Bounded-Curvature Shortest Paths through a Sequence of Points

Xavier Goaoc, Hyo-Sil Kim, Sylvain Lazard

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Bounded-Curvature Shortest Paths through a Sequence of Points

Xavier Goaoc — Hyo-Sil Kim — Sylvain Lazard

N° 7465

November 2010

Rapport de recherche

ISSN 0249-6399 ISRN INRIA/RR--7465--FR+ENG

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE
INRIA
centre de recherche NANCY - GRAND EST

Bounded-Curvature Shortest Paths through a Sequence of Points

Xavier Goaoc* , Hyo-Sil Kim† , Sylvain Lazard*

Thème : Algorithmique, calcul certifié et cryptographie
Équipe-Projet Végas

Rapport de recherche n° 7465 — November 2010 — 50 pages

Abstract: We consider the problem of computing shortest paths having curvature at most one almost everywhere and visiting a sequence of n points in the plane in a given order. This problem arises naturally in path planning for point car-like robots in the presence of polygonal obstacles, and is also a sub-problem of the Dubins Traveling Salesman Problem.

This problem reduces to minimizing the function $F : \mathbb{R}^n \to \mathbb{R}$ that maps $(\theta_1, \ldots, \theta_n)$ to the length of a shortest curvature-constrained path that visits the points $p_1, \ldots, p_n$ in order and whose tangent in p_i makes an angle θ_i with the x-axis. We show that when consecutive points are distance at least 4 apart, all minima of F are realized over at most 2^k disjoint convex polyhedra over which F is strictly convex; each polyhedron is defined by $4n-1$ linear inequalities and k denotes, informally, the number of p_i such that the angle $\angle(p_{i-1}, p_i, p_{i+1})$ is small. A curvature-constrained shortest path visiting a sequence points can therefore be approximated by standard convex optimization methods, which presents an interesting alternative to the known polynomial-time algorithms that can only compute a multiplicative constant factor approximation.

Our technique also opens new perspectives for bounded-curvature path planning among polygonal obstacles. In particular, we show that, under certain conditions, if the sequence of points where a shortest path touches the obstacles is known then "connecting the dots" reduces to a family of convex optimization problems.

Key-words: Path planning, Bounded curvature, Dubins path, Convex optimization.

This work was supported by the INRIA *Équipe Associée* KI, the BK21 project, and Mid-career Researcher Program through NRF grant funded by the MEST (No. R01-2008-000-11607-0).

* INRIA Nancy Grand Est, LORIA laboratory, Nancy, France. `firstname.name@loria.fr`
† Department of Computer Science, KAIST, Daejeon, Korea. `hyosil@tclab.kaist.ac.kr`

Centre de recherche INRIA Nancy – Grand Est
LORIA, Technopôle de Nancy-Brabois, Campus scientifique,
615, rue du Jardin Botanique, BP 101, 54602 Villers-Lès-Nancy
Téléphone : +33 3 83 59 30 00 — Télécopie : +33 3 83 27 83 19

Plus Court Chemin de Courbure Bornée Passant par une Séquence de Points

Résumé : Nous considérons le problème du calcul d'un plus court chemin de courbure au plus 1 presque partout et visitant une séquence de points dans un ordre donné. Cette question apparaît naturellement en planification de trajectoire pour des véhicules de type voiture en présence d'obstacles polygonaux ainsi que dans le *problème du voyageur de commerce de Dubins*.

Ce problème se ramène à minimiser une fonction $F : \mathbb{R}^n \to \mathbb{R}$ qui à $(\theta_1, \ldots, \theta_n)$ associe la longueur du plus court chemin de courbure bornée qui visite les points $p_1, \ldots, p_n$ dans cet ordre et dont la tangente en p_i fait un angle θ_i avec l'axe des abscisses. Nous montrons que quand la distance entre points consécutifs est au moins 4, tous les minima de F sont atteints sur un domaine qui est l'union disjointe d'au plus 2^k polyèdres sur lesquels F est strictement convexe; chaque polyèdre est défini par $4n-1$ inégalités linéaires et k représente, informellement, le nombre de points p_i tels que l'angle $\angle(p_{i-1}, p_i, p_{i+1})$ est faible. On peut ainsi approximer un plus court chemin de courbure borné visitant une séquence de points par des techniques standard d'optimisation convexe. Cela représente une alternative intéressante aux méthodes antérieures, qui ne peuvent garantir qu'une approximation à un facteur constant multiplicatif.

Notre technique ouvre aussi de nouvelles perspectives concernant la planification de trajectoire à courbure bornée en présence d'obstacles polygonaux. En particulier, nous montrons que sous certaines conditions, si la séquence des points en lesquels le chemin touche les obstacles est connue alors relier ces points se ramène à une famille de probl'emes d'optimisation convexe.

Mots-clés : Planification de trajectoire, Courbure bornée, Chemins de Dubins, Optimisation convexe

1 Introduction

Path-planning problems involve computing feasible paths, possibly optimal for some criterion such as time or length, for a robot moving among obstacles. These problems are central in robotics and they have been widely studied; see, for instance, the books and survey papers [15, 18, 19, 29]. In its simplest form, path planning focuses on collision-free paths. However, robots generally come with physical limitations, such as bounds on the velocity, acceleration or curvature. Such differential constraints, called *nonholonomic*, restrict the geometry of the paths it can follow. Although there has been a considerable amount of work on nonholonomic motion planning in the robotics and control communities, relatively little work has been done, in comparison, from an algorithmic perspective.

In this paper, we study the path-planning problem for a car-like robot. The robot *configuration* is specified by both its location, a point p in $\mathbb{R}^2$ (typically, the midpoint of the rear axle), and its direction of travel which we represent by its polar angle θ in $\mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z}$. The robot is constrained to move in the forward direction, and its turning radius is bounded from below by a positive constant, which can be assumed to be equal to one by scaling the space. In this context, the robot follows *bounded-curvature* paths, that is, differentiable curves whose curvature is constrained to be at most one almost everywhere. Furthermore, for any robot configuration (p, θ), the oriented tangent to the curve at p has a polar angle θ.

The first results on curvature-constrained shortest paths go back to Dubins [12] in 1957. He proved that, in the plane without obstacles, bounded-curvature shortest paths consist of arcs of unit radius circles (C-segments) and straight line segments (S-segments); moreover, such shortest paths are of type CCC or CSC, or a substring thereof. These types of paths are generally referred to as *Dubins paths.*

We consider the problem of computing bounded-curvature shortest paths that visit, in order, a given sequence of n points in the plane (with no obstacles). This question is related to the problem of path planning in the presence of polygonal obstacles, because, roughly speaking, such a shortest path is also a locally shortest path through a sequence of points in the absence of obstacles (see below for details). This problem is also a sub-problem of the Dubins Traveling Salesman Problem which has been substantially studied in the robotics literature, for instance, in the context of UAV (unmanned air vehicles) path planning.

Our results. Let $p_1, \ldots, p_n$ be a sequence of points in the plane. For $1 < i < n$, we say that p_i is a *sharp turn* if the angle $\angle(p_{i-1}, p_i, p_{i+1})$ is acute, and one of p_i's neighbors, p_{i-1} or p_{i+1}, is within distance 4 from the segment joining p_i to its other neighbor. Let $F : (\mathbb{S}^1)^n \to \mathbb{R}$ map a sequence $(\theta_1, \ldots, \theta_n)$ of angles to the length of a shortest curvature-constrained path visiting the configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order. Since $\mathbb{S}^1$ can be *lifted to* (that is, represented by) some interval of length 2π in $\mathbb{R}$, we can also see F as a function from (a subset of) $\mathbb{R}^n$ to $\mathbb{R}$. Dubins' characterization implies that a curvature-constrained shortest path between two configurations can be computed in constant time. Computing a curvature-constrained shortest path visiting the points in order is thus equivalent, from a computational point of view, to finding a minimum of the function F. Our main result is the following:

Theorem 1. *Let $p_1, \ldots, p_n$ be a sequence of points in the plane that has k sharp turns and such that any two consecutive points are at least distance 4 apart. All global minima of F are realized in an open domain of $(\mathbb{S}^1)^n$ that can be lifted to a union of up to 2^k disjoint convex polyhedra in $\mathbb{R}^n$, each defined by at most $4n - 1$ linear inequalities, in two variables each.*[1] *Moreover, through this lifting, F is strictly convex over each of these polyhedra.*

Theorem 1 implies that a curvature-constrained shortest path visiting the points in order can be approximated by standard methods for convex optimization problems. Since convex optimization methods are known to be efficient in practice, this appears to be an interesting alternative to the known polynomial-time algorithms, which can only compute a *multiplicative* constant factor approximation. It should also be stressed that Theorem 1 easily extends to the case where the initial and final directions θ_1 and θ_n are prescribed (Theorem 28).

The technique we develop to prove Theorem 1 opens the way to a new approach to path planning among polygonal obstacles. It is known that a bounded-curvature shortest path between *two* configurations in the presence of polygonal obstacles is a concatenation of Dubins paths whose extremities are

[1] Note that, through the lifting, F remains continuous and strictly convex on the closure of each polyhedron.

extremal configurations or contact points on the boundary of the obstacles [14, 17]. We show that, given the sequence of contact points and knowing which one lie on *anchored* circular arcs (i.e., arcs of the path that touch the obstacle more than once), the reconstruction of the whole path reduces, under certain conditions, to a family of convex optimization problems (Theorem 31); in other words, "connecting the dots" is now easy, and the difficult task appears to be the discrete subproblem of computing the contact points and the anchored circular arcs.

We also establish several new properties of shortest paths of bounded curvature. In particular, we prove two fundamental results on Dubins paths of type CSC. First, if a CSC-path joining two configurations is such that its two circular arcs are shorter than π, then it is necessarily the shortest of all the CSC-paths (Proposition 5). Also, given any two points p_1 and p_2, the length $F(\theta_1, \theta_2)$ of the shortest CSC-path from (p_1, θ_1) to (p_2, θ_2) is C^2 and locally strictly convex over the domain where the two circular arcs are shorter than π (Theorem 6). We also prove that the length function F may have 2^{n-2} local minima (Corollary 10), each of them corresponding to a distinct *locally shortest path* in which the length of the circular arcs preceding and following every point p_i, $1 < i < n$, are equal or sum up to 2π (Proposition 9).

Our results say little on the theoretical complexity of approximating the shortest curvature-constrained path through a sequence of points since most convex optimization methods are notoriously hard to analyze. Nevertheless, in the absence of sharp turn ($k = 0$), and if consecutive points in the input are distance at least ς apart for some $\varsigma > 2 + \sqrt{5} \sim 4.24$, the ellipsoid method, although known to be inefficient both in theory and in practice, can be analyzed: it computes in $O\big(n^4 \log \frac{n}{\varepsilon}\big)$ time (in an extended real RAM model) a curvature-constrained path that visits $p_1, \ldots, p_n$ in order and whose length exceeds that of the (unique) optimum by at most ε (Corollary 24).

Previous work. As mentioned above, the study of bounded-curvature path planning started with Dubins' [12] first characterization of the geometry of shortest paths of bounded curvature, in the plane without obstacles. A more direct proof of this result, using ideas from control theory, was presented later by Boissonnat et al. [6] and Sussmann and Tang [32], independently.

The problem of computing shortest paths of bounded curvature through an ordered sequence of points was, to our knowledge, first considered by Bui [9] in 1994. Bui's high-level approach was similar to ours, that is to argue that a shortest path corresponds to a minimum of a convex function; unfortunately, several of the proofs from [9] have serious gaps. Bui also showed that a path of minimal length corresponds to a solution of one of 2^n algebraic systems, each consisting of $O(n)$ equations of bounded degree; even though this is totally unpractical, this solution illustrates well the difficulty of the problem.

An approximate solution can easily be obtained by considering the points as given on-line and greedily concatenating shortest paths from configuration (p_i, θ_i) to point p_{i+1} (where θ_1 is the polar angle of $\overrightarrow{p_1p_2}$ and $\theta_{i>1}$ is the polar angle at p_i of the path from (p_{i-1}, θ_{i-1}) to p_i); if p_i and p_{i+1} are at least distance $d \geqslant 4$ apart they are connected by a path of type CS whose length is at most $d + 2\pi - 2\arctan d$. When any two consecutive points are at least distance 4 apart, this greedy approach yields a path whose length is less than 1.91 times that of the optimum. Without lower bound on the distance between consecutive points the approximation factor of the greedy algorithm cannot be bounded. This was addressed in 2000 by Lee et al. [20] who presented a linear-time approximation algorithm for computing a path that is at most 5.03 times longer than the optimal one; we note that when the distance between any two consecutive points is at least $d > 2$, the guarantee on the approximation factor improves to $1 + \frac{2\pi}{d}$ which is less than 2.58 for $d = 4$.

The problem of computing shortest paths of bounded curvature through an *unordered* set of points, referred to as *Dubins TSP*, has also been studied [21, 23, 25, 28]; see also [24] for a short survey. Surprisingly, it is only recently that the corresponding decision problem was shown to be NP-hard [23]. All proposed approximation algorithms are based on a discretization of the directions at the via-points. The discretizations are, however, very rough: in essentially all cases, only one direction is chosen at each point. The stochastic version of this problem, in which the n targets are randomly distributed, has also been studied; see, e.g., [13, 16, 28].

Boissonnat and Lazard [7] also considered the related problem of computing the convex hull of bounded curvature of a set of points, that is the shortest bounded-curvature closed curve that encloses all the points. Here the path does not necessarily pass through every point. This simplifies the problem because it then reduces to computing the polygon of shortest perimeter whose vertices lie, in order, *inside* the unit disks centered at the vertices of the (regular) convex hull of the input points. Furthermore, the length of this polygon, defined over the Cartesian product of these disks, is shown to be a convex function, thus the minimum is unique and it can be computed by convex optimization.

Our problem is also related to the problem of computing bounded-curvature shortest paths in the presence of polygonal obstacles. Jacobs and Canny [17] proved the existence of a shortest path when there exists a feasible one.[2] They also proved, in parallel with Fortune and Wilfong [14], that such a shortest path consists of a concatenation of Dubins paths joined at points on the boundary of the obstacles. Fortune and Wilfong [14] also presented an exponential-time and space algorithm for deciding the existence of a feasible path between two configurations. A few years later, Reif and Wang [27] showed that the decision problem corresponding to finding a shortest path is NP-hard. Several approximation algorithms were proposed [17, 30, 33]; in particular, Wang and Agarwal [33] presented a $O(\frac{n^2}{\varepsilon^2}\log n)$-time algorithm for computing a $(1+\varepsilon)$-approximation of a shortest ε-*robust* path (informally, a path is ε-robust if it remains feasible after an ε-perturbation of the configurations touching the obstacles). It is only very recently that Backer and Kirkpatrick [3] presented the first polynomial-time algorithm that computes a $(1+\varepsilon)$-approximation of a shortest path, or reports that there is no path shorter than a given constant ℓ (the complexity is polynomial in terms of the number of polygon vertices, the number of bits of precision used to specify them, ε^{-1}, and ℓ). Shortest or feasible bounded-curvature paths have also been studied inside convex polygons [1], narrow corridors [5], and among obstacles of bounded curvature [2, 8].

Note finally that other models of car-like robots have also been studied. In particular, the Reeds and Shepp model [26], in which both forward and backward motions are allowed, has been extensively studied. Note also that other, and more general, dynamic constraints have been considered, and that Dubins paths have been generalized to the three-dimensional case [31]. We refer to [19] for a recent overview of such path planning problems.

Paper organization. After some preliminaries, the next section outlines the proof of our main result (Theorem 1). We then proceed to prove the local convexity of the length function, first between two configurations, in Section 3, then in $(\mathbb{S}^1)^n$ over a domain called *lemon* (due to its evocative shape in 2D), in Section 4. In Section 5, we prove that any global minimum of F belongs to a subset of the lemon region, and in Section 6, we show that this subset admits a "bounding box" that remains inside the lemon region and whose connected components are all convex. We then show, in Section 7, that many of these connected components can often be discarded by considering sharp turns, and conclude the proof of our main result. We analyze the complexity of the ellipsoid method applied to our problem in Section 8. In the first part of Section 9, we extend our results to the case where the initial and the final directions are prescribed, and in the second part of the section, we discuss the connection with the problem of computing curvature-constrained shortest paths in the presence of polygonal obstacles.

2 Preliminaries and proof outline

We give here some terminology and present a high-level sketch of the proof of our main result. We say that a sequence of points $p_1, \ldots, p_n$ satisfies the (D_d) *condition* if every two consecutive points p_i and p_{i+1} are at least distance d apart (the constant d is specified where needed).

Recall that a bounded-curvature path is a differentiable path whose curvature is defined almost everywhere, and is at most one in our case. A configuration is defined as a pair $(p, \theta) \in \mathbb{R}^2 \times \mathbb{S}^1$. A path goes from a configuration (p_s, θ_s) to a configuration (p_f, θ_f) if it starts in p_s and ends in p_f with the polar angles of its starting and final tangents being θ_s and θ_f, respectively. A shortest bounded-curvature path between two configurations, without obstacles, is a concatenation of at most three arcs that are circular arcs of unit radius (C-segments) or line segments (S-segments) [12]. Between any two configurations, there are at most six such paths, called Dubins paths, and their types are defined as CSC or CCC (or a subsequence thereof). We sometimes need to specify whether the path is turning right (clockwise) or left (counterclockwise) on a circle; we then refer to paths of types LSR, etc. Note that there is always a unique path of type, say LSR, between two configurations, because the circular arcs are considered to be shorter than 2π. We may also specify the length of an arc as an index; for instance, a path of type $L_\pi SR$ refer to a path that consists of a circular arc of length π turning left, a line segment, and a circular

[2] Note that this is not trivial and actually not true when reversals are allowed [11], that is for the model of Reeds and Shepp [26].

arc turning right. Finally, if the points p_s and p_f are at least distance 4 apart, then a shortest Dubins path between any two configurations (p_s, θ_s) and (p_f, θ_f) is of type CSC [10].[3] Moreover, we have the following straightforward property.

Lemma 2. *If p_s and p_f are at distance $d \geqslant 4$, then the shortest Dubins path between any configurations (p_s, θ_s) and (p_f, θ_f) has type CSC and the line segment has length at least $\sqrt{(d-2)^2 - 4}$.*

Proof. Let s be the length of the straight line segment and λ be the distance between the center of the circles supporting the two circular arcs of the CSC path. Either the line segment is an outer tangent to these two circles, in which case $s = \lambda$, or it is an inner tangent, in which case we get that $(\frac{s}{2})^2 + 1 = (\frac{\lambda}{2})^2$ by considering one of the two triangles induced by the non-simple quadrilateral formed by the two circle centers and the two segment endpoints. Hence, $s = \lambda$ or $s = \sqrt{\lambda^2 - 4}$, thus $s \geqslant \sqrt{\lambda^2 - 4}$ in all cases. The result follows since $\lambda + 2 \geqslant d$. □

We denote by $\mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z}$ the space of angles. In some cases, it will be more convenient to consider angles in $\mathbb{R}$, that is to lift $\mathbb{S}^1$ to some interval of length 2π in $\mathbb{R}$.

Proof outline. We outline here our proof of Theorem 1. Let $\mathcal{L}(\alpha)$ denote the set of points in $(\mathbb{S}^1)^n$ whose associated shortest path has all its circular arcs, between every consecutive pair of points p_i and p_{i+1}, of length less than α (we consider all possible shortest paths, if it is not unique). Recall that $F : (\mathbb{S}^1)^n \to \mathbb{R}$ is the function that maps a sequence $(\theta_1, \ldots, \theta_n)$ of angles to the length of a shortest curvature-constrained path visiting the configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order. Our first result, the cornerstone of this paper, is that, under the (D_4) condition, F is locally convex over $\mathcal{L}(\pi)$. We first show this for the case of two points (Theorem 6). The generalization to n points is then straightforward since F decomposes into

$$F(\theta_1, \ldots, \theta_n) = F_1^2(\theta_1, \theta_2) + \ldots + F_{n-1}^n(\theta_{n-1}, \theta_n),$$

where $F_i^{i+1}(\theta_i, \theta_{i+1})$ is the length of the shortest path from p_i to p_{i+1} and whose tangents in those points have polar angles θ_i and θ_{i+1}, respectively. We also prove that any global minimum of F must belong to $\mathcal{L}\left(\frac{3\pi}{4}\right)$ under the (D_4) condition (Lemma 11).

At this point, since F is locally convex on an open region containing all its global minima, one could hope that it has a unique global minimum which can be computed by convex optimization. However, the geometry of $\mathcal{L}(\alpha)$ turns out to be quite complicated. Two main issues are that $\mathcal{L}(\alpha)$ is, in general, not connected, and that the connected components are not convex, when lifted to $\mathbb{R}^n$. In particular, we give an example where $\mathcal{L}(\pi)$ has 2^{n-2} connected components, each of which contains a local minimum of F (Corollary 10).

To overcome these issues, we first show that there exists a "nice" region $\mathcal{D}$, called a *diamond*, such that $\mathcal{L}\left(\frac{3\pi}{4}\right) \subset \mathcal{D} \subset \mathcal{L}(\pi)$ (Lemma 16) and whose connected components are, once lifted to $\mathbb{R}^n$, convex polyhedra defined by $O(n)$ inequalities each (Lemma 17). The inclusion $\mathcal{D} \subset \mathcal{L}(\pi)$ and the convexity of these components guarantee that the function F is (globally) convex on each of these polyhedra. The other inclusion $\mathcal{L}\left(\frac{3\pi}{4}\right) \subset \mathcal{D}$ ensures that $\mathcal{D}$ contains all the global minima of F. Hence, all global minima of F can be computed by minimizing F over every connected component of $\mathcal{D}$ over which F is convex (Proposition 18). Unfortunately, there might still be $\Theta(2^n)$ such components.

To reduce the search space, consider the map $\sigma : (\mathbb{S}^1)^n \to \{-1, +1\}^n$ such that $\left(\sigma(\theta_1, \ldots, \theta_n)\right)_i$ is -1 if the vector with polar angle θ_i belongs to the positive cone of $\overrightarrow{p_i p_{i-1}}$ and $\overrightarrow{p_{i+1} p_i}$, and $+1$ otherwise. The map σ characterizes the connected components of $\mathcal{D}$, and we show that if p_i is not a sharp turn then $\left(\sigma(\theta_1, \ldots, \theta_n)\right)_i = 1$ for every global minimum $(\theta_1, \ldots, \theta_n)$ of F (Lemma 20). This narrows the search to 2^k connected components of $\mathcal{D}$, where k is the number of sharp turns among $p_1, \ldots, p_n$, and Theorem 1 follows.

If $k = 0$, we have that all the global minima of F belong to a single convex polyhedron on which F is strictly convex. It follows that F has a unique global minimum which corresponds to the unique shortest path (by Proposition 5). Furthermore, an analysis of the ellipsoid method yields that we can compute a path whose length is at most ε plus the length of the globally shortest path, in $O\left(n^4 \log \frac{n}{\varepsilon}\right)$ time (Corollary 24).

[3] Although this is not stated explicitly in [10], Bui et al. show (in §4.3) that if a shortest Dubins path is of type CCC, then p_f lies inside a disk of radius 2 (denoted $\mathcal{C}_H$) which contains p_s. Furthermore, if p_f lies on the boundary of that disk, then the path is also of type CSC with the line segment of length zero.

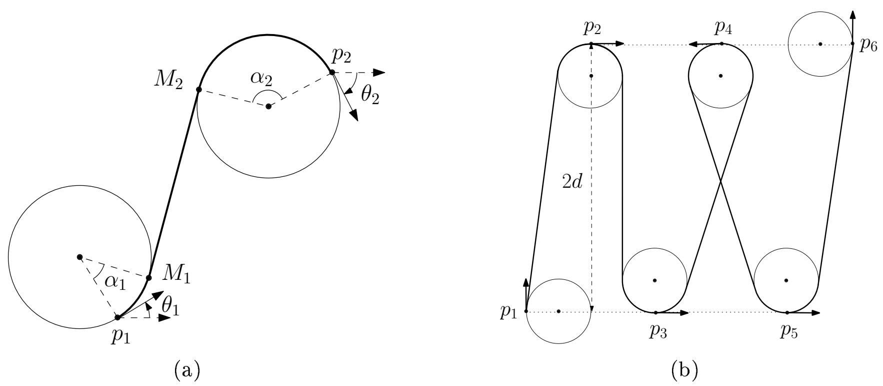

Figure 1: (a) LSR-path from (p_1, θ_1) to (p_2, θ_2). (b) Every sequence of polar angles $\theta_2, \ldots, \theta_{n-1}$ in $\{0, \pi\}$ defines a path that is arbitrarily close to a local optimum, for d large enough.

3 Local convexity of the length function of a CSC-path

In this section, we consider the case of two points, and let $F_{csc}(\theta_1, \theta_2)$ denote the length of a shortest CSC-path from a configuration (p_1, θ_1) to a configuration (p_2, θ_2).[4] We prove that F_{csc} is locally strictly convex[5] at any point (θ_1, θ_2) such that both circular arcs of the corresponding path are shorter than π. Figure 2 shows an example of such domain, and the graph of $F_{csc}(\theta_1, \theta_2)$ over that domain. We emphasize that the local convexity of F_{csc} holds without any distance assumption although we ignore here CCC paths which may be the shortest when p_1 and p_2 are distance less than 4 apart.

We start by showing that for a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the path of type T (T-path) from (p_1, θ_1) to (p_2, θ_2) is such a locally strictly convex function (Proposition 4). We prove this by computing the Hessian of the length function (Proposition 3). We then prove that this local convexity extends to the length $F_{csc}(\theta_1, \theta_2) = \min_{T \in \{LSR, RSL, LSL, RSR\}} F_T(\theta_1, \theta_2)$ of a shortest CSC-path from (p_1, θ_1) to (p_2, θ_2) (Theorem 6). We prove this by first showing the interesting property that, if both circular arcs of a CSC-path are shorter than π, then this path is *the* shortest CSC-path (Proposition 5).

The proofs are rather technical, and, for clarity, we provide here proof sketches and postpone the complete proofs to Appendix A.

Notation. Refer to Figure 1(a). For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, let $F_T(\theta_1, \theta_2)$ denote the length of the T-path from (p_1, θ_1) to (p_2, θ_2). For a given CSC-path, let α_i be the length of its i-th circular arc, let M_1 and M_2 be the first and last endpoint of its line segment, and let M_1M_2 denote the Euclidean distance from M_1 to M_2. Let μ_B be equal to 1 if B is true and to -1 otherwise. In particular, for a type of path $T \in \{LSR, RSL, LSL, RSR\}$, $\mu_{C_j=R}$ ($j = 1, 2$) is equal to 1 if the type of the j-th circular arc in T is R and it is equal to -1 otherwise. Finally, let $\delta_{i,j}$ equal to 1 if $i = j$ and to 0 otherwise.

We start by computing the first and second derivatives of the length function F_T, and its Hessian. The proof is calculatory and technical but a careful observation of the expressions involved leads to simple expressions of the derivatives.

Proposition 3. *For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is twice differentiable at any point (θ_1, θ_2) such that the corresponding T-path exists and none of its arcs vanishes. Furthermore,*

$$\frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} = \mu_{i=1}\mu_{C_i=R}(1 - \cos\alpha_i) \qquad \frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} = \delta_{i,j} \ \sin\alpha_i + \frac{\sin\alpha_i \sin\alpha_j}{M_1M_2}$$

[4]For two points ($n = 2$), the length function F was defined in Section 1 as the length of a shortest curvature-constrained path from (p_1, θ_1) to (p_2, θ_2). When d is sufficiently small, F_{csc} and F differ at points (θ_1, θ_2) such that the shortest path from (p_1, θ_1) to (p_2, θ_2) is (only) realized by a CCC-path. However, for $d \geqslant 4$, the shortest paths are always of type CSC, and thus F_{csc} and F coincide.

[5]The local convexity of F_{csc} naturally refers to the local convexity of $F_{csc} \circ \tau$, where τ is the quotient map from $\mathbb{R}^2$ to $\mathbb{S}^1 \times \mathbb{S}^1$; in other words, the angles are seen in $\mathbb{R}$.

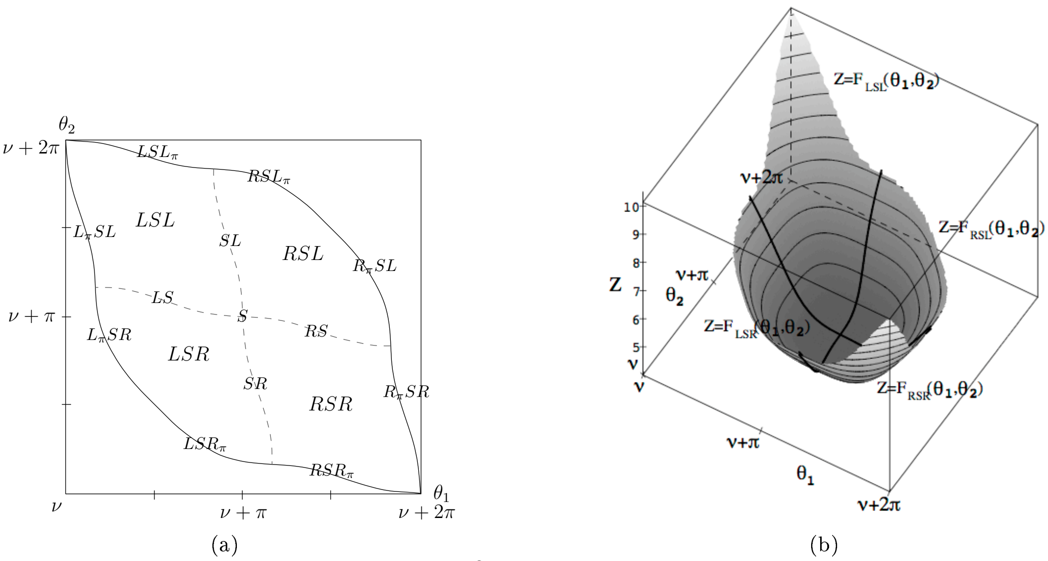

Figure 2: (a) Domain of (θ_1, θ_2) in $[\nu, \nu + 2\pi]^2$ such that both circular arcs of the shortest CSC-path from (p_1, θ_1) to (p_2, θ_2) are shorter than π, where $\|\overrightarrow{p_1 p_2}\| = 4$ and ν is the polar angle of $\overrightarrow{p_2 p_1}$ (for $\theta_1 = \theta_2 = \nu + \pi$, the path is reduced to a line segment). (b) Graph of F_{csc} over that domain.

and the determinant of the Hessian of F_T is

$$\sin\alpha_1 \sin\alpha_2 \left(1 + \frac{\sin\alpha_1 + \sin\alpha_2}{M_1 M_2}\right).$$

If (θ_1, θ_2) is such that the lengths α_1 and α_2 of the circular arcs of the T-path are in $(0, \pi)$, then $\frac{\partial^2 F_T}{\partial \theta_1^2}$ and the determinant of the Hessian of F_T are positive. This implies that F_T is positive definite (by Sylvester's criterion) and thus locally strictly convex at (θ_1, θ_2). As a consequence, we get the following result.

Proposition 4. *For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is locally strictly convex at any point (θ_1, θ_2) such that the corresponding T-path exists, none of its arcs vanishes, and its two circular arcs have length less than π.*

Note that we get trivially a similar result if we consider the length function in terms of only one angle, say θ_1 with θ_2 fixed: then $F_T(\theta_1)$ is locally strictly convex at any θ_1 such that $\alpha_1 \in (0, \pi)$.

Proposition 5. *If both circular arcs of a CSC-path from (p_1, θ_1) to (p_2, θ_2) are strictly shorter than π, then all the other distinct CSC-paths are strictly longer.*

Sketch of proof. We consider two geometrically distinct paths of type T and T' in $\{LSR, RSL, LSL, RSR\}$, from (p_1, θ_1) to (p_2, θ_2), such that both circular arcs of the T-path are shorter than π. We consider all possible types of T and T' in turn and we show using geometric arguments that, in every case, the line segment of the T-path is shorter than the one of the T'-path, and that the same holds for the total length of the circular arcs.

□

We can now prove the local convexity of the length function $F_{csc}(\theta_1, \theta_2)$ of the shortest CSC-paths from (p_1, θ_1) to (p_2, θ_2) on the domain of (θ_1, θ_2) such that both circular arcs are shorter than π. Figure 2 shows an example of such domain and of the graph of $F_{csc}(\theta_1, \theta_2)$ over that domain.

Graph of F_{csc} over that domain. ($|p_1 p_2| = 4$ and ν is the polar angle of $\overrightarrow{p_2 p_1}$).

Theorem 6. *The length $F_{csc}(\theta_1, \theta_2)$ of the shortest CSC-path from (p_1, θ_1) to (p_2, θ_2) is locally strictly convex at any point (θ_1, θ_2) such that both circular arcs of the corresponding path are strictly shorter than π. Furthermore, F_{csc} is C^2 at such a point.*

Sketch of proof. Propositions 4 and 5 yield that $F_{csc}(\theta_1, \theta_2)$ is locally strictly convex at any point such that both circular arcs of the shortest CSC-path have length in $(0, \pi)$. When one circular arc vanishes, two T and T'-paths coincide ($T \neq T'$ in $\{LSR, \ldots\}$), and Proposition 3 yields that F_{csc} is locally C^2 and thus locally convex at this point. Finally, we prove the *strict* local convexity at such a point by considering the third derivatives of the length functions of the T and T'-paths. □

4 Local convexity of F and lemon regions

Theorem 6 suggests that the length function $F(\theta_1, \ldots, \theta_n)$ is well-behaved on a certain domain of $(\mathbb{S}^1)^n$. In this section, we define this domain and analyze its geometric structure.

Definition of the lemons. Assume that condition (D_4) holds, and let $\alpha \in (0, \pi]$. Let $L_i^{i+1}(\alpha)$ denote the set of angles (θ_i, θ_{i+1}) in $(\mathbb{S}^1)^2$ such that both circular arcs of the shortest path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) have length (strictly) less than α. This set is well-defined because Lemma 2 ensures that the shortest path is of type CSC and Proposition 5 guarantees it is unique (since $\alpha \leqslant \pi$). Note that $L_i^{i+1}(\alpha)$ is an open set, since the circular arcs must be *strictly* shorter than α. Theorem 6 now simply asserts that the length function F_i^{i+1} is locally strictly convex on $L_i^{i+1}(\pi)$. We call $L_i^{i+1}(\alpha)$ a *lemon* region due to its evocative shape (see Figure 2(a)).

We now define the *n-dimensional lemon* region $\mathcal{L}(\alpha) \subset (\mathbb{S}^1)^n$ as the set of tuples $(\theta_1, \ldots, \theta_n)$ such that the shortest path visiting the configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ has all its circular arcs, between any two consecutive points p_i and p_{i+1}, of length less than α. The shortest path through a sequence of configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ is the concatenation of the shortest paths from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) for $i = 1, \ldots, n-1$. This ensures, with Proposition 5, that $\mathcal{L}(\alpha)$ is well-defined for any $\alpha \in (0, \pi]$. That also implies that a point $(\theta_1, \ldots, \theta_n)$ is in $\mathcal{L}(\alpha)$ if and only if, for $i = 1, \ldots, n-1$, the shortest path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) uses circular arcs of length less than α, that is $(\theta_i, \theta_{i+1}) \in L_i^{i+1}(\alpha)$. This rewrites as

$$\mathcal{L}(\alpha) = \bigcap_{i=1}^{n-1} (\mathbb{S}^1)^{i-1} \times L_i^{i+1}(\alpha) \times (\mathbb{S}^1)^{n-i-1}, \tag{1}$$

with the convention that $(\mathbb{S}^1)^0 \times A = A \times (\mathbb{S}^1)^0 = A$. Now, since $F(\theta_1, \ldots, \theta_n) = \sum_{i=1}^{n-1} F_i^{i+1}(\theta_i, \theta_{i+1})$, and a sum of locally convex functions is locally convex, Theorem 6 yields that F is locally convex over $\mathcal{L}(\pi)$. Furthermore, since $F_i^{i+1}(\theta_i, \theta_{i+1})$ is locally *strictly* convex over $L_i^{i+1}(\pi)$, it is easy to show that F is locally *strictly* convex over $\mathcal{L}(\pi)$, even though F_i^{i+1} is not locally strictly convex as a function of $(\theta_1, \ldots, \theta_n)$:

Proposition 7. *If the (D_4) condition holds, the length function $F(\theta_1, \ldots, \theta_n)$ is C^2 and locally strictly convex over $\mathcal{L}(\pi)$.*

Proof. First recall that, as shown in the proof of Theorem 6, the function F is C^2 over $\mathcal{L}(\pi)$. The Hessian of F is thus defined over $\mathcal{L}(\pi)$ and we prove that it is positive definite over $\mathcal{L}(\pi)$, which implies the result. The Hessian of $F = \sum_{i=1}^{n-1} F_i^{i+1}(\theta_i, \theta_{i+1})$ is

$$H = \begin{pmatrix}
\frac{\partial^2 F_1^2}{\partial \theta_1^2} & \frac{\partial^2 F_1^2}{\partial \theta_1 \partial \theta_2} & 0 & \cdots & & & & 0 \\
\frac{\partial^2 F_1^2}{\partial \theta_1 \partial \theta_2} & \frac{\partial^2 F_1^2}{\partial \theta_2^2} + \frac{\partial^2 F_2^3}{\partial \theta_2^2} & \frac{\partial^2 F_2^3}{\partial \theta_2 \partial \theta_3} & 0 & \cdots & & & 0 \\
0 & \frac{\partial^2 F_2^3}{\partial \theta_2 \partial \theta_3} & \frac{\partial^2 F_2^3}{\partial \theta_3^2} + \frac{\partial^2 F_3^4}{\partial \theta_3^2} & \frac{\partial^2 F_3^4}{\partial \theta_3 \partial \theta_4} & 0 & \cdots & & 0 \\
0 & 0 & \frac{\partial^2 F_3^4}{\partial \theta_3 \partial \theta_4} & \frac{\partial^2 F_3^4}{\partial \theta_4^2} + \frac{\partial^2 F_4^5}{\partial \theta_4^2} & \frac{\partial^2 F_4^5}{\partial \theta_4 \partial \theta_5} & 0 & \cdots & 0 \\
\vdots & & & & & & & \vdots \\
0 & \cdots & & & \cdots & 0 & \frac{\partial^2 F_{n-1}^n}{\partial \theta_{n-1} \partial \theta_n} & \frac{\partial^2 F_{n-1}^n}{\partial \theta_n^2}
\end{pmatrix}.$$

Hence, for any $\Theta = (\theta_1, \ldots, \theta_n)$,

$$\Theta^T H \Theta = \begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix}^T \begin{pmatrix}\frac{\partial^2 F_1^2}{\partial\theta_1^2} & \frac{\partial^2 F_1^2}{\partial\theta_1\partial\theta_2}\\ \frac{\partial^2 F_1^2}{\partial\theta_1\partial\theta_2} & \frac{\partial^2 F_1^2}{\partial\theta_2^2}\end{pmatrix}\begin{pmatrix}\theta_1\\ \theta_2\end{pmatrix} + \begin{pmatrix}\theta_2\\ \theta_3\end{pmatrix}^T \begin{pmatrix}\frac{\partial^2 F_2^3}{\partial\theta_2^2} & \frac{\partial^2 F_2^3}{\partial\theta_2\partial\theta_3}\\ \frac{\partial^2 F_2^3}{\partial\theta_2\partial\theta_3} & \frac{\partial^2 F_2^3}{\partial\theta_3^2}\end{pmatrix}\begin{pmatrix}\theta_2\\ \theta_3\end{pmatrix} + \cdots$$
$$+ \begin{pmatrix}\theta_{n-1}\\ \theta_n\end{pmatrix}^T \begin{pmatrix}\frac{\partial^2 F_{n-1}^n}{\partial\theta_{n-1}^2} & \frac{\partial^2 F_{n-1}^n}{\partial\theta_{n-1}\partial\theta_n}\\ \frac{\partial^2 F_{n-1}^n}{\partial\theta_{n-1}\partial\theta_n} & \frac{\partial^2 F_{n-1}^n}{\partial\theta_n^2}\end{pmatrix}\begin{pmatrix}\theta_{n-1}\\ \theta_n\end{pmatrix}.$$

For any Θ in $\mathcal{L}(\pi)$, (θ_i, θ_{i+1}) belongs to the lemon $L_i^{i+1}(\pi)$, for all $i = 1, \ldots, n-1$, and Theorem 6 implies that every term of the above sum is strictly positive. Hence H is positive definite over $\mathcal{L}(\pi)$ which concludes the proof. □

Lifting. Before we proceed to use the convexity of F over $\mathcal{L}(\pi)$, let us give an explicit lifting of $(\mathbb{S}^1)^n$ to some hypercube in $\mathbb{R}^n$ that preserves the connected components of $\mathcal{L}(\pi)$. Let ν_{i+1}^i denote the polar angles of $\overrightarrow{p_{i+1}p_i}$ and consider the family $\mathcal{H}$ of all the hyperplanes

$$H_i^+ = \{(\theta_1, \ldots, \theta_n) \in (\mathbb{S}^1)^n \mid \theta_i = \nu_{i+1}^i\}, \quad i = 1, \ldots, n-1,$$
$$H_i^- = \{(\theta_1, \ldots, \theta_n) \in (\mathbb{S}^1)^n \mid \theta_i = \nu_i^{i-1}\}, \quad i = 2, \ldots, n.$$

Lemma 8. *$\mathcal{L}(\pi)$ does not intersect any hyperplane from $\mathcal{H}$.*

Proof. In any CSC-path from (p_i, ν_{i+1}^i) to (p_{i+1}, θ_{i+1}), for any $\theta_{i+1} \in \mathbb{S}^1$, the circular arc following p_i has length at least π if $|p_ip_{i+1}| \geqslant 2$. Thus, $\mathcal{L}(\pi)$ does not intersect any hyperplane H_i^+, and similarly for H_i^-. □

For $i = 1, \ldots, n-1$, let Λ_i be the interval (closed on its left side and open on its right side) of length 2π that contains 0 and has its endpoints in $\nu_{i+1}^i + 2\pi\mathbb{Z}$, and let $\Lambda_n = \Lambda_{n-1}$. Now, let

$$\Lambda = \prod_{1 \leqslant i \leqslant n} \Lambda_i \subset \mathbb{R}^n.$$

For $1 < i < n$, Λ_i contains one point from $\nu_i^{i-1} + 2\pi\mathbb{Z}$ which splits it into two intervals; we denote the larger of these intervals by Λ_i^+ and the smaller by Λ_i^- (if the two intervals have the same length the names have no importance); by convention we let $\Lambda_1^+ = \Lambda_1^- = \Lambda_1$ and $\Lambda_n^+ = \Lambda_n^- = \Lambda_n$. Now, in the lifting of $(\mathbb{S}^1)^n$ to Λ, each cell of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ is lifted to a box $\prod_{1 \leqslant i \leqslant n} \Lambda_i^{\varepsilon_i}$ where $\varepsilon_i \in \{-, +\}$. An immediate consequence of Lemma 8 is that the image, through this lifting of every connected component of $\mathcal{L}(\pi)$ is connected.

Three issues. With Proposition 7, one could hope to use convex optimization methods to find the minimum of the length function F. This can only work if $\mathcal{L}(\pi)$ contains a global minimum of F. We show in Section 5 that under the (D_4) condition *any* global minimum of F lies in $\mathcal{L}(\frac{3\pi}{4}) \subset \mathcal{L}(\pi)$ (Lemma 11). This still leaves us with two issues.

On the one hand, it is clear from the example of Figure 2(a) that even for two points, $\mathcal{L}(\pi)$ may be non-convex; this means that there could be many (local) minima of F in every connected component of $\mathcal{L}(\pi)$. We handle this in Section 6 by describing a simple region $\mathcal{D}$, which we call a *diamond*, such that $\mathcal{L}(\frac{3\pi}{4}) \subset \mathcal{D} \subset \mathcal{L}(\pi)$ and the lifting $(\mathbb{S}^1)^n \to \Lambda$ maps each connected component of $\mathcal{D}$ to a convex polyhedron.

On the other hand, the regions $\mathcal{L}(\pi)$, $\mathcal{L}(\frac{3\pi}{4})$ and $\mathcal{D}$ may have exponentially many connected components. Indeed, by Lemma 8, $\mathcal{L}(\pi)$ is contained in the interior of the cells of the arrangement of $\mathcal{H}$, and Figure 1(b) shows an example of via-points for which (for d large enough) every choice of $(\theta_1, \ldots, \theta_n)$ in $\{\frac{\pi}{2}\} \times \{0, \pi\}^{n-2} \times \{\frac{\pi}{2}\}$ belongs to $\mathcal{L}(\frac{3\pi}{4})$ and to a distinct cell of that arrangement. In section 7 we give a geometric condition on (p_{i-1}, p_i, p_{i+1}) that implies that for $1 < i < n$, the polar angle of the tangent at p_i to any globally shortest path belongs to Λ_i^+(Lemma 20); this will avoid searching in each of the 2^{n-2} cells of $(\mathbb{S}^1)^n \setminus \mathcal{H}$.

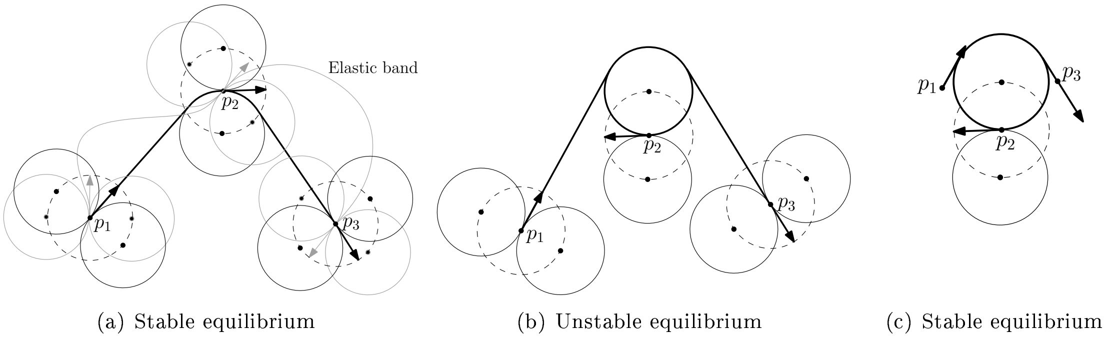

Figure 3: Bounded-curvature paths seen as equilibriums of a mechanical devices consisting of freely-rotating pulleys attached at p_1, p_2, p_3 and an elastic band.

5 Any global minimum of F belongs to $\mathcal{L}(\frac{3\pi}{4})$

We prove here that not only does any global minimum of the length function F belong to the region $\mathcal{L}(\pi)$, but it also belongs to the smaller region $\mathcal{L}(\frac{3\pi}{4})$. Let γ be a shortest bounded-curvature path through a sequence of configurations $(p_1, \tilde{\theta}_1), \ldots, (p_n, \tilde{\theta}_n)$. We assume that the points $p_1, \ldots, p_n$ satisfy condition (D_4) so that γ is of type CSC between any two consecutive configurations (by Lemma 2).

We first consider locally shortest paths, where γ is *locally shortest* if it cannot be shortened by perturbing the θ_i, that is if $\tilde{\Theta} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_n)$ is a local minimum of F. Locally shortest paths can be nearly characterized as follows (we postpone the complete proof to section 5.1):

Proposition 9. *If γ is a locally shortest path, then (i) its initial and final circular arcs vanish, (ii) the two circular arcs preceding and following every point p_i, $1 < i < n$, have the same orientation (R or L), and (iii) their lengths are either equal or sum up to 2π.*

Conversely, if γ satisfies (i) and (ii), and if (iii)' the lengths of the two circular arcs incident to every p_i, $1 < i < n$, are equal and strictly less than π, then γ is a locally shortest path; on the other hand, if this length is strictly larger than π for some p_i, then γ is not locally shortest if $(D_{2+\sqrt{5}})$ holds, and γ might be locally shortest otherwise.[6]

Sketch of proof. The first two claims follow from the expression of the first derivatives of the length of a CSC-path obtained in Proposition 3, and from the local convexity of the length function F of γ (Proposition 7). Concerning the last claim, Figures 3(b) and 3(c) show two configurations where the path has circular arcs of length more than π; in the former the path is not locally shortest, whereas it is in the latter. This follows from considerations on the length of the segments. If $(D_{2+\sqrt{5}})$ holds, then the segments are longer than 1 (by Lemma 2), and Proposition 3 yields that $\frac{\partial^2 F(\tilde{\Theta})}{\partial \theta_2^2}$ is strictly negative, which shows that $\tilde{\Theta}$ is not a local minimum of F. Conversely, a third-order Taylor expansion of the length function reveals that when one of the line segments is sufficiently short, γ is locally shortest. $\square$

Note that these properties are intuitively obvious if one consider a mechanical model where the path is modeled by an elastic band passing through the points $p_1, \ldots, p_n$ at which freely rotating double-pulleys (i.e., two tangent unit disks) are attached (see Figure 3 and [9]). It is, however, interesting to note the limitation of this intuition: if the two circular arcs before and after every point p_i have the same length, then the mechanical model is at an equilibrium; if these arcs are shorter than π, it seems clear that this equilibrium is stable, implying that the path is locally shortest, and Proposition 9 indeed proves it; however, if these circular arcs are strictly longer than π, it also seems fairly clear that the equilibrium is unstable, implying that the path is not locally shortest, but Proposition 9 shows that this is not necessary true and it depends on the length of the line segments.

[6] We do not claim that the bound $2 + \sqrt{5}$ is sharp in the constraint $(D_{2+\sqrt{5}})$.

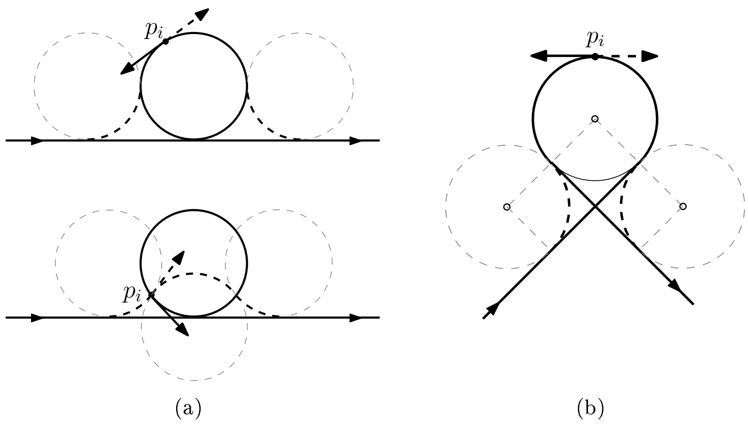

Figure 4: A globally shortest path has all its circular arcs of length at most $\frac{3\pi}{4}$ (if($D_{2+2\sqrt{2}}$) holds).

Corollary 10. *There might be at least 2^{n-2} locally shortest paths of bounded curvature through $p_1, \ldots, p_n$.*

Proof. This lower bound follows from the second statement of Proposition 9. Indeed, refer to Figure 1(b), and consider n points with coordinates $(2i, (-1)^i d)$ for some constant d, and the polar angles $\theta_1 = \frac{\pi}{2}$, $\theta_i = 0$ or π for $1 < i < n$, and $\theta_n = (-1)^n \frac{\pi}{2}$. If d is sufficiently large, then any of the 2^{n-2} choices of paths through $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ depicted in the figure are arbitrarily close to locally shortest paths. □

We can now prove the main result of this section. We first give a simple proof of this result under the $(D_{2+2\sqrt{2}})$ condition, and postpone the complete proof, under the (D_4) condition, to Section 5.2.

Lemma 11. *Under condition (D_4), in any globally shortest path γ the circular arcs preceding and following each via-point have length less than $\frac{3\pi}{4}$.*

Sketch of proof. We give here a simple proof of the lemma under the $(D_{2+2\sqrt{2}})$ condition. Let C_i^- and C_i^+ denote the circular arcs of γ that precede and follow p_i, respectively. By Proposition 9, since γ is a locally shortest path (and $2 + 2\sqrt{2} > 2 + \sqrt{5}$), the length of C_i^- and C_i^+ are equal and less than π, or their lengths sum up to 2π. By Lemma 2, $(D_{2+2\sqrt{2}})$ implies that the line segments preceding and following p_i have length at least 2; the path can thus be trivially shortened if C_i^- and C_i^+ sum up to 2π (see Figure 4(a)), contradicting the global optimality of γ. On the other hand, if the circular arcs have equal length in $[\frac{3\pi}{4}, \pi]$ (see Figure 4(b)), the fact that the segments preceding and following p_i have length at least 2 again allows to shorten the path (by convexity of the two circular shortcuts), and the statement follows. □

5.1 Proof of Proposition 9

Recall that γ is a locally shortest path that visit the configurations $(p_1, \tilde{\theta}_1), \ldots, (p_n, \tilde{\theta}_n)$ in order, i.e., $\tilde{\Theta} = (\tilde{\theta}_1, \ldots, \tilde{\theta}_n)$ is a local minimum of F, and γ is a shortest Dubins path between every two consecutive configurations. We prove Proposition 9, which partially characterizes the geometry of γ under the (D_4) condition (which ensures that γ is a concatenation of type CSC-paths). This characterization is, unfortunately, not complete even in the case where condition $(D_{2+\sqrt{5}})$ holds. To clarify the subtlety of the statement, we restate Proposition 9 in terms of the following properties:

(a) The initial and final circular arcs of γ vanish.
(b) The circular arcs that precede and follow p_i have the same orientation (R or L), for $1 < i < n$.
(c_i) The circular arcs that precede and follow p_i have same length.
(c_i^+) The circular arcs that precede and follow p_i have same length, which is strictly less than π.
(d_i) The lengths of the circular arcs that precede and follow p_i sum up to exactly 2π.

Then, Proposition 9 can be rewritten as follows:

$$
\begin{array}{rcll}
\gamma \text{ is locally shortest} & \Rightarrow & (a) \text{ and } (b) \text{ and, } \forall\, 1 < i < n, & (c_i) \text{ or } (d_i) \\
\gamma \text{ is locally shortest} & \Leftarrow & (a) \text{ and } (b) \text{ and, } \forall\, 1 < i < n, & (c_i^+) \\
(D_{2+\sqrt{5}}) \text{ and } \gamma \text{ is locally shortest} & \Rightarrow & (a) \text{ and } (b) \text{ and, } \forall\, 1 < i < n, & (c_i^+) \text{ or } (d_i) \\
\gamma \text{ is locally shortest} & \not\Rightarrow & (a) \text{ and } (b) \text{ and, } \forall\, 1 < i < n, & (c_i^+) \text{ or } (d_i).
\end{array}
$$

The following three lemmas prove these four statements. In the proofs, we denote by C_i^- (resp. C_i^+) the circular arc preceding (resp. following) p_i, and by α_i^- (resp. α_i^+) the length of this arc.

Lemma 12. *If γ is a locally shortest path, then (a) its initial and final circular arcs vanish, (b) the two circular arcs preceding and following every point p_i $(1 < i < n)$ have the same orientation (R or L), and their lengths are either (c_i) equal or (d_i) sum up to 2π.*

Proof. Since F is the minimum of several length functions (associated with different path types), it is difficult to determine where F is differentiable (we only know that F is differentiable over $\mathcal{L}(\pi)$, by Theorem 6). We thus consider, in the proof, the length function associated with the path type of γ, instead of F.

Recall that Proposition 3 states that the length of a CSC-path of type T from (p_1, θ_1) to (p_2, θ_2) is differentiable at any (θ_1, θ_2) such that the circular arcs of the corresponding T-path do not vanish, and for $i = 1, 2$:

$$\frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} = \mu_{i=1}\mu_{C_i=R}(1 - \cos\alpha_i).$$

This first implies Statement (a). Indeed, consider the subpath of γ between p_1 and p_2, let T be its type, and suppose for a contradiction that the circular arc at p_1 does not vanish. If the circular arc at p_2 does not vanish either, then F_T is differentiable at $\tilde{\theta}_1$ and its derivative, $\mu_{C_1=R}(1 - \cos\alpha_1)$, is nonzero; thus γ is not locally shortest, a contradiction. On the other hand, if the circular arc at p_2 vanishes, the type T is not uniquely defined, but it can be chosen so that the path changes continuously if θ_1 increases from $\tilde{\theta}_1$ (and, similarly, if θ_1 decreases); the length of the path thus changes continuously and may decrease since the derivative defined by continuity at $\tilde{\theta}_1$ is nonzero and does not depend on the orientation L or R at p_2. Hence, the initial arc of γ vanishes, and similarly for its final arc.

We now prove the rest of the lemma. Consider any nonterminal point p_i, and the subpath of γ between p_{i-1} and p_{i+1}; denote γ_i this subpath. Let $F_{\mathcal{T}}(\theta_i)$ be the length of the path from $(p_{i-1}, \tilde{\theta}_{i-1})$, through (p_i, θ_i), and to $(p_{i+1}, \tilde{\theta}_{i+1})$, whose type before and after p_i is that of γ_i (these types are not uniquely defined if some circular arc vanishes). Proposition 3 then yields that if the circular arcs of the subpath of γ do not vanish, then $F_{\mathcal{T}}(\theta_i)$ is differentiable at $\tilde{\theta}_i$, and

$$F'_{\mathcal{T}}(\tilde{\theta}_i) = -\mu_{C_i^-=R}(1 - \cos\alpha_i^-) + \mu_{C_i^+=R}(1 - \cos\alpha_i^+).$$

Since $\tilde{\theta}_i$ is a local minimum of $F_{\mathcal{T}}$, either $F'_{\mathcal{T}}(\tilde{\theta}_i) = 0$ or $F_{\mathcal{T}}$ is not differentiable at $\tilde{\theta}_i$. In the latter case, some circular arcs of γ_i vanishes, and the types of the CSC-path before and after p_i can then be chosen so that the corresponding path from $(p_{i-1}, \tilde{\theta}_{i-1})$, through (p_i, θ_i), and to $(p_{i+1}, \tilde{\theta}_{i+1})$ changes continuously, and so its length, when θ_i increases from $\tilde{\theta}_i$ (and, similarly, if θ_i decreases). Furthermore, the value of the derivative of $F_{\mathcal{T}}$ defined by continuity at $\tilde{\theta}_i$ is independent of that choice of type (since $\mu_{C_i^\pm=R}(1 - \cos\alpha_i^\pm) = 0$ when $C_i^\pm$ vanishes). Hence, if the derivative is negative, the length of the path decreases when θ_i increases from $\tilde{\theta}_i$, contradicting its optimality (and, similarly, if the derivative is positive). Therefore, $F'_{\mathcal{T}}(\tilde{\theta}_i) = 0$ in all cases.

Now, if the orientations (R or L) of the two circular arcs C_i^- and C_i^+ differ, $F'_{\mathcal{T}}(\tilde{\theta}_i) = \mu_{C_i^+=R}(2 - \cos\alpha_i^- - \cos\alpha_i^+)$ which is zero only if $\alpha_i^+ = \alpha_i^- = 0$; in that case, the arcs may be considered to have the same orientation, which implies Statement (b). It follows that $F'_{\mathcal{T}}(\theta_i) = \mu_{C_i^+=R}(\cos\alpha_i^- - \cos\alpha_i^+)$, which is zero only if $\alpha_i^- = \alpha_i^+$ or $\alpha_i^- + \alpha_i^+ = 2\pi$ modulo 2π; moreover these equalities are true not modulo 2π since $0 \leqslant \alpha_i^\pm < 2\pi$, which proves Statement (c_i-d_i). □

Lemma 13. *γ is a locally shortest path if (a) its initial and final circular arcs vanish, (b) the two circular arcs preceding and following every point p_i $(1 < i < n)$ have the same orientation (R or L), and (c_i^+) they have equal length, which is strictly shorter than π.*

Proof. From the expression of the first-order derivative of F (Proposition 3), we have that:

$$\left|\frac{\partial F(\Theta)}{\partial \theta_i}\right| = \begin{cases} 1-\cos\alpha_1^+ & \text{for } i=1, \\ \cos\alpha_i^- - \cos\alpha_i^+ & \text{for } 1<i<n, \\ 1-\cos\alpha_n^- & \text{for } i=n. \end{cases}$$

Since $\alpha_1^+ = \alpha_n^- = 0$ and $\alpha_i^- = \alpha_i^+$ for $1<i<n$, the gradient of F vanishes at any point Θ for which the path γ satisfies the hypotheses of the lemma. These hypotheses also imply, by Proposition 7, that the length function is locally strictly convex in Θ, which concludes the proof. □

Lemma 14. *If $(D_{2+\sqrt{5}})$ holds and γ is a locally shortest path then, for $1<i<n$, the total length of the two circular arcs preceding and following p_i is smaller or equal to 2π. This can be false if $(D_{2+\sqrt{5}})$ does not hold.*

Proof. We start by proving the first statement of the lemma. Consider a locally shortest path γ. By Lemma 12, the lengths α_i^+ and α_i^- of the two circular arcs incident to p_i $(1<i<n)$ are either equal or sum up to 2π. Assume, for a contradiction, that $\alpha_i^+ = \alpha_i^- \in (\pi, 2\pi)$ and that the condition $(D_{2+\sqrt{5}})$ holds (see Figure 3(b)). Proposition 3 yields that for $1<i<n$,

$$\frac{\partial^2 F(\Theta)}{\partial \theta_i^2} = \sin\alpha_i^- + \frac{\sin^2\alpha_i^-}{s_{i-1}^i} + \sin\alpha_i^+ + \frac{\sin^2\alpha_i^+}{s_i^{i+1}}, \tag{2}$$

where s_i^{i+1} denote the length of the line segment between p_i and p_{i+1} in γ. The condition $(D_{2+\sqrt{5}})$ ensures that $s_i^{i+1} \geqslant 1$ (Lemma 2), thus $\alpha_i^\pm \in (\pi, 2\pi)$ implies that $\frac{\partial^2 F(\Theta)}{\partial \theta_i^2} < 0$, contradicting the local optimality of γ. This proves the first statement.

We now prove the second part of the lemma. We consider a path γ through three points p_1, p_2, p_3 that consists, as in Figure 3(c), of a sufficiently short line segment, two circular circular arcs of equal length in $(\pi, 2\pi)$ (and same orientation L or R), and another sufficiently short line segment. Let $\tilde{\Theta} = (\tilde{\theta_1}, \tilde{\theta_2}, \tilde{\theta_3})$ be the sequence of polar angles of γ at p_1, p_2, p_3.

First note that, as in the proof of Lemma 13, the gradient of F vanishes at $\tilde{\Theta}$. Since the circular arcs at p_1 and p_3 vanish, Proposition 3 also yields that all partial second derivatives of F are zero at $\tilde{\Theta}$ except for $\frac{\partial^2 F(\tilde{\Theta})}{\partial \theta_2^2}$ which is strictly positive if the line segments of γ are sufficiently short (see Equation (2)). Hence, the second-order Taylor expansion of F at $\tilde{\Theta}$ is $F(\tilde{\Theta}+\Theta) = F(\tilde{\Theta}) + \frac{1}{2}\frac{\partial^2 F(\tilde{\Theta})}{\partial \theta_2^2}\theta_2^2 + o(\|\Theta\|^2)$, and it is strictly greater than $F(\tilde{\Theta})$ when $\theta_2 \neq 0$ and $\|\Theta\|$ is small enough.

When $\theta_2 = 0$, we consider the third-order Taylor expansion. However, since the length function F is not three-times differentiable at $\tilde{\Theta}$, we consider the length function $F_{T_{12}T_{23}} = F_{T_{12}} + F_{T_{23}}$ where $T_{i,i+1}$ is one of the two CSC-types of γ between p_i and p_{i+1} (there are two types since one circular arc vanishes), and $F_{T_{i,i+1}}$ is the length of the shortest path of type $T_{i,i+1}$ from (p_i, θ_i) to (p_{i+1}, θ_{i+1}). Then, when $\theta_2 = 0$, the third-order Taylor expansion of $F_{T_{12}T_{23}}$ at $\tilde{\Theta}$ is (see Equation (21)) $F_{T_{12}T_{23}}(\tilde{\Theta}+\Theta) = F_{T_{12}T_{23}}(\tilde{\Theta}) + \frac{1}{6}\frac{\partial^3 F_{T_{12}}(\tilde{\Theta})}{\partial \theta_1^3}\theta_1^3 + \frac{1}{6}\frac{\partial^3 F_{T_{23}}(\tilde{\Theta})}{\partial \theta_3^3}\theta_3^3 + o(\|\Theta\|^3)$. We have proved in the proof of Theorem 37 that, for $\|\Theta\|$ small enough, $\frac{\partial^3 F_{T_{12}}(\tilde{\Theta})}{\partial \theta_1^3}\theta_1^3 = \mu_{C_1=R}\,\theta_1^3$, which is strictly positive unless $\theta_1 = 0$, and similarly for $\frac{\partial^3 F_{T_{23}}(\tilde{\Theta})}{\partial \theta_3^3}\theta_3^3$. Hence, $F_{T_{12}T_{23}}(\tilde{\Theta}+\Theta) > F_{T_{12}T_{23}}(\tilde{\Theta})$ for $\|\Theta\| > 0$ small enough. Hence, independently of whether θ_2 vanishes, $F(\tilde{\Theta}+\Theta) > F(\tilde{\Theta})$ for all $\|\Theta\| > 0$ sufficiently small, and thus γ is locally shortest. We conclude the proof by noting that our proof of the second statement only requires that one of the two line segment of γ is sufficiently short. □

5.2 Proof of Lemma 11

We give here a complete proof of Lemma 11 which states that, under condition (D_4), in any globally shortest path γ the circular arcs preceding and following each via-point have length less than $\frac{3\pi}{4}$. We argue that if the length of a circular arc preceding or following p_i is at least $\frac{3\pi}{4}$ then the part of the path γ between p_{i-1} and p_{i+1} can be transformed into a shorter path that still visits p_{i-1}, p_i and p_{i+1} in that order, contradicting its global optimality.

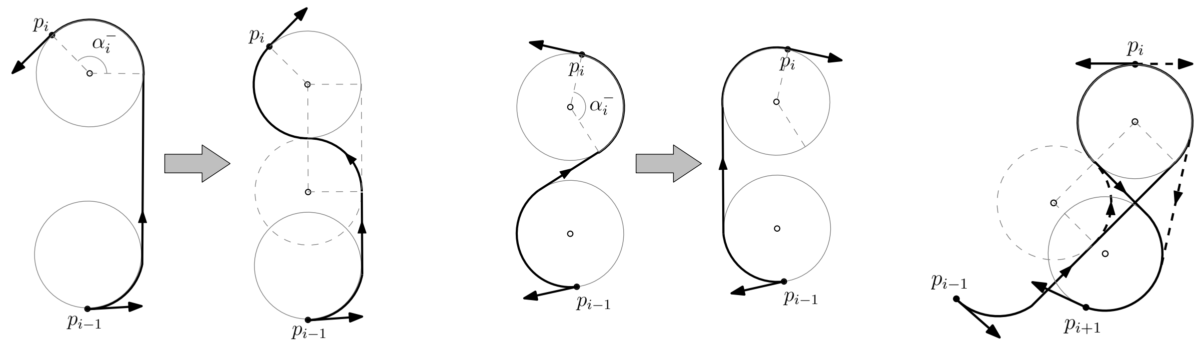

Figure 5: The transformations A (left) and B (center) and a combination of both (right).

The transformations. Let C_i^- and C_i^+ denote the circular arcs of γ that precede and follow p_i, respectively, let α_i^- and α_i^+ denote their lengths, and let s^- and s^+ denote the lengths of the line segments of the path γ that precedes and follows p_i respectively. We define two elementary transformations on the path from p_{i-1} to p_i; symmetric transformations between p_i and p_{i+1} are obtained by considering the path from p_{i+1} to p_i (traced backward).

Both transformations are illustrated in Figure 5. If $s^- \geqslant 2$, *transformation A* replaces the CSC path from (p_{i-1}, θ_{i-1}) to (p_i, θ_i) by a $CSCC$ from (p_{i-1}, θ_{i-1}) to $(p_i, \theta_i + \pi)$. If $s^- < 2$, the (D_4) condition implies that the path from p_{i-1} to p_i is of type LSR or RSL; *transformation B* then replaces it by a path of type LSL or RSR, respectively, using the same circles.

Both transformations change the configuration at p_i from (p_i, θ_i) into $(p_i, \theta_i + \pi)$, while leaving the configurations in p_{i-1} and p_{i+1} unchanged. We make three claims:

(i) If $s^- \geqslant 2$ and $\frac{3\pi}{4} \leqslant \alpha_i^- \leqslant \frac{3\pi}{2}$ then transformation A shortens the path from p_{i-1} to p_i.

(ii) If $s^- < 2$ and $\frac{3\pi}{4} \leqslant \alpha_i^-$ then transformation B shortens the path from p_{i-1} to p_i.

(iii) If $s^- < 2$ then $\alpha_i^- \leqslant \frac{3\pi}{2}$.

Symmetric claims apply to the transforms between p_{i+1} and p_i, under similar conditions on α_i^+ and s^+. Note that Claim (i) is straightforward, because, if $\alpha_i^- \leqslant \frac{3\pi}{2}$, the difference between the lengths of the original subpath and the new one is $(2 + \alpha_i^-) - (2\pi - \alpha_i^-) = 2(\alpha_i^- - \pi + 1)$, which is positive if $\alpha_i^- \geqslant \frac{3\pi}{4}$. Claims (ii) and (iii) are less straightforward and will be proved later.

Now, since γ is a locally shortest path, it must satisfy condition (iii) of Proposition 9, that is $\alpha_i^- = \alpha_i^+$ or $\alpha_i^- + \alpha_i^+ = 2\pi$. We treat the two cases separately.

Case where $\alpha_i^- = \alpha_i^+$. We can assume that the lengths of *both* circular arcs are at least $\frac{3\pi}{4}$. If s^- is at least 2 then we apply transformation A between p_{i-1} and p_i, otherwise we apply transformation B; we proceed similarly between p_{i+1} and p_i, obtaining a path from (p_{i-1}, θ_{i-1}) to (p_{i+1}, θ_{i+1}) through $(p_i, \theta_i + \pi)$.

If s^- and s^+ are smaller than 2, then Claim (ii) implies that the two transformations shorten the path γ. If exactly one of the two line segments is shorter than 2, then the circular arc, say C_i^-, of the corresponding CSC subpath has length at most $\frac{3\pi}{2}$ by Claim (iii), and thus the other arc C_i^+ is also shorter than or equal to $\frac{3\pi}{2}$. Hence, the two transformations also shorten the path γ. The remaining case is when both s^- and s^+ are larger than or equal to 2. The proof of Lemma 14 implies that if both segments are longer than 1 then $\alpha_i^- = \alpha_i^+ \leqslant \pi$; that $(D_{2+\sqrt{5}})$ does not hold is not an issue here, as its only purpose in the proof of Lemma 14 was to bound from below the length of the segments. Hence, $\alpha_i^- = \alpha_i^+ < \frac{3\pi}{2}$ and the above transformations again shorten the path γ, and thus concluding the proof when $\alpha_i^- = \alpha_i^+$

Case where $\alpha_\mathbf{i}^- + \alpha_\mathbf{i}^+ = \mathbf{2}\pi$. Let λ^- (resp. λ^+) denote the distance between the centers of the two circles supporting the circular arcs of the path γ between p_{i-1} and p_i (resp. p_i and p_{i+1}).

First, if both s^- and s^+ are larger than or equal to 2, the path can be trivially shortened as shown in Figure 4(a). Second, if both line segments are shorter than 2, we apply transformation B between p_{i-1}

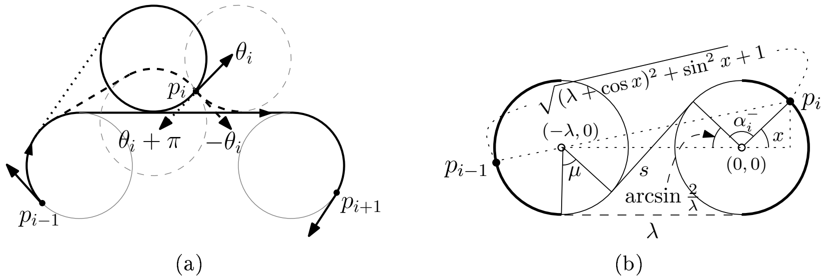

Figure 6: (a) The case where $\alpha_i^- + \alpha_i^+ = 2\pi$, $s^- < 2$, $s^+ \geqslant 2$, and $\alpha_i^+ > \frac{3\pi}{2}$. The dashed curve shortens the original path. (b) For the proof of Claims (ii) and (iii) in Section 5.2.

and p_i and between p_i and p_{i+1}. The difference between the lengths of the original subpath and the new one is $(2\alpha_i^- - \pi + s^- - \lambda^-) + (2\alpha_i^+ - \pi + s^+ - \lambda^+)$, as shown below in the proof of Claim (ii). Since $\alpha_i^- + \alpha_i^+ = 2\pi$, this difference is equal to $(\pi + s^- - \lambda^-) + (\pi + s^+ - \lambda^+)$, and $s^2 = \lambda^2 - 4$ implies that $\pi + s^\pm > \lambda^\pm$ which concludes the proof when both line segments are shorter than 2.

We are thus left with the case where exactly one line segment has length at least 2; assume, without loss of generality, that $s^- < 2$ and $s^+ \geqslant 2$. If $\alpha_i^+ \leqslant \frac{3\pi}{2}$, then applying transformation B from p_{i-1} to p_i and transformation A from p_i to p_{i+1} shorten the path, even if α_i^- or α_i^+ is shorter than $\frac{3\pi}{4}$; indeed, the difference between the lengths of the original subpath and the new one is $(2\alpha_i^- - \pi + s^- - \lambda^-) + (2\alpha_i^+ - 2\pi + 2)$. Since $\alpha_i^- + \alpha_i^+ = 2\pi$, this difference is equal to $\pi + s^- - \lambda^- + 2$ which is positive since, as we just noticed, $\pi + s^- > \lambda^-$.

We can thus assume that $s^- < 2$, $s^+ \geqslant 2$, and $\alpha_i^+ > \frac{3\pi}{2}$ (Figure 6(a)); furthermore, assume without loss of generality, that the line segments lie on the x-axis, that the arcs $C_i^\pm$ are oriented L. Then, $\theta_i \in [0, \frac{\pi}{2}]$ (since C_i^+ is longer than $\frac{3\pi}{2}$), and the path from p_{i-1} to p_i is of type RSL (since its segment is shorter than 2). As shown in Figure 6(a), we replace the RSL path from (p_{i-1}, θ_{i-1}) to (p_i, θ_i) by a RSR path (dashed) from (p_{i-1}, θ_{i-1}) to $(p_i, -\theta_i)$, and we replace the LSC path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) by a RSC path from $(p_i, -\theta_i)$ to (p_{i+1}, θ_{i+1}) that uses the same final circular arc. Between p_i and p_{i+1}, the difference between the lengths of the original subpath and the new one is larger than $\alpha_i^+ - (2\pi - \alpha_i^+) = 2\alpha_i^+ - 2\pi$. Between p_{i-1} and p_i, the (dashed) transformed path is shorter than the (dotted) RSR path from (p_{i-1}, θ_{i-1}) to $(p_i, \theta_i + \pi)$; call this latter path the intermediate path.[7] Hence, the difference between the lengths of the original path and the transformed path is larger than the difference between the lengths of the original path and the intermediate path, which is $2\alpha_i^- - \pi + s^- - \lambda^-$ (see the proof of Claim (ii)). Thus, the difference between the lengths of the original subpath from p_{i-1} to p_{i+1} and the transformed path is $(2\alpha_i^- - \pi + s^- - \lambda^-) + (2\alpha_i^+ - 2\pi) = \pi + s^- - \lambda^-$, which is positive as noted above. It remains to prove Claims (ii) and (iii).

Proof of Claims (ii) and (iii). We first argue that the (D_4) condition and $s^- < 2$ imply that p_{i-1} and p_i lie respectively on the two half-circles shown in bold in Figure 6(b). Suppose, without loss of generality, that p_i lies on the unit circle $\mathcal{C}_i$ centered at $(0,0)$, and p_{i-1} lies on the unit circle $\mathcal{C}_{i-1}$ centered at $(-\lambda, 0)$. Let x denote the polar angle of p_i, or more precisely a measure in $[-\pi, \pi]$ of this angle. When p_i and $\mathcal{C}_{i-1}$ are fixed, the maximum of the distance $\|p_{i-1}p_i\|$ is $\sqrt{(\lambda + \cos x)^2 + \sin^2 x} + 1$ since it is realized when the segment $p_i p_{i-1}$ contains the center of $\mathcal{C}_{i-1}$. See Figure 6(b). By the (D_4) assumption, this distance is at least 4, which gives $\cos x \geqslant \frac{8-\lambda^2}{2\lambda}$. Since $s^- < 2$, $\lambda = \sqrt{s^{-2} + 4} \in [2, 2\sqrt{2}]$, which implies $\frac{8-\lambda^2}{2\lambda} \in [0,1]$. Thus $x \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, that is p_i lies on the bold half-circle of $\mathcal{C}_i$ in Figure 6(b), and similarly for p_{i-1}. Furthermore, $|x| \leqslant \arccos \frac{8-\lambda^2}{2\lambda}$.

It follows that $\alpha_i^- \leqslant \frac{3\pi}{2}$ (Claim (iii)) and also that the difference between the lengths of the original subpath and the new one is, with the notation of Figure 6(b), $(\mu + s^- + \alpha_i^-) - (\lambda + \pi + \mu - \alpha_i^-) =$

[7] This is straightforward by a convexity argument. Indeed, the (dotted) RSR path from (p_{i-1}, θ_{i-1}) to $(p_i, \theta_i + \pi)$ lies on the boundary of its convex hull (this follows from the fact, shown in the proof of Claim (iii), that p_{i-1} and p_i lie on the opposite half-circles of the circle supporting the original path from p_{i-1} to p_i, as shown in Figure 6(b)). The (dashed) RSR path (p_{i-1}, θ_{i-1}) to $(p_i, -\theta_i)$ has the same property, and its convex hull is included in the other convex hull. It follows that latter path is shorter than the former one.

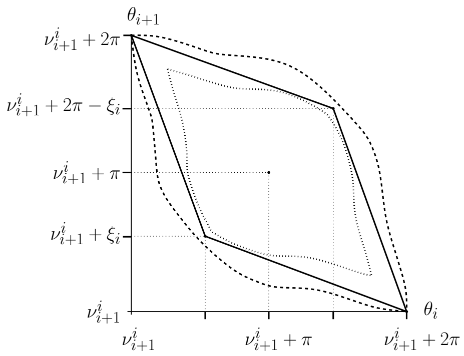

Figure 7: Diamond D_i^{i+1}, contained in $L_i^{i+1}(\pi)$ (dashed) and containing $L_i^{i+1}(\frac{3\pi}{4})$ (dotted) for $|p_ip_{i+1}| = 4$.

$2\alpha_i^- - \pi + s^- - \lambda$. We prove that this difference, denoted G, is nonnegative. If $s^- \geqslant \frac{4}{\pi} - \frac{\pi}{4}$, then $\pi s^- + \left(\frac{\pi}{2}\right)^2 \geqslant 4$, and $(s^- + \frac{\pi}{2})^2 \geqslant {s^-}^2 + 4 = \lambda^2$; thus $s^- - \lambda \geqslant -\frac{\pi}{2}$ and $G \geqslant 2\alpha_i^- - \frac{3\pi}{2}$ which is nonnegative when $\alpha_i^- \geqslant \frac{3\pi}{4}$.

Now, suppose that $s^- < \frac{4}{\pi} - \frac{\pi}{4}$, which implies that $\lambda = \sqrt{{s^-}^2 + 4} \in [2, \frac{4}{\pi} + \frac{\pi}{4}]$. Consider first the case where $x \geqslant 0$ and refer to Figure 6(b). Then $\alpha_i^- = \pi - \arccos\frac{2}{\lambda} - x \geqslant \pi - \arccos\frac{2}{\lambda} - \arccos\frac{8-\lambda^2}{2\lambda}$, and thus $G \geqslant \pi - 2\arccos\frac{2}{\lambda} - 2\arccos\frac{8-\lambda^2}{2\lambda} + \sqrt{\lambda^2 - 4} - \lambda$. G is positive for $\lambda \in [2, \frac{4}{\pi} + \frac{\pi}{4}]$ because the right-hand side of this inequality, denoted E, is positive on that interval; indeed, the derivative of E with respect to λ is always negative[8] and E is positive[9] for $\lambda = \frac{4}{\pi} + \frac{\pi}{4}$.

For $x < 0$, $\alpha_i^- \geqslant \pi - \arccos\frac{2}{\lambda}$, and, similarly as above, $G \geqslant \pi - 2\arccos\frac{2}{\lambda} + \sqrt{\lambda^2 - 4} - \lambda$. We proved above that the derivative of the right-hand side is always negative, and we prove, similarly as before, that the right-hand side expression is positive[10] for $\lambda = \frac{4}{\pi} + \frac{\pi}{4}$. Hence, G is again positive, which concludes the proof of Claim (ii), and of the lemma.

6 Convex bounding boxes of the connected components of $\mathcal{L}(\frac{3\pi}{4})$

In this section we prove that there exists a simple shape, which we call the *diamond* $\mathcal{D}$, that contains $\mathcal{L}(\frac{3\pi}{4})$ and is contained in $\mathcal{L}(\pi)$. It will then follow that the lift of $\mathcal{D}$ to the hypercube Λ defined in Section 4 consists of $O(2^n)$ *convex* components, over which F is convex, and which contain the global minimum of F.

Consider two consecutive points p_i and p_{i+1}, let $d_i = |p_ip_{i+1}|$ and let $\xi_i = 2\pi/(d_i - 1/d_i)$. Recall that ν_{i+1}^i denotes the polar angle of the vector $\overrightarrow{p_{i+1}p_i}$. We define D_i^{i+1} as the image of the open quadrilateral with vertices $(0, 2\pi)$, (ξ_i, ξ_i), $(2\pi, 0)$, and $(2\pi - \xi_i, 2\pi - \xi_i)$ under the translation of vector $(\nu_{i+1}^i, \nu_{i+1}^i)$ (see Figure 7). For two points, this shape has the following property (the proof is technical and is postponed to Appendix B, for clarity):

[8] On one hand, the derivative of $-2\arccos\frac{8-\lambda^2}{2\lambda}$ is negative (recall that the derivative of $\arccos f(\lambda)$ is $-\frac{f'(\lambda)}{\sqrt{1-f^2(\lambda)}}$). On the other hand, the derivative of the other terms of E is equal to $-\frac{4}{\lambda\sqrt{\lambda^2-4}} + \frac{\lambda}{\sqrt{\lambda^2-4}} - 1$ which is negative or zero if and only if $\lambda - \frac{4}{\lambda} \leqslant \sqrt{\lambda^2 - 4}$; the left-hand side is nonnegative because $\lambda \geqslant 2$ (by the D_4 assumption), and the inequality is thus equivalent (after squaring) to $\lambda \geqslant 2$, which proves that the derivative of E is always negative.

[9] The fact that E is positive for $\lambda = \frac{4}{\pi} + \frac{\pi}{4}$ can be observed by evaluating numerically the function (which gives roughly 0.25). More formally, this can be proved by first replacing $-2\arccos z$ by $-\pi + 2\arcsin z$ and considering the terms of degree up to 9 in the power series of arcsin. Since, all the terms in the power series are positive, the resulting expression bounds E from below. This expression is (up to the positive factor $\frac{1}{2705829396480\,\pi^9\,(16+\pi^2)^9}$) a polynomial expression of degree 36 in π, with integer coefficients. The expression can be shown to be positive by regarding it as a polynomial in x, then by applying the sequence of change of variable $y = 10\,x$, $z = y - 31$, $u = 1/z$, and $v = u - 1$, which transforms the interval $(3.1, 3.2)$ in x into the interval $(0, +\infty)$ in v. All the coefficients of the resulting polynomial are positive, which implies that the initial polynomial in x is positive over $(3.1, 3.2)$, and thus is positive for $x = \pi$.

[10] By first replacing $-2\arccos z$ by $-\pi + 2\arcsin z$ and using that $\arcsin z > z$ for $z > 0$, we get that $G \geqslant \pi + (-\pi + 2\frac{2}{\lambda}) + \sqrt{\lambda^2 - 4} - \lambda$, which is equal, for $\lambda = \frac{4}{\pi} + \frac{\pi}{4} = \frac{16+\pi^2}{4\pi}$, to $\frac{16\pi}{16+\pi^2} + (\frac{4}{\pi} - \frac{\pi}{4}) - (\frac{4}{\pi} + \frac{\pi}{4}) = \frac{\pi(16-\pi^2)}{2(16+\pi^2)} > 0$.

Lemma 15. *If* $|p_i p_{i+1}| \geqslant 4$ *then* $L_i^{i+1}\left(\frac{3\pi}{4}\right) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$.

Sketch of proof. We first observe that $L_i^{i+1}(\alpha)$ is symmetric with respect to the lines $y = x$ and $y = 2\nu_{i+1}^i + 2\pi - x$, so it suffices to prove the inclusions in one quadrant. We focus on the quadrant $\nu_{i+1}^i \leqslant x \leqslant y \leqslant 2\nu_{i+1}^i + 2\pi - x$, and show that in this quadrant, the boundary of $L_i^{i+1}(\alpha)$ consists of points whose corresponding shortest path has type $L_\alpha SL$ and $L_\alpha SR$. We then obtain analytical expressions for these two arcs and prove that the segment bounding D_i^{i+1} in that quadrant separates the arcs for $\alpha = \frac{3\pi}{4}$ and $\alpha = \pi$. □

We then extend this construction to an arbitrary number of points by defining the *diamond of* $p_1, \dots, p_n$ as

$$\mathcal{D} = \bigcap_{i=1}^{n-1} (\mathbb{S}^1)^{i-1} \times D_i^{i+1} \times (\mathbb{S}^1)^{n-i-1}, \tag{3}$$

where D_i^{i+1} denotes the 2-dimensional diamond of p_i and p_{i+1} defined above (with the convention that $(\mathbb{S}^1)^0 \times A = A \times (\mathbb{S}^1)^0 = A$). Notice the similarity with Equation (1):

$$\mathcal{L}(\alpha) = \bigcap_{i=1}^{n} (\mathbb{S}^1)^{i-1} \times L_i^{i+1}(\alpha) \times (\mathbb{S}^1)^{n-i-1}.$$

The inclusions $L_i^{i+1}\left(\frac{3\pi}{4}\right) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$ from Lemma 15 then yield:

Corollary 16. *If* $p_1, \dots, p_n$ *satisfy the* (D_4) *condition then* $\mathcal{L}\left(\frac{3\pi}{4}\right) \subset \mathcal{D} \subset \mathcal{L}(\pi)$.

We now describe the image of $\mathcal{D}$ through the lifting $(\mathbb{S}^1)^n \to \Lambda$, defined in Section 4 from the family of hyperplanes $\mathcal{H}$.

Lemma 17. *$\mathcal{D}$ does not intersects the hyperplanes of $\mathcal{H}$, and the lifting $(\mathbb{S}^1)^n \to \Lambda$ maps the intersection of $\mathcal{D}$ with each cell of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ to a convex polyhedron defined by at most $4(n-1)$ linear inequalities in $\mathbb{R}^n$, each involving 2 variables.*

Proof. By definition, D_i^{i+1} intersects none of the lines $\{(\theta_i, \theta_{i+1}) \in (\mathbb{S}^1)^2 \mid \theta_i = \nu_{i+1}^i\}$ and $\{(\theta_i, \theta_{i+1}) \in (\mathbb{S}^1)^2 \mid \theta_{i+1} = \nu_{i+1}^i\}$, thus $\mathcal{D}$ intersects none of the hyperplanes of $\mathcal{H}$.

If I_1 and I_2 are two intervals in $\mathbb{R}$ whose interiors avoid $\nu_{i+1}^i + 2\pi\mathbb{Z}$, then the image of D_i^{i+1} through the (partial) lift $\mathbb{S}^1 \times \mathbb{S}^1 \to I_1 \times I_2$ is connected, and thus convex by definition of D_i^{i+1}. Now, every cell of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ is mapped through the lift $(\mathbb{S}^1)^n \to \Lambda$ to a box of the type $\Xi = \prod_i \Lambda_i^{\varepsilon_i}$ where $\varepsilon_i \in \{-, +\}$, and the image of $\mathcal{D}$ through the (partial) lift $(\mathbb{S}^1)^n \to \Xi$ is convex. In other words, the lift $(\mathbb{S}^1)^n \to \Lambda$ maps each connected component of $\mathcal{D}$ to a convex polyhedron in $\mathbb{R}^n$.

Each diamond D_i^{i+1} is defined by 4 linear inequalities involving θ_i and θ_{i+1}. Every such inequality, defined over $(\mathbb{S}^1)^2$, can be lifted through $(\mathbb{S}^1)^2 \to \Lambda_i^{\varepsilon_i} \times \Lambda_{i+1}^{\varepsilon_{i+1}}$ to at most *one* inequality in $\mathbb{R}^2$, since the image of D_i^{i+1} through $(\mathbb{S}^1)^2 \to \Lambda_i^{\varepsilon_i} \times \Lambda_{i+1}^{\varepsilon_{i+1}}$ is convex. Hence, the image of $\mathcal{D}$ through the (partial) lift $(\mathbb{S}^1)^n \to \Xi$ is a convex polyhedron; this polyhedron can be defined by at most $4(n-1)$ linear inequalities in $\mathbb{R}^n$, each involving 2 variables. The result follows. □

We can already prove that computing a global shortest path through a sequence of points reduces to solving $O(2^n)$ convex optimizations problems in n variables and with $O(n)$ constraints each.

Proposition 18. *If $p_1, \dots, p_n$ satisfy the (D_4) condition, any global minimum of F is contained in the image of $\mathcal{D}$ through the lifting $(\mathbb{S}^1)^n \to \Lambda \subseteq \mathbb{R}^n$, which consists of at most 2^{n-2} convex polyhedra, each defined by at most $4(n-1)$ linear inequalities in 2 variables each. Furthermore, F is strictly convex over each of these convex polyhedra.*

Proof. The global minimum of F is realized over $\mathcal{D}$ as it is realized over $\mathcal{L}\left(\frac{3\pi}{4}\right)$ (by Lemma 11) and $\mathcal{L}\left(\frac{3\pi}{4}\right) \subset \mathcal{D}$ (by Corollary 16). For each cell c in $(\mathbb{S}^1)^n \setminus \mathcal{H}$, the lifting $(\mathbb{S}^1)^n \to \Lambda$ maps $\mathcal{D} \cap c$ to a convex polyhedron in $\mathbb{R}^n$ defined by $O(n)$ linear inequalities (by Lemma 17). Moreover, the length function F is convex on each such polyhedron since it is convex on $\mathcal{L}(\pi)$ (by Proposition 7) and $\mathcal{D} \subset \mathcal{L}(\pi)$ (by Corollary 16). Since $(\mathbb{S}^1)^n \setminus \mathcal{H}$ has at most 2^{n-2} cells, the statement follows. □

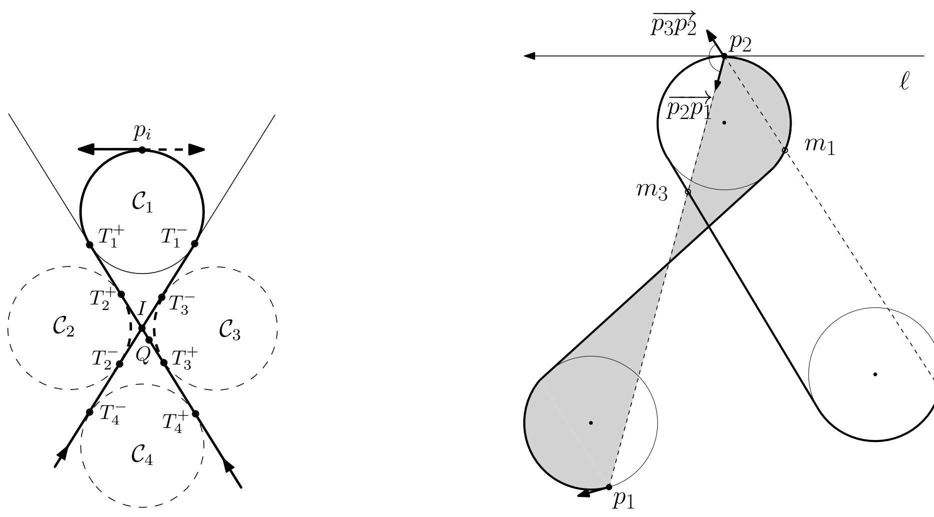

Figure 8: (a) Shortening a path that self-intersects. (b) If $\theta_2 \in \Lambda_2^-$ and p_2 is not a sharp turn, the path self-intersects. (The vectors are not drawn to scale for clarity.)

7 Pruning the connected components of $\mathcal{D}$

With Proposition 18, the computation of a global minimum of F reduces to solving a convex optimization problem on each cell of $(\mathbb{S}^1)^n \setminus \mathcal{H}$, where each of the cell is mapped to a product $\prod_i \Lambda_i^{\varepsilon_i}$ through the lifting $(\mathbb{S}^1)^n \to \Lambda$; in this section we give a condition on the respective positions of p_{i-1}, p_i and p_{i+1} that implies that only the cells where $\varepsilon_i = +$ can contain a global minimum of F. This will reduce the search from 2^{n-2} to 2^k cells, where k is the number of triples of consecutive points that violate our condition.

Sharp turns and self-intersections. Let φ_i denote the angle between $\overrightarrow{p_ip_{i-1}}$ and $\overrightarrow{p_ip_{i+1}}$ that is smaller or equal to π. We say that p_{i-1}, p_i, p_{i+1} form a *sharp turn*, or for simplicity that p_i is a *sharp turn*, if $|\varphi_i| \leqslant \frac{\pi}{2}$ and if one of its neighbors, p_{i-1} and p_{i+1}, is within distance 4 from the segment formed by p_i and its other neighbor. Note that, if (p_{i-1}, p_i, p_{i+1}) satisfies the (D_4) condition, then p_i is a sharp turn if and only if the latter condition is satisfied, that is, $|\sin\varphi_i| \leqslant \frac{4}{\min(|p_{i-1}p_i|,|p_ip_{i+1}|)}$.

We start by observing that "local" self-intersection of the globally shortest path through the via-points can be traced back to sharp turns.

Lemma 19. *If $p_1, \ldots, p_n$ satisfy the (D_4) condition and p_i is not a sharp turn then the portion, from p_{i-1} to p_{i+1}, of a globally shortest path has no self-intersection.*

Proof. Let γ denote the portion, from p_{i-1} to p_{i+1}, of a globally shortest path. We assume that γ self-intersects, and prove that p_i is a sharp turn. Let C_t^- and C_t^+ denote the circular arcs of γ that precede and follow p_t, respectively, and S^- and S^+ the line segments that precede and follow p_i. The self-intersection of γ is an intersection between two elements in $\{C_{i-1}^+, S^-, C_i^-, C_i^+, S^+, C_{i+1}^-\}$. We discuss the various situations in turn.

Assume first that two circular arcs intersect. It can be neither C_{i-1}^+ and $C_i^\pm$ (since $|p_{i-1}p_i| > 4$), nor C_{i+1}^- and $C_i^\pm$ (since $|p_ip_{i+1}| > 4$), nor C_i^- and C_i^+ (since their length is at most $\frac{3\pi}{4}$ by Lemma 11). Then C_{i-1}^+ intersects C_{i+1}^-, which implies that $|p_{i-1}p_{i+1}| \leqslant 4$ and p_i is a sharp turn.

Any point in S^+ is within distance at most 2 from the segment $p_{i+1}p_i$. Thus, if C_{i-1}^+ intersects S^+ the distance from p_{i-1} to the segment $p_{i+1}p_i$ is at most 4, and p_i is a sharp turn. A similar argument handles the symmetric case where C_{i+1}^- intersects S^-. Since both S^+ and S^- are tangent to the circle supporting C_i^- and C_i^+ and to, respectively, C_{i+1}^- and C_{i-1}^+, no other intersection between a circular arc and a segment is possible.

Now assume that S^- and S^+ intersect in some point I. Let L^- and L^+ denote the two lines supporting S^- and S^+. We consider the four unit circles $\mathcal{C}_1, \ldots, \mathcal{C}_4$ tangent to both lines, and label the line/circle intersections as shown in Figure 8(a). We assume that T_1^+, T_2^+, T_3^+, and T_4^+ appear in this order on L^+; this is without loss of generality because the arcs $C_i^\pm$ are shorter than $\frac{3\pi}{4}$, by Lemma 11. If T_3^+ lies on S^+ and T_2^- lies on S^- then we can shorten γ using arcs of $\mathcal{C}_2$ and $\mathcal{C}_3$ (see Figure 8(a)), which contradicts the assumption that γ is part of a global shortest path. So, assume that $T_3^+ \notin S^+$ (the other case is handled similarly) and let Q denote the endpoint of S^+ other than T_1^+. Since Q lies between T_3^+ and I, each of the two unit circles tangent to L^+ at Q intersect L^-. The ray starting at T_1^- and containing S^- then intersects each of the disks of radius 2 centered in p_{i-1} and p_{i+1}. If that ray meets the disk centered in p_{i-1} first, then p_{i-1} is distance at most 4 from the segment $p_i p_{i+1}$, and a similar argument holds in the symmetric case. Hence, p_i is a sharp turn, which concludes the proof. □

Sharp turns and tangents at p_i. Let $1 < i < n$, $(\varepsilon_1, \ldots, \varepsilon_n) \in \{-,+\}^n$ and consider a point $(\theta_1, \ldots, \theta_n) \in \prod_i \Lambda_i^{\varepsilon_i}$. Let γ denote a shortest curvature-constrained path visiting the configurations (p_i, θ_i) in order. Then $\varepsilon_i = -$ if and only if the tangent to the path at p_i lies in the positive cone of the vectors $\overrightarrow{p_i p_{i-1}}$ and $\overrightarrow{p_{i+1} p_i}$ (see Figure 8(b)). We now show that if p_i is not a sharp turn then the tangent at p_i to any globally shortest path must lie outside this positive cone, and thus only the cells $\prod_i \Lambda_i^{\varepsilon_i}$ with $\varepsilon_i = +$ need to be considered when searching a global minimum of F.

Lemma 20. *If $p_1, \ldots, p_n$ satisfy the (D_4) condition, and p_i is not a sharp turn, then for any globally shortest path, the polar angle θ_i of the tangent at p_i belongs to Λ_i^+.*

Proof. Let γ be a globally shortest path and assume, for a contradiction, that the polar angle θ_i of its tangent vector at p_i is in Λ_i^-, and that p_i is not a sharp turn. We show that the arc of γ from p_{i-1} to p_{i+1} has a self-intersection, a contradiction with Lemma 19.

For simplicity, we consider $i = 2$ and refer to Figure 8(b). Let ℓ be the oriented line tangent to the (oriented) path γ at p_2, and, for any two distinct points a and b, let (ab) denote the oriented line from a to b. Without loss of generality, we assume that p_1 is to the left[11] of ℓ; since $\theta_i \in \Lambda_i^-$, p_3 is to the left of ℓ and to the right of (p_1p_2). Moreover, if p_3 is on (p_1p_2) then p_2 is a sharp turn, so p_3 is *strictly* to the right of (p_1p_2) and, similarly, p_1 is strictly to the left of (p_3p_2).

For $j = 1, 3$ let γ_j denote the portion of γ between p_j and p_2. By Lemma 11, the circular arcs of γ_j have length at most $\frac{3\pi}{4}$, which is strictly less than π. Thus, p_1 and p_3 are strictly to the left of ℓ, and γ_1 and γ_3 are also entirely strictly to the left of ℓ, except for p_2.

We now argue that γ_1 intersects the line (p_2p_3) in p_2 and exactly one other point, denoted m_1, at which γ_1 traverses (p_2p_3). Let c be the number of intersection points between $\gamma_1 \setminus p_2$ and (p_3p_2), counted with multiplicity. We first observe that γ crosses the line (p_3p_2) from right to left in p_2, and p_1 is strictly to the left of (p_3p_2), so c must be odd. Next, γ_1 intersects any line other than the line supporting "its" segment, in at most three points, counted with multiplicity. Indeed, since the circular arcs have length at most π, if a circular arc meets the line in two points (possibly identical), then the segment does not intersect the line in another point, and the second circular arc intersects the line in at most one point (counted with multiplicity). Since p_2 contributes at least one to this count, c must be at most two. Since c must also be odd, $c = 1$ and γ_1 intersects (p_2p_3) in p_2 and exactly one other point m_1, at which it traverses this line.

We furthermore prove that m_1 belongs to the open segment $[p_2p_3]$. First, since γ_1 is to the left of ℓ, so is m_1. As p_2 is on ℓ and p_3 is strictly to the left of ℓ, it follows that either m_1 belongs to the segment $[p_2p_3]$, or p_3 belongs to the segment $[p_2m_1]$. In the latter case, p_3 lies in the convex hull of γ_1 and is thus within distance at most 2 from the segment of γ_1; this is impossible, as it would imply that p_3 is at distance at most 4 from the segment p_1p_2, i.e. that p_2 is a sharp turn. Hence m_1 belongs to the closed segment $[p_2p_3]$. Finally, $m_1 \neq p_2$ by definition, and $m_1 \neq p_3$ because otherwise p_3 lies in the convex hull of γ_1, again requiring that p_2 be a sharp turn.

Similarly, γ_3 intersects (p_1p_2) in p_2 and exactly one other point, denoted m_3 which belongs to the open segment $[p_1p_2]$.

Consider now the curve ρ obtained as the union of γ_1 and the segment $[p_1p_2]$. It is closed, and thus delimits a bounded region $\mathcal{R}$ (not necessarily connected, see Figure 8(b)). Note that the line (p_2p_3) meets ρ in exactly $\{m_1, p_2\}$, and m_1 lies strictly in between p_2 and p_3, thus p_3 lies strictly outside the region $\mathcal{R}$.

[11] Unless specified otherwise, the constraint to be to the left (or to the right) of an oriented line is considered non-strict.

Consider finally the intersection between γ_3 and ρ. Let γ'_j be γ_j minus its endpoint p_2. γ'_3 intersects the line (p_1p_2) in exactly one point, m_3, at which it traverses (p_2p_3). Since m_3 lies on the open segment $[p_1p_2]$, either m_3 is a (the) point of self-intersection of ρ, or γ'_3 intersects the interior of $\mathcal{R}$ in a neighborhood of m_3. In the former case, m_3 then lies on γ'_1 and thus γ is self-intersecting between p_1 and p_3, contradicting Lemma 19. In the latter case, when γ'_3 intersects the interior of $\mathcal{R}$, γ'_3 must intersect ρ in some other point because γ'_3 does not intersect $\mathcal{R}$ in some neighborhoods of p_2 and p_3. Since γ_3 is simple and intersects line (p_2p_3) only at p_2 and m_3, γ'_3 must intersect γ'_1. Then, again, γ is self-intersecting between p_1 and p_3, concluding the proof. □

Proof of Theorem 1. We can now complete the proof of our main result (Theorem 1) which states the following:

> *Let $p_1, \ldots, p_n$ be a sequence of points in the plane that satisfy the (D_4) condition and has k sharp turns. All global minima of F are realized in a domain of $(\mathbb{S}^1)^n$ that can be lifted to a union of up to 2^k disjoint convex polyhedra, each defined by $O(n)$ linear inequalities in $\mathbb{R}^n$. Moreover, through this lifting, F is strictly convex over each of these polyhedra.*

Proof of Theorem 1. By Proposition 18, under the (D_4) condition, the global minima of the length function F are contained, through the lifting, in 2^{n-2} disjoint convex polyhedra, each defined by $O(n)$ linear inequalities, over which F is strictly convex. If p_i is a non-sharp turn then any global minimum of F has its θ_i-coordinate in Λ_i^+. The number of cells of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ that remains to explore is thus 2^k, and for each of them it suffices to solve a convex optimization problem on a n-dimensional polyhedron defined by $O(n)$ linear inequalities. □

8 Convex optimization

There are essentially two general methods that solve convex optimization problems with guaranteed complexity, the ellipsoid method and the class of interior point methods. While the latter is usually more efficient, it only works if one can compute so-called *self-concordant barriers* for the function to be minimized and its constraints. In our problem, the function F to be minimized is defined as the minimum of four functions which use inverse trigonometric functions, and computing *self-concordant barriers* currently seems out of reach. We thus focus on the ellipsoid method. We consider an extended real RAM model where arithmetic operations, trigonometric and inverse trigonometric functions, and the min function can be evaluated in constant time over the reals.

Overview of the ellipsoid method. Let us recall briefly its principles as described in the book of Ben-Tal and Nemirovski [4, Lecture 5] (see also the book of Gärtner and Matoušek [22, Chapter 7] for a quick overview). Given $\varepsilon > 0$, a convex function $f : \mathbb{R}^n \to \mathbb{R}$, and a closed non-empty convex domain $X \subset \mathbb{R}^n$, we want to compute $\hat{x} \in X$ such that $f(\hat{x}) \leqslant \min_{x \in X} f(x) + \varepsilon$. Here, X is one of the polyhedra over which we want to minimize F, and f is a convex extension over $\mathbb{R}^n$ of the restriction of F to X (see Section 8.1 for details). The method starts with some ellipsoid $\mathcal{E}_0$ (typically a ball) containing X, and constructs a sequence of ellipsoids $(\mathcal{E}_i)_{i \geqslant 1}$ such that $\mathcal{E}_i$ contains the minimum of f and the volume of $\mathcal{E}_i$ decreases exponentially fast. Given the ellipsoid $\mathcal{E}_i$ with center c_i, $\mathcal{E}_{i+1}$ is constructed as follows. A *separation oracle* is first used to test if c_i is contained in the domain X and, if not, to obtain a hyperplane H separating it from the domain. If c_i belongs to the domain, a *first-order oracle* is used to compute $f(c_i)$ and the gradient $\overrightarrow{n}$ of f in c_i; let H denote the hyperplane through c_i with normal $\overrightarrow{n}$. Then, $\mathcal{E}_{i+1}$ is defined as the ellipsoid with minimum volume that contains the portion of $\mathcal{E}_i$ bounded by H and containing the minimum of f. At step i, an approximation $\hat{x}$ is given by the center c_j of $\mathcal{E}_j$, $j \leqslant i$, that lies in X and minimizes f. The cost of one iteration of this construction is quadratic in the number of constraints.

Lemma 21. *$\mathcal{E}_{i+1}$ can be computed from $\mathcal{E}_i$ in $O(n^2)$ time.*

Proof. Testing if c_i is contained in the domain X takes $O(n)$ time, since the domain is defined by $O(n)$ linear inequalities, each involving only two variables (by Lemma 17). If c_i violates one inequality, we can obtain a separating hyperplane H from that inequality in $O(1)$ time. If c_i is in the domain, we can compute $f(c_i) = F(c_i)$ by adding up the length of the shortest CSC path joining each pair of consecutive configurations, and the gradient of F in c_i by using Proposition 3; altogether this takes $O(n)$ time. The complexity of the computation of $\mathcal{E}_{i+1}$ from $\mathcal{E}_i$ and H is dominated by the multiplication of a $n \times n$ matrix and a n-dimensional vector [4, Lecture 5]; it can thus be done in $O(n^2)$ time. □

Number of iterations. The number of steps needed to achieve an additive error of ε on the solution depends on the geometry of the domain:

Theorem 22 ([4, Theorem 5.2.1]). *Let the convex set X of the problem contain a Euclidean ball of a given radius $r > 0$ and be contained in the ball $\mathcal{E}_0 = \{||x||^2 < R\}$ of a given radius R. For every input accuracy $\varepsilon > 0$, the ellipsoid method terminates after no more than*

$$N(\varepsilon) = Ceil\left(2n^2\left(\ln\left(\frac{R}{r}\right) + \ln\left(\frac{\varepsilon + Var_R(f)}{\varepsilon}\right)\right)\right) + 1$$

steps, where

$$Var_R(f) = \max_{x\in\mathcal{E}_0} f(x) - \min_{x\in\mathcal{E}_0} f(x).$$

Moreover the result $\hat{x}$ generated by the method belongs to X and satisfies $f(\hat{x}) \leqslant \min_{x\in X} f(x) + \varepsilon$.

In our setting, bounding $\frac{R}{r}$ and $\mathrm{Var}_R(f)$ requires an analysis of the geometry of the domain. If we consider any cell in $(\mathbb{S}^1)^n \setminus \mathcal{H}$ then already for $n = 3$ the ratio $\frac{R}{r}$ can be arbitrarily large[12] There is, however, one cell for which we can bound $N(\varepsilon)$ in terms of n and ε (we postpone the rather technical proof to Section 8.1):

Lemma 23. *For any $\varsigma > 2+\sqrt{5}$, if condition (D_ς) holds and the considered domain X lies in $\prod_{1\leqslant i\leqslant n} \Lambda_i^+$, then $N(\varepsilon) = O(n^2 \ln \frac{n}{\varepsilon})$.*

Recall that, if none of the p_i is a sharp turn, then, by Lemmas 17 and 20, the minimum of F lies in a convex polyhedron contained in $\prod_{1\leqslant i\leqslant n} \Lambda_i^+$. Thus, this cell is the only one we need to consider, and the convexity of F ensures that there is a unique globally shortest path. Moreover, we have:

Corollary 24. *Let $\varsigma > 2 + \sqrt{5}$. If $p_1, \ldots, p_n$ is a sequence of points in the plane that satisfy the (D_ς) condition and has no sharp turn then we can compute in time $O\big(n^4 \log \frac{n}{\varepsilon}\big)$ a path of curvature at most 1 that visits the p_i in order and whose length is at most that of the globally shortest path plus ε.*

8.1 Proof of Lemma 23

We give bounds on r, R and $\mathrm{Var}_R(f)$, and the statement will follow from the expression of $N(\varepsilon)$ from Theorem 22. Bounding R is easy, as the domain is contained in Λ which is, by construction, contained in $[-2\pi, 2\pi]^n$; we then have that $R^2 \leqslant n(2\pi)^2$ and $R = O(\sqrt{n})$.

Bounding r is less straightforward. Recall that the diamond D_i^{i+1} is defined as the image of the quadrilateral with vertices $(0, 2\pi)$, (ξ_i, ξ_i), $(2\pi, 0)$, and $(2\pi - \xi_i, 2\pi - \xi_i)$ under the translation of vector $(\nu_{i+1}^i, \nu_{i+1}^i)$ where $\xi_i = \frac{2\pi}{d_i - \frac{1}{d_i}}$ (see Section 6). Let $\square_i^{i+1}$ be the square in $\Lambda_i \times \Lambda_{i+1}$ of side-length $2\pi - 2\xi_i$ whose bottom-left corner is $(\nu_{i+1}^i + \xi_i, \nu_{i+1}^i + \xi_i)$; see Figure 9. Just like we proceeded to define $\mathcal{D}$ from the D_i^{i+1}, we extend each $\square_i^{i+1}$ into a cylinder and intersect them all:

$$\square = \bigcap_{i=1}^{n}\left(\prod_{1\leqslant j<i} \Lambda_j\right) \times \square_i^{i+1} \times \left(\prod_{i+1<j\leqslant n} \Lambda_j\right),$$

with the convention that $A \times \prod_{j\in J} \Lambda_j = A$ whenever $J = \emptyset$. Since $\square_i^{i+1}$ is contained in the lift of D_i^{i+1}, it follows that $\square$ is contained in the lift of $\mathcal{D}$. Let $\Lambda^+ = \prod_{1\leqslant i\leqslant n} \Lambda_i^+$. Now, any ball contained in the box $\square \cap \Lambda^+$ is contained in our domain. It thus remains to bound the side-length of $\square \cap \Lambda^+$ from below to get a lower bound on r.

[12] Consider three points at $p_1 = (-d, 0)$, $p_2 = (0, 2)$ and $p_3 = (d, 0)$, and consider the cell $\Lambda_1 \times \Lambda_2^- \times \Lambda_3$ (see Section 4) as d goes to infinity. Since the length of Λ_1 and Λ_3 is equal to 2π, independently of d, R remains larger than π. However, since the length of Λ_2^- goes to 0, so does r and the ratio R/r is therefore unbounded.

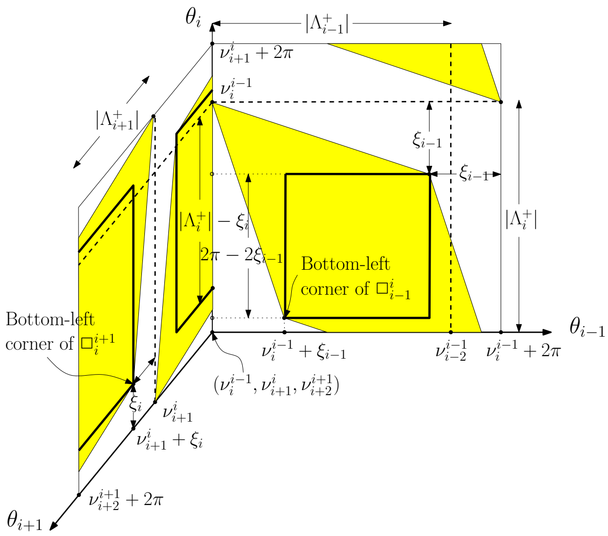

Figure 9: Determining a non-empty hypercube in the cell $X \subseteq \prod_{1 \leqslant i \leqslant n} \Lambda^+_i$.

The projection of $\square \cap \Lambda^+$ on the i^{th}-coordinate axis is $A \cap B$, where A and B denote, respectively, the projections of $\square^i_{i-1} \cap (\Lambda^+_{i-1} \times \Lambda^+_i)$ and $\square^{i+1}_i \cap (\Lambda^+_i \times \Lambda^+_{i+1})$ on the i^{th}-coordinate axis. The length of A is either $2\pi - 2\xi_{i-1}$ or $|\Lambda^+_i| - \xi_{i-1}$ and that of B is either $2\pi - 2\xi_i$ or $|\Lambda^+_i| - \xi_i$ (see Figure 9). Let $\xi = \max_i \xi_i$. Since $|A| + |B| - |A \cap B| \leqslant |\Lambda^+_i|$, we are in one of the three cases:

$$|A \cap B| \geqslant 4\pi - 4\xi - |\Lambda^+_i|, \quad \text{or} \quad |A \cap B| \geqslant 2\pi - 3\xi, \quad \text{or } |A \cap B| \geqslant |\Lambda^+_i| - 2\xi.$$

Since $2\pi \geqslant |\Lambda^+_i| \geqslant \pi$ we get in all three cases that

$$|A \cap B| \geqslant \min(2\pi - 4\xi, 2\pi - 3\xi, \pi - 2\xi) = \pi - 2\xi \quad \text{if } \xi \leqslant \frac{\pi}{2}.$$

By solving a degree-two equation, we get that $\xi_i = \frac{2\pi}{d_i - \frac{1}{d_i}} < \frac{\pi}{2}$ is equivalent to $d_i > 2 + \sqrt{5}$. Thus, for any constant $\varsigma > 2 + \sqrt{5}$ the (D_ς) condition ensures that $\xi < \frac{\pi}{2} - \varsigma'$ for some positive constant ς', and thus that $|A \cap B|$ and the side-length of $\square \cap \Lambda^+$ is bounded from below by some positive constant.[13] This box contains a ball of radius $\Omega(1)$, and thus $r = \Omega(1)$.

We now turn our attention to $\mathrm{Var}_R(f)$. We start by showing how F can easily be extended into a convex function, f, over $\mathbb{R}^n$. Let H^{i+1}_i denote the family of all the tangent planes to the graph of F^{i+1}_i over the intersection of $\Lambda^+_i \times \Lambda^+_{i+1}$ with the lift of D^{i+1}_i to $\Lambda_i \times \Lambda_{i+1}$. We let $f^{i+1}_i(x_i, x_{i+1})$ denote the function over $\mathbb{R}^2$ whose graph is the lower envelope of H^{i+1}_i, let $f = \sum_{i=1}^{n-1} f^{i+1}_i$, and note that f is convex over $\mathbb{R}^n$ and coincides with F over the domain.

We now bound $\mathrm{Var}_R(f)$. First, we have that $F^{i+1}_i(\theta_i, \theta_{i+1}) \geqslant |p_i p_{i+1}|$ and since the domain contains the global minimum of F in its interior, H^{i+1}_i contains a horizontal hyperplane and thus the same lower bound holds for $f^{i+1}_i(\theta_i, \theta_{i+1})$. It follows that:

$$\min_{x \in \mathcal{E}_0} f(x) \geqslant \sum_{i=1}^{n-1} |p_i p_{i+1}|.$$

Next, since the slope of any plane in H^{i+1}_i is at most 2 by Proposition 3, we have that

$$\max_{x \in \mathcal{E}_0} f(x) \leqslant \sum_{i=1}^{n-1} (|p_i p_{i+1}| + O(1)) \leqslant \min_{x \in \mathcal{E}_0} f(x) + O(n),$$

and $\mathrm{Var}_R(f) = O(n)$.

Now, injecting $R = O(\sqrt{n})$, $r = \Omega(1)$ and $\mathrm{Var}_R(f) = O(n)$ into the statement of Theorem 22, we obtain that $N(\varepsilon) = O\big(n^2 (\log \sqrt{n} + \log \frac{n}{\varepsilon})\big) = O\big(n^2 \log \frac{n}{\varepsilon}\big)$.

[13] Note that if $(D_{2+\sqrt{5}})$ is violated, then ξ_i may be larger than $\frac{\pi}{2}$ and $A \cap B$, and therefore $\square \cap \Lambda^+$, may be empty.

9 Extensions

In many situations, the initial and/or the final directions of the path are fixed and given. We first show, in Section 9.1, how our results can easily be extended to this variant of the problem. Next, we discuss the connections between the problems of finding curvature-constrained shortest paths (i) in the presence of polygonal obstacles, and (ii) through a given sequence of points, without obstacles.

9.1 Fixed initial and/or final directions

Our approach essentially relies on four properties (holding under various distance assumptions): (i) F is locally strictly convex over $\mathcal{L}(\pi)$, (ii) any global minimum belongs to $\mathcal{L}(\frac{3\pi}{4})$, (iii) there exists a shape $\mathcal{D}$ sandwiched in-between $\mathcal{L}(\frac{3\pi}{4})$ and $\mathcal{L}(\pi)$ whose connected components are lifted to convex polyhedra, and (iv) every non-sharp turn reduces the number of connected components of $\mathcal{D}$ to consider. We now explain how these properties can be extended when the initial and final directions (i.e., the values of θ_1 and θ_n) are fixed.

Property (i). The proof of Proposition 4 (local convexity of $F_T(\theta_1, \theta_2)$) immediately yields that, when θ_1 is fixed, for any path type T, the map $\theta_2 \mapsto F_T(\theta_1, \theta_2)$ is locally strictly convex whenever the path exists, the length of the circular arc at p_2 is in $(0, \pi)$, and the other circular arc does not vanish. Note that, the length of the circular arc at p_1 may be larger than π and thus Proposition 5 (uniqueness of the shortest CSC-path when both arcs are shorter than π) cannot be used directly. It follows that it is unclear whether Theorem 6 can be extended to showing that $\theta_2 \mapsto \min_{T \in \{LSL, LSR, RSL, RSR\}} F_T(\theta_1, T_2)$ is locally strictly convex. Nonetheless, we easily get that each of

$$\theta_2 \mapsto \min_{T \in \{LSL, LSR\}} F_T(\theta_1, \theta_2) \quad \text{and} \quad \theta_2 \mapsto \min_{T \in \{RSL, RSR\}} F_T(\theta_1, \theta_2) \tag{4}$$

is locally strictly convex whenever the length of the circular arc at p_2 is less than π, and the other circular arc does not vanish.[14]

Since the minimum of F_T over all possible T is not known to be locally convex, we consider $F_\sigma(\theta_2, \ldots, \theta_{n-1})$, $\sigma = 1, \ldots, 4$, the length function of a shortest curvature-constrained path that goes through the configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order, and such that the circular arcs incident to p_1 and p_n have a fixed orientation L or R determined by σ; $\sigma = 1, \ldots, 4$ corresponds to each of the four choices of orientations. Let $\mathcal{L}_\sigma(\alpha)$ denote the set of $(\theta_2, \ldots, \theta_{n-1})$ such that the shortest path visiting $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order and whose type in p_1 and p_n is prescribed by σ has all its circular arcs of length less than α, except possibly the first and the last arc. The local convexity of the functions in (4), as well as the similar functions when θ_2 is fixed, yields, with the proof of Proposition 7 (F is locally convex over $\mathcal{L}(\pi)$), the following analogue to Property (i):

Proposition 25. *For any fixed θ_1, θ_n, and $1 \leqslant \sigma \leqslant 4$, F_σ is locally strictly convex over $\mathcal{L}_\sigma(\pi)$.*

Property (ii). The necessary conditions (ii) and (iii) on locally shortest paths of Proposition 9 and the proof of Lemma 11 ($\mathcal{L}(\pi)$ contains the global minima) rely on finding (local or global) shortcuts that would contradict the optimality of the path. These arguments can still be used in our setting as any locally/globally shortest path that starts and ends with prescribed types of circular arcs is indeed a local/global minimum of F_σ for the corresponding σ. Since these shortcuts never consider the circular arcs incident to p_1 or p_n, our property (ii) extends immediately:

Lemma 26. *For any fixed θ_1, θ_n, and $1 \leqslant \sigma \leqslant 4$, any global minimum of F_σ belongs to $\mathcal{L}_\sigma(\frac{3\pi}{4})$.*

[14] It is sufficient to show that the minimum of the two functions F_{LSL} and F_{LSR} is C^1 at any point where their graphs intersect (and similarly for F_{RSL} and F_{RSR}). If two paths of types LSL and LSR have their second circular arc shorter than π, then exactly one of these two paths has its first circular arc shorter than π (indeed, by Proposition 5, both paths cannot have both arcs shorter than π, nor can they have both arcs longer than π, since otherwise both paths can be shortened by half a circle, yielding that two paths with both arcs are shorter than π). Thus F_{LSL} and F_{LSR} take distinct values at any point where the length of the second circular arc of their corresponding paths is in $(0, \pi)$. When the second circular arc vanishes, the paths of type LSL and LSR coincide, and the minimum of F_{LSL} and F_{LSR} is smooth (as in the proof of Theorem 6).

Property (iii). For any $1 < i < n-1$, the lemon $L_i^{i+1}(\alpha)$ depends on neither θ_1 nor θ_n, thus the diamond D_i^{i+1} defined in Section 6, still satisfies $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$. For $i = 1$ (and similarly for $i = n-1$), the lemon $L_1^2(\pi)$ is one-dimensional since θ_1 is fixed and, moreover, it is a segment.[15]

The domain $\mathcal{D}$ can then be adapted into a new domain $\mathcal{D}_\sigma$ by simply setting D_1^2 to be that segment (and, similarly, for D_{n-1}^n). Lemma 17 then immediately extends:

Lemma 27. *For any fixed θ_1, θ_n, and $1 \leqslant \sigma \leqslant 4$, $\mathcal{D}_\sigma$ does not intersect the hyperplanes of $\mathcal{H}$, and the lifting $(\mathbb{S}^1)^n \to \Lambda$ maps the intersection of $\mathcal{D}_\sigma$ with each cell of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ to a convex polyhedron defined by at most $4(n-1)$ linear inequalities in $\mathbb{R}^n$, each involving 2 variables.*

Property (iv). Finally, the proof of Lemma 20 uses the fact that the circular arcs at p_{i-1} and p_{i+1} are shorter than π which may be an issue when analyzing p_2 and p_{n-1}. It does, however, generalize immediately for $p_3, \ldots, p_{n-2}$ and allows to reduce the search to 2^{k+2} connected components where k is the number of sharp turns in $\{p_3, \ldots, p_{n-2}\}$. From there, the results of Section 8 extend immediately.

We finally obtain the following analogue of Theorem 1:

Theorem 28. *Let $p_1, \ldots, p_n$ be a sequence of points in the plane that satisfy the (D_4) condition and let k denote the number of sharp turns at p_i, $3 \leqslant i \leqslant n-2$. For any fixed θ_1 and θ_n, all curvature-constrained shortest paths from (p_1, θ_1) to (p_n, θ_n) that visit $p_1, \ldots, p_n$ in order can be computed by minimizing four functions over 2^{k+2} polyhedra in $\mathbb{R}^n$ over which they are strictly convex, each polyhedron being defined by at most $4n-1$ linear inequalities in two variables each.*

9.2 Shortest paths among polygonal obstacles

Any shortest path of bounded curvature between two given points *in the presence of polygonal obstacles* is a concatenation of Dubins paths whose extremities are the starting point, the ending point, and some *contact points* on the boundary of the obstacles [14, 17]. We now discuss how our results imply that, under some conditions, once the sequence of contact points is known, “filling the dots” reduces to computing locally shortest curvature-constrained paths visiting certain sequences of contact points *in the absence of obstacles* (Proposition 29), and admits a convex optimization formulation similar to Theorem 1 (Theorem 31). This suggests that the real difficulty in curvature-constrained path planning among polygonal obstacles resides in the subproblem of finding the sequence of contact points, a question that we do not address here.

Problem statement. Let ρ be a bounded-curvature shortest path between two configurations in the presence of polygonal obstacles. We call a circular arc of ρ *anchored* if it touches the obstacles in at least two points. The question we consider is the following: we assume that we are given the sequence $p_1, \ldots, p_n$ of contact points of ρ with the obstacles as well as the anchored circular arcs contained in ρ and we want to reconstruct the whole geometry of the path. As in the previous sections, we work under the (restrictive) condition that ρ is a concatenation of CSC-paths with non-vanishing line segments; in other words, *we rule out shortest paths that contain two consecutive non-vanishing circular arcs.*

To keep the presentation simple, and without loss of generality, we assume that ρ touches the obstacles in a finite number of points (since the case where it travels along the boundary of an obstacle can be handled easily) and that we know the orientation of the circular arcs following the first point and preceding the last point of ρ (we discuss after Theorem 31 how to manage without this information).

The length-reducing perturbations of SCS-paths of Dubins [12, Lemma 1] imply that for any endpoint p_i of an anchored circular arc of ρ, the tangent to ρ in p_i and the orientation, L or R, of the circular arcs preceding and following p_i are prescribed by the geometry of the obstacle in p_i. Splitting ρ at the endpoints of circular arcs, we obtain a sequence of subpaths that contain no anchored circular arcs and join known initial and final configurations. We can thus assume, again without loss of generality, that ρ *contains no anchored circular arc.*

[15] If a path from (p_1, θ_1) to (p_2, θ_2) of type LSL or LSR uses an arc of length exactly π ending in p_2 then it uses a segment that is tangent to both the L circle of (p_1, θ_1) and the circle of radius 2 centered in p_2. Since there can be at most two such segments (because of orientation of circle L), it follows that for any fixed θ_1 and σ the corresponding lemon $L_1^2(\pi)$ has at most two endpoints and thus is a segment.

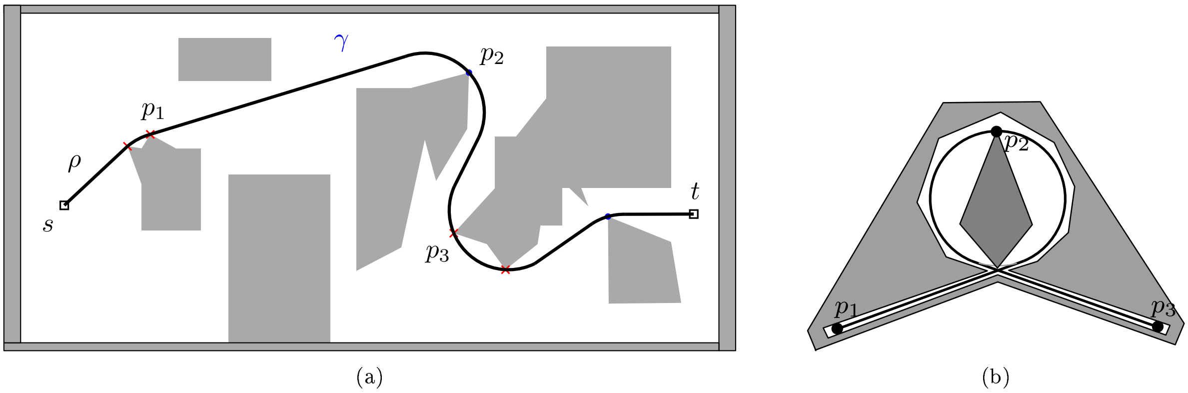

Figure 10: (a) A subpath γ of the shortest path ρ from s to t; γ has one via-point, p_2. (b) A shortest path in the presence of obstacles may have, in general, circular arcs of length arbitrarily close to π, preceding and following a via-point.

Local geometry of ρ. Let $\tilde{\theta}_i$ denote the polar angle of the tangent to ρ in p_i; we are given $\tilde{\theta}_1$ and $\tilde{\theta}_n$ and our first step is to characterize $\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1}$. We use a separate set of "free" variables $\theta_2, \ldots, \theta_{n-1}$ and, for the sake of the presentation, we also consider θ_1 and θ_n which are fixed equal to $\tilde{\theta}_1$ and $\tilde{\theta}_n$, respectively. As in Section 9.1, we let $F_\sigma(\theta_2, \ldots, \theta_{n-1})$ denote the length of a shortest curvature-constrained path that goes, in the absence of obstacles, through the configurations $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order and such that the circular arcs incident to p_1 and p_n have the prescribed orientation.[16] Note that here σ is fixed and corresponds to the prescribed orientation at p_1 and p_n of *the* optimal path ρ. We can now state our first result.

Proposition 29. *$(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ is a local minimum of F_σ. Moreover, for $i = 2, \ldots, n-2$, the part of ρ between p_i and p_{i+1} coincide with* the shortest *CSC-path from $(p_i, \tilde{\theta}_i)$ to $(p_{i+1}, \tilde{\theta}_{i+1})$ in the absence of obstacles.*

We remark that for $i = 1$ and $i = n-1$, the shortest CSC-path in the absence of obstacle may collide with the obstacles.

Proof. We can assume that $n > 2$ since otherwise there is nothing to prove. Thus, ρ cannot be reduced to a single circular arc, and the obstacles touch every circular arc of ρ exactly once.

Let us first assume that ρ has no vanishing circular arc. Let $T_i \in \{LSR, RSL, LSL, RSR\}$ be the type of the CSC-subpath of ρ from $(p_i, \tilde{\theta}_i)$ to $(p_{i+1}, \tilde{\theta}_{i+1})$, $i = 1, \ldots, n-1$. We define $F_{\mathcal{T}}(\theta_1, \ldots, \theta_n)$ as the length of the path through $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ whose type between the i-th and $(i+1)$-th configurations is T_i. We note that $F_{\mathcal{T}}(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ measures the length of ρ.

The length-reducing perturbations of SCS-paths of Dubins [12, Lemma 1] imply that at each contact point $p_2, \ldots, p_{n-1}$ the obstacle is locally inside the disk supporting the circular arcs. Thus, each non-terminal circular arc of ρ can rotate slightly around the contact point it contains without creating any new intersection with the obstacles. Since the path the length of which is measured by $F_{\mathcal{T}}(\theta_2, \ldots, \theta_{n-1})$ deforms continuously as $(\theta_2, \ldots, \theta_{n-1})$ changes, we have that in a neighborhood of $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ this path does not cross the obstacles properly. It follows that $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ is a local minimum of $F_{\mathcal{T}}$.

The proof of Lemma 12 now implies that the two circular arcs preceding and following every point p_i $(1 < i < n)$ have the same orientation (R or L), and their lengths are either equal or sum up to 2π.[17] Since ρ is a globally shortest path, the total length of the two arcs incident to p_i is strictly less than 2π.

[16] If the initial (resp. final) circular arc of ρ vanishes then F_σ makes no requirement on the type of the initial (resp. final) circular arc.

[17] The statement in Lemma 12 concerning the nonterminal points p_i $(1 < i < n)$ was proved by considering paths from $(p_{i-1}, \tilde{\theta}_{i-1})$, through (p_i, θ_i), and to $(p_{i+1}, \tilde{\theta}_{i+1})$ (with $\tilde{\theta}_{i-1}, \tilde{\theta}_{i+1}$ fixed); the fact that the direction of the path at p_1 and p_n was not constrained in Lemma 12 has thus no impact here. Note also that, formally, the path γ we considered in the statement of Lemma 12 was, for ease of exposition, a *shortest* Dubins path between every two consecutive configurations. This is not the case here, since the Dubins path between any two consecutive configurations is not known to be the shortest. However, the proof of Lemma 12 does not use the hypothesis that the Dubins subpaths of γ are the *shortest*; it considers the type of the paths incident to each p_i, regardless of its optimality.

It follows that the arcs incident to every p_i, $1 < i < n$, have the same orientation and the same length, which is less than π. Hence, by Proposition 5, the path ρ between every two nonterminal consecutive configurations is the shortest CSC-path. It follows that in a neighborhood of $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$, F_σ and $F_\mathcal{T}$ coincide. This proves both results.

If some circular arc of ρ vanishes then we relax the definition of $F_\mathcal{T}$ by allowing the corresponding circular arc to have an arbitrary type. This ensures that the path the length of which is measured by $F_\mathcal{T}$ still deforms continuously when $(\theta_2, \ldots, \theta_{n-1})$ changes so $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ is still a local minimum of $F_\mathcal{T}$. We thus get that the circular arcs preceding and following each contact point have equal length and are shorter than π; from there, Proposition 5 implies both statements. $\square$

We stress that Proposition 29 states that, in the presence of obstacles, a bounded-curvature shortest path that contains no CC subpath and no anchored circular arc is such that every nonterminal CSC-subpath connecting two configurations touching the obstacles is *the* shortest CSC-path in the *absence* of obstacles. We believe that this property of shortest paths among obstacles is new; in particular, it strengthens the result of Fortune and Wilfong [14], and Jacobs and Canny [17] that a shortest path in the presence of obstacles is a concatenation of *Dubins* paths.

Reduction to convex optimization. Proposition 29 suggests that the problem of computing a bounded-curvature shortest path in the presence of obstacles, given the sequence of its contact points with the obstacles and its initial and final orientation, can be broken down into computing minima of length functions F_σ and therefore may be amenable to the machinery we developed in Sections 2–7.

As described in Section 9.1, and with the same definition of $\mathcal{L}_\sigma(\alpha)$,[18] our approach relies on four main properties : (i) F_σ is locally strictly convex over $\mathcal{L}_\sigma(\pi)$, (ii) any global minimum belongs to $\mathcal{L}_\sigma(\frac{3\pi}{4})$, (iii) there exists a shape $\mathcal{D}_\sigma$ sandwiched in-between $\mathcal{L}_\sigma(\frac{3\pi}{4})$ and $\mathcal{L}_\sigma(\pi)$ whose connected components are lifted to convex polyhedra, and (iv) every non-sharp turn reduces the number of connected components of $\mathcal{D}_\sigma$ to consider.

In our setting, Properties (i) and (iii) hold exactly as in Section 9.1. However, Property (ii) does not directly hold for several reasons. First, the minima we are seeking are only guaranteed to be *local minima* (Proposition 29), whereas our machinery was developed for *global minimum* (and the globality of the minimum was used for showing that it belongs to $\mathcal{L}_\sigma(\frac{3\pi}{4})$). Second, the local minima we are seeking may actually not belong to $\mathcal{L}_\sigma(\frac{3\pi}{4})$: Figure 10(b) shows an example in which a bounded-curvature shortest path from p_1 to p_3 in the presence of obstacles is, in the absence of obstacles, a locally shortest path through p_1, p_2, p_3 where the length of the circular arcs preceding and following p_2 can be made arbitrarily close to π.

Such situations can, however, be circumvented by assumptions that are reasonable in many contexts. Say that the path ρ is *locally simple* in p_i if the portion of ρ from p_{i-1} to p_{i+1} has no self-intersection.

Lemma 30. *If $p_1, \ldots, p_n$ satisfy the $(D_{2+\sqrt{5}})$ condition and if each p_i in which ρ is not locally simple is contained in an obstacle of diameter larger than $1+\sqrt{2}$, then $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ belongs to $\mathcal{L}_\sigma(\frac{3\pi}{4})$.*

Proof. Assume, for a contradiction, that at least one of the circular arcs incident to some p_i, $1 < i < n$, has length at least $\frac{3\pi}{4}$. By Proposition 9, since $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ is a local minimum of F_σ, both circular arcs at p_i, $1 < i < n$, have same length or their lengths sum up to 2π. The latter case contradicts the optimality of ρ. In the former case, condition $(D_{2+\sqrt{5}})$ and Lemma 2 imply that the incident line segments are longer than 1 and the path therefore self-intersects. Moreover, the closed region formed by the two circular arcs and the two segments clipped at their intersection point is of diameter at most $1+\sqrt{2}$. Thus, if the obstacles have larger diameters, they cannot be contained in that region. Hence, the concatenation of the two circular arcs at p_i can only touch the obstacles on its concave side, and it must touch the obstacles in at least two points (including p_i) since ρ is a shortest path, contradicting the hypothesis that ρ contains no anchored circular arc. Therefore, $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ belongs to $\mathcal{L}_\sigma(\frac{3\pi}{4})$. $\square$

We consider finally Property (iv). The path corresponding to a point in $\mathcal{L}_\sigma(\frac{3\pi}{4})$ has all its circular arc shorter than $\frac{3\pi}{4}$, except possibly the arc incident to p_1 and p_n. For such a path, the proof of Lemma 20 shows that, under the (D_4) condition (and thus also under $(D_{2+\sqrt{5}})$), and for $3 \leqslant i \leqslant n-2$, if p_i is

[18] Recall, for clarity, that $\mathcal{L}_\sigma(\alpha)$ denote the set of $(\theta_2, \ldots, \theta_{n-1})$ such that the shortest path (without obstacles) visiting $(p_1, \theta_1), \ldots, (p_n, \theta_n)$ in order satisfies two conditions: (i) the circular arcs following p_1 and preceding p_n have their orientations prescribed by σ, that is the same as in ρ, and (ii) all the other circular arcs have length less than α.

not a sharp turn and θ_i is in Λ_i^-, then ρ is not locally simple in p_i; thus if ρ is locally simple in p_i, and p_i is not a sharp turn, then $\theta_i \in \Lambda_i^+$. Putting everything together, we obtain the following analogue of Theorems 1 and 28.

Theorem 31. *Assume that $p_1, \ldots, p_n$ satisfy the $(D_{2+\sqrt{5}})$ condition and that every p_i in which ρ is not known to be locally simple is contained in an obstacle of diameter larger than $1+\sqrt{2}$. Let k be the number of p_i, $3 \leqslant i \leqslant n-2$, that are sharp turns or in which ρ is not known to be locally simple.*

(i) $(\tilde{\theta}_2, \ldots, \tilde{\theta}_{n-1})$ is one of the local minima of F_σ over a collection of at most 2^{k+2} convex polyhedra in $\mathbb{R}^n$, over which F_σ is strictly convex; every polyhedron is defined by at most $4n-1$ linear inequalities in two variables each.

(ii) For each of these at most 2^{k+2} local minima $(\theta_2, \ldots, \theta_{n-1})$, consider the concatenation of the shortest CSC-paths from (p_i, θ_i) to (p_{i+1}, θ_{i+1}), for $i = 1, \ldots, n-1$ (with the circular arcs following p_1 and preceding p_n of prescribed orientation). The path ρ is the (or a) shortest of these paths that avoids the obstacles.

Note that, since ρ contains no anchored circular arc, if ρ is not locally simple at p_i then the loop of ρ passing through p_i, and joining the self-intersection to itself, must enclose an obstacle. In particular, if ρ is a bounded-curvature shortest path in a *simple polygon*, it is locally simple in every point, and k is the number of sharp turns among $p_3, \ldots, p_{n-2}$.

Unknown initial and final orientations. If the orientation of the circular arcs of ρ following p_1 and preceding p_n are unknown, a natural approach is to consider four distinct subproblems, one for each choice of orientations of these arcs. More precisely, we want to solve the problem (i) of finding a shortest path ρ between $(p_1, \tilde{\theta}_1)$ to $(p_n, \tilde{\theta}_n)$ and through $p_2, \ldots, p_{n-1}$ that avoids the obstacles. We have showed how to solve the problem (ii) where the orientation (L or R) of ρ is *known* at p_1 and p_n; we now consider the problem (iii) of finding a shortest path ρ_σ where the orientation of ρ_σ is *prescribed* at p_1 and p_n.

A difficulty with Problem (iii) is that if we enforce the path to turn, say, right at p_1, then it is possible that the circular arc of ρ_σ at p_1 vanishes which implies that θ_2 is constrained; this situation can propagate and nothing ensures that the two circular arcs incident to a p_i have the same length. In other words the proof of Proposition 29 does not hold anymore. However, we note that this is not a problem in the end because the set of candidate shortest paths obtained by applying Theorem 31 to each length function F_σ is guaranteed to contain ρ as Proposition 29 is true for the choice of orientations that correspond to ρ.

10 Concluding remarks

We conclude with a few remarks on possible extensions of our results.

Distance condition. The requirement in Theorems 1 and 28 that the points $p_1, \ldots, p_n$ satisfy the (D_4) condition is used in three places. First, it excludes the occurrence of CCC-paths. Then, it is used to argue that the global minima of the length function F belong to the lemon $\mathcal{L}(\frac{3\pi}{4})$ (Lemma 11). Finally, it is instrumental for proving that the minima of F can be searched for in convex components over which F is convex (Lemma 15). Further relaxing the distance condition in these theorems therefore seems a considerable task. In particular, this would require to study the convexity properties of the length function of CCC paths, task we did not undertake here since this paper is already substantial.

Corollary 24 uses a (D_ς) condition, with $\varsigma > 2+\sqrt{5}$, to bound the complexity of the ellipsoid method for computing the globally shortest path in the absence of sharp turn. While the reduction to a single convex optimization problem holds under the weaker distance condition (D_4), we do not see how the complexity analysis could be extended for $\varsigma \leqslant 2+\sqrt{5}$.

Last, Theorem 31, our extension of Theorem 1 in the presence of obstacles, considers the stronger distance assumption $(D_{2+\sqrt{5}})$. This assumption is only used to ensure, in Lemma 30, that the mimima of the length function belong to $\mathcal{L}(\frac{3\pi}{4})$. This is done in a straightforward way and is likely to be improved, at least in specific settings such as inside a simple polygon.

More efficient pruning. The idea underlying Theorem 1 is that finding a bounded-curvature globally shortest path visiting a sequence of points in a prescribed order can be decomposed into two problems: (a) identifying which, among 2^{n-2} convex polyhedra, contains a global minimum of the length function,[19] and (b) minimizing the length function over that polyhedron. We essentially solved problem (b) by showing that it admits a convex optimization formulation. Regarding problem (a), we observed that excluding certain self-intersections in the path reduces (in fact, halves) the set of candidate polyhedra; we encapsulated this condition in the coarser, but simpler and more intuitive, notion of sharp turns; exploring other conditions that reduce the combinatorial explosion of problem (a) is a natural follow-up question.

Obstacles. A natural question raised by our results of Section 9.2 is whether the sequence of points where the shortest path between two configurations in the presence of obstacles can be computed, or even approximated. A candidate setting in which this question could perhaps be tackled is inside a simple polygon, where the homotopy class of the shortest path is trivial.

Acknowledgment

The authors would like to thank Jean-Daniel Boissonnat for his implication at a very early stage of this work, during the PhD of the third author, and Otfried Cheong for helpful discussions.

[19] More technically, this amounts to deciding at each via-point p_i, $2 \leqslant i \leqslant n-1$, whether the polar angle of the tangent to a globally shortest path at that point belongs to Λ_i^+ or Λ_i^-.

A Local convexity of the length function of a CSC-path

We present here complete proofs of the results claimed in Section 3. Recall that we prove here that the length $F_{csc}(\theta_1, \theta_2)$ of a shortest CSC-path from configuration (p_1, θ_1) to (p_2, θ_2) is a locally strictly convex function at at any point (θ_1, θ_2) such that both circular arcs of the corresponding path are shorter than π. For simplicity, we refer in this section to F_{csc} as F. The angles θ_1 and θ_2 are considered, throughout this section, in $\mathbb{R}$ rather than in $\mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z}$. In other words, the function F is defined over $\mathbb{R}^2$, seen as the standard universal covering of the torus $\mathbb{S}^1 \times \mathbb{S}^1$. The (local) convexity of the function is thus defined over $\mathbb{R}^2$ in the standard way.

We start by showing that for a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is such a locally strictly convex function (Proposition 35). We prove this by first computing the first derivatives and the Hessian of the length function (Propositions 33 and 34).

We then prove that the local convexity extends to the length $F(\theta_1, \theta_2) = \min_{T\in\{LSR,RSL,LSL,RSR\}} F_T(\theta_1, \theta_2)$ of a shortest CSC-path (Theorem 37). We prove this by first showing the interesting property that, if both circular arcs of a CSC-path are shorter than π, then this path is *the* shortest CSC-path (Proposition 36).

Notation. We introduce here some new notation; we also recall, for clarity, the notation introduced in Section 3. Refer to Figure 11. For a given CSC-path, let O_i, $i = 1, 2$, denote the center of the unit circle supporting the i-th circular arc (defined by continuity if one of the circular arcs vanish), let α_i be the length of this i-th circular arc, and let M_1 and M_2 be the first and last endpoint of the line segment of the path. Let $\overrightarrow{U_{12}}$ denote the vector $\overrightarrow{M_1M_2}/M_1M_2$ where M_1M_2 denotes the Euclidean distance from M_1 to M_2. Let μ_B be equal to 1 if B is true and to -1 otherwise. In particular, for a type of path $T \in \{LSR, RSL, LSL, RSR\}$, $\mu_{C_j=R}$ ($j = 1, 2$) is equal to 1 if the type of the j-th circular arc in T is R and it is equal to -1 otherwise. Let $\delta_{i,j}$ be equal to 1 if $i = j$ and to 0 otherwise. Finally, if $\vec{u}$ and $\vec{v}$ are two vectors of $\mathbb{R}^2$, $\vec{u} \times \vec{v}$ denote their determinant (or, equivalently, the nonzero coordinates of their cross product, seen as vectors of $\mathbb{R}^3$).

First derivatives of $\mathbf{F_T}(\theta_1, \theta_2)$

We start by a straightforward preliminary lemma.

Lemma 32. *If $\vec{u}$, $\vec{v}$ and $\vec{w}$ are three vectors of $\mathbb{R}^2$ such that $\|\vec{v}\| = 1$, and $\vec{u} \cdot \vec{v} = 0$ or $\vec{v} \cdot \vec{w} = 0$, then*

$$\vec{u} \cdot \vec{w} = -(\vec{u} \times \vec{v})(\vec{v} \times \vec{w}).$$

In particular, if $\vec{u}$ and $\vec{v}$ are two vectors such that $\|\vec{v}\| = 1$, then

$$\vec{u} \cdot \frac{\partial \vec{v}}{\partial \theta_i} = -(\vec{u} \times \vec{v}) \left(\vec{v} \times \frac{\partial \vec{v}}{\partial \theta_i}\right)$$

Proof. The proof of the first equality is straightforward:

$$\begin{aligned}
\vec{u} \cdot \vec{w} &= \|\vec{u}\|.\|\vec{w}\| \cos(\angle(\vec{u}, \vec{w})) \\
&= \|\vec{u}\|.\|\vec{w}\| \cos(\angle(\vec{u}, \vec{v}) + \angle(\vec{v}, \vec{w})) \\
&= 0 - \|\vec{u}\|.\|\vec{w}\| \sin(\angle(\vec{u}, \vec{v})) \sin(\angle(\vec{v}, \vec{w})) \\
&= -(\vec{u} \times \vec{v})(\vec{v} \times \vec{w})
\end{aligned}$$

The second equality of the lemma follows by considering $\vec{w} = \frac{\partial \vec{v}}{\partial \theta_i}$ because, since $\vec{v}$ has norm one, its scalar product with its derivative is zero. □

Proposition 33. *For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is differentiable at any point (θ_1, θ_2) such that the corresponding T-path exists and both its circular arcs do not vanish. Furthermore,*

$$\frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} = \mu_{i=1}\mu_{C_i=R}(1 - \cos \alpha_i). \tag{5}$$

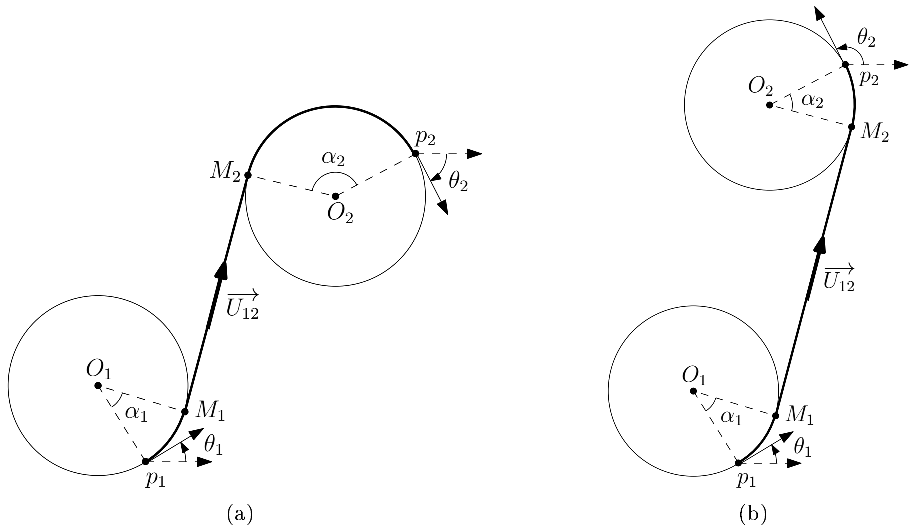

Figure 11: LSR- and LSL-paths from (p_1, θ_1) to (p_2, θ_2).

Proof. Consider a T-path from (p_1, θ_1) to (p_2, θ_2), with $T \in \{LSR, RSL, LSL, RSR\}$ and refer to Figure 11. We determine, in turn, the derivatives with respect to θ_i of the length α_i of its i-th circular arc, and of the length g of its line segment. We then prove

$$\frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} = \mu_{i=1} \left(\mu_{C_i=R} + \overrightarrow{U_{12}} \times \overrightarrow{p_i O_i} \right) \tag{6}$$

which is equivalent to Eq. (5) because $\overrightarrow{U_{12}} \times \overrightarrow{p_i O_i} = \mu_{C_i=L} \cos \alpha_i$ (see Figure 11(a)). Note that this yields that F_T is differentiable at any point (θ_1, θ_2) such that the corresponding T-path exists and none of its circular arcs vanishes, because, at such a point, F_T and its partial derivatives are defined and continuous.

If the length of the i-th $(i = 1, 2)$ circular arc is smaller than π, then $\alpha_i = \arccos(\overrightarrow{O_i p_i} \cdot \overrightarrow{O_i M_i})$, otherwise, $\alpha_i = 2\pi - \arccos(\overrightarrow{O_i p_i} \cdot \overrightarrow{O_i M_i})$. It follows that

$$\alpha_i(\theta_1, \theta_2) = \mu_{(\alpha_i < \pi)} \arccos(\overrightarrow{O_i p_i} \cdot \overrightarrow{O_i M_i}) \mod 2\pi$$

Its derivative with respect to θ_j, $j = 1, 2$ is thus:

$$\frac{\partial \alpha_i}{\partial \theta_j} = \mu_{(\alpha_i < \pi)} \frac{-\left(\overrightarrow{O_i M_i} \cdot \frac{\partial \overrightarrow{O_i p_i}}{\partial \theta_j} + \overrightarrow{O_i p_i} \cdot \frac{\partial \overrightarrow{O_i M_i}}{\partial \theta_j} \right)}{\sqrt{1 - (\overrightarrow{O_i p_i} \cdot \overrightarrow{O_i M_i})^2}}$$

Since $\overrightarrow{O_i p_i}$ and $\overrightarrow{O_i M_i}$ have norm 1, the denominator can be simplified into $|\overrightarrow{O_i p_i} \times \overrightarrow{O_i M_i}|$. On the other hand, by Lemma 32, the terms of the numerator can be rewritten as follows:

$$\overrightarrow{O_i M_i} \cdot \frac{\partial \overrightarrow{O_i p_i}}{\partial \theta_j} = -(\overrightarrow{O_i M_i} \times \overrightarrow{O_i p_i})(\overrightarrow{O_i p_i} \times \frac{\partial \overrightarrow{O_i p_i}}{\partial \theta_j}),$$

$$\overrightarrow{O_i p_i} \cdot \frac{\partial \overrightarrow{O_i M_i}}{\partial \theta_j} = -(\overrightarrow{O_i p_i} \times \overrightarrow{O_i M_i})(\overrightarrow{O_i M_i} \times \frac{\partial \overrightarrow{O_i M_i}}{\partial \theta_j}).$$

Furthermore, $\overrightarrow{O_i p_i} = (\cos(\theta_i \pm \pi/2), \sin(\theta_i \pm \pi/2))$, thus

$$\overrightarrow{O_i p_i} \times \frac{\partial \overrightarrow{O_i p_i}}{\partial \theta_j} = \delta_{i,j} \tag{7}$$

and

$$\overrightarrow{O_iM_i} \cdot \frac{\partial \overrightarrow{O_ip_i}}{\partial \theta_j} = \delta_{(i,j)} \ (\overrightarrow{O_ip_i} \times \overrightarrow{O_iM_i}).$$

Thus, if ε_i denotes the sign of $\overrightarrow{O_ip_i} \times \overrightarrow{O_iM_i}$, we have

$$\frac{\partial \alpha_i}{\partial \theta_j} = -\mu_{(\alpha_i<\pi)}\varepsilon_i \left(\delta_{(i,j)} - \overrightarrow{O_iM_i} \times \frac{\partial \overrightarrow{O_iM_i}}{\partial \theta_j} \right).$$

Now, one can easily observe (see Figure 11) that, for any $i = 1, 2$, we have $\mu_{(\alpha_i<\pi)} \ \varepsilon_i = \mu_{i=1} \ \mu_{C_i=L}$, and since also $\overrightarrow{O_2M_2} = \pm\overrightarrow{O_1M_1}$, we obtain

$$\frac{\partial \alpha_i}{\partial \theta_j} = -\mu_{i=1} \ \mu_{C_i=L} \left(\delta_{(i,j)} - \overrightarrow{O_jM_j} \times \frac{\partial \overrightarrow{O_jM_j}}{\partial \theta_j} \right).$$

Thus, if the path is of type LSL or RSR, then $\mu_{i=1} \ \mu_{C_i=L}$ changes sign for $i = 1, 2$, and

$$\frac{\partial \alpha_1 + \alpha_2}{\partial \theta_j} = -\mu_{C_i=L} \ (\delta_{1,j} - \delta_{2,j}) = \mu_{C_j=R} \ \mu_{j=1}. \tag{8}$$

Otherwise, if the path is of type LSR or RSL, then $\mu_{i=1} \ \mu_{C_i=L}$ is equal to $\mu_{C_2=R}$, and

$$\frac{\partial \alpha_1 + \alpha_2}{\partial \theta_j} = -\mu_{C_2=R} \left(1 - 2 \ \overrightarrow{O_jM_j} \times \frac{\partial \overrightarrow{O_jM_j}}{\partial \theta_j} \right). \tag{9}$$

It remains to determine the derivative of g, the length of segment M_1M_2. Since $g^2 = \overrightarrow{M_1M_2} \cdot \overrightarrow{M_1M_2}$, we have $\frac{\partial g^2}{\partial \theta_j} = 2g \frac{\partial g}{\partial \theta_j} = 2\overrightarrow{M_1M_2} \cdot \frac{\partial \overrightarrow{M_1M_2}}{\partial \theta_j}$. Using the notation $\overrightarrow{U_{12}} = \frac{\overrightarrow{M_1M_2}}{M_1M_2}$, we thus have

$$\frac{\partial g}{\partial \theta_j} = \overrightarrow{U_{12}} \cdot \frac{\partial \overrightarrow{M_1M_2}}{\partial \theta_j}.$$

We decompose $\overrightarrow{M_1M_2}$ into $\overrightarrow{M_1O_1} + \overrightarrow{O_1p_1} + \overrightarrow{p_1p_2} + \overrightarrow{p_2O_2} + \overrightarrow{O_2M_2}$. Similarly as before, by Lemma 32,

$$\overrightarrow{U_{12}} \cdot \frac{\partial \overrightarrow{O_ip_i}}{\partial \theta_j} = -\left(\overrightarrow{U_{12}} \times \overrightarrow{O_ip_i} \right) \left(\overrightarrow{O_ip_i} \times \frac{\partial \overrightarrow{O_ip_i}}{\partial \theta_j} \right),$$

and, as noted above the last term is equal to $\delta_{i,j}$ (Eq. 7) so

$$\overrightarrow{U_{12}} \cdot \frac{\partial \overrightarrow{O_ip_i}}{\partial \theta_j} = \delta_{i,j}(\overrightarrow{U_{12}} \times \overrightarrow{p_iO_i}). \tag{10}$$

Hence,

$$\begin{aligned} \overrightarrow{U_{12}} \cdot \frac{\partial \overrightarrow{O_1p_1} - \overrightarrow{O_2p_2}}{\partial \theta_j} &= \delta_{1,j}(\overrightarrow{U_{12}} \times \overrightarrow{p_1O_1}) - \delta_{2,j}(\overrightarrow{U_{12}} \times \overrightarrow{p_2O_2}) \\ &= \mu_{j=1}(\overrightarrow{U_{12}} \times \overrightarrow{p_jO_j}). \end{aligned}$$

Now, if the path is of type LSL or RSR then, $\overrightarrow{M_1O_1} + \overrightarrow{O_2M_2} = 0$, thus

$$\frac{\partial g}{\partial \theta_j} = \mu_{j=1}(\overrightarrow{U_{12}} \times \overrightarrow{p_jO_j}),$$

which, together with Eq. (8), concludes the lemma for paths of type LSL and RSR.

On the other hand, if the path is of type LSR or RSL, then $\overrightarrow{M_1O_1} = \overrightarrow{O_2M_2}$, and, by Lemma 32,

$$\overrightarrow{U_{12}} \cdot \frac{\partial \overrightarrow{O_2M_2}}{\partial \theta_j} = -\left(\overrightarrow{U_{12}} \times \overrightarrow{O_2M_2} \right) \left(\overrightarrow{O_2M_2} \times \frac{\partial \overrightarrow{O_2M_2}}{\partial \theta_j} \right).$$

Furthermore, $\overrightarrow{U_{12}} \times \overrightarrow{O_2M_2}$ is equal to $\mu_{C_2=R}$ (see Figure 11(a)). Hence,

$$\frac{\partial g}{\partial \theta_j} = \mu_{j=1}(\overrightarrow{U_{12}} \times \overrightarrow{p_jO_j}) - 2\ \mu_{C_2=R}\left(\overrightarrow{O_2M_2} \times \frac{\partial \overrightarrow{O_2M_2}}{\partial \theta_j}\right).$$

Finally, since $\overrightarrow{O_2M_2} \times \frac{\partial \overrightarrow{O_2M_2}}{\partial \theta_j} = \overrightarrow{O_jM_j} \times \frac{\partial \overrightarrow{O_jM_j}}{\partial \theta_j}$, Eq. (9) yields that

$$\frac{\partial F_T}{\partial \theta_j} = \mu_{j=1}(\overrightarrow{U_{12}} \times \overrightarrow{p_jO_j}) - \mu_{C_2=R}.$$

This concludes the proof of the lemma since, for paths of types LSR and RSL, $\mu_{C_2=R}$ is equal to $-\mu_{j=1}\ \mu_{C_j=R}$. □

Second derivatives of $\mathbf{F_T}(\theta_1, \theta_2)$

We start by computing the second derivatives of the length $F_T(\theta_1, \theta_2)$ of a path of given type $T \in \{LSR, RSL, LSL, RSR\}$. We then compute its Hessian and show that it is positive definite, and thus that $F_T(\theta_1, \theta_2)$ is locally convex, for any (θ_1, θ_2) such that both circular arcs of the corresponding path are shorter than π.

Proposition 34. *For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is twice differentiable at any point (θ_1, θ_2) such that the corresponding T-path exists and none of its arcs vanishes. Furthermore,*

$$\frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} = \delta_{i,j}\ \sin\alpha_i + \frac{\sin\alpha_i \sin\alpha_j}{M_1M_2} \tag{11}$$

and the determinant of the Hessian of F_T is

$$\sin\alpha_1 \sin\alpha_2 \left(1 + \frac{\sin\alpha_1 + \sin\alpha_2}{M_1M_2}\right).$$

Proof. We first prove

$$\frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} = \delta_{i,j}\ \mu_{i=1}\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i} + \mu_{i=j}\frac{(\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i})(\overrightarrow{U_{12}} \cdot \overrightarrow{p_jO_j})}{M_1M_2} \tag{12}$$

which is equivalent to Eq. (11) because $\mu_{i=1}\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i} = \sin\alpha_i$ (see Figure 11). Note that, this will yield that, F_T is twice differentiable at any point (θ_1, θ_2) such that the corresponding T-path exists and none of its arcs vanishes, because at such a point F_T and all its first and second derivatives are defined and continuous.

By Proposition 33 (see Eq. (6)), we have

$$\frac{\partial F_T}{\partial \theta_i} = \mu_{i=1}\ \overrightarrow{U_{12}} \times \overrightarrow{p_iO_i} \pm 1,$$

and thus,

$$\frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} = \mu_{i=1}\left(\overrightarrow{U_{12}} \times \frac{\partial \overrightarrow{p_iO_i}}{\partial \theta_j} - \overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j}\right) \tag{13}$$

Notice, on the other hand that, given any three vectors $\vec{u}$, $\vec{v}$ and $\vec{w}$ of $\mathbb{R}^2$ such that $\|\vec{v}\| = 1$,

$$\vec{u} \times \vec{w} = (\vec{u} \times \vec{v})(\vec{v} \cdot \vec{w}) + (\vec{u} \cdot \vec{v})(\vec{v} \times \vec{w}). \tag{14}$$

Indeed,

$$\begin{aligned} \vec{u} \times \vec{w} &= \|\vec{u}\|.\|\vec{w}\| \sin(\angle(\vec{u}, \vec{v})) \cos(\angle(\vec{v}, \vec{w})) + \|\vec{u}\|.\|\vec{w}\| \cos(\angle(\vec{u}, \vec{v})) \sin(\angle(\vec{v}, \vec{w})) \\ &= (\vec{u} \times \vec{v})(\vec{v} \cdot \vec{w}) + (\vec{u} \cdot \vec{v})(\vec{v} \times \vec{w}) \end{aligned}$$

We can thus rewrite the terms of (13) using (14) as follows. Recall that since $\overrightarrow{p_iO_i}$ has norm 1, its scalar product with its derivative is zero.

$$\overrightarrow{U_{12}} \times \frac{\partial \overrightarrow{p_iO_i}}{\partial \theta_j} = (\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i}) \left(\overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{p_iO_i}}{\partial \theta_j} \right).$$

Furthermore, since $\overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{p_iO_i}}{\partial \theta_j} = \delta_{i,j}$ (Eq. (7)), we get

$$\mu_{i=1}\ \overrightarrow{U_{12}} \times \frac{\partial \overrightarrow{p_iO_i}}{\partial \theta_j} = \mu_{i=1}\ \delta_{i,j}\ \overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i}, \tag{15}$$

which is the first term of (12).

Consider now the second term of (13). Again, using (14) and since $\overrightarrow{U_{12}}$ has norm 1, its cross product with $\overrightarrow{O_1M_1}$ is zero, thus

$$\overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} = (\overrightarrow{p_iO_i} \times \overrightarrow{O_1M_1}) \left(\overrightarrow{O_1M_1} \cdot \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} \right).$$

Now, since $\overrightarrow{U_{12}} = \overrightarrow{M_1M_2}/M_1M_2$, its derivative is

$$\frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} = \frac{1}{M_1M_2} \frac{\partial \overrightarrow{M_1M_2}}{\partial \theta_j} - \frac{1}{M_1M_2^2} \frac{\partial M_1M_2}{\partial \theta_j} \overrightarrow{M_1M_2}, \tag{16}$$

and, since $\overrightarrow{O_1M_1}$ and $\overrightarrow{M_1M_2}$ are orthogonal, we get:

$$\overrightarrow{O_1M_1} \cdot \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} = \frac{1}{M_1M_2} \frac{\partial \overrightarrow{M_1M_2}}{\partial \theta_j} \cdot \overrightarrow{O_1M_1}.$$

We decompose $\overrightarrow{M_1M_2}$ into $\overrightarrow{M_1O_1} + \overrightarrow{O_1p_1} + \overrightarrow{p_1p_2} + \overrightarrow{p_2O_2} + \overrightarrow{O_2M_2}$, and since the derivatives of $\overrightarrow{M_1O_1}$ and $\overrightarrow{O_2M_2}$ are orthogonal to $\overrightarrow{O_1M_1}$, we obtain:

$$\begin{aligned}\overrightarrow{O_1M_1} \cdot \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} &= \frac{1}{M_1M_2} \left(\frac{\partial \overrightarrow{O_1p_1}}{\partial \theta_j} - \frac{\partial \overrightarrow{O_2p_2}}{\partial \theta_j} \right) \cdot \overrightarrow{O_1M_1} \\ &= \mu_{j=1} \frac{1}{M_1M_2} \frac{\partial \overrightarrow{O_jp_j}}{\partial \theta_j} \cdot \overrightarrow{O_1M_1}\end{aligned} \tag{17}$$

Furthermore, similarly as before, $\overrightarrow{O_1M_1} \cdot \frac{\partial \overrightarrow{O_jp_j}}{\partial \theta_j} = -\overrightarrow{O_1M_1} \times \overrightarrow{O_jp_j}$, by Lemma 32 and Eq. (7). Hence,

$$\overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} = -\mu_{j=1} \frac{(\overrightarrow{p_iO_i} \times \overrightarrow{O_1M_1})(\overrightarrow{p_jO_j} \times \overrightarrow{O_1M_1})}{M_1M_2}.$$

Finally, remark that, by Lemma 32,

$$\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i} = -(\overrightarrow{U_{12}} \times \overrightarrow{O_1M_1})(\overrightarrow{O_1M_1} \times \overrightarrow{p_iO_i}),$$

and, since $\overrightarrow{U_{12}} \times \overrightarrow{O_1M_1} = \pm 1$,

$$\overrightarrow{p_iO_i} \times \frac{\partial \overrightarrow{U_{12}}}{\partial \theta_j} = -\mu_{j=1} \frac{(\overrightarrow{U_{12}} \cdot \overrightarrow{p_iO_i})(\overrightarrow{U_{12}} \cdot \overrightarrow{p_jO_j})}{M_1M_2}.$$

Hence the second terms of (12) and (13) are equal, which concludes the proof of Eq. (12) (and thus of Eq. (11)) since, as we have seen in (15), the first terms of (11) and (13) are also equal.

This also concludes the proof of the proposition because the expression of the determinant of the Hessian follows directly from the expression of the second derivative (11):

$$\left(\sin\alpha_1 + \frac{\sin^2\alpha_1}{M_1M_2}\right)\left(\sin\alpha_2 + \frac{\sin^2\alpha_2}{M_1M_2}\right) - \left(\frac{\sin\alpha_1\sin\alpha_2}{M_1M_2}\right)^2 = \sin\alpha_1\sin\alpha_2\left(1 + \frac{\sin\alpha_1 + \sin\alpha_2}{M_1M_2}\right).$$

Note finally that one can easily prove that this determinant is also equal to $\frac{-(\overrightarrow{U_{12}}\cdot\overrightarrow{p_1O_1})(\overrightarrow{U_{12}}\cdot\overrightarrow{p_2O_2})(\overrightarrow{U_{12}}\cdot\overrightarrow{p_1p_2})}{M_1M_2}$.

□

Local convexity of $\mathbf{F_T}(\theta_1, \theta_2)$

As we have already seen in Section 3, Proposition 34 yield the local convexity of the length function F_T. We recall it for ease of reading.

Proposition 35. *For a given path type $T \in \{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is locally strictly convex at any point (θ_1, θ_2) such that the corresponding T-path exists, none of its arcs vanishes, and its two circular arcs have length less than π.*

Proof. Proposition 34 and the assumption that the length α_i of each circular arc is in $(0, \pi)$ imply that $\frac{\partial^2 F_T}{\partial \theta_1^2}$ and the determinant of the Hessian of F_T are positive. Thus, F_T is positive definite, by Sylvester's criterion, and thus locally strictly convex at any point (θ_1, θ_2) that satisfies the hypotheses. □

Geometric properties of CSC-paths

We proved so far that, for a given path type T in $\{LSR, RSL, LSL, RSR\}$, the length $F_T(\theta_1, \theta_2)$ of the T-path from (p_1, θ_1) to (p_2, θ_2) is locally convex at any (θ_1, θ_2) such that both circular arcs are shorter than π. We now prove that this property is also true for the length $F(\theta_1, \theta_2)$ of a shortest CSC-path, that is for the function $\min_{T \in \{LSR, RSL, LSL, RSR\}} F_T(\theta_1, \theta_2)$. For that purpose, we prove the following geometric property of CSC-paths, which is interesting in its own rights.

Proposition 36. *If both circular arcs of a CSC-path from (p_1, θ_1) to (p_2, θ_2) are strictly shorter than π, then all the other distinct CSC-paths are strictly longer.*

Proof. Consider two geometrically distinct paths of type T and T' in $\{LSR, RSL, LSL, RSR\}$, from (p_1, θ_1) to (p_2, θ_2), such that both circular arcs of the T-path are strictly shorter than π. Similarly as for the T-path, let O'_i be the center of the unit circle supporting the i-th circular arc of the T'-path (defined by continuity if one of the circular arcs vanish), let α'_i be its length, and let M'_1 and M'_2 be the first and last endpoint of the line segment of the T'-path.

In the first part of the proof, we show that the length of the line segment of the T-path is smaller than the one of the T'-path. In the second part, we show the same property about the circular arcs.

Length of the line segments. Note first that, the hypothesis that length of both circular arcs of the T-path are smaller than π implies that (see Figure 11(a)):

$$\begin{cases} \overrightarrow{p_1O_1} \cdot \overrightarrow{M_1M_2} \geqslant 0 \\ \overrightarrow{O_2p_2} \cdot \overrightarrow{M_1M_2} \geqslant 0 \end{cases} \tag{18}$$

By considering if necessary, the reverse paths possibly up to a symmetry, we can assume, without loss of generality, that the first circular arcs of the T and T'-paths have different orientations, R and L, respectively.

We prove that $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M'_1M'_2}\|$ by considering, in turn, the case where (i) $T' = LSL$, and otherwise, that is if $T' = LSR$, the cases where (ii) $T = RSR$ or (iii) $T = RSL$.

Case $\mathbf{T' = LSL}$. Refer, for instance, to Figure 12. In this case, $\|\overrightarrow{M'_1M'_2}\| = \|\overrightarrow{O'_1O'_2}\|$ and $\overrightarrow{O'_1O'_2}$ can be decomposed as follow:

$$\overrightarrow{O'_1O'_2} = \overrightarrow{O'_1p_1} + \overrightarrow{p_1O_1} + \overrightarrow{O_1M_1} + \overrightarrow{M_1M_2} + \overrightarrow{M_2O_2} + \overrightarrow{O_2p_2} + \overrightarrow{p_2O'_2}.$$

This implies that

$$\overrightarrow{O'_1O'_2} \cdot \overrightarrow{M_1M_2} \geqslant \overrightarrow{M_1M_2} \cdot \overrightarrow{M_1M_2};$$

Indeed, first, depending on T and T', $\overrightarrow{O'_1p_1} + \overrightarrow{p_1O_1}$ is equal either to $\vec{0}$ or $2\overrightarrow{p_1O_1}$ and $\overrightarrow{p_1O_1} \cdot \overrightarrow{M_1M_2} \geqslant 0$ by (18). Thus $(\overrightarrow{O'_1p_1} + \overrightarrow{p_1O_1}) \cdot \overrightarrow{M_1M_2} \geqslant 0$ and similarly for $\overrightarrow{O_2p_2} + \overrightarrow{p_2O'_2}$. This yields the above inequality since $\overrightarrow{O_1M_1}$ and $\overrightarrow{M_2O_2}$ are orthogonal to $\overrightarrow{M_1M_2}$.

We thus get

$$\|\overrightarrow{O'_1O'_2}\|\|\overrightarrow{M_1M_2}\| \geqslant \overrightarrow{O'_1O'_2} \cdot \overrightarrow{M_1M_2} \geqslant \|\overrightarrow{M_1M_2}\|^2$$

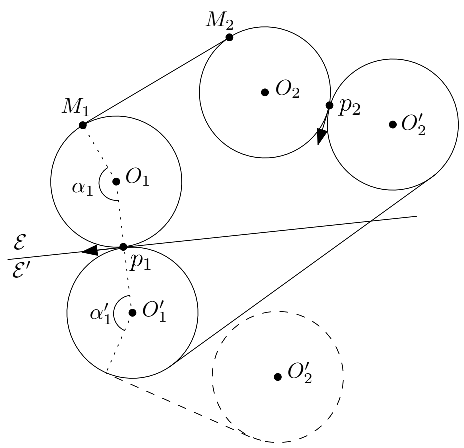

Figure 12: For the proof of Prop. 36: $(T, T') = (RSR, LSL)$, and a contradiction if $\alpha_1, \alpha'_1 < \pi$.

and $\|\overrightarrow{M'_1M'_2}\| = \|\overrightarrow{O'_1O'_2}\| \geqslant \|\overrightarrow{M_1M_2}\|$.

Case $(\mathbf{T}, \mathbf{T}') = (\mathbf{RSR}, \mathbf{LSR})$. Refer to Figure 14. In this case, $O_2 = O'_2$ and

$$\overrightarrow{O'_1O'_2} = \overrightarrow{O'_1O_2} = 2\overrightarrow{p_1O_1} + \overrightarrow{O_1O_2},$$
$$\|\overrightarrow{O'_1O'_2}\|^2 = 4 + \|\overrightarrow{O_1O_2}\|^2 + 4\overrightarrow{p_1O_1} \cdot \overrightarrow{O_1O_2}.$$

Since $T' = LSR$, $\|\frac{\overrightarrow{M'_1M'_2}}{2}\|^2 = \|\frac{\overrightarrow{O'_1O'_2}}{2}\|^2 - 1$ and

$$\|\overrightarrow{M'_1M'_2}\|^2 = \|\overrightarrow{O_1O_2}\|^2 + 4\overrightarrow{p_1O_1} \cdot \overrightarrow{O_1O_2}.$$

Since $T = RSR$, $\overrightarrow{O_1O_2} = \overrightarrow{M_1M_2}$ and

$$\|\overrightarrow{M'_1M'_2}\|^2 = \|\overrightarrow{M_1M_2}\|^2 + 4\overrightarrow{p_1O_1} \cdot \overrightarrow{M_1M_2}.$$

Hence by (18), we get that $\|\overrightarrow{M'_1M'_2}\| \geqslant \|\overrightarrow{M_1M_2}\|$.

Case $(\mathbf{T}, \mathbf{T}') = (\mathbf{RSL}, \mathbf{LSR})$. The proof is more subtle than in the two previous cases. Suppose, without loss of generality, that $p_1 = (0, 0)$ and $p_2 = (d, 0)$ in an Euclidean coordinate system. Since $T = RSL$, $\|\frac{\overrightarrow{M_1M_2}}{2}\|^2 = \|\frac{\overrightarrow{O_1O_2}}{2}\|^2 - 1$ and similarly for the T'-path. Thus $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M'_1M'_2}\|$ is equivalent to $\|\overrightarrow{O_1O_2}\| \leqslant \|\overrightarrow{O'_1O'_2}\|$. The coordinates of the circle centers O_1 and O'_1 are (see Figure 11)

$$O_1 = (\cos(\theta_1 - \frac{\pi}{2}), \sin(\theta_1 - \frac{\pi}{2})) = (\sin\theta_1, -\cos\theta_1),$$
$$O'_1 = (\cos(\theta_1 + \frac{\pi}{2}), \sin(\theta_1 + \frac{\pi}{2})) = (-\sin\theta_1, \cos\theta_1),$$

and, similarly,

$$O_2 = (d - \sin\theta_2, \cos\theta_2),$$
$$O'_2 = (d + \sin\theta_2, -\cos\theta_2).$$

Thus, $\|\overrightarrow{O_1O_2}\|^2 \leqslant \|\overrightarrow{O'_1O'_2}\|^2$ if and only if

$$(d - \sin\theta_2 - \sin\theta_1)^2 \leqslant (d + \sin\theta_2 + \sin\theta_1)^2$$

which simplifies into $-4\,d\,(\sin\theta_1 + \sin\theta_2) \leqslant 0$, that is,

$$\sin\theta_1 + \sin\theta_2 \geqslant 0. \tag{19}$$

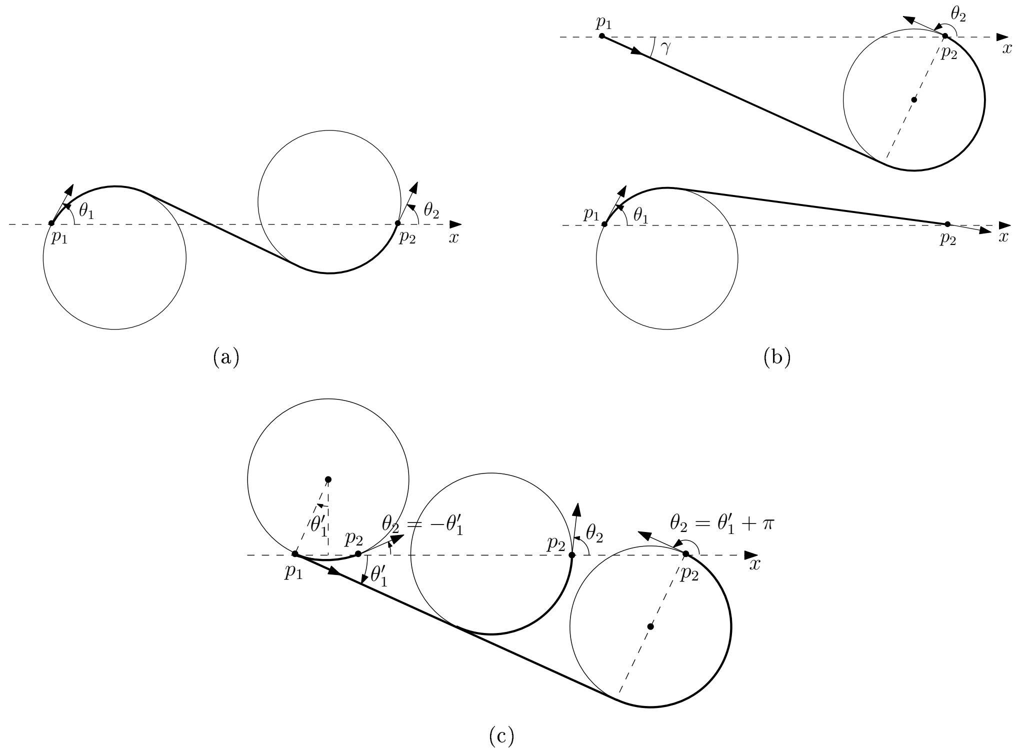

Figure 13: For the proof of Prop. 36 in the case where $T = RSL$.

We now prove that (19) is satisfied which implies that $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M_1'M_2'}\|$. Refer to Figure 13(a). The hypothesis that $T = RSL$ implies that the two circular arcs lie on opposite sides of the line segment. Furthermore, since the second circular arc is shorter than π, if θ_2 decreases continuously with θ_1 fixed, both circular arcs shorten until one of them vanishes. Hence, there exists a path of type RS or SL from (p_1, θ_1) to p_2 whose circular arc is shorter than π (see Figure 13(b)). If the resulting modified path is of type RS, then the segment lies above the x-axis and, since the first circular arc is shorter than π, θ_1 belongs to $[0, \pi]$. Similarly, if the modified path is of type SL, the segment lies below the x-axis and, since the second circular arc is shorter than π, θ_1 belongs to $[-\gamma, 0]$, where $\gamma = \arcsin \frac{2}{d}$ ($d \geqslant 2$ since, otherwise, the circular arc is greater than π). Similarly, if we fix θ_2 and continuously decrease θ_1, both circular arcs shorten until the path is of type RS or SL, and we get, as before, that θ_2 belongs to $[-\gamma, \pi]$. Hence, both θ_1 and θ_2 are in $[-\gamma, \pi]$.

Now, if both θ_1 and θ_2 are in $[0, \pi]$, (19) is satisfied, and thus $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M_1'M_2'}\|$. If θ_1 belong to $[-\gamma, 0]$, we decrease θ_1 with θ_2 fixed, until until one circular arc vanishes. Note that, while θ_1 decreases in $[-\gamma, 0] \subseteq [-\pi/2, 0]$, $\sin\theta_1 + \sin\theta_2$ decreases as well. The resulting modified path, from (p_1, θ_1') to (p_2, θ_2), is of type SL because, if it was of type RS, then $\theta_1 \geqslant \theta_1' > 0$, contradicting the assumption that $\theta_1 \in [-\gamma, 0]$ (if $\theta_1' = 0$, the path is a segment and the path is also of type SL). We now argue that θ_2 is in $[-\theta_1', \theta_1' + \pi,]$ and thus that $\sin\theta_1 + \sin\theta_2 \geqslant \sin\theta_1' + \sin\theta_2 \geqslant 0$. Refer to Figure 13(c) and consider a path of type SL from (p_1, θ_1'), fixed, to a moving configuration (p_2, θ_2) where d increases continuously from $2\sin|\theta_1'|$ to $\frac{2}{\sin|\theta_1'|}$ (outside of this range, there is no path of type SL from (p_1, θ_1') to p_2 with a circular arc shorter than π). When d increases continuously from $2\sin|\theta_1'|$ to $\frac{2}{\sin|\theta_1'|}$, θ_2 increases monotonically from $-\theta_1'$ to $\theta_1' - \pi$. Hence (19) is satisfied, and thus $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M_1'M_2'}\|$.

Finally, if θ_2 belong to $[-\gamma, 0]$, we proceed similarly. We decrease θ_2 with θ_1 fixed until we get a path of type RS from (p_1, θ_1) to (p_2, θ_2') (see Figure 13(b)). A symmetry about p_2 yields a path of type SL from (p_2, θ_2') to (p_1', θ_1) and thus, as above we get that θ_1 is in $[-\theta_2', \theta_2' + \pi]$ and thus that $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M_1'M_2'}\|$.

This concludes the proof that $\|\overrightarrow{M_1M_2}\| \leqslant \|\overrightarrow{M_1'M_2'}\|$ in all cases.

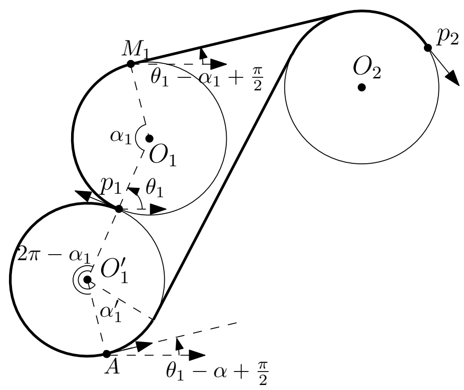

Figure 14: For the proof of Prop. 36, when $\alpha_1 + \alpha'_1 > 2\pi$.

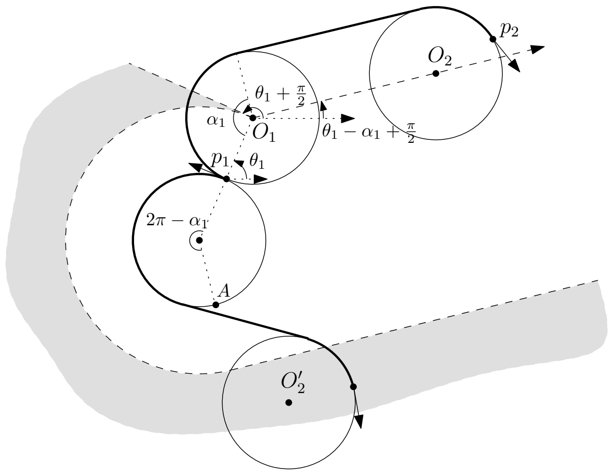

Figure 15: For the proof of Prop. 36, when $T = RSR$ and $T' = LSR$.

Length of the circular arcs. We now show that the sum of the lengths of the circular arcs is strictly smaller on the T-path than on the T'-path. Recall that α_i and α'_i denote the length of the i-th circular arc of the T-path and T'-path, respectively.

Assume first that the first circular arcs of the T and T'-paths have different orientations (R or L), and that the sum $\alpha_1 + \alpha'_1$ of their lengths is larger than 2π. Refer to Figure 14. Since the first circular arc of the T-path has length $\alpha_1 < \pi$, we can define the point A on the first circular arc of the T'-path such that its sub-path from p_1 to A has length $2\pi - \alpha_1$. The sub-path from M_1 to p_2 on the T-path consists of a straight line segment and a circular arc which is shorter than π. Thus, the angular distance from M_1 to p_2 on the T-path is minimal, and the angular distance from A to p_2 on the T'-path is necessarily larger or equal (since the polar angles of the oriented tangents to the two paths at M_1 and A are equal). In other words, the total length of the circular arcs from A to p_2 on the T'-path is larger than or equal to the length of the second circular arc of the T-path. On the other hand, the length of the circular arc from p_1 to A on the T'-path is $2\pi - \alpha_1$ which is greater than α_1 since $\alpha_1 < \pi$.

Hence, if the first circular arcs of the T and T'-paths have different orientations and $\alpha_1 + \alpha'_1 > 2\pi$, the total length of the circular arcs on the T-path is less than on the T'-path, i.e., $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$. We get the same result (by considering the reverse paths), if the second circular arcs of the T and T'-paths have different orientations and if $\alpha_2 + \alpha'_2 > 2\pi$.

Now, assume, without loss of generality, that the first circular arcs of the T and T'-paths have different orientations, R and L, respectively. Three cases occur: either the T and T' have the same orientation on the second circular arc, or (T, T') is equal to (RSL, LSR) or to (RSR, LSL). We consider these three cases in turn, and show that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$.

T and T$'$ have the same orientation (R or L) on the second circular arc. Assume first that $T = RSR$ and $T' = LSR$, and refer to Figure 15.

We argue that the sum $\alpha_1 + \alpha'_1$ of the lengths of the first circular arcs of the two paths is larger than 2π, which, as shown above, yields the result that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$.

Since $T = RSR$, O_2 lies on the ray starting at O_1 with polar angle $\theta_1 - \alpha_1 + \frac{\pi}{2}$. If the length α'_1 of the first circular arc of the LSR-path is smaller than or equal to $2\pi - \alpha_1$, then O'_2 lies in the (closed) grey region of Figure 15. Since $O_2 = O'_2$ lies in the intersection of this region and the ray, either $\alpha_1 = 0$ or the line segment of the T-path vanishes. In both cases, the T-path is of type SR or R which are sub-types of LSR, and thus the T and T'-paths coincide, contradicting our hypothesis. Hence, $\alpha'_1 > 2\pi - \alpha_1$.

Note that we did not use here that the circular arcs of the T-path are shorter than π. Thus, if $T = RSL$ and $T' = LSL$, we can exchange the roles of T and T' and we get, up to a symmetry, paths of type $T = RSR$ and $T' = LSR$. Thus $\alpha_1 + \alpha'_1 > 2\pi$ on the resulting (symmetric-exchanged) paths, and thus also on the initial path.

Case $(\mathbf{T}, \mathbf{T'}) = (\mathbf{RSL}, \mathbf{LSR})$. If $\alpha'_1 \geqslant \pi$ and $\alpha'_2 \geqslant \pi$, then, since $\alpha_1 < \pi$ and $\alpha_2 < \pi$, we directly get the result that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$. We can thus assume that one of the circular arcs of the T'-path is shorter than π. Assume, without loss of generality, that this is its first arc, that is, $\alpha'_1 < \pi$.

We increase θ_1 continuously until α_1 or α_2 reaches π, or α'_1 or α'_2 reaches 0. Consider the T'-path: the fact that $T' = LSR$ and that its first circular arc is shorter than π implies that, while θ_1 increases, both circular arcs shorten, that is α'_1 and α'_2 decrease (see Figure 13(a)). Similarly, for $T = RSL$, the fact that the first circular arc is shorter than π implies that, while θ_1 increases, both circular arcs get longer, that is α_1 and α_2 increase. Hence, while θ_1 increases, $\alpha_1 + \alpha_2$ increases and $\alpha'_1 + \alpha'_2$ decreases. It is thus sufficient to show that the resulting modified paths verify that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$.

If α'_2 reaches 0, the T'-path becomes of type LS; this is a sub-type of LSL and we already proved, in the case where $T = RSL$ and $T' = LSL$, that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$. The case where α'_1 reaches 0 is similar (by considering the reverse paths). On the other hand, if α_1 reaches π, then $\alpha'_1 > \pi$. Indeed, otherwise, O'_2 lies in the intersection of the two grey regions of Figure 16, that is $O'_2 = O_1$, and the two paths are identical and of type R; thus $\alpha_2 = 0$ and, since α_2 increases during the motion, it is equal to 0 during the whole motion, and the initial T-path is of type RS; similarly as above, this is a sub-type of RSR, and the case where $T = RSR$ and $T' = LSR$ was treated before. Thus, if α_1 reaches π, then $\alpha'_1 > \pi$. Hence, $\alpha_1 + \alpha'_1 > 2\pi$, which implies that $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$. We get the same result if α_2 reaches π, by considering, for instance, the reverse path.

Case $(\mathbf{T}, \mathbf{T'}) = (\mathbf{RSR}, \mathbf{LSL})$. Refer to Figure 12. Let $\mathcal{E}$ (resp. $\mathcal{E}'$) denote the (closed) half-plane delimited by the line tangent to the two paths at p_1, and containing the first circular arc of the T-path (resp. T'-path).

If α_1 and α'_1 are smaller than π, then the circles supporting the second circular arcs lie in $\mathcal{E}$ and $\mathcal{E}'$ respectively. Since these two circles have the point p_2 in common, both paths are either identical and reduced to a segment or both paths have their two circular arcs of length π, contradicting the hypothesis in both cases. Hence, $\alpha'_1 \geqslant \pi$ and $\alpha_1 < \alpha'_1$. Similarly, $\alpha_2 < \alpha'_2$ and thus $\alpha_1 + \alpha_2 < \alpha'_1 + \alpha'_2$.

We thus have shown that if both circular arcs of a CSC-path are shorter than π, this path is strictly shorter than the other distinct CSC-paths, which concludes the proof of the proposition. □

Local convexity of $\mathbf{F}(\theta_1, \theta_2)$

We can now prove the local convexity of the length function $F(\theta_1, \theta_2)$ of the shortest CSC-paths from (p_1, θ_1) to (p_2, θ_2) on the domain of (θ_1, θ_2) such that both circular arcs are shorter than π. Figures 2(a) and 2(b) show an example of such domain and of the graph of $F(\theta_1, \theta_2)$ over that domain.

Theorem 37. *The length $F(\theta_1, \theta_2)$ of the shortest CSC-path from (p_1, θ_1) to (p_2, θ_2) is locally strictly convex at any point (θ_1, θ_2) such that both circular arcs of the corresponding path are strictly shorter than π.*

Proof. Theorem 37 essentially follows from Propositions 35 and 36, and from the fact that for any (θ_1, θ_2) such that two T and T'-paths coincide ($T \neq T'$ in $\{LSR, RSL, LSL, RSR\}$), that is when a circular arc vanishes, the first and second derivatives of F_T and $F_{T'}$ are equal and thus F is locally C^2 at this point.

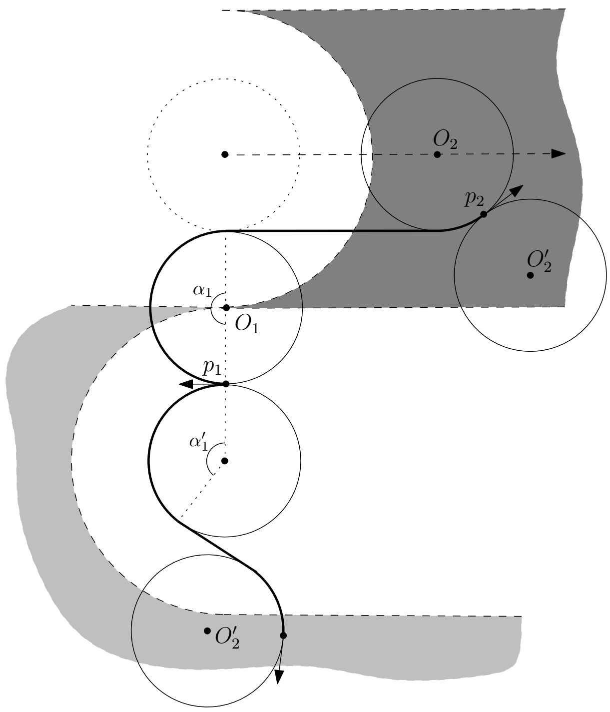

Figure 16: For the proof of Prop. 36: contradiction when $T = RSL$, $T' = LSR$, $\alpha_1 = \pi$ and $\alpha_1' < \pi$.

More precisely, consider a shortest CSC-path such that both its circular arcs are strictly shorter than π, and let T in $\{LSR, RSL, LSL, RSR\}$ denote its type. (All the paths we consider here are from (p_1, θ_1) to (p_2, θ_2).)

If both its circular arcs have nonzero length, it is geometrically distinct from all the other T'-paths, for $T' \neq T$ in $\{LSR, RSL, LSL, RSR\}$. Furthermore, by Propositions 35 and 36, this T-path is strictly shorter than all other T'-paths, and its length function F_T is locally strictly convex at (θ_1, θ_2). Hence, $F = \min_{T \in \{LSR,RSL,LSL,RSR\}} F_T$ is also locally strictly convex at (θ_1, θ_2).

Now, if one circular arc of the T-path has zero length, this path is geometrically equal to another path of type $T' \neq T$. We first show that F is locally convex at such point by showing that F is locally C^2. However, note that this does not yet show the *strict* local convexity at such point. Recall that (see Propositions 33 and 34)

$$\begin{aligned} \frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} &= \mu_{i=1} \mu_{C_i=R} (1 - \cos \alpha_i) \\ \frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} &= \delta_{i,j} \ \sin \alpha_i + \frac{\sin \alpha_i \sin \alpha_j}{M_1 M_2}. \end{aligned} \tag{20}$$

Since the two paths are geometrically equal, the only values that differ in expressions of the partial derivatives of F_T and F_T' is the $\mu_{C_i=R}$ for which the C_i vanish(es). But then, $\alpha_i = 0$ and thus $\frac{\partial F_T(\theta_1,\theta_2)}{\partial \theta_i} = \frac{\partial F_{T'}(\theta_1,\theta_2)}{\partial \theta_i} = 0$. Thus, the function F is locally C^2 at such point, and thus locally convex.

We now show that the function F is locally strictly convex at any point $\tilde{\Theta} = (\tilde{\theta}_1, \tilde{\theta}_2)$ such that two T and T'-paths are geometrically identical with a vanishing first circular arc and a second circular arc of length in $(0, \pi)$. We show that for $\Theta = (\theta_1, \theta_2)$ small enough (in norm), $F(\tilde{\Theta} + \Theta)$ is strictly above the tangent plane to the graph of F at $(\tilde{\Theta}, F(\tilde{\Theta}))$. Consider first the second-order Taylor expansion of F at $\tilde{\Theta}$

$$F(\tilde{\Theta} + \Theta) = F(\tilde{\Theta}) + \Theta \cdot \nabla F + \frac{1}{2} \Theta^T H \ \Theta + o(\|\Theta\|^2)$$

where ∇F and H denote the gradient and the Hessian of F at $\tilde{\Theta}$. Since F is C^2 at $\tilde{\Theta}$, (20) yields that all the partial second derivatives of F at $\tilde{\Theta}$ are zero except for $\frac{\partial^2 F(\tilde{\Theta})}{\partial \theta_2^2}$ which is strictly positive. We thus get that $\Theta^T H\ \Theta = \frac{\partial^2 F(\tilde\Theta)}{\partial \theta_2^2}\theta_2^2$ which is strictly positive unless $\theta_2 = 0$. When $\theta_2 = 0$, we consider the third-order Taylor expansion of F_T (the terms in $\theta_1^2\theta_2, \theta_1\theta_2^2$, and θ_2^3 are zero since $\theta_2 = 0$):

$$F_T(\tilde\Theta + \Theta) = F_T(\tilde\Theta) + \Theta \cdot \nabla F_T + \frac{1}{6}\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3}\theta_1^3 + o(\|\Theta\|^3) \tag{21}$$

where $\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3}$ denotes the third derivative of F_T with respect to θ_1 at $\tilde\Theta$. Equation (20) gives that

$$\frac{\partial^3 F_T}{\partial \theta_1^3} = \frac{\partial \sin\alpha_1}{\partial\theta_1}\left(1 + \frac{\sin\alpha_1}{M_1M_2}\right) + \sin\alpha_1 \frac{\partial\left(1 + \frac{\sin\alpha_1}{M_1M_2}\right)}{\partial\theta_1}.$$

However, since the first circular arc vanishes, i.e., $\alpha_1 = 0$, at $\tilde\Theta$, we get that

$$\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3} = \frac{\partial \sin\alpha_1}{\partial\theta_1}.$$

Now, $\sin\alpha_1 = \overrightarrow{U_{12}} \cdot \overrightarrow{p_1O_1}$, thus

$$\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3} = \frac{\partial\overrightarrow{U_{12}}}{\partial\theta_1}\cdot\overrightarrow{p_1O_1} + \overrightarrow{U_{12}}\cdot\frac{\partial\overrightarrow{p_1O_1}}{\partial\theta_1}.$$

The second term is equal, by (10), to $-\overrightarrow{U_{12}} \times \overrightarrow{p_1O_1} = \mu_{C_1=R}$. Thus, by (16), we get that

$$\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3} = \left(\frac{1}{M_1M_2}\frac{\partial\overrightarrow{M_1M_2}}{\partial\theta_j} - \frac{1}{M_1M_2^2}\frac{\partial M_1M_2}{\partial\theta_j}\overrightarrow{M_1M_2}\right)\cdot\overrightarrow{p_1O_1} + \mu_{C_1=R}.$$

Furthermore, since $\overrightarrow{M_1M_2} \cdot \overrightarrow{p_1O_1} = 0$ at $\tilde\Theta$,

$$\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3} = \frac{1}{M_1M_2}\frac{\partial\overrightarrow{M_1M_2}}{\partial\theta_j}\cdot\overrightarrow{p_1O_1} + \mu_{C_1=R}.$$

Now, we decompose $\overrightarrow{M_1M_2}$ into $\overrightarrow{M_1O_1} + \overrightarrow{O_1p_1} + \overrightarrow{p_1O_2} + \overrightarrow{O_2M_2}$. The derivative of $\overrightarrow{M_1O_1} = \pm\overrightarrow{O_2M_2}$ is orthogonal to $\overrightarrow{M_1O_1}$ which is equal to $\overrightarrow{p_1O_1}$ at $\tilde\Theta$, the derivative of $\overrightarrow{O_1p_1}$ is orthogonal to $\overrightarrow{O_1p_1}$, and the derivative of $\overrightarrow{p_1O_2}$ with respect to θ_1 is zero. Hence,

$$\frac{\partial^3 F_T(\tilde\Theta)}{\partial \theta_1^3} = \mu_{C_1=R}.$$

Finally, since the first circular arc vanishes in the shortest path with polar angle $\tilde\Theta$, we observe that the shortest path with polar angle $\tilde\Theta + \Theta$, with $\theta_2 = 0$ and $|\theta_1|$ small enough, is such that C_1 is oriented R if $\theta_1 > 0$, and L if $\theta_1 < 0$. Thus $\mu_{C_1=R}\ \theta_1^3 > 0$ for $\theta_1 \neq 0$. Therefore, we have proved that for Θ small enough (in norm), $F(\tilde\Theta + \Theta)$ is strictly above the tangent plane to the graph F at $\tilde\Theta$.

Hence, we have proved that F is locally strictly convex at any point such that one circular arc has length in $[0, \pi)$ and the other has length in $(0, \pi)$. By continuity, F is also locally strictly convex at the point such that both circular arcs vanishes, which concludes the proof. □

B Diamonds and lemons

Let $\alpha \in (0, \pi]$ and recall that the lemon $L_i^{i+1}(\alpha)$ is defined as the set of angles (θ_i, θ_{i+1}) in $(\mathbb{S}^1)^2$ such that both circular arcs of the shortest path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) have length *strictly* less than α. Also recall that the diamond D_i^{i+1} is defined as the image of the open quadrilateral with vertices $(0, 2\pi)$, (ξ_i, ξ_i), $(2\pi, 0)$, and $(2\pi - \xi_i, 2\pi - \xi_i)$ under the translation of vector $(\nu_{i+1}^i, \nu_{i+1}^i)$, where $\xi_i = \frac{2\pi}{d_i - \frac{1}{d_i}}$ and $d_i = |p_ip_{i+1}|$. Here, we give a complete proof for Lemma 15, which states:

If $|p_ip_{i+1}| \geqslant 4$ *then* $L_i^{i+1}\left(\frac{3\pi}{4}\right) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$.

For the rest of the section, let d denote the distance $d_i = |p_ip_{i+1}|$.

B.1 Translations and symmetries

We first identify certain symmetries of the regions $L_i^{i+1}(\alpha)$ that will simplify the discussion.

Translations of the lemon $\mathbf{L_i^{i+1}}(\alpha)$. First, we can assume that p_{i+1} lies at the origin, because translating $\{p_i, p_{i+1}\}$ changes neither $L_i^{i+1}(\alpha)$ nor D_i^{i+1}. Now, rotating p_i about p_{i+1} by any angle μ, increases θ_i and θ_{i+1} by μ, that is, translates $L_i^{i+1}(\alpha)$ and D_i^{i+1} by vector (μ, μ). This does not change their relative positions. We can thus assume that $p_{i+1} = (0,0)$, $p_i = (d, 0)$, and $\nu_{i+1}^i = 0$. Then, the lemon $L_i^{i+1}(\alpha)$ and diamond D_i^{i+1} do not intersect the two hyperplanes $\{(\theta_i, \theta_{i+1}) \in (\mathbb{S}^1)^2 \mid \theta_i = \nu_{i+1}^i = 0\}$ and $\{(\theta_i, \theta_{i+1}) \in (\mathbb{S}^1)^2 \mid \theta_{i+1} = \nu_{i+1}^i = 0\}$ (as noted in the proofs of Lemmas 8 and 17). It follows that we can lift $L_i^{i+1}(\alpha)$ and D_i^{i+1} from $(\mathbb{S}^1)^2$ onto the square $(0, 2\pi)^2$ of $\mathbb{R}^2$, and work in a Cartesian coordinate system (x, y). Refer to Figure 7 with $\nu_{i+1}^i = 0$ (or similarly to Figure 2(a)). Note that this lift is different from the lift $\mathbb{S}^1 \times \mathbb{S}^1 \to \Lambda_i \times \Lambda_{i+1}$ considered in Section 4, because, even though $\Lambda_i = [0, 2\pi)$, the endpoints of Λ_{i+1} are in $\nu_{i+2}^{i+1} + 2\pi\mathbb{Z}$ which does not contain 0.

Symmetries of $\mathbf{L_i^{i+1}}(\alpha)$. We first observe some symmetries of the region $L_i^{i+1}(\alpha)$. We define the following transformations on Dubins paths:

s is the symmetry with respect to the x-axis. It maps the L_aSL_b (resp. L_aSR_b, R_aSL_b, R_aSR_b) path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) to the R_aSR_b (resp. R_aSL_b, L_aSR_b, L_aSL_b) path from $(p_i, 2\pi - \theta_i)$ to $(p_{i+1}, 2\pi - \theta_{i+1})$.

r is composed of three transforms. We first reverse the path, turning the L_aSL_b (resp. L_aSR_b, R_aSL_b, R_aSR_b) path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) into the R_bSR_a (resp. L_bSR_a, R_bSL_a, L_bSL_a) path from $(p_{i+1}, \theta_{i+1} - \pi)$ to $(p_i, \theta_i - \pi)$. We then apply a symmetry with respect to the bisecting line of p_i and p_{i+1}, turning the R_bSR_a (resp. L_bSR_a, R_bSL_a, L_bSL_a) path from $(p_{i+1}, \pi + \theta_{i+1})$ to $(p_i, \pi + \theta_i)$ into the L_bSL_a (resp. R_bSL_a, L_bSR_a, R_bSR_a) path from $(p_i, 2\pi - \theta_{i+1})$ to $(p_{i+1}, 2\pi - \theta_i)$. We finally apply the symmetry s, which turns the path into the R_bSR_a (resp. L_bSR_a, R_bSL_a, L_bSL_a) path from (p_i, θ_{i+1}) to (p_{i+1}, θ_i).

The symmetry s implies that the shortest Dubins path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) and the one from $(p_i, 2\pi - \theta_i)$ to $(p_{i+1}, 2\pi - \theta_{i+1})$ have equal length and that their longest circular arcs also have equal length. This implies that $L_i^{i+1}(\alpha)$ is symmetric with respect to point (π, π). Similarly, the transform r implies that the shortest Dubins path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) and the one from (p_i, θ_{i+1}) to (p_{i+1}, θ_i) have equal length and that their longest circular arcs also have equal length. Thus, $L_i^{i+1}(\alpha)$ is symmetric with respect to the line $y = x$. By composing this symmetry and the symmetry about (π, π), we get that $L_i^{i+1}(\alpha)$ is also symmetric with respect to the line $y = 2\pi - x$ (the line through (π, π) and orthogonal to $y = x$).

B.2 Boundary of $\mathbf{L_i^{i+1}}(\alpha)$

We now give an analytical description of the boundary of $L_i^{i+1}(\alpha)$. As mentioned above, we consider $L_i^{i+1}(\alpha)$ through the lift from $(\mathbb{S}^1)^2$ to $(0, 2\pi)^2$. Its boundary thus lies in the closed square $[0, 2\pi]^2$. For simplicity, for any $(x, y) \in [0, 2\pi]^2$, the *path corresponding to* (x, y) refers to the shortest CSC-path from (p_i, x) to (p_{i+1}, y) whose longest circular arc has minimal length (in the case where several shortest CSC paths with minimal longest circular arc exist, we pick any of them).

We obtain our analytical description in three steps. We first show that any point on the boundary of $L_i^{i+1}(\alpha)$ has the property that the longest circular arc of its corresponding path has length α. We then give an analytical description of arcs of curves that are guaranteed to contain any point with that property. Let us emphasize that at this point, we will only have showed inclusions and do not claim that any point with the above property is on the boundary of $L_i^{i+1}(\alpha)$ nor that any point on the arcs of curves have that property. Instead of trying to prove these reverse inclusions directly, in the third step, we argue that the union of the arcs forms a simple closed curve σ in $[0, 2\pi]^2$; since $L_i^{i+1}(\alpha)$ has nonempty interior and exterior regions, its boundary must disconnect $\mathbb{S}^1 \times \mathbb{S}^1$ and therefore cannot be a proper subset of a simple closed curve.

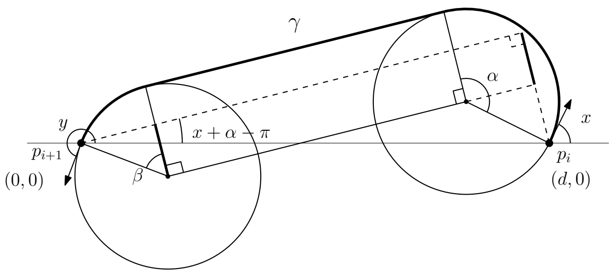

Figure 17: For the expression of β when (x, y) is on $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$.

Step 1. We first claim that the path corresponding to any point (x, y) on the boundary of $L_i^{i+1}(\alpha)$ has its longest circular arc of length exactly α. This statement may seem obvious, but it is not: one has to account for possible discontinuities in the function that maps (x, y) to the length of the longest circular arc of its corresponding path. If (x, y) is such a discontinuity then there must exist two shortest path from (p_i, x) to (p_{i+1}, y) with longest circular arcs of distinct lengths; this cannot occur in $L_i^{i+1}(\pi)$ by Proposition 5. As a consequence we get that the path corresponding to any point (x, y) on the boundary of $L_i^{i+1}(\alpha)$ has its longest circular arc of length either equal to α or strictly larger than π. In the latter case, our definition of "corresponding path" implies that *any* shortest CSC-path from (p_i, x) to (p_{i+1}, y) has its longest circular arc of length strictly more than π; it follows that in a neighborhood of (x, y), the longest circular arc of the corresponding path remains strictly larger than π, and such a point (x, y) cannot be on the boundary of $L_i^{i+1}(\alpha)$. The claim follows.

Step 2. Let $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ be the set of points $(x, y) \in [0, 2\pi]^2$ whose corresponding path, γ, has type $L_\alpha SL_\beta$ and $L_\alpha SR_\beta$, respectively, with $\beta \leqslant \alpha$. For technical reasons, it is convenient to consider (x, y) in the *open* square $(0, 2\pi)^2$, and to actually define $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ by continuity over the closed square $[0, 2\pi]^2$. [20] We now describe analytically these curves.

Let us first consider $(x, y) \in \mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$. Referring to Figure 17, we have:

$$y = x + \alpha + \beta. \tag{22}$$

Note that this equality is not modulo 2π. Indeed, since γ is of type LSL, $\beta \leqslant \alpha \leqslant \pi$, and p_{i+1} is left of p_i on the x-axis, we have that x is in $(0, \pi]$ and y is in $[\pi, 2\pi)$. Since $d \geqslant 4$, the two circles supporting the circular arcs of γ are separated by a vertical line, which is crossed (from right to left) by the line segment of γ; hence $x + \alpha \in [\frac{\pi}{2}, \frac{3\pi}{2}]$. Furthermore, since $\beta \leqslant \alpha$, the oriented line segment of γ must point down, that is $x + \alpha \in [\pi, \frac{3\pi}{2}]$. Since $\beta \leqslant \alpha \leqslant \pi$, we thus have $x + \alpha + \beta \in [\pi, \frac{5\pi}{2}]$, and the fact that $y \in [\pi, 2\pi)$ yields the claim.

Furthermore, since $x + \alpha \in [\pi, \frac{3\pi}{2}]$, we obtain the expression of β in terms of α, by considering the length of the short bold segments in Figure 17:

$$\cos\beta \;=\; d\sin(x + \alpha - \pi) - \sin\Big(\alpha - \frac{\pi}{2}\Big) \;=\; \cos\alpha - d\sin(x + \alpha). \tag{23}$$

Plugging this expression in (22) we get that the curve $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ has the following equation; the upper bound on the domain of x corresponds to the condition $\cos\alpha - d\sin(x + \alpha) \leqslant 1$ with $\pi \leqslant x + \alpha \leqslant \frac{3\pi}{2}$, and the lower bound comes from $\pi \leqslant x + \alpha$.

$$\begin{aligned} \mathcal{C}_{L_\alpha SL_{\leqslant\alpha}} : \quad & y = x + \alpha + \arccos(\cos\alpha - d\sin(x + \alpha)) \\ & \text{for} \quad \pi - \alpha \leqslant x \leqslant \pi - \alpha + \arcsin\Big(\frac{1 - \cos\alpha}{d}\Big). \end{aligned} \tag{24}$$

[20] When $\alpha < \pi$ it is straightforward that $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ are contained in $(0, 2\pi)^2$. When $\alpha = \pi$, there are exactly four points of the boundary of the square $[0, 2\pi]^2$, namely its corners, that have corresponding paths of type $L_\pi SL_\beta$ or $L_\pi SR_\beta$ with $\beta \leqslant \pi$. Although these four corners correspond to the same configurations when seen in $(\mathbb{S}^1)^2$, $(0, 2\pi)$ and $(2\pi, 0)$ are the only ones that belong, as points of $[0, 2\pi]^2$, to the closure of the parts of $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ contained in $(0, 2\pi)^2$; we omit the other two points, $(0, 0)$ and $(2\pi, 2\pi)$, since they are isolated and play no role in the boundary of $L_i^{i+1}(\pi)$ in $[0, 2\pi]^2$.

Note that the lower bound on x corresponds to a LSL path whose circular arcs both have length α (and whose segment has polar angle π). The upper bound on x corresponds to $\cos\beta = 1$, that is, to a path whose second circular arc vanishes. A similar argument yields that the curve $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ has equation:

$$\begin{aligned}\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}:\quad & y = x+\alpha-\arccos(2-\cos\alpha+d\sin(x+\alpha))\\ & \text{for}\quad \pi-\alpha+\arcsin\left(\frac{1-\cos\alpha}{d}\right)\leqslant x\leqslant \pi-\alpha+\arcsin\left(\frac{2-2\cos\alpha}{d}\right).\end{aligned}\tag{25}$$

The curves $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ meet in one of their endpoints (at $x = \pi-\alpha+\arcsin\left(\frac{1-\cos\alpha}{d}\right)$ and $y = x+\alpha$), which corresponds to a path of type $L_\alpha S$. Their union is thus a connected curve τ whose endpoints are $(\pi-\alpha, \pi+\alpha)$ on the line $y = 2\pi - x$, and $(\tilde{x},\tilde{x})$ on the line $y = x$, with $\tilde{x} = \pi-\alpha+\arcsin\left(\frac{2-2\cos\alpha}{d}\right)$. Moreover, we prove below (from (26) and (29)) that the slope of τ is everywhere less than -1, which implies that τ is simple and lies entirely in the quadrant $x\leqslant y\leqslant 2\pi - x$.

Step 3. Since any boundary point of $L_i^{i+1}(\alpha)$ with corresponding path of type $L_\alpha SL_{\leqslant\alpha}$ or $L_\alpha SR_{\leqslant\alpha}$ must belong to τ, the symmetries r and s yield that any boundary point of $L_i^{i+1}(\alpha)$ must belong to τ or one of its symmetric copies with respect to point (π,π), to line $y = x$, and their combinations. We already know that τ is simple and lies in the quadrant $x\leqslant y\leqslant 2\pi-x$ with each endpoint on each of the two lines $y = x$ and $y = 2\pi - x$. Thus, the union of τ and its three symmetric copies forms a simple closed curve σ in $[0, 2\pi]^2$.

Let $\hat{\sigma}$ denote the projection of σ on $\mathbb{S}^1\times\mathbb{S}^1$. For $\alpha<\pi$, τ lies in the open square $(0,2\pi)^2$ and $\hat{\sigma}$ is therefore a simple closed curve in $\mathbb{S}^1\times\mathbb{S}^1$. Since both $L_i^{i+1}(\alpha)$ and its complement have interior points, the boundary of $L_i^{i+1}(\alpha)$ cannot be a proper subset of a simple closed curve and thus it must be equal to $\hat{\sigma}$. The situation for $\alpha = \pi$ requires more care, as τ lies in the open square $(0, 2\pi)^2$ except for its endpoint with coordinates $(0, 2\pi)$. Thus, σ contains two points, $(0, 2\pi)$ and $(2\pi, 0)$, which correspond to the same point in $\mathbb{S}^1\times\mathbb{S}^1$. More precisely, σ consists of two curves that are simple, disjoint, and lie in $(0, 2\pi)^2$ except for their endpoints which are equal to $(0, 2\pi)$ and $(2\pi, 0)$. It follows that $\hat{\sigma}$ is a closed curve in $\mathbb{S}^1\times\mathbb{S}^1$ consisting of two simple loops meeting in exactly one point: $(0, 0)$. We claim that each loop contains some point, other than $(0, 0)$, that belongs to the boundary of $L_i^{i+1}(\pi)$. Indeed, as shown previously, the junction between $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}$ and $\mathcal{C}_{L_\pi SR_{\leqslant\pi}}$ is an interior point of τ with corresponding path of type $L_\pi S$, and thus lies on the boundary of $L_i^{i+1}(\pi)$. By symmetry, the other (top) curve also contains a point on the boundary of $L_i^{i+1}(\pi)$. Together with the observation that both $L_i^{i+1}(\pi)$ and its complement have interior points, this implies that the boundary of $L_i^{i+1}(\pi)$ cannot be a proper subset of $\hat{\sigma}$ but is equal to the whole curve.

We thus obtained a complete description of $L_i^{i+1}(\alpha)$. Note finally that the point (π,π), which clearly belongs to any lemon, indicates on which side of this curve the lemon lies.

B.3 Position of $\mathbf{D_i^{i+1}}$

To complete the proof, it suffices to show that the segment bounding D_i^{i+1} in the quadrant $x\leqslant y\leqslant 2\pi-x$ lies above the two curves $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}$ and $\mathcal{C}_{L_\pi SR_{\leqslant\pi}}$, and below the two curves $\mathcal{C}_{L_{\frac{3\pi}{4}}SL_{\leqslant\frac{3\pi}{4}}}$ and $\mathcal{C}_{L_{\frac{3\pi}{4}}SR_{\leqslant\frac{3\pi}{4}}}$ (as in Figure 7). Before completing the proof of the lemma, we first provide a very simple proof of a similar result in the case where $d > 8.6$.

B.3.1 Simple case where $\mathbf{d > 8.6}$

Our simple proof uses a different diamond D_i^{i+1}, defined independently of d. Specifically, we let D_i^{i+1} be the image of the quadrilateral with vertices $(0, 2\pi)$, $(\frac{\pi}{4},\frac{\pi}{4})$, $(2\pi, 0)$, and $(\frac{7\pi}{4},\frac{7\pi}{4})$ under the translation of vector $(\nu_{i+1}^i,\nu_{i+1}^i)$. The segment that bounds D_i^{i+1} in the quadrant $x\leqslant y\leqslant 2\pi-x$ lies on the line $y = 2\pi-7x$ with x ranging from 0 to $\frac{\pi}{4}$. Consider the functions of x whose graphs are the curves $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$. A simple calculation shows that, for any α, the derivative of these functions are less than -7.[21] We have seen in Step 2 of Section B.2 that the leftmost point of $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}\cup\mathcal{C}_{L_\pi SR_{\leqslant\pi}}$ is the point

[21]More precisely, the derivative of the function of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ minus -7 is $8+\frac{d\cos(x+\alpha)}{\sqrt{1-(\cos(\alpha)-d\sin(x+\alpha))^2}}$. Noting that the term in the square root is non-negative (by (23)) and that $\cos(x+\alpha)\leqslant 0$ (since $\pi\leqslant x+\alpha\leqslant\frac{3\pi}{2}$), the expression

$(0, 2\pi)$ which lies on the line $y = 2\pi - 7x$. It follows that this line is strictly above $\mathcal{C}_{L_\pi SL_{\leqslant\pi}} \cup \mathcal{C}_{L_\pi SR_{\leqslant\pi}}$, except for the endpoint $(0, 2\pi)$. On the other hand, a simple calculation also shows that the rightmost point of $\mathcal{C}_{L_{\frac{3\pi}{4}} SL_{\leqslant\frac{3\pi}{4}}} \cup \mathcal{C}_{L_{\frac{3\pi}{4}} SR_{\leqslant\frac{3\pi}{4}}}$ is strictly above the line $y = 2\pi - 7x$, which implies that the line is strictly below this curve, and concludes the proof.

B.3.2 Case where $\mathbf{d \geqslant 4}$

Our proof consists of two parts considering the curves $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$, independently. Let ℓ be the line through points $(0, 2\pi)$ and (ξ_i, ξ_i), that is the line that contains the boundary segment of D_i^{i+1} lying in $x \leqslant y \leqslant 2\pi - x$. The idea of the proof is to bound the slope of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ from above by the slope of ℓ and to to bound the values of the local extrema of $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ using partial sums of its power series. In the following, we show how the proof is conducted but we do not detail all the calculations (which are done with Maple).

Curves $\mathcal{C}_{\mathbf{L}_\alpha \mathbf{SL}_{\leqslant\alpha}}$. We show here that for any x, the slope of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ is less than the slope of ℓ. This will imply that $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}$ is below ℓ because the leftmost point of $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}$ is $(0, 2\pi)$, which lies on ℓ. This will also imply that $\mathcal{C}_{L_{\frac{3\pi}{4}} SL_{\leqslant\frac{3\pi}{4}}}$ lies above ℓ once we proved that its rightmost point lies above ℓ, which will come from the second part of the proof where we show that $\mathcal{C}_{L_{\frac{3\pi}{4}} SR_{\leqslant\frac{3\pi}{4}}}$ lies above ℓ.

The slope of ℓ is $-(d - \frac{1}{d} - 1)$, and that of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ is

$$1 + \frac{d\cos(x+\alpha)}{\sqrt{1 - (\cos(\alpha) - d\sin(x+\alpha))^2}}. \tag{26}$$

Noting that $\pi \leqslant x + \alpha \leqslant \frac{3\pi}{2}$, for any x, the slope of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ is less than the slope of ℓ if and only if

$$d^2\cos^2(x+\alpha) - \left(d - \frac{1}{d}\right)^2 \left(1 - (\cos(\alpha) - d\sin(x+\alpha))^2\right) > 0,$$

which is equivalent, after multiplying by d^2 and replacing $\sin(x+\alpha)$ by z, to

$$d^2\left(d^4 - 3d^2 + 1\right) z^2 - 2\cos(\alpha)\, d\left(d^2 - 1\right)^2 z + d^4 - \left(1 - \cos^2(\alpha)\right)\left(d^2 - 1\right)^2 > 0. \tag{27}$$

Consider this expression as a degree-two polynomial in z. The values of d for which its leading coefficient vanishes are all smaller than 2, and thus the leading coefficient is positive for any $d \geqslant 2$. On the other hand, the discriminant is a polynomial in d, which, after simplifying by $4d^2$ and substituting d by $u + 4$, is

$$\begin{aligned}(\cos^2\alpha - 2)u^6 &+ (24\cos^2\alpha - 48)u^5 + (238\cos^2\alpha - 473)u^4 + (-2448 + 1248\cos^2\alpha)u^3 \\ &+ (-7013 + 3649\cos^2\alpha)u^2 + (-10536 + 5640\cos^2\alpha)u + 3600\cos^2\alpha - 6479.\end{aligned}$$

Since all its coefficients are negative, it is negative for all $u \geqslant 0$, and thus the discriminant is negative for all $d \geqslant 4$. Therefore, (27) is satisfied, which completes the proof that the slope of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ is less than the slope of ℓ, for any x.

Curve $\mathcal{C}_{\mathbf{L}_\alpha \mathbf{SR}_{\leqslant\alpha}}$. Unlike in the proof for $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$, it is not true that the slope of $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ is less than that of ℓ, for any x and $d \geqslant 4$. Thus, here we take a different approach, that is, to bound the values of local extrema.

Recall that we need to prove that (i) $\mathcal{C}_{L_\pi SR_{\leqslant\pi}}$ is below ℓ and that (ii) $\mathcal{C}_{L_{\frac{3\pi}{4}} SR_{\leqslant\frac{3\pi}{4}}}$ is above ℓ. Let F_α be the function of x whose graph is $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$ minus the function whose graph is ℓ (see (25)):

$$\begin{aligned} F_\alpha : &\left[\pi - \alpha + \arcsin\left(\frac{1-\cos\alpha}{d}\right), \pi - \alpha + \arcsin\left(\frac{2 - 2\cos\alpha}{d}\right)\right] \to \mathbb{R} \\ &x \mapsto x + \alpha - \arccos(2 - \cos\alpha + d\sin(x+\alpha)) - 2\pi + \left(d - \frac{1}{d} - 1\right)x. \end{aligned} \tag{28}$$

is negative if and only if $64\left(1 - (\cos(\alpha) + d\sin(x+\alpha))^2\right) - d^2\cos^2(x+\alpha) < 0$. This is a degree-two polynomial in $\sin(x+\alpha)$ whose leading coefficient is $-63\,d^2 < 0$ and discriminant is $4\,d^2\left(-63\,d^2 + 4032 + 64\cos^2\alpha\right)$, which is negative for any $d \geqslant 8.6 > \sqrt{\frac{4032+64}{63}} \geqslant \sqrt{\frac{4032+64\cos^2\alpha}{63}}$. Hence, for any x and $d \geqslant 8.6$, the slope of $\mathcal{C}_{L_\alpha SL_{\leqslant\alpha}}$ is less than that of $y = 2\pi - 7x$. The calculation is similar for $\mathcal{C}_{L_\alpha SR_{\leqslant\alpha}}$.

Let x_{min} and x_{max} denote the leftmost and rightmost endpoints of the domain of definition of F_α. We want to prove that $F_\pi(x) < 0$ and $F_{\frac{3\pi}{4}}(x) > 0$ for all x in $[x_{min}, x_{max}]$. For that purpose, we first prove that F_α has at most two local extrema, x_- and x_+, other than x_{min} and x_{max}, and then that the above inequalities are satisfied for x_{min}, x_{max}, x_-, and x_+, for any $d \geqslant 4$.

Local extrema of F. We prove here that F_α has at most two local extrema, $x_\pm$, other than its endpoints x_{min} and x_{max}. Since F_α is continuous, we prove this by showing that (i) the denominator of F'_α does not vanish on (x_{min}, x_{max}), and that (ii) the numerator of F'_α admits two real roots (we do not need to argue if and when these roots belong to (x_{min}, x_{max})).

The derivative of F_α is

$$F'_\alpha(x) = \frac{d\cos(x+\alpha)}{\sqrt{1-(-2+\cos(\alpha)-d\sin(x+\alpha))^2}} + d - \frac{1}{d}. \tag{29}$$

The denominator vanishes when the square root term in $F'_\alpha(x)$ vanishes, that is when $-2+\cos(\alpha) - d\sin(x+\alpha) = \varsigma = \pm 1$. When $\varsigma = -1$, this is equivalent to $\sin(x+\alpha) = -\frac{1-\cos(\alpha)}{d}$, and to $x = x_{min} = \pi - \alpha + \arcsin\left(\frac{1-\cos\alpha}{d}\right)$ since $\pi \leqslant x+\alpha \leqslant \frac{3\pi}{2}$. When $\varsigma = +1$, the equality is equivalent to $x = \pi - \alpha + \arcsin\left(\frac{3-\cos\alpha}{d}\right)$, which is larger or equal to x_{max}; indeed, this is equivalent to $\arcsin\left(\frac{3-\cos\alpha}{d}\right) \geqslant \arcsin\left(\frac{2-2\cos\alpha}{d}\right)$, that is to $3-\cos\alpha \geqslant 2-2\cos\alpha$, or $\cos\alpha \geqslant -1$. Hence, the denominator of $F'_\alpha(x)$ does not vanish on (x_{min}, x_{max}).

The numerator of F'_α is

$$d^2\cos(x+\alpha) + \left(d^2-1\right)\sqrt{\delta} \quad \text{with} \quad \delta = 1-(-2+\cos(\alpha)-d\sin(x+\alpha))^2.$$

Since $\cos(x+\alpha) \leqslant 0$ (by (25)) and $d^2-1 > 0$ (for $d \geqslant 4$), the numerator vanishes if and only if

$$\delta\left(d^2-1\right)^2 - d^4\cos^2(x+\alpha) = 0. \tag{30}$$

This is a degree-two equation in $\sin(x+\alpha)$. Set $z = \sin(x+\alpha)$ and denote by $z_\pm$ its two roots. Then,

$$z_\pm = \frac{\left(d^2-1\right)^2(2-\cos\alpha) \pm \sqrt{\Delta}}{d\left(d^2-\left(d^2-1\right)^2\right)}, \tag{31}$$

$$\text{with} \quad \Delta = \left(d^2-1\right)^4(2-\cos\alpha)^2 - \left(d^2-\left(d^2-1\right)^2\right)\left(\left(d^2-1\right)^2\left(1-(2-\cos\alpha)^2\right) - d^4\right).$$

Hence F_α has at most two local extrema in the open interval (x_{min}, x_{max}), namely the values $x_\pm = \pi - \arcsin(z_\pm) - \alpha$ (recall that $x+\alpha \in [\pi, \frac{3\pi}{2}]$ by (25)).

Sign of $\mathbf{F}_\alpha(\mathbf{x}_\pm)$. Using $\arccos u = \frac{\pi}{2} - \arcsin u$ and $x = \pi - \arcsin(z) - \alpha$, we can rewrite F_α as

$$F_\alpha(x) = -\frac{5}{2}\pi + \alpha + \arcsin(2-\cos\alpha + d\,z) + \left(d - \frac{1}{d}\right)(\pi - \alpha - \arcsin z). \tag{32}$$

For studying the sign of $F_\alpha(x_\pm)$, we bound arcsin by partial sums of its power series (we need to use partial sums of degree at least 5 for the proof to work). We also need to determine the sign of the operands of arcsin, that is, of $z_\pm$ and $2-\cos(\alpha) + d\,z_\pm$.

Bounds on arcsin. Let $\Gamma_1(z) = z + \frac{1}{6}z^3 + \frac{3}{40}z^5$ and $\Gamma_2(z) = z + \frac{1}{6}z^3 + (\frac{\pi}{2} - \frac{7}{6})z^5$; note that $\Gamma_1(z)$ is the partial sum, up to degree five, of the power series of arcsin, and that the degree-five coefficient of $\Gamma_2(z)$ is such that $\Gamma_2(1) = \frac{\pi}{2} = \arcsin 1$. We prove here that

$$\begin{aligned} \Gamma_1(z) \leqslant \arcsin z \leqslant \Gamma_2(z) \quad &\text{for} \quad z \in [0,1], \\ \Gamma_2(z) \leqslant \arcsin z \leqslant \Gamma_1(z) \quad &\text{for} \quad z \in [-1,0]. \end{aligned} \tag{33}$$

Note first that the second set of inequalities follows from the first set, since all three functions are odd. Also the inequality $\Gamma_1(z) \leqslant \arcsin z$, for $z \geqslant 0$, follows from the fact that all the terms of the power series of arcsin are positive. We thus only have to prove that $\arcsin z \leqslant \Gamma_2(z)$, for $z \in [0,1]$.

The derivative of $\Gamma_2(z)-\arcsin z$ is $1+\frac{1}{2}z^2+5(\frac{\pi}{2}-\frac{7}{6})z^4-\frac{1}{\sqrt{1-z^2}}$, which is zero only if $\left(1+\frac{1}{2}z^2+5(\frac{\pi}{2}-\frac{7}{6})z^4\right)^2(1-z^2)-1=0$. The changes of variable $z=1/t$ and $t=u+1$ transform the z-interval $(0,1)$ into the u-interval $(0,+\infty)$, and the polynomial into

$$(-447+180\,\pi)\,u^6+(-2682+1080\,\pi)\,u^5+(-6504+2610\,\pi)\,u^4+(-8136+3240\,\pi)\,u^3 \\ +\left(-4064+1020\,\pi+225\,\pi^2\right)u^2+\left(-1560\,\pi+992+450\,\pi^2\right)u-36.$$

One can easily check that all the coefficients are positive except for the constant, thus by Descartes' rule of signs, the derivative of $\Gamma_2(z)-\arcsin z$ vanishes exactly once in $(0,+\infty)$. Since $\Gamma_2(z)$ and $\arcsin z$ are equal for $z=0$ and $z=1$, they cannot be equal for $z\in(0,1)$. The result follows because the inequality is verified for z sufficiently small since $\frac{\pi}{2}-\frac{7}{6}>\frac{3}{40}$ and all the terms of higher degree of the power series are negligible for z small enough.

Signs of $\mathbf{z}_\pm$ and $(\mathbf{2}-\cos\alpha+\mathbf{d}\,\mathbf{z}_\pm)$. We prove here that $z_\pm\leqslant 0$, $(2-\cos\alpha+d\,z_+)\leqslant 0$, and $(2-\cos\alpha+d\,z_-)>0$. First note that $z_\pm\leqslant 0$, since $z_\pm=\sin(x_\pm+\alpha)$ and $\pi\leqslant x+\alpha\leqslant\frac{3\pi}{2}$ (by (25)). By (31), we have

$$-dz_\pm=\frac{(d^2-1)^2}{(d^2-1)^2-d^2}(2-\cos\alpha)\pm\frac{\sqrt{\Delta}}{(d^2-1)^2-d^2}.$$

For $d\geqslant 4$, the coefficient of $(2-\cos\alpha)$ is larger than 1 and the denominator $(d^2-1)^2-d^2$ is positive, thus $-dz_+\geqslant 2-\cos\alpha$. Concerning z_-, since the denominator of $-dz_-$ is positive, we easily get that $-dz_-<2-\cos\alpha$ if and only if $d^2(2-\cos\alpha)<\sqrt{\Delta}$. Since both terms are positive, this is equivalent to

$$d^4(2-\cos\alpha)^2-\Delta<0.$$

Making the change of variable $d=u+4$, and setting $G=(2-\cos\alpha)^2-2$, the left-hand side of the inequality is equal to

$$-G\,u^6-24\,G\,u^5-(237\,G+1)\,u^4-(1232\,G+16)\,u^3-(3553\,G+93)\,u^2-(5384\,G+232)\,u-3344\,G-209.$$

Since $G>0$ for $\alpha\in[\frac{\pi}{2},\pi]$, the above expression is negative for all $u\geqslant 0$, which concludes the proof that $2-\cos\alpha+d\,z_->0$.

$\mathbf{F}_\pi(\mathbf{x}_\pm)<\mathbf{0}$ and $\mathbf{F}_{\frac{3\pi}{4}}(\mathbf{x}_\pm)>\mathbf{0}$. Applying the bounds on arcsin and the signs of $z_\pm$, we can now prove this claim. We first prove that $F_\pi(x_+)<0$. Since $2-\cos\alpha+d\,z_+\leqslant 0$ and $z_+\leqslant 0$, (32) gives

$$\begin{aligned} F_\pi(x_+) &= -\frac{5}{2}\pi+\pi+\arcsin(2-\cos\pi+d\,z_+)+\left(d-\frac{1}{d}\right)(\pi-\pi-\arcsin z_+) \\ &\leqslant -\frac{3}{2}\pi+\Gamma_1(3+d\,z_+)-\left(d-\frac{1}{d}\right)\Gamma_2(z_+). \end{aligned} \tag{34}$$

We replace z_+ by its value which is of the form $\frac{-b+\sqrt{\Delta}}{a}$. As noted before, the denominator of $z_\pm$ is negative for $d\geqslant 4$, and Γ_1 and Γ_2 are odd functions, thus the right hand side of (34) is negative if and only if its numerator is positive. This numerator is of the form $A+B\sqrt{\Delta}$. Substituting $d=u+4$ in A (resp. B), we obtain a polynomial of degree 26 (resp. 20), whose coefficients are all positive (resp. negative). Thus $A>0$ and $B<0$ for $d\geqslant 4$, and $A+B\sqrt{\Delta}>0$ if $A^2-B^2\Delta>0$. We substitute, again, $d=u+4$ in $A^2-B^2\Delta$, which gives a polynomial of degree 52 whose coefficients are all positive. Hence $A^2-B^2\Delta>0$ for $d\geqslant 4$, which concludes the proof that $F_\pi(x_+)<0$.

The proofs for the three other cases are similar. First, since $2-\cos\alpha+dz_->0$ and $z_-\leqslant 0$,

$$F_\pi(x_-)\leqslant-\frac{3}{2}\pi+\Gamma_2(3+d\,z_-)-\left(d-\frac{1}{d}\right)\Gamma_2(z_-).$$

Again, the right-hand side of the inequality is of the form $A+B\sqrt{\Delta}$ over a negative denominator (for $d\geqslant 4$), and we prove similarly as above that both A and B are positive for $d\geqslant 4$. Hence $F_\pi(x_-)<0$ for $d\geqslant 4$. Second,

$$\begin{aligned} F_{\frac{3\pi}{4}}(x_-) &= -\frac{5}{2}\pi+\frac{3\pi}{4}+\arcsin\left(2-\cos\frac{3\pi}{4}+d\,z_-\right)+\left(d-\frac{1}{d}\right)\left(\pi-\frac{3\pi}{4}-\arcsin z_-\right) \\ &\geqslant -\frac{7}{4}\pi+\Gamma_1\left(2+\frac{\sqrt{2}}{2}+d\,z_-\right)+\left(d-\frac{1}{d}\right)\left(\frac{\pi}{4}-\Gamma_1(z_-)\right). \end{aligned}$$

Again, the right-hand side is of the form $A+B\sqrt{\Delta}$ over a negative denominator and, similarly as above, $A<0$ and $B>0$ for $d \geqslant 4$, and substituting $d=u+4$ in $B^2\Delta - A^2$, we get a polynomial of degree 54 whose coefficients are all negative. Hence $F_{\frac{3\pi}{4}}(x_-) > 0$ for $d \geqslant 4$. Third,

$$F_{\frac{3\pi}{4}}(x_+) \geqslant -\frac{7}{4}\pi + \Gamma_2\left(2+\frac{\sqrt{2}}{2}+d\,z_+\right) + \left(d-\frac{1}{d}\right)\left(\frac{\pi}{4}-\Gamma_1(z_+)\right).$$

Here the right-hand side of the inequality is of form $A+B\sqrt{\Delta}$ over a negative denominator, and both A and B are negative for $d \geqslant 4$. Hence $F_{\frac{3\pi}{4}}(x_+) > 0$ for $d \geqslant 4$.

$\mathbf{F_\pi(x) < 0}$ **and** $\mathbf{F_{\frac{3\pi}{4}}(x) > 0}$ **for** $\mathbf{x = x_{min}}$ **and** $\mathbf{x_{max}}$**.** Recall that we already showed that $\mathcal{C}_{L_\pi SL_{\leqslant\pi}}$ is below ℓ and that its rightmost point is equal to the leftmost point of $\mathcal{C}_{L_\pi SR_{\leqslant\pi}}$ This implies that $F_\pi(x_{min}) < 0$.

By (25), $x_{max} = \pi - \alpha + \arcsin\left(\frac{2-2\cos\alpha}{d}\right)$ and so, $2-\cos\alpha + d\,\sin(x_{max}+\alpha) = \cos\alpha$. Thus,

$$F_\alpha(x_{max}) = -2\pi + \alpha - \arccos\cos\alpha + \left(d-\frac{1}{d}\right)x_{max} = -2\pi + \left(d-\frac{1}{d}\right)x_{max}.$$

For $\alpha=\pi$, $x_{max} = \arcsin(4/d) \leqslant \Gamma_2(4/d)$, which yields that $F_\pi(x_{max})$ is smaller or equal than a degree 6 polynomial in d over $3\,d^6$. The change of variable $d=u+4$ gives a polynomial whose coefficients are all negative which implies that $F_\pi(x_{max}) < 0$ for $d \geqslant 4$.

Similarly, for $\alpha = \frac{3\pi}{4}$, $x_{max} = \frac{\pi}{4} + \arcsin\left(\frac{2+\sqrt{2}}{d}\right) \geqslant \frac{\pi}{4} + \Gamma_1\left(\frac{2+\sqrt{2}}{d}\right)$. This yields that $F_{\frac{3\pi}{4}}(x_{max})$ is larger or equal to a degree 7 polynomial in d over $120\,d^6$. As before, the change of variable $d=u+4$ gives a polynomial whose coefficients are all positive; hence $F_{\frac{3\pi}{4}}(x_{max}) > 0$ for $d \geqslant 4$.

Finally, we prove that $F_{\frac{3\pi}{4}}(x_{min}) < 0$. Recall that $x_{min} = \pi - \alpha + \arcsin\left(\frac{1-\cos\alpha}{d}\right)$; thus $2-\cos\alpha + d\,\sin(x_{min}+\alpha) = 1$, and

$$F_\alpha(x_{min}) = -2\pi + \alpha + \left(d-\frac{1}{d}\right)x_{min}.$$

For $\alpha = \frac{3\pi}{4}$, $x_{min} = \frac{\pi}{4} + \arcsin\left(\frac{1+\sqrt{2}/2}{d}\right) \geqslant \frac{\pi}{4} + \Gamma_1\left(\frac{1+\sqrt{2}/2}{d}\right)$. As before, this yields that $F_{\frac{3\pi}{4}}(x_{min})$ is larger or equal to a degree 7 polynomial in d over $3840\,d^6$, and the change of variable $d=u+4$ gives a polynomial whose coefficients are all positive. Hence $F_{\frac{3\pi}{4}}(x_{min}) > 0$ for $d \geqslant 4$, which concludes the proof of this lemma.

References

[1] P. K. Agarwal, T. Biedl, S. Lazard, S. Robbins, S. Suri, and S. Whitesides. Curvature-constrained shortest paths in a convex polygon. *SIAM Journal on Computing*, 31(6):1814–1851, 2002.

[2] P. K. Agarwal, P. Raghavan, and H. Tamaki. Motion planning for a steering-constrained robot through moderate obstacles. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*, pages 343–352, New York, NY, USA, 1995. ACM Press.

[3] J. Backer and D. Kirkpatrick. A complete approximation algorithm for shortest bounded-curvature paths. In *Proceedings of the nineteenth International Symposium on Algorithms and Computation*, pages 628–643, Dec 2008.

[4] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization.* MPS/SIAM Series on Optimization, 2008.

[5] S. Bereg and D. Kirkpatrick. Curvature-bounded traversals of narrow corridors. In *Proceedings of the twenty-first annual symposium on Computational Geometry*, pages 278–287, New York, NY, USA, 2005. ACM Press.

[6] J.-D. Boissonnat, A. Cérézo, and J. Leblond. Shortest paths of bounded curvature in the plane. *Journal of Intelligent and Robotic Systems*, 11(1-2):5–20, 1994.

[7] J-D. Boissonnat and S. Lazard. Convex hulls of bounded curvature. In *Proc. 8th Canad. Conf. Comput. Geom.*, pages 14–19, 1996.

[8] J.-D. Boissonnat and S. Lazard. A polynomial-time algorithm for computing a shortest path of bounded curvature amidst moderate obstacles. *International Journal of Computational Geometry and Applications*, 13(3):189–229, June 2003.

[9] X.-N. Bui. *Planification de trajectoire pour un robot polygonal non-holonome dans un environnement polygonal*. PhD thesis, École des Mines de Paris, 1994.

[10] X.-N. Bui, P. Souères, J.-D. Boissonnat, and J.-P. Laumond. The shortest paths synthesis for non-holonomic robots moving forwards. In *Proceedings of the IEEE International conference on Robotics and Automation*, pages 2–7, San Diego, CA, 1994.

[11] G. Desaulniers. On shortest paths for a car-like robot maneuvering around obstacles. Les Cahiers du GERAD G-94-35, Ecole des Hautes Etudes Commerciales, Montreal, Canada, 1994.

[12] L. E. Dubins. On curves of minimal length with a constraint on average curvature, and prescribed initial and terminal positions and tangents. *American Journal of Mathematics*, 79(3):497–516, 1957.

[13] J. Enright and E. Frazzoli. UAV routing in a stochastic time-varying environment. In *IFAC World Congress*, 2005.

[14] S. Fortune and G. Wilfong. Planning constrained motion. *Annals of Mathematics and Artificial Intelligence*, 3(1):21–82, 1991.

[15] D. Halperin, L. E. Kavraki, and J.-C. Latombe. Robotics. In *CRC Handbook of Discrete and Computational Geometry*, chapter 41, pages 1065–1093. CRC Press, 2004.

[16] S. Itani and M. A. Dahleh. On the stochastic TSP for the Dubins vehicle. In *Proceedings of the 2007 American Control Conference*, New York, USA, 2007.

[17] P. Jacobs and J. Canny. Planning smooth paths for mobile robots. In Z. Li and J.F. Canny, editors, *Nonholonomic Motion Planning*, pages 271–342. Kluwer Academic, 1992.

[18] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.

[19] S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006.

[20] J.-H. Lee, O. Cheong, W.-C. Kwon, S.-Y. Shin, and K.-Y. Chwa. Approximation of curvature-constrained shortest paths through a sequence of points. In *Algorithms - ESA 2000*, pages 314–325, 2000.

[21] X. Ma and D. A. Castañón. Receding horizon planning for Dubins traveling salesman problems. In *45th IEEE Conference on Decision and Control*, Dec 2006.

[22] J. Matoušek and B. Gärtner. *Understanding and Using Linear Programming*. Springer, 2007.

[23] J. Le Ny, E. Feron, and E. Frazzoli. On the curvature-constrained traveling salesman problem. Manuscript. Conditionally accepted for publication in *IEEE Transactions on Automatic Control*, 2009.

[24] J. Le Ny, E. Feron, and E. Frazzoli. The curvature-constrained traveling salesman problem for high point densities. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 5985–5990, 2007.

[25] S. Rathinam, R. Sengupta, and S. Darbha. A resource allocation algorithm for multivehicle systems with nonholonomic constraints. *IEEE Transactions on Automation Science and Engineering*, 4, Jan 2007.

[26] J. A. Reeds and L. A. Shepp. Optimal paths for a car that goes both forwards and backwards. *Pacific J. Math*, 145(2):367–393, 1990.

[27] J. Reif and H. Wang. The complexity of the two dimensional curvature-constrained shortest-path problem. In *Proceedings of the third Workshop on the Algorithmic Foundations of Robotics on Robotics: the Algorithmic Perspective*, pages 49–57, Natick, MA, USA, 1998. A. K. Peters, Ltd.

[28] K. Savla, E. Frazzoli, and F. Bullo. Traveling salesperson problems for the Dubins vehicle. *IEEE Transactions on Automatic Control*, 53, July 2008.

[29] J. T. Schwartz and M. Sharir. Algorithmic motion planning in robotics. In J. van Leeuwen, editor, *Algorithms and Complexity*, volume A of *Handbook of Theoretical Computer Science*, pages 391–430. Elsevier, Amsterdam, 1990.

[30] J. Sellen. Approximation and decision algorithms for curvature-constrained path planning: a state-space approach. In *Proceedings of the third workshop on the algorithmic foundations of robotics on Robotics : the algorithmic perspective*, pages 59–67, Houston, Texas, USA, 1998.

[31] H. J. Sussmann. Shortest 3-dimentional paths with a prescribed curvature bound. In *Proceedings of the 34th IEEE Conference on Decision and Control*, pages 3306–3311, 1995.

[32] H. J. Sussmann and G. Tang. Shortest paths for the reeds-shepp car: A worked out example of the use of geometric techniques in nonlinear optimal control. Report SYCON 91-10, Rutgers University, New Brunswick, NJ, 1991.

[33] H. Wang and P.K. Agarwal. Approximation algorithms for curvature-constrained shortest paths. *SIAM Journal on Computing*, 30(6):1739–1772, 2001.

Centre de recherche INRIA Nancy – Grand Est
LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399