free space is given by $\mathcal{X}_d \setminus \mathcal{O} \subset \mathbb{R}^n$ for a set of obstacles $\mathcal{O} = \cup_i \mathcal{O}_i$ where \mathcal{O}_i is convex polytope.

To begin, N points are uniformly sampled $\mathbf{v}_i \sim U(\mathcal{X}_d)$ and the matrices \mathbf{F} and \mathbf{G} from Theorem 1 are generated.

Algorithm 1 Bézier Graphs

- 1: **hyperparameters:** $(N, T) \triangleright$ (Node count and time interval)
 - 2: $\mathcal{G} \leftarrow \text{buildGraph}(\mathcal{X}_d, \mathcal{U})$
 - 3: $\mathcal{C} \leftarrow \text{cutGraph}(\mathcal{G}, \mathcal{O})$
 - 4: $\mathbf{v}_k^* \leftarrow \text{findPath}(\mathcal{C})$
 - 5: $\mathbf{x}_d(\cdot) \leftarrow \text{refineWithMpc}(\mathbf{v}_k^*, \mathcal{O})$
 - 6: **return** $\mathbf{x}_d(\cdot)$
-

Algorithm 2 cutHeuristic

```

1: if any( $\mathbf{A}\mathbf{v}_i \leq \mathbf{b}$ ) then                                ▷ Figure 3a)
2:   return 0
3: end if
4:  $\mathbf{v}^o \leftarrow \text{closestPoint}(\mathbf{v}_i)$ 
5:  $(\mathbf{h}, \mathbf{l}) \leftarrow \text{adjacentHyperplane}(\mathbf{v}^o)$ 
6: if all( $\mathbf{h}\mathbf{v}_i \leq \mathbf{l}$ ) then                                ▷ Figure 3b)
7:   return 1
8: end if
9: Solve QP                                                ▷ Figure 3c)
10: return  $\|\delta^*\| > 0$ 

```

Two points \mathbf{v}_i and \mathbf{v}_j are connected with an edge $e_{i,j}$ if they satisfy $\mathbf{F} \begin{bmatrix} \mathbf{v}_i^\top & \mathbf{v}_j^\top \end{bmatrix}^\top \leq \mathbf{G}$. As such, every edge in the graph implies the existence of a dynamically feasible Bézier curve connecting them. This initial graph will be fixed throughout the remainder of the Algorithm.

B. Cut Graph

Given a graph of Bézier curves \mathcal{G} , we remove the edges that could result in collisions with any obstacles that are present. Given a set of obstacles \mathcal{O} , Property 1 states that a Bézier curve is guaranteed to be collision-free if the convex hull of its control points does not intersect with any obstacles, i.e. $\text{conv}\{\mathbf{P}\} \cap \mathcal{O} = \emptyset$. For each edge e with control points \mathbf{P} and convex obstacle \mathcal{O}_i characterized by n_c hyperplane constraints $\mathbf{A}^{\mathcal{O}_i} \in \mathbb{R}^{n \times n_c}$ and $\mathbf{b}^{\mathcal{O}_i} \in \mathbb{R}^{n_c}$, the problem of determining if the curve is collision-free can be formulated as a linear program feasibility check via Property (1), which we solve using the following quadratic program:

$$\begin{aligned}
 \min_{\lambda, \delta} \quad & \delta^\top \delta && \text{(Cut-QP)} \\
 \text{s.t.} \quad & \mathbf{A}^{\mathcal{O}_i} \mathbf{P} \lambda \leq \mathbf{b}^{\mathcal{O}_i} + \delta \\
 & \lambda_j \geq 0, \quad \sum_j \lambda_j = 1
 \end{aligned}$$

where $\lambda \in \mathbb{R}^{p+1}$ is the convex interpolation variable and $\delta \in \mathbb{R}^{n_c}$ is a slack variable. Given the solution δ^* , if $\|\delta^*\|_2 > 0$ then the curve is obstacle free. This program can easily be vectorized to solve all edges simultaneously.

Although this is a quadratic program and therefore can be solved efficiently, solving it for every edge and obstacle becomes too computationally expensive for real-time operation (even when parallelized), especially when the edge and number of obstacle count is large. To address this, we introduce the heuristic outlined in Algorithm 2, which consists of three checks. First, we determine if any of the control points of \mathbf{P} lie in the obstacle – if so, the curve may intersect the obstacle. Next, we attempt to find a separating hyperplane between the obstacle and the control points – if one exists, then the curve will be collision free. If the points do not lie on one side of the hyperplane, then the heuristic is indeterminate, and the original quadratic program must be solved. In practice, this heuristic typically removes $> 99\%$ of the edges that need to be checked; furthermore, this heuristic is easily implemented on GPU as it only requires linear algebra operations, which significantly improves execution speed (as reported in Section IV).




Fig. 3: The heuristic employed to check if the Bézier curve is collision free. First, obstacle membership is checked. If not satisfied, a single proposed separating hyperplane is checked. Finally, the Quadratic Program (Cut-QP) is solved.

C. Find Path

Due to the random sampling of the points in the graph, the starting and goal positions are likely not vertices. Therefore, the start and goal nodes are chosen as the closest normed-distance points to the desired start and goal positions. The edge cost is taken to be $\sum_i \|\mathbf{P}_{i+1} - \mathbf{P}_i\|$, the upper bound on the path length of the Bézier curve from Property 2. We then use Dijkstra's algorithm to solve Problem 2:

Problem 2. In the context of Problem 1, consider a collision-free graph of Bézier curves $\mathcal{C} = (V, E)$ produced by Algorithm 1. Find a path $\{\mathbf{v}_k^*\}_{k=0}^K$ in the graph connecting \mathbf{v}_0^* to $\mathbf{v}_K^* = \Pi(x_G)$ with \mathbf{v}_0^* satisfying $\Pi(\mathbf{x}(0)) \in \mathbf{v}_0^* \oplus \bar{\mathcal{E}}$.

D. Refine the Path

If the graph solve \mathbf{v}_k^* exists, it represents a feasible path for the system to safely traverse the cluttered environment; however, it may be significantly suboptimal. To improve its optimality, we can solve a finite time optimal control program, which balances tracking the graph solution with short-horizon optimality. Note that by Lemma 1, a solution Problem 2 can be converted into a dynamically feasible curve $\mathbf{x}_d(\cdot)$ for the reduced order model (2). Then, we sample this curve at the optimal control time discretization to produce a sequence of reference points \mathbf{r}_k , which can be optimized:

$$\begin{aligned}
 \min_{\mathbf{x}, \mathbf{u}} \quad & \sum_{k=0}^{N-1} \mathbf{x}_k^\top \mathbf{Q} \mathbf{x}_k + \mathbf{F} \mathbf{r}_k + \mathbf{u}_k^\top \mathbf{R} \mathbf{u}_k + \mathbf{x}_N^\top \mathbf{V} \mathbf{x}_N && \text{(MPC)} \\
 \text{s.t.} \quad & \mathbf{x}_{k+1} = \mathbf{A} \mathbf{x}_k + \mathbf{B} \mathbf{u}_k && (7a) \\
 & \mathbf{x}_0 \in \Pi(\mathbf{x}(0)) \oplus \bar{\mathcal{E}} && (7b) \\
 & \mathbf{D} \begin{bmatrix} \mathbf{x}_k & \mathbf{x}_{k+1} \end{bmatrix} \in \mathcal{X}_d \setminus (\mathcal{O} \oplus \mathcal{E}) && (7c) \\
 & \mathbf{F} \begin{bmatrix} \mathbf{x}_k & \mathbf{x}_{k+1} \end{bmatrix} \leq \mathbf{G} && (7d) \\
 & \mathbf{x}_N = \mathbf{r}_N && (7e)
 \end{aligned}$$

where $\mathbf{Q}, \mathbf{F} \in \mathbb{R}^{n \times n}$ are symmetric positive definite matrices weighting distance to reference as well as path length, $\mathbf{R} \in \mathbb{R}^{m \times m}$ is a positive definite input scaling matrix, and $\mathbf{V} \in \mathbb{R}^{n \times n}$ is a terminal cost. The matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$ and $\mathbf{B} \in \mathbb{R}^{n \times m}$ are the exact discretization of the integrator dynamics in (2). The constraint $\mathcal{X}_d \setminus (\mathcal{O} \oplus \mathcal{E})$ is enforced by fining a separating hyperplane between the node \mathbf{x}_k and the closest obstacle (where the same adjacentHyperplane call form the heuristic can be made). In order to improve the quality of solutions, (MPC) can be iteratively solved in an SQP fashion.




Fig. 4: 500 randomly generated obstacles with graph replanning at 10 Hz and (MPC) at 200 Hz. Cyan indicates the Bézier graph solve \mathbf{v}_k^* and blue the MPC Bézier solution $\mathbf{x}_d(\cdot)$.

Importantly, the graph solve represents a feasible warm start for the MPC program, and can be thought of as providing a corridor and a dynamically feasible path which MPC can refine. The main motivation for the structures used in this paper is the notion of guaranteed feasibility:

Theorem 2. *If Problem 2 is feasible, then Problem 1 is solved by applying (MPC) in closed loop.*

Proof. First, we must show that at time $t = 0$, the solution \mathbf{v}_k^* from Problem 2 provides a feasible solution for (MPC). Let \mathbf{r}_k denote the refined sequence of \mathbf{v}_k^* , which is constructed to produce the reference for (MPC) using Property 3. As this is associated with a Bézier curve for the reduced-order model, it satisfies the discrete dynamics (7a). Next, by the definition of Problem 2, we know $\Pi(\mathbf{x}_0) \in \mathbf{v}_0^* \oplus \bar{\mathcal{E}}$, which implies that (7b) is satisfied by Mikowski addition properties. As \mathbf{v}_k^* is in the collision-free graph \mathcal{C} , it satisfies (7d) and (7c). By construction, (7e) is satisfied. Therefore, \mathbf{r}_k^* represents a feasible solution for (MPC).

Next, take \mathbf{x}_k^* to denote the MPC solution and consider the time interval $I = [0 \ h]$ for MPC discretization $h > 0$. From Lemma (1), we know that there exists a unique curve $\mathbf{x}_d(\cdot)$ defined over I connecting \mathbf{x}_0^* to \mathbf{x}_1^* . By Theorem 1, the closed loop system tracking this curve will satisfy (4b) and (4c) for all $t \in I$. Furthermore, by Definition (1), we have that $\Pi(\mathbf{x}(h)) \in \mathbf{x}_1^* \oplus \bar{\mathcal{E}}$. Therefore, we can appeal to standard Robust tube MPC theory [31] to claim recursive feasibility and robust stability of the closed loop system tracking the MPC solution, meaning (4a) is achieved and (4b) and (4c) are satisfied for all time. \square

We have shown that if there is a solution to the graph problem in Algorithm 1, then the closed loop system tracking MPC will satisfy state and input constraints $\mathcal{X}_d \setminus \mathcal{O}$ and \mathcal{U} for all time, and will approach a neighborhood of the goal \mathbf{x}_G .

TABLE I: Comparison of functions on CPU, parallelized CPU, and GPU. Any CPU implementation would violate real time constraints.

	cutHeuristic	adjacentHyperplane
CPU	5920 \pm 36 ms	6.24 \pm 0.08 ms
Parallel CPU	3640 \pm 31 ms	2.52 \pm 0.33 ms
GPU	72.0 \pm 1.5 ms	0.61 \pm 0.14 ms




Fig. 5: A long-horizon maze with 600 cells, 300 obstacles and 50,000 edges. The graph cut is solved at 10 Hz, and graph solve and MPC both run at 200 Hz. Blue represents the graph solve, and red the closed loop system behavior.

IV. RESULTS

We deploy Algorithm 1 on the 3D hopping robot, ARCHER [26] – a video can be found at [32] and the code for this project is available at [33]. 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 all of the above states, as well as foot and flywheel positions and velocities. Although planning dynamically feasible paths on the full hybrid system dynamics is possible [34], the added complexity of planning collision free paths on the full order state directly quickly becomes intractable. The most challenging problem preventing the nonconvex problem from being directly solved in real time is appropriately choosing a collection of corridors to traverse. By combining problems at various time scales, the graph solve is able to provide a coarse estimate of the keep in constraints while foregoing path optimality, and the short horizon MPC is able to refine it into a more optimal path. Therefore, we choose to plan desired center of mass trajectories $\mathbf{x}_d \in \mathbb{R}^4$ with virtual force inputs $\mathbf{u}_d \in \mathbb{R}^2$, which are assumed to have double-integrator dynamics.

The projection map $\Pi : \mathcal{X} \rightarrow \mathbb{R}^4$ is taken to be the restriction of the full order state to the center of mass x and y positions and velocities. The tracking controller \mathbf{k} is a Raibert-style controller, which takes in desired center of mass state and input trajectories and produces desired orientation quaternions as:

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

and desired angular rates $\boldsymbol{\omega}_d \equiv \mathbf{0}$. This desired quaternion is then tracked by a low-level controller:

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

which runs at 1 kHz. For all experiments, the layered controller was running on an Ubuntu 20.04 machine with an AMD Ryzen 5950x @ 3.4 GHz, an NVIDIA GeForce RTX 4090, and 64 Gb RAM.

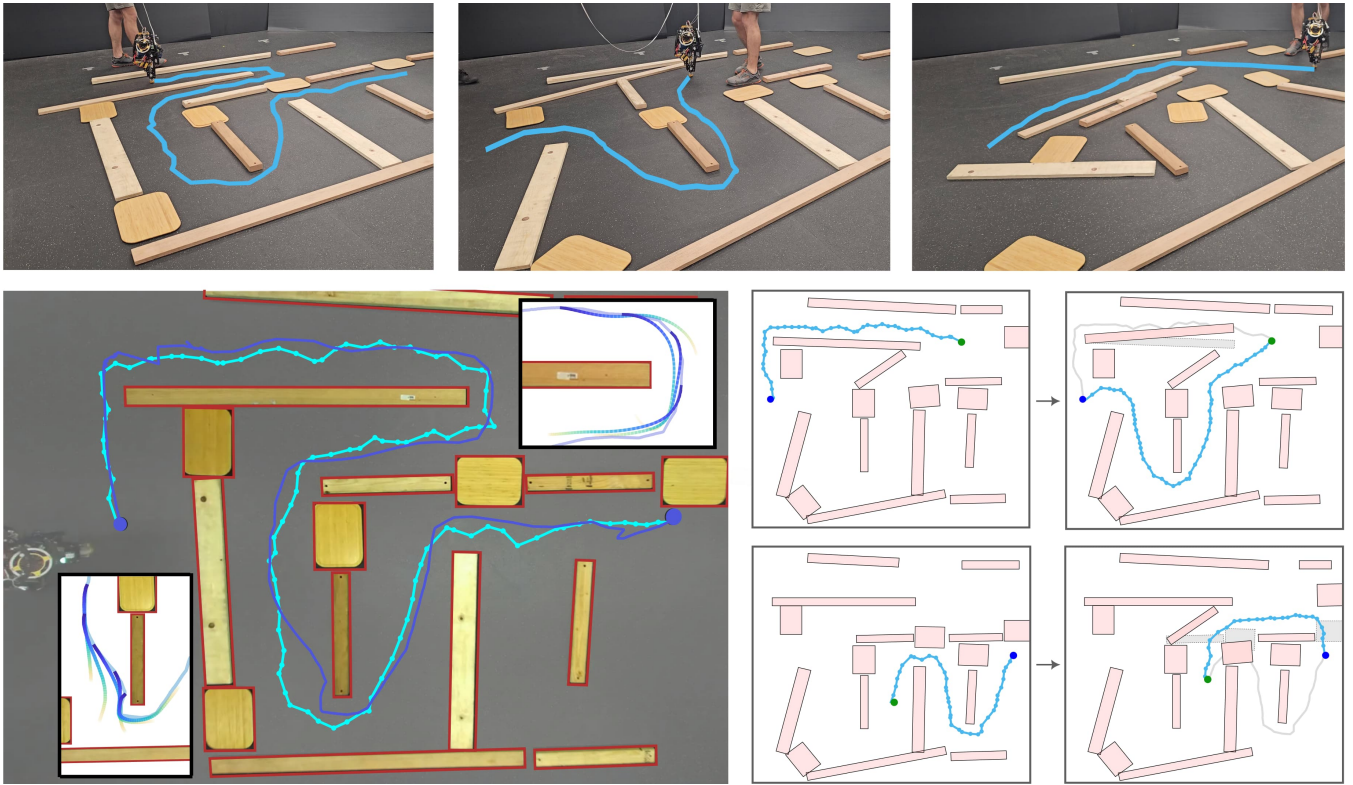


Fig. 6: Experiments run on the ARCHER hardware Platform. (Top) 3 snapshots of the graph solve for various obstacle configurations. (Left) An overhead image of the graph solve and the closed loop system trajectory. In the upper right and bottom left corner, the MPC refinement around the corner can be seen. (Right) Freeze frames of the graph updating online when the obstacle locations are moved.

In order to solve (Cut-QP) and (MPC), the OSQP library [35] was used with maximum iterations limited to 50. The heuristic for cutting the graph and getting the separating hyperplanes was running on a custom written CUDA kernel. Using the GPU to perform these operations was critical to real-time performance – as summarized in Table I, the CUDA kernel provided a significant speedup as compared to multithreaded CPU implementation. A nominal graph with 0.5 second Bézier curves and a 50 node MPC horizon with a timestep of 0.1 s were used. Each iteration of (MPC) was solved at 100 Hz, and in practice 1 SQP iteration was taken. Finally, path length cost was scaled significantly higher than tracking cost in order to incentivize shorter paths than the graph solve could provide.

A. Simulation and Hardware

As can be seen in Figure 4 and Figure 5, the proposed framework is able to solve extremely long horizon tasks with numerous obstacles in real time. Figure 4 shows how the MPC solution $x_d(\cdot)$ can refine the graph solve v_k^* to improve the optimality of the path while maintaining collision avoidance. In Figure 5, a grid was used instead of random sampling due to the structured nature of the problem. In both of these settings, about 5000 nodes were sampled for the graph, leading to 50,000 edges to have to be processed every graph solve time step. In this setup, the (Cut-QP) takes 2.7 seconds to run per obstacle, even when limited to 50 solver iterations. When the heuristic runs, it is able to remove > 99% of the edges and leads to solve times of 10-20 ms.

For the hardware setup seen in Figure 6, in order to estimate obstacle locations, we had an overhead ZED 2 camera and used the SAM2 segmenting repository [36]. Once the experiment started, the segmenter was initialized with a single click per obstacle, and was able to parse obstacle locations for the remainder of the experiment, which were streamed over ROS to the hierarchical controller. SAM2 segmented 20 obstacles at 4 Hz, and the graph cut and MPC ran at 50 Hz and 200 Hz, respectively. We provided various cluttered environments to ARCHER, which was able to solve for feasible paths and traverse them in real time. In the bottom right of Figure 6, the real-time replanning of the graph can be seen as the obstacles were kicked around in the environment.

V. CONCLUSION

We proposed a layered control architecture for performing path planning through cluttered environments. By leveraging the properties of Bézier curves, a kinodynamically feasible graph for a planning model could be efficiently constructed, and could provide curves that were guaranteed to satisfy state and input constraints of the full order system when tracked. By leveraging the CPU and the GPU together, Algorithm (1) could efficiently be run in real time for extremely long-horizon tasks such as maze solving both in simulation and on hardware. Future work includes finding automated ways of determining the graph size and Bézier horizon length and using adaptive time scales for combining high performance and high precision tasks.

REFERENCES

- [1] N. J. Nilsson, "A Mobile Automaton: An Application of Artificial Intelligence Techniques.," Defense Technical Information Center, Fort Belvoir, VA, Tech. Rep., Jan. 1969. [Online]. Available: <http://www.dtic.mil/docs/citations/ADA459660>
- [2] J. Wang, T. Zhang, N. Ma, Z. Li, H. Ma, F. Meng, and M. Q.-H. Meng, "A survey of learning-based robot motion planning," *IET Cyber-Systems and Robotics*, vol. 3, no. 4, pp. 302–314, 2021, _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1049/csy2.12020>. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1049/csy2.12020>
- [3] A. H. Qureshi, Y. Miao, A. Simeonov, and M. C. Yip, "Motion Planning Networks: Bridging the Gap Between Learning-based and Classical Motion Planners," Jun. 2020, arXiv:1907.06013 [cs]. [Online]. Available: <http://arxiv.org/abs/1907.06013>
- [4] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, and C. Cadena, "From Perception to Decision: A Data-driven Approach to End-to-end Motion Planning for Autonomous Ground Robots," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, May 2017, pp. 1527–1533, arXiv:1609.07910 [cs]. [Online]. Available: <http://arxiv.org/abs/1609.07910>
- [5] M. J. Bency, A. H. Qureshi, and M. C. Yip, "Neural Path Planning: Fixed Time, Near-Optimal Path Generation via Oracle Imitation," Apr. 2019, arXiv:1904.11102 [cs]. [Online]. Available: <http://arxiv.org/abs/1904.11102>
- [6] E. Schmerling and M. Pavone, "Kinodynamic Planning," in *Encyclopedia of Robotics*, M. H. Ang, O. Khatib, and B. Siciliano, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2021, pp. 1–9.
- [7] 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.
- [8] 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.
- [9] 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. [Online]. Available: <http://ieeexplore.ieee.org/document/6631299/>
- [10] S. M. LaValle, "Rapidly-exploring random trees : A new tool for path planning," *The annual research report*, 1998.
- [11] A. Shkolnik, M. Walter, and R. Tedrake, "Reachability-guided sampling for planning under differential constraints," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 2859–2865.
- [12] J. A. Reeds and L. A. Shepp, "Optimal paths for a car that goes both forwards and backwards," *Pacific Journal of Mathematics*, Oct. 1990. [Online]. Available: <https://www.scinap.se/papers/1971998222>
- [13] D. Fridovich-Keil, S. L. Herbert, J. F. Fisac, S. Deglurkar, and C. J. Tomlin, "Planning, Fast and Slow: A Framework for Adaptive Real-Time Safe Trajectory Planning," Mar. 2018, arXiv:1710.04731 [cs]. [Online]. Available: <http://arxiv.org/abs/1710.04731>
- [14] A. Shkolnik, M. Walter, and R. Tedrake, "Reachability-guided sampling for planning under differential constraints," in *2009 IEEE International Conference on Robotics and Automation*, ser. 1, May 2009, pp. 2859–2865, iSSN: 1050-4729.
- [15] 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)*, ser. 1, May 2020, pp. 4245–4251, iSSN: 2577-087X.
- [16] D. Ioan, I. Prodan, S. Oлару, F. Stoican, and S.-I. Niculescu, "Mixed-integer programming in motion planning," *Annual Reviews in Control*, vol. 51, pp. 65–87, 2021.
- [17] R. Deits and R. Tedrake, "Efficient mixed-integer planning for UAVs in cluttered environments," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, ser. 1. Seattle, WA, USA: IEEE, May 2015, pp. 42–49. [Online]. Available: <http://ieeexplore.ieee.org/document/7138978/>
- [18] T. Marcucci, J. Umenberger, P. A. Parrilo, and R. Tedrake, "Shortest paths in graphs of convex sets," *arXiv:2101.11565*, 2021.
- [19] T. Marcucci, M. Petersen, D. von Wrangel, and R. Tedrake, "Motion Planning around Obstacles with Convex Optimization," May 2022, arXiv:2205.04422 [cs]. [Online]. Available: <http://arxiv.org/abs/2205.04422>
- [20] J. Lee, E. Bakolas, and L. Sontag, "An efficient and direct method for trajectory optimization of robots constrained by contact kinematics and forces," *Autonomous Robots*, vol. 45, pp. 135–153, 2021.
- [21] S. Stoneman and R. Lampariello, "Embedding nonlinear optimization in rrt for optimal kinodynamic planning," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 3737–3744.
- [22] P. Fernbach, S. Tonneau, A. Del Prete, and M. Taix, "A kinodynamic steering-method for legged multi-contact locomotion," in *2017 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 3701–3707.
- [23] 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.
- [24] 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]. [Online]. Available: <http://arxiv.org/abs/2305.01072>
- [25] 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," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, Dec. 2022, pp. 3732–3739, iSSN: 2576-2370.
- [26] E. R. Ambrose, "Creating ARCHER: A 3D Hopping Robot with Flywheels for Attitude Control," Ph.D. dissertation, California Institute of Technology, 2022.
- [27] T. Apgar, P. Clary, K. Green, A. Fern, and J. W. Hurst, "Fast online trajectory optimization for the bipedal robot cassie," in *Robotics: Science and Systems*, vol. 101, 2018, p. 14.
- [28] E. D. Sontag and Y. Wang, "On characterizations of the input-to-state stability property," *Systems & Control Letters*, vol. 24, no. 5, pp. 351–359, Apr. 1995. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/0167691194000506>
- [29] M. Kamermans, "A primer on bézier curves," (*online book*), 2020. [Online]. Available: <https://pomax.github.io/bezierinfo/>
- [30] N. Csomay-Shanklin and A. D. Ames, "Reachable bézier polytopes: A primitive for layered motion planning," *preprint*, 2024. [Online]. Available: https://noelc-s.github.io/website/research/2024/09/30/ACC_Bezier/
- [31] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017. [Online]. Available: <https://books.google.com/books?id=MrJctAEACAAJ>
- [32] "Supplemental video." [Online]. Available: <https://vimeo.com/1009702220>
- [33] "Code," 2024. [Online]. Available: <https://github.com/noelc-s/PathPlanner>
- [34] N. Csomay-Shanklin, V. D. Dorobantu, and A. D. Ames, "Nonlinear model predictive control of a 3d hopping robot: Leveraging lie group integrators for dynamically stable behaviors," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 12 106–12 112.
- [35] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020. [Online]. Available: <https://doi.org/10.1007/s12532-020-00179-2>
- [36] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," *arXiv:2304.02643*, 2023.