arXiv:2205.04422v1 [cs.RO] 9 May 2022

Motion Planning around Obstacles with Convex Optimization

Tobia Marcucci[1,*] Mark Petersen[2,*] David von Wrangel[1] Russ Tedrake[1,*]

[1]Massachusetts Institute of Technology, Cambridge MA, USA
{tobiam,wrangel,russt}@mit.edu
[2]Harvard University, Cambridge MA, USA
markpetersen@g.harvard.edu

Abstract

Trajectory optimization offers mature tools for motion planning in high-dimensional spaces under dynamic constraints. However, when facing complex configuration spaces, cluttered with obstacles, roboticists typically fall back to sampling-based planners that struggle in very high dimensions and with continuous differential constraints. Indeed, obstacles are the source of many textbook examples of problematic nonconvexities in the trajectory-optimization problem. Here we show that convex optimization can, in fact, be used to reliably plan trajectories around obstacles. Specifically, we consider planning problems with collision-avoidance constraints, as well as cost penalties and hard constraints on the shape, the duration, and the velocity of the trajectory. Combining the properties of Bézier curves with a recently-proposed framework for finding shortest paths in Graphs of Convex Sets (GCS), we formulate the planning problem as a compact mixed-integer optimization. In stark contrast with existing mixed-integer planners, the convex relaxation of our programs is very tight, and a cheap rounding of its solution is typically sufficient to design globally-optimal trajectories. This reduces the mixed-integer program back to a simple convex optimization, and automatically provides optimality bounds for the planned trajectories. We name the proposed planner GCS, after its underlying optimization framework. We demonstrate GCS in simulation on a variety of robotic platforms, including a quadrotor flying through buildings and a dual-arm manipulator (with fourteen degrees of freedom) moving in a confined space. Using numerical experiments on a seven-degree-of-freedom manipulator, we show that GCS can outperform widely-used sampling-based planners by finding higher-quality trajectories in less time.

1 Introduction

In this paper we consider the problem of designing continuous collision-free trajectories for robots moving in environments with obstacles. A wide array of techniques can be found in the literature to tackle this long-standing problem in robotics [14], and selecting the right one requires compromising between multiple features of the problem at hand: dimensionality and complexity of the environment, dynamic constraints, computation limits, completeness and optimality requirements.

*Contributed equally to this work.

Methods based on direct trajectory optimization [7, 2, 35, 25, 41] can design trajectories in high-dimensional spaces while taking into account the robot kinematics and dynamics. However, by transcribing the planning problem as a nonconvex program, and by relying on local optimization, these techniques can fail in finding a collision-free trajectory, especially if the robot configuration space is cluttered. In these scenarios, roboticists typically fall back to sampling-based planners [18, 22] (see also the more recent review [9]). These algorithms are *probabilistically complete*, meaning that, if a feasible path exists, they will eventually find one, regardless of the complexity of the configuration space [21, Chapter 5]. This guarantee, however, comes at a cost. Although many sampling-based planners support "kinodynamic" constraints [39, 12, 40], continuous differential constraints are difficult to impose on discrete samples, making the kinodynamic versions of the classical sampling-based algorithms much less successful in practice. In addition, even using asymptotically-optimal sampling-based planners [17, 11, 16], the trajectories we design can be considerably suboptimal in practice, where only a finite number of samples can be taken. For certain classes of dynamical systems, hybrid approaches, where a trajectory-optimization planner is driven by a higher-level graph search, have been shown to overcome part of these difficulties [29]. Still, these multi-layer architectures do not offer a unified formulation of the planning problem as a single optimization problem.

The promise of the planners based on Mixed-Integer Convex Programming (MICP) [33, 32, 28, 39, 6] is to take the best of the two worlds above: the completeness of sampling-based algorithms, and the ease with which trajectory optimization handles the robot kinematics and dynamics; with the added bonus of global optimality and within a single optimization framework. The spread of MICP techniques, however, is strongly limited by their runtimes: even for small-scale problems, these methods can require several minutes to design a trajectory. Only recently, collision-free planners entirely based on convex optimization have been proposed [8], but their application is currently limited to purely-geometric path planning in low-dimensional spaces.

In this paper, we focus on a limited but important class of motion-planning problems with differential constraints, and we present a planner that, although being based on MICP, reliably solves very high-dimensional problems in a few seconds, through a single convex program.

1.1 Contribution

We consider a formulation of the collision-free planning problem similar to the one from [6]. In particular, we assume the robot configuration space to be partitioned into a collection of "safe" convex regions, i.e., regions that do not intersect with any of the obstacles. In the special case of polygonal obstacles, this partition can be constructed exactly. More generally, approximate decompositions can be efficiently obtained using existing algorithms [3, 5], as well as newly-developed techniques tailored to the complex configuration spaces of multi-link kinematic trees [1]. Our goal is then to design a continuous trajectory that is entirely contained in the union of the safe regions. The optimality criterion and the additional constraints are allowed to depend on the shape, the duration, and the velocity of the trajectory.

The main technical contribution of this work is showing that the trajectory-design problem just described can be formulated as a shortest-path problem in Graphs of Convex Sets (GCS): a recently-studied class of optimizations that lends itself to very efficient mixed-integer programming [26]. Existing MICP planners parameterize a single trajectory and use binary variables to

assign each of its segments to a safe region. Conversely, with the proposed planner, which we name GCS, the safe regions are connected through an adjacency graph and are each assigned a trajectory segment. The optimal probabilities of transitioning between the regions are then computed via an efficient blending of convex and graph optimization. We show that the MICPs constructed in this way have very tight convex relaxations and, in the great majority of practical cases, a single convex program, together with a cheap rounding step, is sufficient to identify a globally-optimal collision-free trajectory. Furthermore, by comparing the costs of the convex relaxation and the rounded trajectory, GCS automatically provides a tight bound on the optimality of the motion plan.

To parameterize trajectories we use Bézier curves: a relatively common tool in motion planning (see e.g. [10, 20, 4]) whose properties are very well suited for mixed-integer programming [19]. This parameterization enables simple convex formulations of the collision-avoidance constraints and, when incorporated in our workflow, leads to very tractable convex optimizations; typically Second-Order-Cone Programs (SOCPs). This is in contrast with existing MICP planners, which require expensive semidefinite constraints to design trajectories that are differentiable more than three times [6]. (Note that the requirement of smooth trajectories is of practical nature: to exploit the differential-flatness properties of quadrotors, for example, it is necessary to design trajectories that are differentiable at least four times [27].)

We demonstrate GCS on a variety of planning problems, ranging from an intricate maze to a quadrotor flying through buildings and a fourteen-dimensional dual-arm manipulation task. The numerical results show that, besides significantly improving on state-of-the-art MICP planners, our relatively unoptimized implementation of GCS can also outperform widely-used sampling-based planners by finding higher-quality trajectories in lower, and consistent, runtimes.

2 Problem Statement

In this section we state the motion-planning problem addressed in this paper in abstract terms, as an optimization over the infinite-dimensional space of trajectories. It will be the goal of Section 5 to present our finite-dimensional transcription of this optimization, which will then be tackled using practical convex programming.

As in [6], we look at the problem of planning around obstacles as the problem of navigating within a collection of "safe" regions. More precisely, we assume the set $\mathcal{Q} \subset \mathbb{R}^n$ of collision-free robot configurations to be decomposed into a family of (possibly overlapping) bounded convex sets $\mathcal{Q}_i \subseteq \mathcal{Q}$, with i in a finite index set $\mathcal{I}$. For polyhedral obstacles this decomposition can be exact, i.e. $\bigcup_{i\in\mathcal{I}} \mathcal{Q}_i = \mathcal{Q}$, while more complex configuration spaces can be decomposed approximately using efficient existing algorithms [5, 1]. Given the regions $\mathcal{Q}_i$, our goal is to find a time $T \in \mathbb{R}_{>0}$ and a trajectory $q : [0, T] \to \mathcal{Q}$ that are a solution of the following optimization problem:[1]

$$\begin{aligned} \text{minimize} \quad & aT + bL(q, T) + cE(\dot{q}, T) && \text{(1a)} \\ \text{subject to} \quad & q \in \mathcal{C}^\eta, && \text{(1b)} \end{aligned}$$

[1] In Section 6 we show how penalties on the second and higher derivatives of q can be approximately integrated in our problem formulation. Further costs and constraints are discussed in Section 8.1.

$$q(t) \in \bigcup_{i \in \mathcal{I}} \mathcal{Q}_i, \qquad \forall t \in [0, T], \tag{1c}$$

$$\dot{q}(t) \in \mathcal{D}, \qquad \forall t \in [0, T], \tag{1d}$$

$$T \in [T_{\min}, T_{\max}], \tag{1e}$$

$$q(0) = q_0,\ q(T) = q_T, \tag{1f}$$

$$\dot{q}(0) = \dot{q}_0,\ \dot{q}(T) = \dot{q}_T. \tag{1g}$$

The objective is a weighted sum, with user-specified weights $a, b, c \in \mathbb{R}_{\geq 0}$, of the trajectory duration T, the length $L(q, T)$ of the trajectory, and the energy $E(\dot{q}, T)$ of the time derivative of the trajectory. Specifically, the latter two quantities are defined as

$$L(q, T) := \int_0^T \|\dot{q}(t)\|_2 dt \quad \text{and} \quad E(\dot{q}, T) := \int_0^T \|\dot{q}(t)\|_2^2 dt. \tag{2}$$

Constraint (1b) asks the trajectory to be continuously differentiable η times. Constraint (1c) ensures that q is contained in the safe sets, and hence is collision free at all times. (Note that this is a stronger constraint than is usual in sampling-based motion planning, where trajectories are typically checked to be collision-free only at a finite number of points.) The set $\mathcal{D}$ in (1d) is required to be convex and can be used to enforce hard limits on the robot velocity. The bounds on the trajectory duration in (1e) are such that $T_{\max} \geq T_{\min} > 0$. Finally, the constraints (1f) and (1g) enforce the boundary conditions on q and its time derivative.

The coupling between the trajectory q and its duration T makes it hard to work with problem (1) directly. Similarly to [38], we break this coupling by introducing the path coordinate $s \in [0, S]$, where S has fixed positive value. We relate the coordinate s to the time variable t via the scaling function $t = h(s)$: the map h is required to be monotonically increasing, and such that $h(0) = 0$ and $h(S) = T$. Expressing the trajectory q as a function of s, we get the curve $r(s) := q(h(s))$. Through a few simple manipulations, we restate problem (1) in terms of the decision variables r and h as

$$\text{minimize} \quad ah(S) + bL(r, S) + cE\left(\dot{r}/\sqrt{\dot{h}}, S\right) \tag{3a}$$

$$\text{subject to} \quad r \circ h^{-1} \in \mathcal{C}^\eta, \tag{3b}$$

$$r(s) \in \bigcup_{i \in \mathcal{I}} \mathcal{Q}_i, \qquad \forall s \in [0, S], \tag{3c}$$

$$\dot{r}(s) \in \dot{h}(s)\mathcal{D},\ \dot{h}(s) > 0, \qquad \forall s \in [0, S], \tag{3d}$$

$$h(0) = 0,\ h(S) \in [T_{\min}, T_{\max}], \tag{3e}$$

$$r(0) = q_0,\ r(S) = q_T, \tag{3f}$$

$$\dot{r}(0) = \dot{h}(0)\dot{q}_0,\ \dot{r}(S) = \dot{h}(S)\dot{q}_T. \tag{3g}$$

In particular, we have used the chain rule to substitute $\dot{q}(t)$ with $\dot{r}(s)/\dot{h}(s)$, and we have changed integration variable in (2) from t to s. This makes $\dot{r}/\sqrt{\dot{h}} : [0, S] \to \mathbb{R}^n$ the argument of the energy function in the objective. The symbol $\circ$ in (3b) denotes the composition operator: notice that the function h is guaranteed to be invertible by the positivity of $\dot{h}$ from (3d). Finally, again by the chain rule, the right-hand sides of the velocity constraints in (3d) and (3g) are multiplied by the derivative $\dot{h}$ of the time scaling.

3 Background on Bézier Curves

In order to tackle problem (3) numerically, it is necessary to parameterize the functions r and h through a finite number of decision variables. To this end, in Section 5, we will employ Bézier curves. The goal of this section is to recall the definition and the basic properties of this family of curves.

A Bézier curve is constructed using Bernstein polynomials. The kth Bernstein polynomial of degree d, with $k = 0, \dots, d$, is defined as

$$\beta_{k,d}(s) := \binom{d}{k} s^k (1-s)^{d-k},$$

where $s \in [0, 1]$. Note that the Bernstein polynomials of degree d are nonnegative and, by the binomial theorem, they sum up to one. Therefore, for each fixed $s \in [0, 1]$, the scalars $\{\beta_{k,d}(s)\}_{k=0}^{d}$ can be thought of as the coefficients of a convex combination. Bézier curves are obtained using these coefficients to combine a given set of $d+1$ *control points* $\gamma_k \in \mathbb{R}^n$:

$$\gamma(s) := \sum_{k=0}^{d} \beta_{k,d}(s)\gamma_k.$$

It is easily verified that Bézier curves enjoy the following properties.

- *Endpoint values.* The curve γ starts at the first control point and ends at the last control point: $\gamma(0) = \gamma_0$ and $\gamma(1) = \gamma_d$.
- *Convex hull.* The curve γ is entirely contained in the convex hull of its control points: $\gamma(s) \in \text{conv}(\{\gamma_k\}_{k=0}^{d})$ for all $s \in [0, 1]$.
- *Derivative.* The derivative $\dot{\gamma}$ of the curve γ is a Bézier curve of degree $d-1$ with control points $\dot{\gamma}_k = d(\gamma_{k+1} - \gamma_k)$ for $k = 0, \dots, d-1$.
- *Integral of convex function.* For a convex function $f : \mathbb{R}^n \to \mathbb{R}$, we have[2]

$$\int_0^1 f(\gamma(s))ds \leq \frac{1}{d+1} \sum_{k=0}^{d} f(\gamma_k). \tag{4}$$

4 The Optimization Framework

Our strategy for solving problem (3) is to first transcribe it as a Shortest-Path Problem (SPP) in GCS, and then use the techniques recently presented in [26] to formulate this SPP as a compact MICP. As we will see in Section 7, the convex relaxation of this MICP is extremely tight in practice, up to the point that a cheap rounding of its solution is almost always sufficient to design a globally-optimal trajectory. In this section, we give a formal statement of the SPP in GCS and we propose a simple randomized rounding for the convex relaxation of our MICP. The latter will effectively reduce the computational cost of the MICP to that of a convex program.

[2]To prove (4), one uses the convexity of f, which gives $f(\gamma(s)) \leq \sum_{k=0}^{d} \beta_{k,d}(s) f(\gamma_k)$, and the formula $\int_0^1 \beta_{k,d}(s)ds = 1/(d+1)$ for the integration of Bernstein polynomials.

4.1 Shortest Paths in Graphs of Convex Sets

The Shortest-Path Problem (SPP) in GCS generalizes the classical SPP with nonnegative edge lengths. We are given a directed graph $G := (\mathcal{V}, \mathcal{E})$ with vertex set $\mathcal{V}$ and edge set $\mathcal{E} \subset \mathcal{V}^2$. Each vertex $v \in \mathcal{V}$ is paired with a bounded convex set $\mathcal{X}_v$, and a point x_v contained in it. In contrast to the classical SPP, where edge lengths are fixed scalars, here the length of an edge $e = (u, v)$ is determined by the continuous values of x_u and x_v via the expression $\ell_e(x_u, x_v)$. The function ℓ_e is assumed to be convex and to take nonnegative values. Convex constraints of the form $(x_u, x_v) \in \mathcal{X}_e$ are allowed to couple the endpoints of edge $e := (u, v)$. A path p in the graph G is defined as a sequence of distinct vertices that connects the source vertex $\sigma \in \mathcal{V}$ to the target vertex $\tau \in \mathcal{V}$. Denoting with $\mathcal{E}_p$ the set of edges traversed by the path p, and with $\mathcal{P}$ the family of all σ-τ paths in the graph G, the SPP in graphs of convex sets is stated as

$$\begin{aligned}
\text{minimize} \quad & \sum_{e:=(u,v)\in\mathcal{E}_p} \ell_e(x_u, x_v) && && \text{(5a)}\\
\text{subject to} \quad & p \in \mathcal{P}, && && \text{(5b)}\\
& x_v \in \mathcal{X}_v, && \forall v \in p, && \text{(5c)}\\
& (x_u, x_v) \in \mathcal{X}_e, && \forall e := (u, v) \in \mathcal{E}_p. && \text{(5d)}
\end{aligned}$$

Here the decision variables are the discrete path p and the continuous values x_v. The objective (5a) minimizes the length of the path p, defined as the sum of the lengths of its edges. Constraint (5b) asks p to be a valid path connecting σ to τ. Importantly, the convex conditions (5c) and (5d) constrain only the continuous variables paired with the vertices visited by the path p, and do not apply to the remaining vertices in the graph. Unlike the classical SPP with nonnegative edge lengths, which is easily solvable in polynomial time, the SPP in GCS can be verified to be NP-hard [26, Theorem 1].

4.2 Rounding the Convex Relaxation of the Shortest-Path Problem

Using recently-developed techniques, problem (5) is formulated as a compact MICP with very tight convex relaxation [26, Equation 21]. In this paper, instead of tackling this MICP with an exact branch-and-bound algorithm, we solve its convex relaxation and we recover an approximate solution via a cheap randomized rounding, that is tailored to the graph structure beneath problem (5). Given the hardness of (5), this approach cannot be guaranteed to work for all instances. Nevertheless, for our planning problems, this strategy turns out to be extremely effective in practice. In addition, this workflow automatically provides us with a bound on the optimality of the approximate solution we identify. In fact, denoting with C_{relax} the cost of the convex relaxation, with C_{opt} the optimal value of (5), and with C_{round} the cost of the rounded solution, we have $C_{\text{relax}} \leq C_{\text{opt}} \leq C_{\text{round}}$. The optimality gap of the rounded solution $\delta_{\text{opt}} := (C_{\text{round}} - C_{\text{opt}})/C_{\text{opt}}$ can be then overestimated as $\delta_{\text{relax}} := (C_{\text{round}} - C_{\text{relax}})/C_{\text{relax}}$ with no additional computation.

For the rounding step we propose a randomized strategy. The MICP from [26] parameterizes a path p by using a binary variable φ_e per edge $e \in \mathcal{E}$, with $\varphi_e = 1$ if and only if $e \in \mathcal{E}_p$. In the convex relaxation, the binary requirement is relaxed to $\varphi_e \in [0, 1]$ and the optimal value of φ_e is naturally interpreted as the probability of the edge e being a part of the shortest path. To round

these probabilities we then run a randomized depth-first search with backtracking. We initialize our candidate path as $p := (\sigma)$, and we denote with $\mathcal{E}_u$ the set of edges $e := (u, v)$ that connect u to a vertex v that the rounding algorithm has not visited yet. At each iteration, calling u the last vertex in the path p, we traverse the edge $e := (u, v) \in \mathcal{E}_u$ with probability $\varphi_e / \sum_{e' \in \mathcal{E}_u} \varphi_{e'}$, and we append a new vertex v to the path p. If a dead end occurs, i.e. if $\varphi_e = 0$ for all $e \in \mathcal{E}_u$, we backtrack to the last vertex in p that admits a way out. The algorithm terminates when $v = \tau$ and the target is reached.[3] Once a path p is identified, its cost, together with the optimal values of the continuous variables x_v, is recovered by solving a small convex program:

$$\text{minimize (5a) subject to (5c) and (5d).} \tag{6}$$

It is easily verified that this rounding strategy always finds a valid path (provided that the convex relaxation of the MICP is feasible). On the other hand, the cost of the path p we find can in principle be infinite, since there might not be an assignment for the continuous variables $\{x_v\}_{v \in p}$ that satisfies the constraints (5c) and (5d).

To increase our chances of finding a high-quality approximate solution, we apply the randomized rounding multiple times. First we run the depth-first search until N distinct paths are identified, or a maximum number M of trials is reached. Then we evaluate the cost of each distinct path by solving a convex program of the form (6), and we return the rounded solution of lowest cost C_{round}.[4] We emphasize that this process is extremely cheap: the runtime of a depth-first search is practically zero (since it is a purely-discrete search in the graph G), while the convex programs (6) are tiny, very sparse, and parallelizable. In this paper we set $N := 10$ and $M := 100$. These values lead to rounding times that are negligible with respect to the solution time of the convex relaxation and, in our experiments, they are typically sufficient to solve the planning problem to global optimality.

Many more details on the MICP formulation of (5) and its convex relaxation can be found in [26]. For the scope of this paper, we will treat the framework from [26] as a modeling language that allows us to formulate, and efficiently solve, an SPP in GCS just by providing the graph G, the edge lengths ℓ_e, and the sets $\mathcal{X}_v$ and $\mathcal{X}_e$.

5 Collision-Free Motion Planning using Graphs of Convex Sets

We now illustrate how problem (3) can be transcribed as an SPP in GCS. As seen in the previous section, to formulate an SPP in GCS we need to: define a graph $G := (\mathcal{V}, \mathcal{E})$, assign a set $\mathcal{X}_v$ to each vertex $v \in \mathcal{V}$, and pair each edge $e \in \mathcal{E}$ with a constraint set $\mathcal{X}_e$ and a length function ℓ_e. Below we describe how each of these components is constructed. At a high level, the plan is to pair each safe region $\mathcal{Q}_i$ with two Bézier curves: a trajectory segment r_i, and a time-scaling

[3] Making this rounding strategy deterministic by, e.g., selecting at each iteration the edge $e \in \mathcal{E}_u$ with larger probability φ_e is, in general, a bad idea. To see this, imagine a graph where multiple paths represent the same underlying decision (e.g. multiple symmetrical solutions). Since the convex relaxation will equally split the probability of this decision being optimal between the edges of these many paths, a greedy deterministic search might end up selecting an alternative path, corresponding to a decision that is overall less likely to be optimal. Conversely, in the same scenario, a randomized rounding correctly weights the two decisions (in expectation).

[4] This sequence of convex optimizations is stopped early if the cost of a path coincides with the cost of the convex relaxation C_{relax}, since this proves the global optimality of the path at hand.

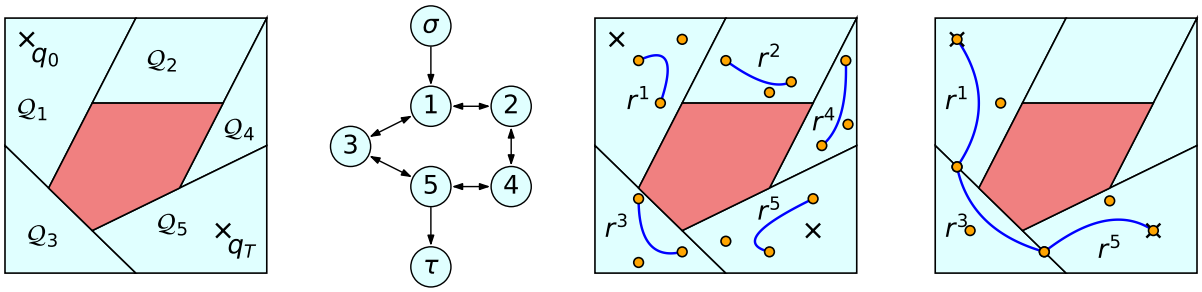

(a) (b) (c) (d)

Figure 1: Formulation of the collision-free motion-planning problem as an SPP in GCS. The red region is the obstacle to be avoided, the light-blue regions $\mathcal{Q}_i$ partition the free space. (a) The collision-free regions $\mathcal{Q}_i$, the starting point q_0, and the ending point q_T. (b) The graph G obtained by connecting intersecting regions, with the source vertex σ and the target vertex τ added to enforce the initial and terminal conditions, respectively. (c) The Bézier curves r^i associated with each region (curves in blue, control points $r_{i,k}$ in orange). (d) A continuous collision-free trajectory r corresponding to the path $p := (\sigma, 1, 3, 5, \tau)$.

function h_i that dictates the speed at which the curve r_i is traveled. The functions r and h in problem (3) will be then reconstructed by sequencing the Bézier curves r_i and h_i paired with the regions $\mathcal{Q}_i$ that are selected by the SPP. Figure 1 provides a visual support to the upcoming discussion.

5.1 The Graph G

We let the vertex set $\mathcal{V}$ contain a vertex i per safe set $\mathcal{Q}_i$ in the decomposition of the configuration space.[5] In addition, we introduce a source vertex σ and a target vertex τ: these will be used to enforce the boundary conditions (3e)–(3g). Overall, we then have $\mathcal{V} := \mathcal{I} \cup \{\sigma, \tau\}$.

We include in the edge set $\mathcal{E}$ all the edges (i, j) such that the intersection of $\mathcal{Q}_i$ and $\mathcal{Q}_j$ is nonempty. Note that, by the symmetry of this condition, $(i, j) \in \mathcal{E}$ implies $(j, i) \in \mathcal{E}$. Similarly, we let $(\sigma, i) \in \mathcal{E}$ and $(i, \tau) \in \mathcal{E}$ if the set $\mathcal{Q}_i$ contains the points q_0 and q_T, respectively. In symbols,

$$\mathcal{E} := \{(i, j) : \mathcal{Q}_i \cap \mathcal{Q}_j \neq \emptyset\} \cup \{(\sigma, i) : q_0 \in \mathcal{Q}_i\} \cup \{(i, \tau) : q_T \in \mathcal{Q}_i\}.$$

Figure 1b shows the graph corresponding to the collision-free regions $\mathcal{Q}_i$, the staring point q_0, and the ending point q_T depicted in Figure 1a.

5.2 The Convex Sets $\mathcal{X}_v$

The source σ and the target τ are auxiliary vertices used to enforce the boundary conditions (3e)–(3g); they require no decision variables and can be safely paired with the empty set $\mathcal{X}_\sigma := \mathcal{X}_\tau := \emptyset$.

[5] As we will see in Section 5.2, the safe set $\mathcal{Q}_i$ does not coincide with the convex set $\mathcal{X}_i$ paired with vertex i in the SPP.

To each of the vertices $i \in \mathcal{I}$, we assign two Bézier curves: $r_i : [0,1] \to \mathcal{Q}_i$ (depicted in Figure 1c) and $h_i : [0,1] \to [0, T_{\max}]$. Both these curves have a user-defined degree $d \geq \eta + 1$, where η is the required degree of differentiability of the overall trajectory q.[6] The convex set $\mathcal{X}_i$ contains the control points of the two curves, i.e. $x_i := (r_{i,0}, \ldots, r_{i,d}, h_{i,0}, \ldots, h_{i,d})$, and is defined by the following conditions:

$$r_{i,k} \in \mathcal{Q}_i, \qquad k = 0, \ldots, d, \tag{7a}$$

$$\dot{h}_{i,k} \geq \dot{h}_{\min}, \qquad k = 0, \ldots, d-1, \tag{7b}$$

$$\dot{r}_{i,k} \in \dot{h}_{i,k} \mathcal{D}, \qquad k = 0, \ldots, d-1, \tag{7c}$$

$$h_{i,0} \geq 0, \; h_{i,d} \leq T_{\max}, \tag{7d}$$

The convex constraint (7a) requires all the control points of r_i to lie in the collision-free set $\mathcal{Q}_i$. By the convex-hull property of the Bézier curves from Section 3, this implies that the whole trajectory segment r_i is contained in $\mathcal{Q}_i$. Again by the properties of Bézier curves, the derivative $\dot{h}_i$ of the time scaling h_i is itself a Bézier curve. Condition (7b) lower bounds each control point of this derivative with a small positive constant $\dot{h}_{\min}$ which, unless differently specified, is set to 10^{-6}. By the convex-hull property, this implies that $\dot{h}_i(s)$ is positive for all $s \in [0,1]$, and hence that h_i is strictly increasing. Since the control points of $\dot{h}_i$ are linear functions of the ones of h_i, constraint (7b) is linear in x_i. Using the definition of convexity and the positivity of $\dot{h}_{i,k}$, condition (7c) can be verified to be convex in $\dot{r}_{i,k}$ and $\dot{h}_{i,k}$: this ensures that $\dot{r}_i(s) \in \dot{h}_i(s)\mathcal{D}$ for all $s \in [0,1]$, since the $(n+1)$-dimensional Bézier curve $(\dot{r}_i, \dot{h}_i)$ is a convex combination of the control points $(\dot{r}_{i,k}, \dot{h}_{i,k})$. Finally, the constraints in (7d) are conservative bounds that ensure the boundedness of $\mathcal{X}_i$, as assumed in Section 4.1.

We remark that asking the control points of a Bézier curve to be in a convex set is only a sufficient condition for the containment of the whole curve. Nonetheless, the conservativeness of the conditions in (7) can be attenuated by increasing the degree of the curves r_i and h_i.

5.3 The Convex Sets $\mathcal{X}_e$

The first role of the edge constraints is to impose the boundary conditions (3e)–(3g). To this end, for all edges $e := (\sigma, i) \in \mathcal{E}$, we define $\mathcal{X}_e$ through the conditions $r_{i,0} = q_0$, $\dot{r}_{i,0} = \dot{h}_{i,0}\dot{q}_0$, and $h_{i,0} = 0$. Given the endpoint property of Bézier curves, these linear constraints on the vector x_i imply $r_i(0) = q_0$, $\dot{r}_i(0) = \dot{h}_i(0)\dot{q}_0$, and $h_i(0) = 0$. Similarly, for all the edges $e := (i, \tau)$, we define $\mathcal{X}_e$ via $r_{i,d} = q_T$, $\dot{r}_{i,d-1} = \dot{h}_{i,d-1}\dot{q}_T$, and $h_{i,d} \in [T_{\min}, T_{\max}]$. For these edges, we then have $r_i(1) = q_T$, $\dot{r}_i(1) = \dot{h}_i(1)\dot{q}_T$, and $h_i(1) \in [T_{\min}, T_{\max}]$.

The second role of the edge constraints is to enforce the differentiability of the overall curves r and h. For all the edges $e := (i,j) \in \mathcal{E} \cap \mathcal{I}^2$, we then define $\mathcal{X}_e$ through the following linear equalities:

$$r_{i,d-l}^{(l)} = r_{j,0}^{(l)} \quad \text{and} \quad h_{i,d-l}^{(l)} = h_{j,0}^{(l)}, \qquad l = 0, \ldots, \eta. \tag{8}$$

[6] The assumption that the curves r and h have the same degree is without loss of generality. The *degree elevation* property of Bézier curves allows us to describe a Bézier curve γ of degree d as a Bézier curve γ' of arbitrary degree $d' \geq d$, with control points that are linear functions of the control points of γ. Any convex cost or constraint on the control points of r and h, that takes advantage of the equal degree of these curves, can then be mapped to an equivalent convex cost or constraint on the control points of curves r and h of different degree.

Here $r_{i,k}^{(l)}$ denotes the kth control point of the lth derivative of r_i. In particular, $r_{i,k}^{(0)} = r_{i,k}$, $r_{i,k}^{(1)} = \dot{r}_{i,k}$, and so on. The same notation is used for h_i.

5.4 The Edge Lengths ℓ_e

The edge lengths ℓ_e must reproduce the cost in (3a) by appropriately weighting the cost of each transition in the graph G. This is achieved by assigning to each edge (σ, i) outgoing the source a length of zero, and to each edge (i, j) or (i, τ) the length

$$a(h_i(1) - h_i(0)) + bL(r_i, 1) + cE\left(\dot{r}_i/\sqrt{\dot{h}_i}, 1\right). \tag{9}$$

While the first term in this sum is immediately restated as the linear cost $a(h_{i,d} - h_{i,0})$, the other two terms require more work to be expressed as convex functions of x_i that are amenable to efficient numerical optimization.

One option is to approximate to arbitrary precision the last two terms in (9) using numerical integration. Since both L and E can be verified to be convex in the functions r_i and h_i, the resulting expression would be convex in x_i, but its numerical implementation would require a large number of second-order-cone constraints, proportional to the density of the integration grid. Instead, we prefer to minimize the following upper bounds of the last two terms in (9):

$$L(r_i, 1) \le \sum_{k=0}^{d-1} \|r_{i,k+1} - r_{i,k}\|_2 \quad \text{and} \quad E\left(\dot{r}_i/\sqrt{\dot{h}_i}, 1\right) \le \sum_{k=0}^{d-1} \frac{\|r_{i,k+1} - r_{i,k}\|_2^2}{h_{i,k+1} - h_{i,k}}. \tag{10}$$

The first inequality overestimates the length of r_i by summing the distances between its control points. The validity of this bound can be verified by applying inequality (4) to the Bézier curve $\dot{r}_i$ and the convex function $\|\dot{r}_i\|_2$. The second inequality does a similar operation with the energy E, and can be checked by applying (4) to the Bézier curve $(\dot{r}_i, \dot{h}_i)$ and the function $\|\dot{r}_i\|_2^2/\dot{h}_i$, which is convex for $\dot{h}_i > 0$.

5.5 Reconstruction of a Collision-Free Trajectory

Once the SPP (5) is solved, the optimal path p determines the sequence of safe regions $\mathcal{Q}_i$ that the robot must traverse. To reconstruct the trajectory r and the time scaling h, we sequence the Bézier curves r_i and h_i associated with these regions, as shown in Figure 1d for $\eta := 0$. Precisely, if the optimal path is $p := (\sigma, i_0, \dots, i_{S-1}, \tau)$, for $\nu = 0, \dots, S-1$, we define

$$r(s) := r_{i_\nu}(s - \nu) \quad \text{and} \quad h(s) := h_{i_\nu}(s - \nu), \qquad \forall s \in [\nu, \nu + 1].$$

Let us verify that the functions r and h just defined form an optimal solution of problem (3), up to the conservativeness of the constraints (7) and the cost bounds (10). The constraints in (8) imply $r, h \in \mathcal{C}^\eta$. This, in turn, gives $h^{-1} \in \mathcal{C}^\eta$ and (3b). That the collision-avoidance constraint (3c) is met, is implied by (7a). The velocity constraint in (3d) is verified thanks to (7c), while (7b) ensures that the function h is monotonically increasing, as required by the second condition in (3d). The boundary conditions (3e)–(3g) are verified because of the constraints on the edges (σ, i) and (i, τ) described in Section 5.3. Finally, summing the edge lengths (9) for all the edges traversed by the path p we get back (3a).

5.6 Class of Optimization Problems

Let us conclude this section by highlighting that, by feeding the SPP in GCS we just constructed to the machinery from [26], we obtain very tractable optimization problems. The framework in [26] relies on *perspective functions* [13, Section IV.2.2] to handle the interplay between the discrete and continuous components of problem (5). These are used to effectively "turn off" the edge costs ℓ_e and the convex constraints $(x_u, x_v) \in \mathcal{X}_e$ and $x_v \in \mathcal{X}_v$ corresponding to the edges e and the vertices v that do not lie along the path p, as required in problem (5). In case of polytopic safe sets $\mathcal{Q}_i$ and a purely-minimum-time objective ($a := 1$ and $b := c := 0$), it can be verified that the perspective operations lead to a mixed-integer Linear Program (LP), which, as discussed in Section 4.2, we tackle as a simple LP followed by a rounding stage. More generally, when the safe sets $\mathcal{Q}_i$ are quadratics or the cost weights b and c are nonzero, we obtain a mixed-integer SOCP, which we solve as a single SOCP plus rounding. In both cases, we then have simple convex optimizations for which very efficient solvers are available (e.g. MOSEK and Gurobi). Conversely, in order to design trajectories that are differentiable more than three times, existing MICP planners formulate prohibitive mixed-integer semidefinite programs that cannot be tackled with common solvers [6].

6 Penalties on the Higher-Order Derivatives of the Trajectory

In many practical applications, we find the need to expand our problem formulation (1) to include convex penalties on the second and higher time derivatives of the trajectory q. These can be used, for example, to indirectly limit the control efforts: for a robot manipulator, in fact, the joint torques needed to execute a trajectory q are proportional to the acceleration $\ddot{q}$ via the inertia matrix; while for a quadrotor the differential-flatness property makes the thrusts a function of the snap $q^{(4)}$ [27]. Unfortunately, even though convex in q, these costs become nonconvex when in problem (3) we optimize jointly over the shape r and the time scaling h of our trajectories. While we are currently working on the design of tight convex approximations of these costs, in this section we show how simple regularization terms can be added to our SPP in GCS to prevent the higher-order derivatives of q from growing excessively in magnitude.

For simplicity, let us first consider the regularization of the second derivative. Using the chain rule, we express the acceleration $\ddot{q}$ in terms of the derivatives of r and h:

$$\ddot{q}(t) = \frac{\ddot{r}(s) - \dot{q}(t)\ddot{h}(s)}{\dot{h}(s)^2},$$

where $\dot{q}(t) = \dot{r}(s)/\dot{h}(s)$ and $s = h^{-1}(t)$. Using this expression, we see that a convex function of $\ddot{q}$ does not, in general, translate into a convex function of r and h, and hence it cannot be directly minimized in our programs. However, provided that we choose a bounded set $\mathcal{D}$ to constrain the velocity $\dot{q}$, the magnitude of $\ddot{q}$ can be kept under control by increasing the minimum value $\dot{h}_{\min}$ of $\dot{h}(s)$ and by penalizing the magnitudes of $\ddot{r}$ and $\ddot{h}$. Letting ε be a small positive scalar, a simple way to achieve the latter is the cost term

$$\varepsilon(E(\ddot{r}, S) + E(\ddot{h}, S)), \tag{11}$$

which can be enforced using the ideas from Section 5.4.

The regularization of the higher-order derivatives of q follows the same logic. Specifically, using Faà di Bruno's formula for differentiating composite functions, we see that the magnitude of $q^{(m)}$ can be regularized by increasing $\dot{h}_{\min}$ and by penalizing the magnitudes of $r^{(l)}$ and $h^{(l)}$, for $l = 2, \dots, m$. The numerical results in the next section show that, even though these regularization terms are not as tight as the velocity bounds in (10), they can sensibly smooth the trajectories we design, while only minimally affecting their cost.

7 Numerical Results

We demonstrate the effectiveness of GCS on a variety of numerical examples. In Section 7.1 we analyze a simple two-dimensional problem, and we illustrate how the different components of problem (1) affect the shape of the trajectories we design. In Section 7.2 we increase the environment complexity and we apply our algorithm to design paths across an intricate maze. In Section 7.3, we run a statistical analysis of the performance of GCS on the task of planning the flight of a quadrotor through randomly-generated buildings. In Section 7.4, we show that, with respect to widely-used sampling-based algorithms, our algorithm is capable of designing higher-quality trajectories in less runtime. Finally, in Section 7.5, we demonstrate the scalability of GCS with a bimanual manipulation problem in a fourteen-dimensional configuration space.

The code necessary to reproduce all the results presented in this section can be found at `https://github.com/mpetersen94/gcs`. It uses an implementation of the SPP in GCS provided by Drake [36]. In addition to the techniques presented in [26], the convex optimizations we solve in this paper feature additional tightening constraints, tailored to the structure of the graphs in our planning problems, and a pre-processing step that eliminates the redundancies in our graphs. These are described in detail in Appendix A. The optimization solver used for the numerical experiments is MOSEK 9.2. All experiments are run on a desktop computer with an Intel Core i7-6950X processor and 64 GB of memory.

7.1 Two-Dimensional Example

The goal of our first numerical example is to illustrate how the different parameters in problem (1) affect the shape of the trajectories we design. To this end, we consider the simple two-dimensional environment depicted in Figure 2a. The initial $q_0 := (0.2, 0.2)$ and final $q_T := (4.8, 4.8)$ configurations are marked with a black cross; the obstacles are the red polygons. The convex decomposition $\{\mathcal{Q}_i\}_{i\in\mathcal{I}}$ of the free space $\mathcal{Q}$ is depicted in light blue in Figure 2b.

The first planning problem we analyze asks to minimize the total Euclidean length of the trajectory. The weights in the objective (1a) are then $a := c := 0$ and $b := 1$. The trajectory q is only required to be continuous ($\eta := 0$), while velocity and time constraints are irrelevant for a minimum-length problem. We let the degree of the Bézier curves r and h be $d := 1$ (i.e. straight lines). Solving the convex relaxation of the SPP in GCS we obtain the cost $C_{\text{relax}} = 10.77$, while the rounding step from Section 4.2 gives us the feasible trajectory depicted in Figure 3a with cost $C_{\text{round}} = 10.96$. By comparing these two numbers, GCS automatically provides the optimality bound $\delta_{\text{relax}} := (C_{\text{round}} - C_{\text{relax}})/C_{\text{relax}} = 1.7\%$ for the rounded solution. However, by actually

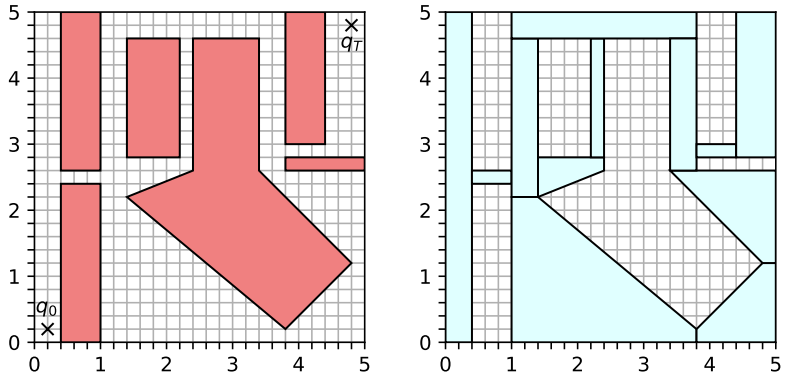

Figure 2: Two-dimensional trajectory-design problem from Section 7.1. (a) Environment with obstacles in red; the initial q_0 and final q_T configurations are marked with crosses. (b) Free space decomposed in convex safe regions $\mathcal{Q}_i$ (in light blue).

running a slightly more expensive mixed-integer solver, it is possible to verify that the rounded solution is indeed the global minimizer: $C_{\text{round}} = C_{\text{opt}}$ and $\delta_{\text{opt}} := (C_{\text{round}} - C_{\text{opt}})/C_{\text{opt}} = 0\%$.

For the second scenario, we consider a minimum-time problem with velocity limits. The weights in problem (1) are set to $a := 1$ and $b := c := 0$. We look for a continuous trajectory ($\eta := 0$), whose velocity $\dot{q}$ is contained in the box $\mathcal{D} := [-1, 1]^2$ for all times t. The bounds on the trajectory duration are set to $T_{\min} \approx 0$ and $T_{\max} \gg 0$, so that they do not affect the optimization problem. For this problem we let the optimizer decide the initial $\dot{q}(0)$ and final $\dot{q}(T)$ velocities by dropping the boundary conditions (1g). Once again, we use Bézier curves of degree $d := 1$. The trajectory generated by GCS is illustrated in Figure 3b. The convex relaxation has cost $C_{\text{relax}} = 9.88$, while the rounded trajectory has duration $C_{\text{round}} = 10.60$ and, therefore, it is certified to be within $\delta_{\text{relax}} = 7.3\%$ of the global minimum. As before, a mixed-integer solver can be used to verify that the trajectory generated by GCS is actually globally optimal ($\delta_{\text{opt}} = 0\%$).

In juxtaposition to the minimum-length case, the minimum-time trajectory in Figure 3b passes below the central obstacle. This is because, although shorter, the trajectory in Figure 3a is everywhere almost horizontal or vertical, and in these directions the speed is limited by the constraint set $\mathcal{D}$ to $\|\dot{q}\|_2 \leq 1$. The route below the obstacle is slightly longer, but it allows diagonal motion with speed $\|\dot{q}\|_2 \leq \sqrt{2}$. In Figure 4a we report the velocity $\dot{q}$ corresponding to the minimum-time trajectory: as expected, the optimal velocity is discontinuous and, at all times t, either the horizontal or the vertical component of $\dot{q}$ reaches the upper bound of 1.

Finally, we show the effects of the regularization strategy discussed in Section 6 on the smoothness of the minimum-time trajectory. We require the curve q to be twice continuously differentiable ($\eta := 2$). The initial and final velocities are forced to be zero, $\dot{q}_0 := \dot{q}_T := 0$, and we set the degree of the Bézier curves to $d := 6$. We increase $\dot{h}_{\min}$ from its default value of 10^{-6} to 10^{-1}, and we add the penalty (11) with weight $\varepsilon = 10^{-1}$. The resulting trajectory is reported in Figures 3c and 4b. As can be seen, the regularization smooths the optimal trajectory significantly and even changes its homotopy class. The costs of the convex relaxation and the rounded solution increase

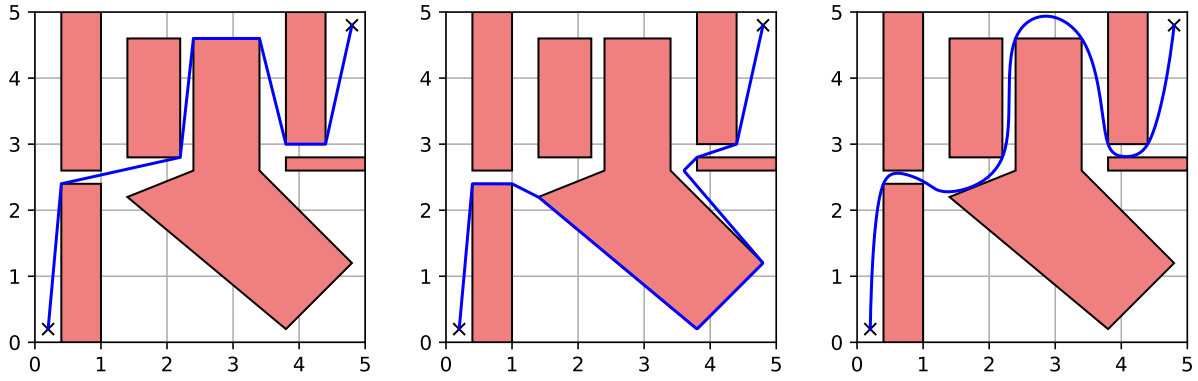

(a) (b) (c)

Figure 3: Trajectories (blue) designed by GCS for the planning problems in Section 7.1. (a) Minimum-length objective. (b) Minimum-time objective with velocity limits $\dot{q} \in [-1, 1]^2$. (c) Minimum-time objective with velocity limits, differentiability constraint $q \in \mathcal{C}^2$, and regularized acceleration. GCS finds the globally-optimal trajectory ($\delta_{\text{opt}} = 0\%$) for each of these tasks by rounding the solution of a single convex program. With no additional computation, it also certifies the following optimality gaps δ_{relax} for the rounded solutions: 1.7% for (a), 7.3% for (b), and 3.0% for (c).

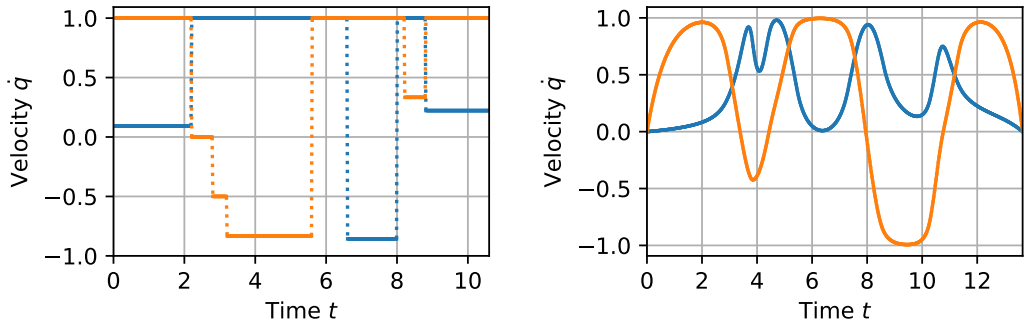

(a) (b)

Figure 4: (a) Velocity profile for the minimum-time trajectory depicted in Figure 3b, with dotted lines representing discontinuities. (b) Velocity profile for the smoothed trajectory in Figure 3c. The horizontal component of $\dot{q}$ is blue, the vertical is orange. In both the problems, the velocity components are constrained to lie in the interval $[-1, 1]$.

to $C_{\text{relax}} = 27.29$ and $C_{\text{round}} = 28.10$, respectively. The optimality gap certified by GCS is hence $\delta_{\text{relax}} = 3.0\%$, but, once again, a mixed-integer solver can be used to verify that the rounded solution is actually globally optimal ($\delta_{\text{opt}} = 0\%$). The duration of the smoothed trajectory is $T = 13.65$.

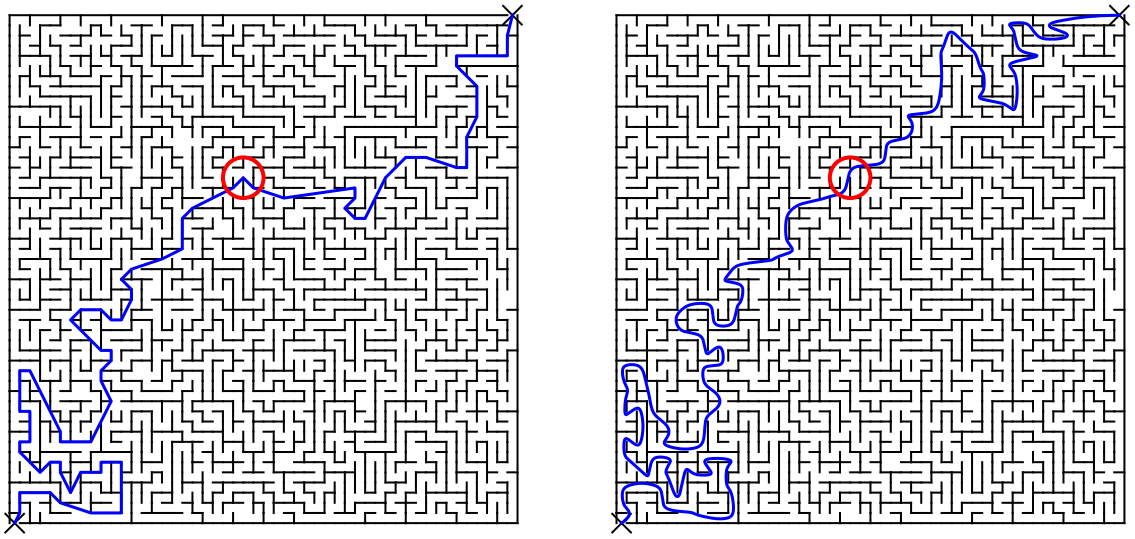

(a) (b)

Figure 5: Solutions of the motion-planning problems through a maze from Section 7.2. (a) Minimum-length trajectory connecting the start (bottom-left cross) and the goal (top-right cross). (b) Solution of the minimum-time problem with velocity constraint $\dot{q} \in [-1, 1]^2$ and regularized acceleration. The solutions of these two problems bifurcate at the red circle, and take different paths across the maze. For both problems, GCS identifies the globally-optimal trajectory via a single SOCP.

7.2 Motion Planning in a Maze

In this example we consider a two-dimensional planning problem of higher complexity than the one just analyzed: we design trajectories through the maze depicted in Figure 5. The maze has $50 \cdot 50 = 2,500$ cells. The starting cell is the one at the bottom left, the goal cell is in the top right. The graph of convex sets is constructed by making each cell into a safe set $\mathcal{Q}_i$. Bidirectional edges are drawn between cells that are not separated by a wall. The maze is generated using random depth-first search. Since mazes constructed using this algorithm have all cells connected to the starting cell by a unique path, to make the planning problem more challenging, we create multiple paths to the goal by randomly selecting and removing 100 walls from the maze.

As in the previous example, we consider a minimum-length problem and a minimum-time problem with regularized acceleration. We set the parameters $(a, b, c, \eta, \mathcal{D}, T_{\min}, T_{\max}, \dot{q}_0, \dot{q}_T, d)$ to the same values we adopted in the corresponding problems in Section 7.1. The optimal trajectories across the maze corresponding to the two objective functions are reported in Figures 5a and 5b. As it can be seen, the two curves visit different sequences of safe sets (cells). In particular, the sharp turn taken by the minimum-length trajectory, circled in red in Figure 5a, would be expensive for the second problem, where we have a penalty on the magnitude of the acceleration. For the curve in Figure 5b, GCS decides then to take a longer but smoother route to the goal. For both

the problems under analysis, the convex relaxation returns a solution with binary probabilities φ_e. Rounding is then unnecessary in this case, and the solution of the convex relaxation is automatically certified to be globally optimal ($\delta_{\text{relax}} = \delta_{\text{opt}} = 0\%$).

We find this example powerful because it highlights the transparency with which GCS blends discrete and continuous optimization. Finding a discrete sequence of cells to traverse the maze in Figure 5 is a trivial graph search. Also finding a path of minimum length, as in Figure 5a, is a relatively simple problem: in fact, in two dimensions, the Euclidean SPP is solvable in polynomial time by constructing a discrete visibility graph [24]. On the other hand, designing a trajectory like the minimum-time one in Figure 5b is a substantially more involved operation. GCS gives us a unified mathematical framework that can tackle all these problems very efficiently, while embracing both the higher-level combinatorial structure and the lower-level convexity of our planning problems.

In conclusion of this example let us illustrate another axis in which GCS significantly improves on existing MICP planners. The worst-case runtime of a mixed-integer solver is typically exponential in the number of binaries in the optimization problem. Previous MICP planners parameterize a single trajectory and subdivide it in a fixed number of segments, then they use a binary variable to assign each segment to each safe region $\mathcal{Q}_i$ [6]. Given that, in the worst case, the optimal trajectory might visit all the safe regions $\mathcal{Q}_i$, this approach requires a total of $|\mathcal{I}|^2$ binary variables. For the maze in Figure 5, we would then have $|\mathcal{I}|^2 = 2,500^2 = 6.25 \cdot 10^6$ binaries: a quantity well beyond the capability of today's solvers. On the contrary, GCS uses only two binaries per pair of intersecting regions, and it yields an MICP with only $5198 \approx 2|\mathcal{I}|$ binaries, which is solved exactly through a single SOCP.

7.3 Statistical Analysis: Quadrotor Flying through Buildings

In this section we present a statistical analysis of the performance of GCS. Taking inspiration from [6], we test our algorithm on the task of planning the motion of a quadrotor through randomly-generated buildings. An example of such a task is illustrated in Figure 6: while moving from the brown to the green block, the quadrotor needs to fly around trees, and through doors and windows. A brief description of how the buildings are generated can be found in Appendix B.

Even though the configuration space of a quadrotor is six dimensional, the differential-flatness property of this system allows us to plan dynamically-feasible trajectories directly in the three-dimensional Cartesian space. In fact, given a four-times-differentiable trajectory of the position of the center of mass, the time evolution of the quadrotor's orientation, together with the necessary control signals, is uniquely defined and easily computed [27]. The space in which we design trajectories is then $\mathcal{Q} \subset \mathbb{R}^3$ and, given that all the obstacles have polyhedral shape (as in Figure 6), the decomposition of this space into convex sets $\mathcal{Q}_i$ can be done exactly. Appendix B provides more details on this decomposition.

In the formulation of the planning problem (1), we penalize with equal weight the duration and the length of the trajectory ($a := b := 1$ and $c := 0$). We parameterize the trajectories using Bézier curves of degree $d := 7$ and, to take advantage of the differential flatness, we require these curves to be continuously differentiable $\eta := 4$ times. The velocity is constrained to be in the

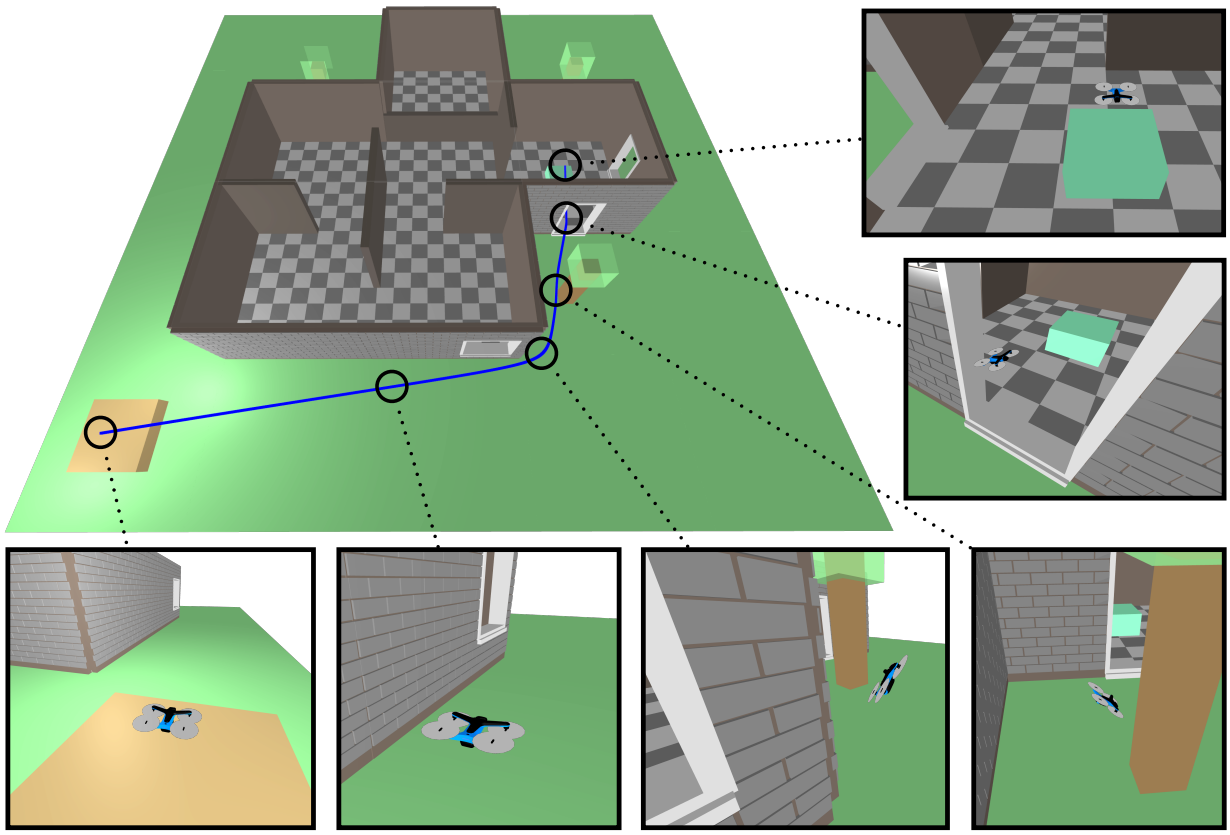

Figure 6: One of the randomly-generated environments for the statistical analysis in Section 7.3. The trajectory generated by GCS for the center of mass of the quadrotor is depicted in blue. The robot orientation is reconstructed taking advantage of the differential flatness of the system dynamics. The snapshots show the starting and ending configurations, as well as the quadrotor flying close to the obstacles in the environment.

box $\mathcal{D} := [-10, 10]^3$ for all times.[7] The limits $T_{\min}$ and $T_{\max}$ on the duration of the trajectory have values that do not affect the optimal solution. As said, the initial q_0 and final q_T positions are above the brown and green boxes, respectively. The boundary values of the velocity are zero $\dot{q}_0 := \dot{q}_T := 0$, as well as the boundary values of the second and the third derivatives of the trajectory.[8] Because of the differential flatness, the latter ensure that the quadrotor starts and ends the motion with horizontal orientation and zero angular velocity. Finally, to regularize the acceleration of the quadrotor, as discussed in Section 6, we set $\dot{h}_{\min} := 10^{-3}$.

We plan the motion of the quadrotor through 100 random buildings. To assess the quality of the trajectories generated by GCS, we look at the optimality gaps δ_{opt} and δ_{relax}. As in

[7] To contextualize the velocity limits, consider that the random environments are squares with sides of length 25, and the collision geometry of the quadrotor is a sphere of radius 0.2 (see also Appendix B).

[8] For $l = 2, \dots, L$, the derivative constraints $q^{(l)}(0) = q^{(l)}(T) = 0$ in problem (1) map to the conditions $r^{(l)}(0) = h^{(l)}(0)\dot{q}_0$ and $r^{(l)}(S) = h^{(l)}(S)\dot{q}_T$ in problem (3). The latter are linear in the decision variables of our SPP in GCS, and can be easily incorporated among the edge constraints listed in Section 5.3.

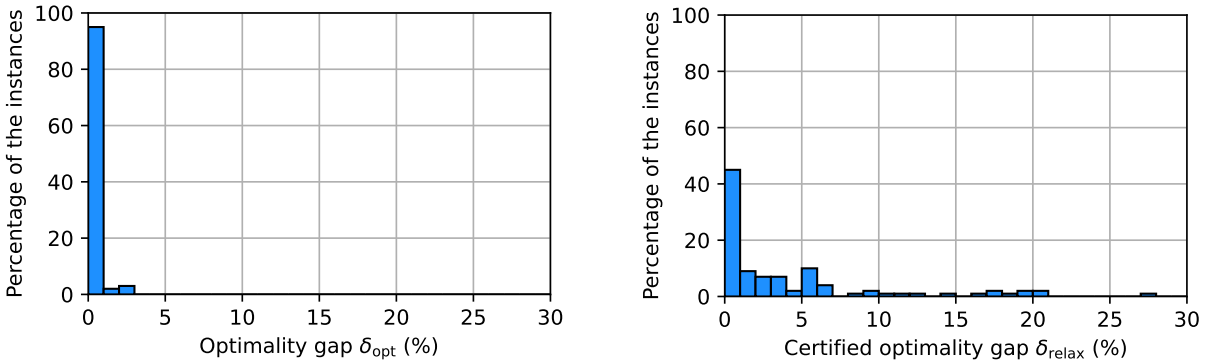

(a) (b)

Figure 7: Histograms of the optimality gaps registered in the statistical analysis in Section 7.3. (a) Optimality gap δ_{opt}: percentage gap between the cost of the solution returned by GCS and the global optimum. On 95% of the environments GCS designs a trajectory with optimality gap smaller than 1%, and, even in the worst case, it finds a solution whose cost is only 2.9% larger than the global minimum. (b) Optimality gap $\delta_{\text{relax}} \geq \delta_{\text{opt}}$ automatically certified by GCS. On 68% (respectively 84%) of the problems GCS certifies that the returned solution has optimality gap smaller than 4% (respectively 7%).

the previous examples, the value of δ_{opt} is computed (just for analysis purposes) by solving the planning problem to global optimality using a mixed-integer algorithm, while δ_{relax} is the upper bound on δ_{opt} that is automatically provided to us by GCS. The histograms of these two quantities across the 100 experiments are reported in Figure 7. Figure 7a shows that on 95% of the environments GCS designs a trajectory whose optimality gap δ_{opt} is smaller than 1%, and, even in the worst case, is only 2.9%. From Figure 7b, we see that on 68% (respectively 84%) of the problems GCS certifies that the returned solution is within 4% (respectively 7%) of the global optimum. The largest optimality gap δ_{relax} certified by GCS is 27.1%, and it corresponds to an environment where we have $\delta_{\text{opt}} = 2.3\%$. Therefore, even for this problem instance, the moderately-large value of δ_{relax} is mostly due to the convex relaxation being slightly loose, rather than the rounded solution being suboptimal.

We report that, for the statistical analysis in this subsection, we set the MOSEK parameter `MSK_IPAR_INTPNT_SOLVE_FORM = 1`, which tells the interior-point solver to interpret our optimizations in standard primal form [23]. Without this, MOSEK encountered numerical issues in the solution of the convex relaxations of the motion-planning problems. This parameter choice has the drawback of sensibly slowing down the planning times: the solve times for the convex relaxations of the 100 motion plans have median 3.7 s, mean 6.4 s, and maximum 31.2 s. However, we are very confident that a deeper analysis of these numerical issues and a tailored pre-solve stage, can reduce these times by at least one order of magnitude.

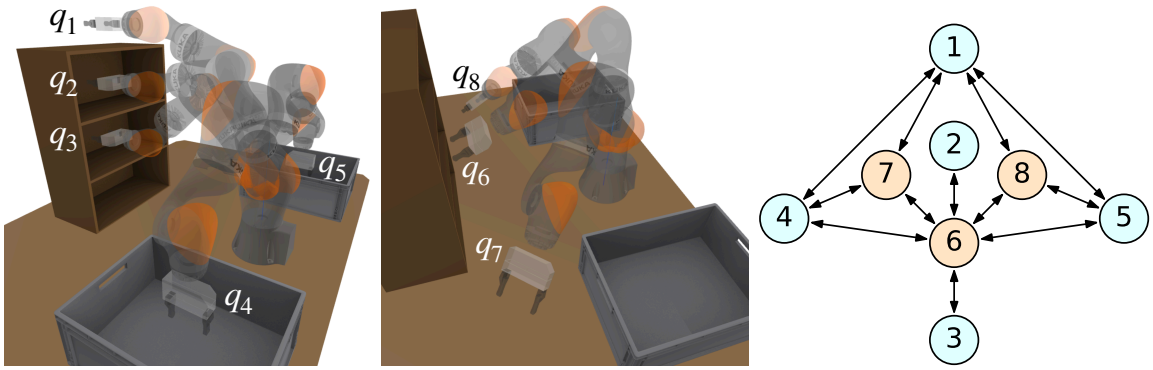

(a) (b) (c)

Figure 8: Construction of the GCS for the motion planning of the robot arm in Section 7.4. (a) Five seed poses $\{q_k\}_{k=1}^{5}$ for the region-inflating algorithm from [5]. These are chosen to fill the space within the rack and the bins. (b) The remaining three seed poses $\{q_k\}_{k=6}^{8}$, chosen to approximately fill the free configuration space. (c) The graph G obtained by processing the safe regions $\mathcal{Q}_i$. The light-blue vertices correspond to the seed poses in (a), the light-brown ones to the additional seeds from (b). Vertices are labeled with the subscripts of the corresponding poses.

7.4 Comparison with PRM: Motion Planning of a Robot Arm

In this subsection we consider the motion planning of a robot arm, and we compare GCS with commonly-used sampling-based planners. GCS is a multiple-query algorithm, meaning that the same data structure (the graph of convex sets) can be used to plan the motion of the robot for many initial and final conditions. Its natural sampling-based comparison is then the Probabilistic-RoadMap (PRM) algorithm [18]. The robot arm we use in this benchmark is the KUKA LBR iiwa with $n = 7$ degrees of freedom: we have chosen a seven-dimensional configuration space $\mathcal{Q}$ since PRM methods can struggle in larger spaces, and both algorithms under analysis can easily design trajectories in lower dimensions.

The robot arm is depicted in Figure 8, and it is required to move within an environment composed of a rack (in front of the robot) and two bins (on the sides). As opposed to the examples considered so far, an exact decomposition of the free configuration space $\mathcal{Q}$ is not feasible in this application. We then adopt the approximate decomposition algorithm, IRIS, from [5]; more precisely, its extension to configuration spaces with nonconvex obstacles, `IrisInConfigurationSpace`, implemented in Drake [36]. Given a "seed pose" of the robot, this algorithm inflates a polytope of robot configurations that are not in collision with the environment. While these polytopes could be rigorously certified to be collision free [1], for the experiments reported here we use a fast implementation based on nonconvex optimization that does not provide a rigorous certification, but that appears to be very reliable in practice.

Automatic seeding of the regions is certainly possible, but we have found that producing seeds manually via inverse kinematics, together with a simple visualization of the graph G to check the connectivity between regions $\mathcal{Q}_i$, is straightforward and highly effective. We use IRIS to construct a total of eight safe polytopes $\mathcal{Q}_i$, whose corresponding seed poses q_i are depicted in

Figures 8. The seed poses $\{q_i\}_{i=1}^5$ in Figures 8a are chosen to create polytopes $\mathcal{Q}_i$ that cover the volume of configuration space for which the end effector is in the vicinity of the rack and the bins. The poses $\{q_i\}_{i=6}^8$ in Figures 8b are picked to approximately fill the rest of the free space. The construction of the safe regions is parallelized, and took us 53 seconds. By processing the safe regions $\mathcal{Q}_i$ as described in Section 5.1, we obtain the graph G depicted in Figure 8c. The vertices $\mathcal{I} = \{1, \dots, 8\}$ are the subscripts of the poses that we use as seeds for the construction of each polytope, i.e., vertex $i \in \mathcal{I}$ is paired with the safe polytope $\mathcal{Q}_i$ obtained from the seed q_i. As can be seen from the connectivity of the graph, the polytopes $\mathcal{Q}_i$ are sufficiently inflated to connect all the seed poses q_i. At runtime, given the initial q_0 and final q_T configuration, the source σ and the target τ vertices are added to the graph and connected to other vertices as described in Section 5.1.

In practice, the plans generated by a PRM can be very suboptimal and are rarely commanded to the robot directly. While asymptotically-optimal versions of the PRM method exist [17], in our experience, in the relatively high-dimensional space we consider here, the increase in performance of these variants is not worth their computational cost. A solution commonly used in practice is then to post-process the plans generated by the PRM with a simple short-cutting algorithm. This algorithm samples pairs of points along the PRM trajectory and connects them via straight segments: if a segment is verified to be collision free the trajectory is successfully shortened. This step can dramatically shorten the PRM trajectories but it requires time-consuming collision checks: for this reason, here we compare GCS with both the regular PRM and the PRM with short-cutting. For both the PRM methods we use the implementation from [31]. More implementation details can be found in Appendix C; here we only mention that our roadmap is composed of $15 \cdot 10^3$ sample configurations and its construction took, with our (not fully optimized) setup, 16 minutes.

The tasks require moving the arm between five waypoint configurations $\rho_i \in \mathcal{Q}$, while avoiding collisions with the rack and the bins. Each waypoint ρ_i is obtained from q_i by perturbing the position of the robot end-effector as shown in Figure 9. We have a total of five tasks: for $i = 1, \dots, 4$, task i asks us to move the robot from ρ_i to ρ_{i+1}; task 5 requires moving the robot from ρ_5 back to ρ_1. The objective is to connect the start and the goal configurations with a continuous ($\eta := 0$) trajectory of minimum Euclidean length ($a := c := 0$ and $b := 1$). Velocity and time constraints are irrelevant given our objective.

As a visual support to the analysis, Figure 9 illustrates the trajectories of the robot end-effector generated by each planner for each task. The blue curves correspond to GCS, the yellow to the regular PRM, and the red to the PRM with short-cutting. Let us emphasize, though, that shorter trajectories in configuration space do not necessarily map to shorter trajectories in task space. The actual configuration-space lengths of these trajectories are reported in Figure 10a, with the same color scheme. The runtimes required by each planner can be found in Figure 10b.[9] In all the tasks, GCS designs trajectories that are shorter than both PRM methods. Moreover, the runtimes of GCS are even smaller than the ones of the regular PRM. The PRM with short-cutting designs higher-quality trajectories than the regular PRM, but its runtimes are significantly larger. The pre-processing described in Appendix A.2 is the reason why our method is extremely fast in

[9]The runtimes of GCS are computed by summing the times necessary for the pre-processing described in Appendix A.2, the solution of the convex relaxation of the SPP in GCS, and the rounding step from Section 4.2.

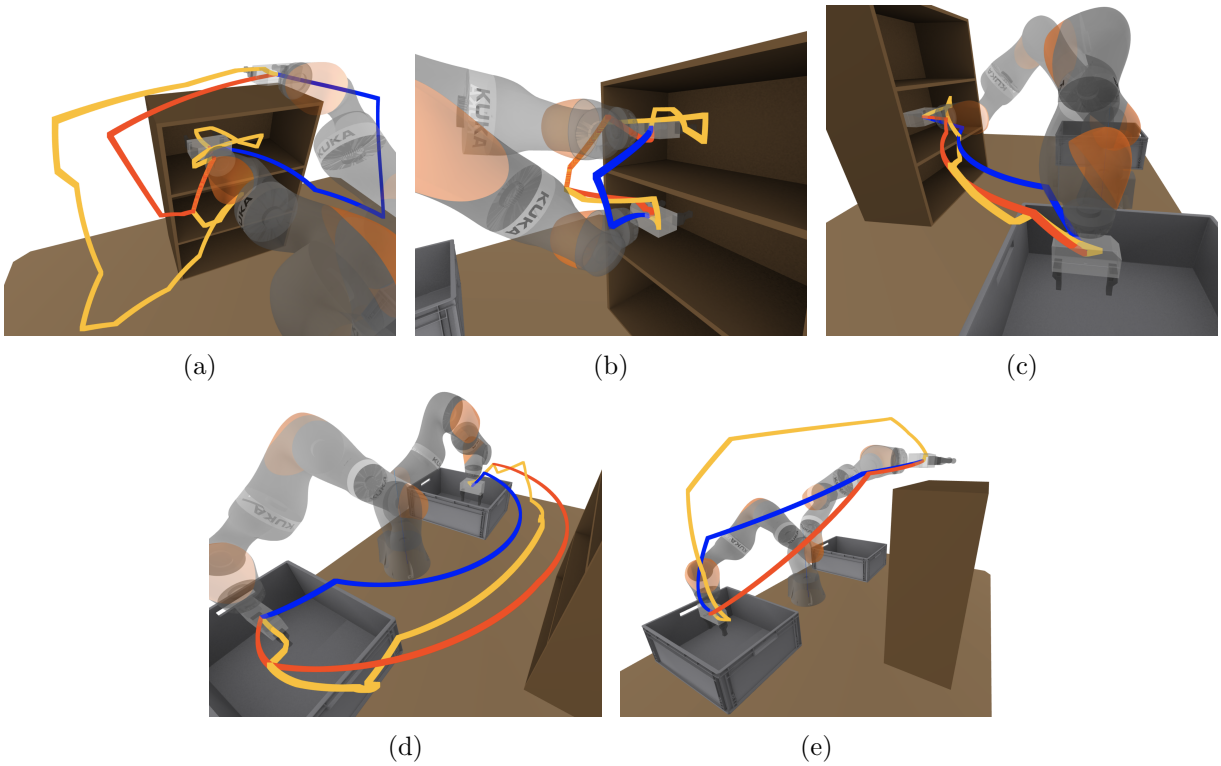

Figure 9: The five motion-planning tasks for the comparison in Section 7.4. End-effector trajectories are depicted in blue for GCS, in yellow for the regular PRM, and in red for the PRM with short-cutting. (a) Task 1: from end-effector above the rack (configuration ρ_1) to end-effector in the upper shelf (configuration ρ_2). (b) Task 2: from upper shelf ρ_2 to lower shelf ρ_3. (c) Task 3: from lower shelf ρ_3 to left bin ρ_4. (d) Task 4: from left bin ρ_4 to right bin ρ_5. (e) Task 5: from right bin ρ_5 to above the rack ρ_1.

solving task 2: in the graph G in Figure 8c there is only one path that connects vertex 2 to vertex 3, and our pre-processing efficiently eliminates all the edges in the graph but $(2, 6)$ and $(6, 3)$.

In conclusion, let us mention that in all the tasks the solution we identify via rounding is the global optimum of the SPP in GCS ($\delta_{\text{opt}} = 0\%$). The certified optimality gap δ_{relax} is 4.1% on average, and achieves a maximum of 13.0% in the first task.

7.5 Coordinated Planning of Two Robot Arms

In the previous subsection we have compared GCS to widely-used PRM methods, choosing a robotic arm with $n = 7$ degrees of freedom because sampling-based algorithms perform poorly in higher dimensions. Here we demonstrate that GCS can tackle planning problems in much higher-dimensional spaces. To this end, we consider the dual-arm manipulator shown in Figure 11, composed of two KUKA LBR iiwa with seven degrees of freedom each, yielding an overall configuration space $\mathcal{Q}$ of $n = 14$ dimensions. The environment is the same as in the previous subsection, but this time, besides the collisions with the rack and the bins, GCS must also prevent

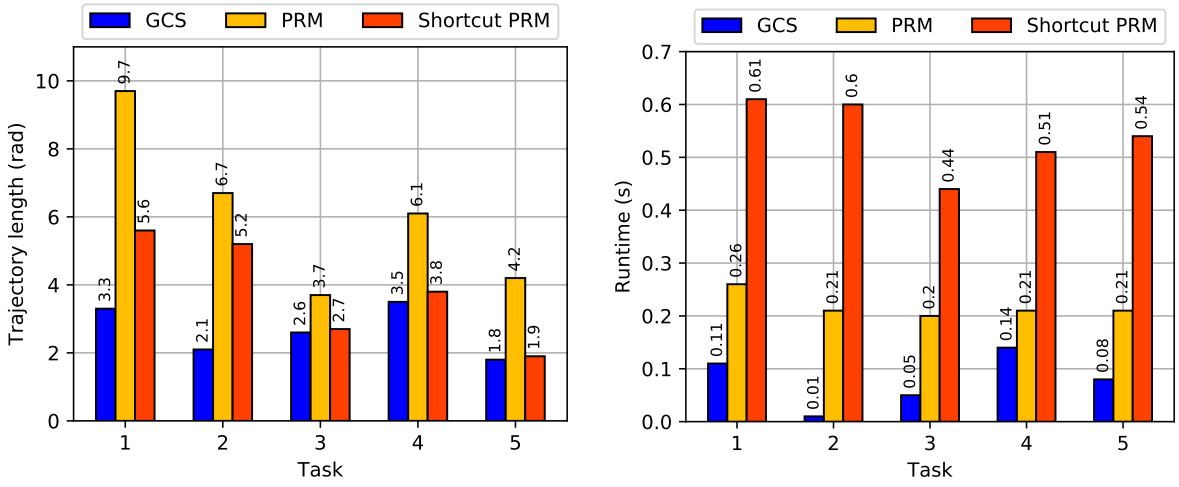

(a) (b)

Figure 10: Comparison of GCS with the PRM method and its version with short-cutting. (a) Length of the trajectories planned by each algorithm for the five tasks depicted in Figure 9. (b) Corresponding runtimes. GCS designs shorter trajectories than the PRM method with short-cutting, and is faster than the regular PRM.

collisions between the arms themselves.

To decompose the configuration space we proceed as in Section 7.4. This time we use a total of 22 seed poses, chosen to approximately cover the workspace around the rack and the bins, as well as the rest of the free space. Also in this case the seeds are produced manually, using inverse kinematics and with the visual support provided by the connectivity of the graph G. We analyze three tasks. In the first task, illustrated in Figure 11a, the arms start in a neutral position and both reach into the top shelf. Task 2, in Figure 11b, asks the arms to cross: the left arm reaches above the rack on the right, and the right arm moves to the left of the bottom shelf. Finally, in Figure 11c, task 3 requires the two arms to reach inside the bins. To make the problem even more challenging, this time we do not limit ourselves to the design of purely-geometric shortest curves as in Section 7.4, but we plan continuously differentiable ($\eta := 1$) trajectories of degree $d := 3$. The weights in the objective (1a) are set to $a := b := 1$ and $c := 0$. The constraint set $\mathcal{D}$ in (1d) ensures that the joint velocities are no greater than 60% of the robot velocity limits. The duration bounds $T_{\min}$ and $T_{\max}$ are set so that they do not affect the optimal trajectory, while the boundary values of the velocity are zero ($\dot{q}_0 := \dot{q}_T := 0$). As described in Section 6, we penalize accelerations via a cost term of the form (11), with weight $\varepsilon = 10^{-3}$. With the same goal, we set $\dot{h}_{\min} := 10^{-3}$.

The trajectories synthesized by GCS for each of the three tasks are represented in Figure 11, with the curves swept by the end-effectors depicted in blue. The optimality gaps δ_{relax} certified by GCS for the three tasks are 3.3%, 2.0%, and 0.6%. Running a mixed-integer solver, we verify that the first two trajectories are, in fact, globally optimal, while the last trajectory has an optimality gap of only $\delta_{\text{opt}} = 0.3\%$. As in Section 7.3, to circumvent numerical issues, we set the MOSEK

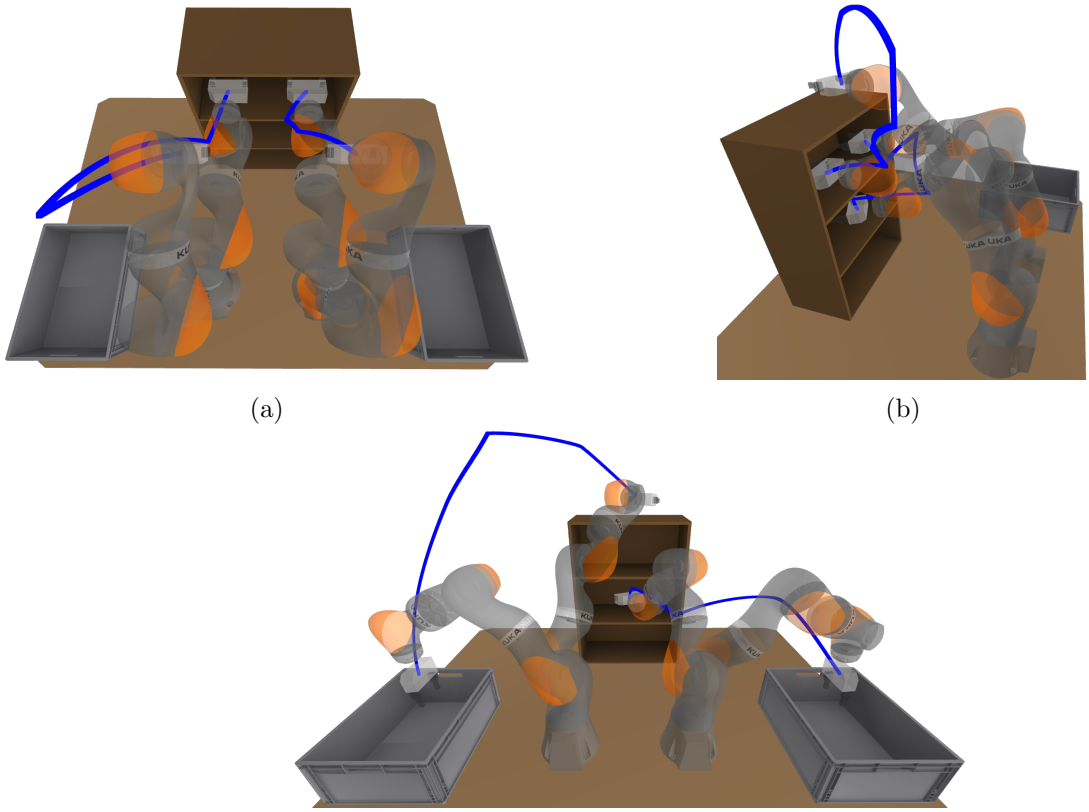

(a) (b)

(c)

Figure 11: Manipulation tasks from Section 7.5. End-effector trajectories are in blue. (a) Task 1: arms from neutral pose to top shelf. (b) Task 2: from top shelf to configuration with crossed arms. (c) Task 3: from crossed arms to lateral bins. Despite the fourteen-dimensional configuration space, the potential collisions between the arms, and the confined environment, GCS can reliably solve the three tasks in a few seconds via convex optimization.

option `MSK_IPAR_INTPNT_SOLVE_FORM = 1` in the solution of the convex relaxations. This leads to the following computation times for the three tasks at hand: 4.0 s, 8.4 s, and 12.9 s. As already mentioned, we are confident that a tailored pre-solve stage can drastically decrease these runtimes.

8 Discussion

On the one hand, transcribing the motion-planning problem as an SPP in GCS allows us to use efficient convex optimization to design trajectories around obstacles. On the other hand, our convexity requirements restrict the class of planning problems we can tackle, and limit the families of trajectories we can parameterize. In this section we comment on the strengths and the limitations of our approach, and we illustrate the pros and the cons of GCS over existing

planning algorithms.

8.1 Additional Costs and Constraints

Besides the derivative penalties discussed in Section 6, there are many additional costs and constraints that our problem statement (1) does not feature but that are relevant in a variety of practical applications. Minimum-distance and minimum-time objectives might lead to unsafe robot trajectories, that do not avoid obstacles with sufficient clearance. A practical workaround in these cases is to discourage the control points of our trajectories to get too close to certain boundaries of the safe regions $\mathcal{Q}_i$. This can be achieved through convex barrier penalties, that are easily included among the edge costs in Section 5.4. Equality constraints that couple the trajectory q to its time derivatives could be used to enforce continuous-time dynamics. However, our choice of optimizing over the shape r and the timing h of the trajectory jointly makes these constraints nonconvex, even for a linear control system. Similarly, the nonlinearity of the kinematics of a robot manipulator makes task-space constraints not directly suitable for our framework. To cope with these nonconvexities, in some applications, it may be practical to post-process the output of GCS with a local nonconvex optimizer.

8.2 Comparison with Existing Mixed-Integer Planners

GCS has three main advantages over existing MICP algorithms for solving problems of the form (1):

1. The tightness of the convex relaxation of our MICPs, demonstrated empirically in the numerical results in Section 7.

2. The reduced number of binary decision variables in our programs, illustrated in the maze example from Section 7.2.

3. The simplicity of the class of optimization problems that our method leads to, discussed in Section 5.6

The first and the second are achieved by leveraging the optimization framework from [26]. The third is partly due to the first (since it is the tightness our MICP formulations that allows us to tackle the motion planning problem as a single convex program, plus rounding), but it is also due to the parameterization of trajectories as Bézier curves.

In Section 5.2, we have leveraged the properties of Bézier curves to enforce infinite families of constraints through a finite number of conditions. For example, in (7a), we have transcribed the safety requirement $r_i(s) \in \mathcal{Q}_i$ for all $s \in [0, 1]$ as a constraint $r_{i,k} \in \mathcal{Q}_i$ per control point $k = 0, \dots, d$. The MICP planner from [6] achieves the same result by using Sums-Of-Squares (SOS) polynomials [30], and semidefinite programming. These approaches are interchangeable and lead to a tradeoff: Bézier curves yield simpler constraints, SOS polynomials parameterize a richer class of trajectories.[10] In the numerical examples analyzed in this paper, we have found this gap to be relatively narrow, and we have then prioritized simpler optimization problems.

[10] Asking a univariate polynomial to be nonnegative by parameterizing it as a Bézier curve with nonnegative control points is more stringent than asking it to be SOS (which, in the univariate case, is equivalent to nonnegativity).

Finally, it is worth mentioning that the problem formulation from [6] features costs and constraints on time derivatives of the trajectory q of any order. These, however, are handled by fixing the duration of each trajectory segment beforehand. A similar result could be achieved with GCS by fixing the time that can be spent in each safe set $\mathcal{Q}_i$.

8.3 Comparison with Sampling-Based Algorithms

As discussed in Section 7.4, among many sampling-based planners, PRM is the natural comparison for GCS. In fact, GCS can be thought of as a generalization of the PRM method, where each collision-free sample is expanded to a collision-free convex region, that is inflated as much as the obstacles allow; reducing in this way a dense roadmap to a compact GCS. In Sections 7.4 and 7.5, we have shown that GCS can outperform PRM in terms of: runtimes, quality of the designed trajectories, scalability with the dimensionality n of the configuration space $\mathcal{Q}$, and variety of objective functions and trajectory constraints. In addition, because of the parallel above, it is reasonable to imagine that many of the techniques developed for PRM to handle, e.g., changes in the environments [15, 37] can be translated to GCS with relatively low effort.

One of the main reasons why sampling-based methods are widely used in academia and industry is their simplicity. Conversely, the implementation of GCS is very involved and requires familiarity with convex-optimization techniques. Nonetheless, we believe that the framework from [26] lends itself to an intuitive mathematical abstraction, and that the programming interface of GCS can be made very easy to use. We have provided a mature implementation of the techniques from [26] within the open-source software Drake [36], and we have developed a simple GCS interface at `https://github.com/mpetersen94/gcs`.

8.4 Comparison with Direct Trajectory Optimization

Direct-trajectory-optimization methods transcribe the motion-planning problem into a nonconvex optimization [7], and can virtually include any sort of cost terms and constraints, including dynamic and task-space constraints. In practice, however, these nonconvex programs can only be tackled with local-optimization algorithms that are slow and unreliable. GCS is different in spirit, as we prioritize low runtimes and the completeness of the planning algorithm over the modelling power.

9 Conclusions and Future Works

In this paper we have introduced GCS: an algorithm based on convex optimization for efficient collision-free motion planning. GCS leverages the framework presented in [26] to design a very tight and lightweight convex relaxation of the planning problem. This convex optimization (typically an SOCP) is quickly solved using commonly-available software, and a cheap randomized rounding of its solution is almost always sufficient to identify a globally-optimal trajectory. We have demonstrated GCS on a variety of scenarios: an intricate maze, a quadrotor flying through buildings, and a manipulation task in a fourteen-dimensional configuration space. Furthermore, we have compared GCS to widely-used PRM methods, showing that our method can find higher-quality trajectories in less time.

This paper presents the first version of a new algorithm, which already compares favorably with widely-used planners that have been optimized over decades. The runtimes of GCS can be drastically reduced (we are currently developing a customized solver for these convex optimizations). We are also highly optimistic that the class of cost functions and constraints that we can handle will expand considerably in the future. In particular, we imagine incorporating task-space constraints, tight penalties on the higher derivatives of the trajectory, as well as dynamic constraints arising from input limits. Furthermore, we wish to extend GCS to problems involving contacts between the robot and the environment. We believe that our planner demonstrates the value of formulating problems as SPPs in GCS, and it can already find multiple real-world applications.

A Further Details on the Implementation of GCS

In this appendix we illustrate two techniques that we employed in the numerical results in Section 7 to tighten and compress the convex relaxations of our planning problems.

A.1 Two-Cycle-Elimination Constraints

The graph G constructed in Section 5.1 connects each pair $\mathcal{Q}_i$ and $\mathcal{Q}_j$ of overlapping safe regions with a two-cycle: $e := (i, j)$ and $f := (j, i)$. Since by traversing both the edges e and f we would visit vertex i twice, and this is not allowed by the definition of a path p, at least one of these edges must be excluded from the shortest path. In other words, the indicator variables φ_e and φ_f cannot be both equal to one. This observation can be used to tighten our convex relaxations, and speed up our planner.

More precisely, for each pair of overlapping regions, we can write the linear constraints

$$\varphi_e + \varphi_f \leq \varphi_i \quad \text{and} \quad \varphi_e + \varphi_f \leq \varphi_j, \tag{12}$$

where φ_i and φ_j represent the total probability flows traversing vertices i and j, respectively. (Note that, since the total flow through a vertex is at most one, these inequalities imply the looser condition $\varphi_e + \varphi_f \leq 1$.) Furthermore, by applying Lemma 1(b) from [26], the two inequalities in (12) can be translated into a pair of convex constraints that tighten the coupling between the flow variables φ_e and the continuous variables x_v in our convex programs. The number of these constraints is linear in the size $|\mathcal{E}|$ of the edge set, and they can substantially increase the tightness of the convex relaxations of our planning problems. They are enforced in all the numerical results presented in Section 7.

A.2 Graph Pre-Processing

The constraints described in Appendix A.1 represent only one of the multiple ways in which we can leverage the knowledge that a path p is allowed to visit a vertex at most once. For example, consider the graph in Figure 8c and task 2 from Section 7.4 of moving the robot arm between the configurations $\rho_2 \in \mathcal{Q}_2$ and $\rho_3 \in \mathcal{Q}_3$. In this case, after connecting the source σ to vertex 2 and vertex 3 to the target τ, we get a graph that admits a single σ-τ path: $p := (\sigma, 2, 6, 3, \tau)$.

Therefore, in this particular case, a pre-processing stage capable of making such an inference would reduce our planning problem to a tiny convex program (exactly).

In general, making the inference just described exactly is infeasible; however, in many practical scenarios, a cheap approximate pre-processing can eliminate most of the redundancies in our graphs G. More precisely, checking if an edge $e := (u, v)$ can be traversed by a σ-τ path is equivalent to solving a vertex-disjoint-paths problem. This problem asks to identify a path p_1 from σ to u and a path p_2 from v to τ such that the overall path $p := (p_1, p_2)$ is a valid path from σ to τ. In other words, the two subpaths p_1 and p_2 are not allowed to share any vertex. This problem is NP-complete [34, Section 70.5], therefore it would not make sense to solve it exactly as a pre-processing for our planner. Nevertheless, the vertex-disjoint-paths problem admits a natural LP relaxation as a fractional multiflow problem [34, Section 70.1] that can be solved very quickly, and can be used as a very-effective sufficient condition to check if an edge is redundant.

We have found this pre-processing to be particularly useful when the graph G is sparse and has small size, and the convex sets $\mathcal{X}_v$ associated to its vertices live in high dimensions. In these cases, the multiflow LPs (which can be tackled in parallel) are solved extremely fast and they can drastically compress and tighten our convex optimizations. We have employed this pre-processing strategy in the numerical examples from Sections 7.3, 7.4, and 7.5: the runtimes of GCS reported in these sections include the time necessary for pre-processing.

B Random Environment Generation for the Quadrotor Example

In this appendix we briefly describe the algorithm we employed for the generation of the random buildings in Section 7.3.

The buildings are constructed over a five-by-five grid, where each cell has sides of length 5. The nine cells at the center of the grid are occupied either by a room, a tree, or obstacle-free grass. The sixteen cells at the boundary of the grid are always occupied by grass. For all the environments, the brown start block is in the cell $(1, 1)$, while the green goal block is in the cell $(4, 3)$ (see Figure 6). To assemble a building we start from the goal cell, which we always require to be a room. Then we mark each adjacent cell either as inside or outside the building, and we repeat this process until the nine inner cells are occupied. For the cells that are marked as outside the building, we decide at random whether to grow a tree or not. Walls divide the rooms from the outside, and are built with either a doorway, a window, two windows, or no openings at all. Walls are also used to divide the rooms; in this case we randomly select a doorway, a vertical half wall, a horizontal half wall, or no wall. The positions of the trees are also drawn at random, while their sizes are taken to be constant.

Given that the walls and the trees have polygonal shape, the decomposition of the configuration space $\mathcal{Q}$ can be done exactly. Specifically, we pair rooms or cells that are occupied by grass with a single box $\mathcal{Q}_i$ of free space, while the space around a tree is decomposed using four non-overlapping boxes. Suitable box-shaped regions are added for each (inner or outer) wall that contains one or more openings. Finally, the safe regions $\mathcal{Q}_i$ are adequately shrunk to take into account the collision geometry of the quadrotor, which is taken to be a sphere of radius 0.2.

C Implementation of the PRM Planner

In this appendix we report the main implementation details for the PRM and the short-cutting algorithm used in the comparison in Section 7.4.

We construct the PRM using the implementation `simple_prm_planner.hpp` from the library [31]. Trying to construct a roadmap just by sampling random robot poses turns out to be infeasible for the application in Section 7.4. In fact, sampling a robot pose $q \in \mathbb{R}^7$ for which the end effector is, e.g., inside one of the shelves in Figure 8 is extremely unlikely: after $3 \cdot 10^5$ samples, and 90 hours of computations, we did not find any such point. As a result, we construct the roadmap in two steps. In the first step, we connect the seed poses $\{q_i\}_{i=1}^8$ from Figure 8 using a collection of bidirectional Rapidly-exploring Random Trees (RRTs) (`simple_rrt_planner.hpp` from [31]). The role of these trees is to form a skeleton for the PRM, and, to keep this skeleton reasonably compact, we mimic the connectivity of our graph G in Figure 5.1. In particular, we connect via RRT only the pairs of seed poses q_i and q_j for which the vertices i and j are connected in G. This process gives us 12 trees, with a total of approximately $2,300$ nodes. In the second step, we fill the rest of the space according to the standard PRM algorithm. We stop the sampling when we reach a total of $15 \cdot 10^3$ nodes in the PRM, included the ones from the RRTs. (In our experience, a larger number of PRM samples would have led to an increase in the runtimes without sensibly improving the quality of the designed trajectories.) During this construction, the collision checks are handled by Drake [36]. With this setup, generating the RRTs took a total of 60 seconds, while the remaining PRM samples required 15 minutes. For the short-cutting algorithm we use the implementation in `path_processing.hpp` from [31].

The numerical parameters we use for the RRT, the PRM, and the short-cutting algorithm are chosen to optimize the tradeoff between the quality of the designed paths and the overall computation times.

Acknowledgements

We would like to thank Greg Izatt for providing us with the environment generator used to benchmark our algorithm in Section 7.3, and for his precious assistance with the visualization of the motion plans in Sections 7.4 and 7.5. We are grateful to Andres Valenzuela for his great help in the initial phases of this project, and to Pablo A. Parrilo, Jack Umenberger, and Calder Phillips-Grafflin for the many insightful suggestions and discussions.

This material is based upon work supported by (in alphabetical ordered): Amazon.com, PO No. 2D-06310236; Department of Defense (DoD) through the National Defense Science & Engineering Graduate (NDSEG) Fellowship Program; Lincoln Laboratory/Air Force, PO No. 7000470769; National Science Foundation, Award No. EFMA-1830901; Office of Naval Research, Award No. N00014-18-1-2210; and Under Secretary of Defense for Research and Engineering, Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

References

[1] Alexandre Amice, Hongkai Dai, Peter Werner, Annan Zhang, and Russ Tedrake. Finding and optimizing certified, collision-free regions in configuration space for robot manipulators. 2022.

[2] Federico Augugliaro, Angela P Schoellig, and Raffaello D'Andrea. Generation of collision-free trajectories for a quadrocopter fleet: A sequential convex programming approach. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1917–1922. IEEE, 2012.

[3] Nora Ayanian and Vijay Kumar. Abstractions and controllers for groups of robots in environments with obstacles. In *2010 IEEE International Conference on Robotics and Automation*, May 2010.

[4] Noel Csomay-Shanklin, Andrew J Taylor, Ugo Rosolia, and Aaron D Ames. Multi-rate planning and control of uncertain nonlinear systems: Model predictive control and control lyapunov functions. *arXiv preprint arXiv:2204.00152*, 2022.

[5] Robin Deits and Russ Tedrake. Computing large convex regions of obstacle-free space through semidefinite programming. In *Algorithmic Foundations of Robotics XI*, pages 109–124. Springer, 2015.

[6] Robin Deits and Russ Tedrake. Efficient mixed-integer planning for UAVs in cluttered environments. In *2015 IEEE International Conference on Robotics and Automation*, pages 42–49. IEEE, 2015.

[7] Moritz Diehl, Hans Georg Bock, Holger Diedam, and P-B Wieber. Fast direct multiple shooting algorithms for optimal robot control. In *Fast Motions in Biomechanics and Robotics*, pages 65–93. Springer, 2006.

[8] Bachir El Khadir, Jean Bernard Lasserre, and Vikas Sindhwani. Piecewise-linear motion planning amidst static, moving, or morphing obstacles. In *2021 IEEE International Conference on Robotics and Automation*, pages 7802–7808. IEEE, 2021.

[9] Mohamed Elbanhawi and Milan Simic. Sampling-based robot motion planning: A review. *IEEE access*, 2:56–77, 2014.

[10] Melvin E Flores. *Real-time trajectory generation for constrained nonlinear dynamical systems using non-uniform rational b-spline basis functions.* California Institute of Technology, 2008.

[11] Jonathan D Gammell, Siddhartha S Srinivasa, and Timothy D Barfoot. Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2997–3004. IEEE, 2014.

[12] Gustavo Goretkin, Alejandro Perez, Robert Platt, and George Konidaris. Optimal sampling-based planning for linear-quadratic kinodynamic systems. In *2013 IEEE International Conference on Robotics and Automation*, pages 2429–2436. IEEE, 2013.

[13] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Convex analysis and minimization algorithms I: Fundamentals*, volume 305. Springer science & business media, 2013.

[14] Michael Hoy, Alexey S Matveev, and Andrey V Savkin. Algorithms for collision-free navigation of mobile robots in complex cluttered environments: a survey. *Robotica*, 33(3):463–497, 2015.

[15] Léonard Jaillet and Thierry Siméon. A PRM-based motion planner for dynamically changing environments. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 2, pages 1606–1611. IEEE, 2004.

[16] Lucas Janson, Edward Schmerling, Ashley Clark, and Marco Pavone. Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions. *The International Journal of Robotics Research*, 34(7):883–921, 2015.

[17] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011.

[18] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

[19] Frans Anton Koolen. *Balance control and locomotion planning for humanoid robots using nonlinear centroidal models*. PhD thesis, Massachusetts Institute of Technology, 2020.

[20] Boris Lau, Christoph Sprunk, and Wolfram Burgard. Kinodynamic motion planning for mobile robots using splines. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2427–2433. IEEE, 2009.

[21] Steven M LaValle. *Planning algorithms*. Cambridge University Press, 2006.

[22] Steven M LaValle et al. Rapidly-exploring random trees: A new tool for path planning. 1998.

[23] Johan Löfberg. Dualize it: software for automatic primal and dual conversions of conic programs. *Optimization Methods & Software*, 24(3):313–325, 2009.

[24] Tomás Lozano-Pérez and Michael A Wesley. An algorithm for planning collision-free paths among polyhedral obstacles. *Communications of the ACM*, 22(10):560–570, 1979.

[25] Anirudha Majumdar and Russ Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research*, 36(8):947–982, 2017.

[26] Tobia Marcucci, Jack Umenberger, Pablo A Parrilo, and Russ Tedrake. Shortest paths in graphs of convex sets. *arXiv preprint arXiv:2101.11565v3*, 2021.

[27] Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *2011 IEEE International Conference on Robotics and Automation*, pages 2520–2525. IEEE, 2011.

[28] Daniel Mellinger, Alex Kushleyev, and Vijay Kumar. Mixed-integer quadratic program trajectory generation for heterogeneous quadrotor teams. In *2012 IEEE International Conference on Robotics and Automation*, pages 477–483. IEEE, 2012.

[29] Ramkumar Natarajan, Howie Choset, and Maxim Likhachev. Interleaving graph search and trajectory optimization for aggressive quadrotor flight. *IEEE Robotics and Automation Letters*, 6(3):5357–5364, 2021.

[30] Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.

[31] Calder Phillips-Grafflin. Common robotics utilities. URL: `https://github.com/calderpg/common_robotics_utilities`.

[32] Arthur Richards and Jonathan P How. Aircraft trajectory planning with collision avoidance using mixed integer linear programming. In *2002 American Control Conference*, volume 3, pages 1936–1941. IEEE, 2002.

[33] Tom Schouwenaars, Bart De Moor, Eric Feron, and Jonathan How. Mixed integer programming for multi-vehicle path planning. In *2001 European Control Conference*, pages 2603–2608. IEEE, 2001.

[34] Alexander Schrijver et al. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

[35] John Schulman, Yan Duan, Jonathan Ho, Alex Lee, Ibrahim Awwal, Henry Bradlow, Jia Pan, Sachin Patil, Ken Goldberg, and Pieter Abbeel. Motion planning with sequential convex optimization and convex collision checking. *The International Journal of Robotics Research*, 33(9):1251–1270, 2014.

[36] Russ Tedrake and the Drake Development Team. Drake: Model-based design and verification for robotics, 2019. URL: `https://drake.mit.edu`.

[37] Jur Van Den Berg, Dave Ferguson, and James Kuffner. Anytime path planning and replanning in dynamic environments. In *2006 IEEE International Conference on Robotics and Automation*, pages 2366–2371. IEEE, 2006.

[38] Diederik Verscheure, Bram Demeulenaere, Jan Swevers, Joris De Schutter, and Moritz Diehl. Time-optimal path tracking for robots: A convex optimization approach. *IEEE Transactions on Automatic Control*, 54(10):2318–2327, 2009.

[39] Dustin J Webb and Jur Van Den Berg. Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In *2013 IEEE International Conference on Robotics and Automation*, pages 5054–5061. IEEE, 2013.

[40] Albert Wu, Sadra Sadraddini, and Russ 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*, pages 4245–4251. IEEE, 2020.

[41] Xiaojing Zhang, Alexander Liniger, and Francesco Borrelli. Optimization-based collision avoidance. *IEEE Transactions on Control Systems Technology*, 29(3):972–983, 2020.