



Bounded-Curvature Shortest Paths through a Sequence of Points using Convex Optimization

Xavier Goaoc, Hyo-Sil Kim, Sylvain Lazard

► **To cite this version:**

Xavier Goaoc, Hyo-Sil Kim, Sylvain Lazard. Bounded-Curvature Shortest Paths through a Sequence of Points using Convex Optimization. SIAM Journal on Computing, Society for Industrial and Applied Mathematics, 2013, 42 (2), pp.662-684. .

HAL Id: hal-00927100

<https://hal.inria.fr/hal-00927100>

Submitted on 10 Jan 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BOUNDED-CURVATURE SHORTEST PATHS THROUGH A SEQUENCE OF POINTS USING CONVEX OPTIMIZATION *

XAVIER GOAOC[†], HYO-SIL KIM[‡], AND SYLVAIN LAZARD[†]

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 is a sub-problem of the Dubins Traveling Salesman Problem and also arises naturally in path planning for point car-like robots in the presence of polygonal obstacles. We show that when consecutive waypoints are distance at least four apart, this question reduces to a family of convex optimization problems over polyhedra in \mathbb{R}^n .

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

AMS subject classification. 68U05

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 [16, 19, 20, 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 constrained to move in the forward direction, and whose 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. The first results on curvature-constrained shortest paths go back to 1957 when Dubins [12] proved that, in the plane without obstacles, bounded-curvature shortest paths consist of arcs of unit radius circles and straight line segments.

We consider the problem of computing a bounded-curvature shortest path that visits, in order, a given sequence p_1, \dots, p_n of n points in the plane (with no obstacles). This problem is 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 (see discussion below). It is also 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 § 8).

Results. We show that a bounded-curvature shortest path through a sequence of n points can be computed by *convex optimization*, provided that any two consecutive points are distance at least 4 apart. More precisely, we show that the problem reduces to a *family* of n -dimensional convex optimization sub-problems, each over a convex polyhedral domain defined by at most $4n$ inequalities. The number of sub-problems is 2^{n-2} in the worst case (Theorem 6.1). This combinatorial complexity is considerable but, in a sense, intrinsic to the problem since there may exist up to 2^{n-2} bounded-curvature *locally* shortest paths visiting the sequence of points (see Figure 5.3). However, the division into sub-problems corresponds to a partition of the space of candidate shortest paths according to some geometric type, and fewer sub-problems can be considered by partially inferring the type of the global solution from the geometry of the point set. In particular, the number of sub-problems can be reduced to 2^k where k is the number of “sharp turns” in the polygonal path $p_1 \dots p_n$ (Theorem 7.2).

The state-of-the-art algorithms for computing bounded-curvature shortest paths through a sequence of points are fast, and often run in time linear in the number of waypoints, but they can only guarantee *multiplicative*-factor approximations of an optimum; the best factor that can currently be achieved is rather large, as it is 1.91 under the same distance assumption as above (see discussion below). Since convex

*This work was supported by the INRIA *Équipe Associée* KI and, in part, by Mid-career Researcher Program through NRF grant (No. R01-2008-000-11607-0) funded by the MEST and, in part, by SRC-GAIA through NRF grant (No. 2011-0030044) funded by the Government of Korea.

[†]Inria, Loria, Villers-lès-Nancy, F-54600, France. `firstname.name@loria.fr`

[‡]Department of Computer Science, KAIST, Daejeon, Korea. `hyosil@tclab.kaist.ac.kr`

optimization methods are known to be efficient in practice and allow for *additive*-factor approximations, our reduction appears to open an interesting alternative. Our results do, however, say little on the theoretical complexity of the problem since most convex optimization methods are notoriously hard to analyze (see § 8).

Bounded curvature path planning problems usually come in two flavors, with *free* or *prescribed* initial and final directions. Even if our presentation focuses on the free-direction setting, our reduction to convex optimization works for both variants. As a consequence, our techniques brings new insight on the problem of path planning in the presence of polygonal obstacles (see § 8). Some of our intermediate results are also interesting in their own rights, in particular our sufficient condition for a path to be a shortest *CSC*-path (Lemma 3.1).

Previous work. The study of bounded-curvature path planning started with Dubins' [12] proof that, in the plane without obstacles, any bounded-curvature shortest path is a concatenation of at most three arcs that are circular arcs of unit radius (*C*-segments) or 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*. 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 impractical, this solution illustrates well the difficulty of the problem.

An approximate solution can easily be obtained by considering the points as given on-line: starting from p_1 by a line segment to p_2 , for $i \geq 2$, the path from p_1 to p_i can be extended by the shortest *CS* path starting from p_i , with the same tangent, and ending in p_{i+1} . 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.¹ 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. [21] 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 [22, 23, 25, 28]; see also [24] for a short survey. It is NP-hard [23] and all proposed approximation algorithms are based on a discretization of the directions at the waypoints. 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 waypoints are randomly distributed, has also been studied; see, e.g., [13, 17, 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 [18] 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

¹Indeed, if p_i and p_{i+1} are at least distance d apart, with $d \geq 4$, they are connected by a path of type *CS* whose length is at most $d + 2\pi - 2 \arctan d$. This method thus builds a path of length at most $n - 1$ times this value, whereas any path has length at least $(n - 1)d$. For $d = 4$, the ratio gives a bound of 1.91.

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

decision problem corresponding to finding a shortest path is NP-hard. Several approximation algorithms were proposed [18, 30, 33]; in particular, Wang and Agarwal [33] presented a $O(\frac{n^2}{\varepsilon} \log n)$ -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). 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 , was presented by Backer and Kirkpatrick [3]; its running time is polynomial in terms of the total bit complexity of the coordinates of the polygon vertices, 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 [20] for a recent overview of such path planning problems.

Proof outline and paper organization. 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 . We consider the function $F : (\mathbb{S}^1)^n \rightarrow \mathbb{R}$ that maps a sequence $(\theta_1, \dots, \theta_n)$ of angles to the length of a shortest curvature-constrained path visiting the configurations $(p_1, \theta_1), \dots, (p_n, \theta_n)$ in order. Dubins' characterization implies that a curvature-constrained shortest path between two configurations can be computed in constant time, so computing a curvature-constrained shortest path visiting the points in order is computationally equivalent to minimizing the function F . Throughout this paper, we assume that consecutive waypoints are distance at least 4 apart, which ensures, in particular, that shortest Dubins paths have type *CSC* between waypoints.

In § 3, we establish that the function F is both twice differentiable and locally strictly convex at any point $(\theta_1, \dots, \theta_n)$ of a domain $\mathcal{L}(\pi)$ such that, in the shortest path visiting the configurations $(p_1, \theta_1), \dots, (p_n, \theta_n)$, all circular arcs between consecutive waypoints have length less than π (Proposition 3.3). We start by considering the case of two points, a situation we handle by using analytical arguments and a new sufficient condition for the optimality of Dubins paths of a given type (Lemma 3.1). The local convexity for multiple waypoints easily follows from the case of two points.

In § 4, we analyze the geometry of *locally* shortest paths (Lemma 4.1) to show that a *globally* shortest path never has arcs of length greater than $\frac{3\pi}{4}$ between waypoints (Proposition 4.2). To keep the presentation simple we give a proof under the assumption that consecutive waypoints are distance at least $2 + 2\sqrt{2}$ apart and outline in § 6 how this condition can be relaxed to distance 4. It follows that F is locally convex on an open domain $\mathcal{L}(\frac{3\pi}{4})$ containing all its global minima. The connected components of this domain are, however, not convex, and they may thus contain several local minima.

To overcome this issue, we construct, in § 5, a “nice” region \mathcal{D} contained in the domain $\mathcal{L}(\pi)$ of local convexity of F , and containing the domain $\mathcal{L}(\frac{3\pi}{4})$ and thus all the global minima of F (Corollary 5.2). The connected components of \mathcal{D} are, once lifted to \mathbb{R}^n , convex polyhedra defined by $4n$ inequalities each (Lemma 5.3). Again, to keep the presentation simple we first construct \mathcal{D} assuming that consecutive waypoints are distance at least 8.6 apart and outline in § 6 how to relax that distance condition. Hence, once lifted to \mathbb{R}^n , F is convex over every connected component of \mathcal{D} and thus its global minimum can be found by convex minimization over each of these components (Proposition 5.4). Unfortunately, there might still be $\Theta(2^n)$ such components.

In § 7, we show that the number of components to be considered is at most exponential in the number of sharp turns (suitably defined) on the polygonal path $p_1 \dots p_n$ (Theorem 7.2). Specifically, we argue that, unless the polygonal path turns sharply at p_i , the polar angle at p_i of any globally shortest path must lie in a restricted range as otherwise the part of the path between p_{i-1} and p_{i+1} would have a self-intersection and therefore admit a global shortcut (Lemma 7.1).

2. Notation and preliminaries. Let p_1, \dots, p_n be a sequence of points in the plane. A *configuration* is defined as a pair $(p, \theta) \in \mathbb{R}^2 \times \mathbb{S}^1$ where $\mathbb{S}^1 = \mathbb{R}/2\pi\mathbb{Z}$ denotes 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} . In particular, when we discuss the local convexity of functions taking angles as input, we implicitly assume that these angles are seen in \mathbb{R} , that is the local convexity of a function $\phi : (\mathbb{S}^1)^n \rightarrow \mathbb{R}$ refers to the local convexity of

$\phi \circ \tau$, where τ is the quotient map from \mathbb{R}^n to $(\mathbb{S}^1)^n$.

We say that a sequence of points p_1, \dots, p_n satisfies the (D_d) condition if every two consecutive points p_i and p_{i+1} are at least distance d apart. Throughout the paper we will consider that at least the (D_4) condition holds so that, in particular, a shortest Dubins path between any two configurations (p_i, θ_i) and (p_{i+1}, θ_{i+1}) is of type *CSC* [10, §4.3]³; furthermore, it is straightforward to bound the length of its line segment as follows.

LEMMA 2.1. *If p_i and p_{i+1} are at distance $d \geq 4$, then the shortest Dubins path between any configurations (p_i, θ_i) and (p_{i+1}, θ_{i+1}) has type *CSC* and the line segment has length at least $\sqrt{(d-2)^2 - 4}$.*

We sometimes need to specify whether a Dubins path is turning right (clockwise) or left (counterclockwise) on a circle; we then refer to paths of types *LSR*, *RSL*, *LSL*, and *RSR*. 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 arc turning right.


3. Local convexity of the length function. In this section we show that the problem can be recast as the minimization of a function from \mathbb{R}^n to \mathbb{R} that is C^2 and locally convex over a certain subset of \mathbb{R}^n .

The parameterization. Let $F : (\mathbb{S}^1)^n \rightarrow \mathbb{R}$ map a sequence $(\theta_1, \dots, \theta_n)$ of angles to the length of a shortest curvature-constrained path visiting the configurations $(p_1, \theta_1), \dots, (p_n, \theta_n)$ 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 .

The function F breaks down into

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

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. All functions F_i^{i+1} behave similarly so we first focus on F_1^2 and define, for any given path type $T \in \{LSR, RSL, LSL, RSR\}$, the function $F_T(\theta_1, \theta_2)$ that denotes the length of the path of type T (T -path) from (p_1, θ_1) to (p_2, θ_2) . Each function 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. The partial derivatives and Hessian of F_T at such a point are



$$\frac{\partial F_T(\theta_1, \theta_2)}{\partial \theta_i} = \mu_{i=1} \mu_{C_i=R} (1 - \cos \alpha_i) \quad (3.1)$$

$$\frac{\partial^2 F_T}{\partial \theta_i \partial \theta_j} = \delta_{i,j} \sin \alpha_i + \frac{\sin \alpha_i \sin \alpha_j}{s} \quad (3.2)$$

$$\det H(F_T) = \sin \alpha_1 \sin \alpha_2 \left(1 + \frac{\sin \alpha_1 + \sin \alpha_2}{s} \right) \quad (3.3)$$

where α_i denotes the length of the i -th circular arc of the *CSC*-path from p_1 to p_2 , s stands for the length of the line segment, μ_B equals 1 if B is true and -1 otherwise, and $\delta_{i,j}$ equals 1 if $i = j$ and 0 otherwise. These proofs are nontrivial, long and technical although they use only elementary calculus and geometry. As they shed no insight on the problem, we omit them here and refer the interested reader to the technical report [15, Proposition 3]. The navigation between the functions F_T and F_1^2 is made possible by the next lemma.

LEMMA 3.1. *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. 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

³Although they do not state it explicitly, Bui et al. [10] show that if a shortest Dubins path is of type *CCC* between p_s and p_f , 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.




Fig. 3.1: (a) Lemon $L_1^2(\pi)$ for $|p_1 p_2| = 4$ (where ν is the polar angle of $\overrightarrow{p_2 p_1}$) and (b) graph of F_1^2 over that domain.

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. Again, since these proofs are not particularly illuminating, we refer the interested reader to the technical report [15, Proposition 5]. \square

Local convexity of F_i^{i+1} . Assume that condition (D_4) holds. For $\alpha \in (0, \pi]$ we 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 2.1 ensures that the shortest path is of type CSC and Lemma 3.1 guarantees it is unique (since $\alpha \leq \pi$). Note that $L_i^{i+1}(\alpha)$ is an open set, since the circular arcs must be *strictly* shorter than α . We call $L_i^{i+1}(\alpha)$ a *lemon* region due to its evocative shape (see Figure 3.1).

LEMMA 3.2. *If the (D_4) condition holds, the length function F_i^{i+1} is C^2 and locally strictly convex over $L_i^{i+1}(\pi)$.*

Proof. Without loss of generality it suffices to consider F_1^2 . As mentioned above, the (D_4) condition and the fact that $(\theta_1, \theta_2) \in L_1^2(\pi)$ ensure that there is a geometrically unique shortest path from (p_1, θ_1) to (p_2, θ_2) and it is of type CSC . Let α_1 and α_2 denote the lengths of its circular arcs.

First, assume that neither α_1 nor α_2 vanishes. Then the type $T \in \{LSR, RSL, LSL, RSR\}$ of the shortest path is uniquely defined. Moreover, the length of a circular arc of the T -path from (p_1, θ_1) to (p_2, θ_2) varies continuously in (θ_1, θ_2) provided it remains in $(0, 2\pi)$. Thus, in a neighborhood of (θ_1, θ_2) the lengths of the circular arcs of the T -path remain in $(0, \pi)$ and by Lemma 3.1 F_1^2 and F_T coincide locally. Finally, since $\alpha_1, \alpha_2 \in (0, \pi)$, Equations (3.1) and (3.3) yield that $\frac{\partial^2 F_T}{\partial \theta_1^2}$ 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 F_T and F_1^2 are locally strictly convex at (θ_1, θ_2) .

If at least one α_i vanishes, at least two T and T' -paths coincide ($T \neq T'$ in $\{LSR, \dots\}$), and the expressions of the derivatives of F_T yield, by continuity, that F_1^2 is locally C^2 and thus locally convex at this point. We prove the *strict* local convexity at such a point by showing, using Taylor expansions of F_T , that the graph of the function is locally strictly above its tangent plane. Specifically, if only α_1 vanishes then the Taylor expansion at order 2 suffices, except for $\theta_2 = 0$ in which case the expansion of order 3 yields the result. The case where only α_2 vanishes is symmetric. The situation where both α_1 and α_2 vanish occur for exactly one point of $L_1^2(\pi)$, so the strict local convexity follows by continuity from the other cases. These computations are straightforward though tedious and we refer to the technical report [15, Theorem 6] for

more details. \square

Local convexity of F . We still assume that condition (D_4) holds and we now define the n -dimensional lemon region $\mathcal{L}(\alpha) \subset (\mathbb{S}^1)^n$ as the set of tuples $(\theta_1, \dots, \theta_n)$ such that the shortest path visiting the configurations $(p_1, \theta_1), \dots, (p_n, \theta_n)$ 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 $(p_1, \theta_1), \dots, (p_n, \theta_n)$ is the concatenation of the shortest paths from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) for $i = 1, \dots, n-1$. This ensures, with Lemmas 2.1 and 3.1, that $\mathcal{L}(\alpha)$ is well-defined for any $\alpha \in (0, \pi]$. This also implies that a point $(\theta_1, \dots, \theta_n)$ is in $\mathcal{L}(\alpha)$ if and only if, for $i = 1, \dots, n-1$, 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 $(\theta_i, \theta_{i+1}) \in L_i^{\alpha}(\alpha)$. This rewrites as

$$\mathcal{L}(\alpha) = \bigcap_{i=1}^{n-1} (\mathbb{S}^1)^{i-1} \times L_i^{\alpha}(\alpha) \times (\mathbb{S}^1)^{n-i-1}, \quad (3.4)$$

with the convention that $(\mathbb{S}^1)^0 \times A = A \times (\mathbb{S}^1)^0 = A$.

PROPOSITION 3.3. *If the (D_4) condition holds, the length function $F(\theta_1, \dots, \theta_n)$ is C^2 and locally strictly convex over $\mathcal{L}(\pi)$.*

Proof. By Lemma 3.2, the function $F = \sum_{i=1}^{n-1} F_i^{\alpha}(\theta_i, \theta_{i+1})$ is C^2 over $\mathcal{L}(\pi)$. Since a sum of locally convex functions is locally convex, F is locally convex over $\mathcal{L}(\pi)$. The functions F_i^{α} are not, however, strictly convex as functions of $(\theta_1, \dots, \theta_n)$. To prove the strict local convexity of F , we consider its Hessian, well defined over $\mathcal{L}(\pi)$:

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

For any $\Theta = (\theta_1, \dots, \theta_n)$,

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

For any Θ in $\mathcal{L}(\pi)$, (θ_i, θ_{i+1}) belongs to the lemon $L_i^{\alpha}(\pi)$, for all $i = 1, \dots, n-1$. Lemma 3.2 implies that every term of the above sum is strictly positive, hence H is positive definite over $\mathcal{L}(\pi)$. \square

4. The domain of local convexity contains all global minima. We now prove that the domain over which we know the length function to be locally convex contains all its global minima. Let γ be a shortest bounded-curvature path through a sequence of configurations $(p_1, \theta_1), \dots, (p_n, \theta_n)$. We assume that the (D_4) condition holds so that γ is of type *CSC* between any two consecutive configurations (by Lemma 2.1). We say that the path γ is *locally shortest* if it cannot be shortened by small deformations (while going through p_1, \dots, p_n and remaining of bounded curvature), that is if $(\theta_1, \dots, \theta_n)$ is a local minimum of F .

LEMMA 4.1. *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 .*⁴

⁴A near reciprocal of this statement actually holds; see [15, Proposition 9].

Proof. Since F is the minimum of several length functions, each associated with a different path type (e.g. $LSR - RSR - LSR\dots$), it is difficult to determine where F is differentiable; we only know that F is differentiable over $\mathcal{L}(\pi)$, by Lemma 3.2. We thus consider, in the proof, the length function associated with the path type of γ , instead of F . For clarity, we assume that γ visits the configurations (p_i, θ_i) , that is we distinguish between the variables θ_i and the value $(\tilde{\theta}_1, \dots, \tilde{\theta}_n)$ of a local minimum of F .

Recall 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 $i = 1, 2$ Equation (3.1) states:

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

This first implies Statement (i). 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 $F_T(\theta_i)$ 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). If the circular arcs of the subpath of γ do not vanish, then $F_T(\theta_i)$ is differentiable at $\tilde{\theta}_i$, and

$$F'_T(\tilde{\theta}_i) = -\mu_{C_i^-=R} (1 - \cos \alpha_i^-) + \mu_{C_i^+=R} (1 - \cos \alpha_i^+),$$

where C_i^- and C_i^+ denote the circular arcs preceding and following p_i , respectively, and α_i^- and α_i^+ denote their lengths.

Since $\tilde{\theta}_i$ is a local minimum of F_T , either $F'_T(\tilde{\theta}_i) = 0$ or F_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_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, $F'_T(\tilde{\theta}_i) = 0$ in all cases.

Now, if the orientations (R or L) of the two circular arcs C_i^- and C_i^+ differ, $F'_T(\tilde{\theta}_i) = \mu_{C_i^+=R} (2 - \cos \alpha_i^- - \cos \alpha_i^+)$ 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 (ii). It follows that $F'_T(\theta_i) = \mu_{C_i^+=R} (\cos \alpha_i^- - \cos \alpha_i^+)$, 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 $0 \leq \alpha_i^\pm < 2\pi$, which proves Statement (iii). \square

We now prove that all the global minima of the length function F belong not only to $\mathcal{L}(\pi)$, but to the smaller lemon $\mathcal{L}(\frac{3\pi}{4})$. This sharper result will be useful later.

PROPOSITION 4.2. *Under condition $(D_{2+2\sqrt{2}})$, in any globally shortest path γ the circular arcs preceding and following each waypoint have length less than $\frac{3\pi}{4}$.*

Proof. As before, let C_i^- and C_i^+ denote the circular arcs of γ that precede and follow p_i , respectively, and α_i^- and α_i^+ be their lengths. By Lemma 4.1, since γ is a locally shortest path, α_i^- and α_i^+ are equal or sum up to 2π . By Lemma 2.1, $(D_{2+2\sqrt{2}})$ implies that the line segments preceding and following p_i have length at least 2. In both cases, the path can be trivially shortened, as illustrated in Figures 4.1(a) and 4.1(b), by using new circular arcs and a new direction of travel at p_i (shown in dashed in the figure). This contradicts the global optimality of γ . \square




Fig. 4.1: Global shortcuts (in dashed) of locally shortest paths, for the proof of Proposition 4.2.

5. Reduction to a family of convex optimization sub-problems. With Propositions 3.3 and 4.2, one could hope to use convex optimization methods to find the minimum of the length function F . It is, however, clear from the example of Figure 3.1(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 $\mathcal{L}(\pi)$. In this section, we describe 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 an adequate lifting $(\mathbb{S}^1)^n \rightarrow \mathbb{R}^n$ maps each connected component of \mathcal{D} to a convex polyhedron, suitable for convex optimization.

Two-dimensional diamonds. Consider two consecutive points p_i and p_{i+1} , and denote by ν_{i+1}^i the polar angle of vector $\overrightarrow{p_{i+1}p_i}$. We define D_i^{i+1} as the image of the open quadrilateral with vertices $(0, 2\pi)$, $(\frac{\pi}{4}, \frac{\pi}{4})$, $(2\pi, 0)$, and $(\frac{7\pi}{4}, \frac{7\pi}{4})$ under the translation of vector $(\nu_{i+1}^i, \nu_{i+1}^i)$ (see Figure 5.1). Two-dimensional diamonds have the following property.

LEMMA 5.1. *If $|p_i p_{i+1}| \geq 8.6$ then $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$.*

Proof. The proof is done in four steps. First, we show that we can assume without loss of generality that $\nu_{i+1}^i = 0$ and thus that $L_i^{i+1}(\alpha)$ and D_i^{i+1} can be lifted to the open square $(0, 2\pi)^2$ of \mathbb{R}^2 without intersecting the square boundary. We thus refer to Figure 5.1 with $\nu_{i+1}^i = 0$. Second, we show that $L_i^{i+1}(\alpha)$ and D_i^{i+1} are symmetric with respect to the two diagonals of the square, which implies that it is enough to prove the inclusion in only one quadrant. Third, we give an analytical expression of the boundary of $L_i^{i+1}(\alpha)$. Fourth, using these analytical expressions, we prove the inclusions $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$ in the considered quadrant.

Step 1. 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 $\nu_{i+1}^i = 0$. In any CSC-path from $(p_i, 0)$ 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 .

The same property holds for paths from (p_i, θ_i) to $(p_{i+1}, 0)$, thus for any $\alpha \leq \pi$ the lemon $L_i^{i+1}(\alpha)$ does not intersect the two hyperplanes $\theta_i = 0$ and $\theta_{i+1} = 0$. Furthermore, D_i^{i+1} does not meet these hyperplanes




Fig. 5.1: Diamond D_i^{i+1} (solid), contained in $L_i^{i+1}(\pi)$ and containing $L_i^{i+1}(\frac{3\pi}{4})$ (both dashed) for $|p_i p_{i+1}| = 8.6$.

either, by construction. It follows that we can lift $L_i^{i+1}(\alpha)$ and D_i^{i+1} from $(\mathbb{S}^1)^2$ onto the *open* square $(0, 2\pi)^2$ of \mathbb{R}^2 .

Step 2. We first prove that $L_i^{i+1}(\alpha)$ is symmetric with respect to point (π, π) and to the lines $y = x$ and $y = 2\pi - x$. The first symmetry is straightforward from the observation that shortest *CSC*-paths from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) and from $(p_i, 2\pi - \theta_i)$ to $(p_{i+1}, 2\pi - \theta_{i+1})$ are symmetric with respect to the x -axis. The second symmetry of $L_i^{i+1}(\alpha)$ is similar though more subtle: it follows from the observation that shortest *CSC*-paths from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) and from (p_i, θ_{i+1}) to (p_{i+1}, θ_i) can be transformed from one another by reversing path and then taking its mirror image with respect to, first, the bisecting line of p_i and p_{i+1} and, second, the x -axis. Finally, the last symmetry of $L_i^{i+1}(\alpha)$ is simply a composition of the two others. It is clear that D_i^{i+1} exhibits the same symmetries, so it suffices to prove the inclusions in the quadrant $Q = \{0 \leq x \leq y \leq 2\pi - x\}$.

Step 3. The proof of the analytical description of the boundary of $L_i^{i+1}(\alpha)$ is quite intricate and goes in three steps, as follows. The *path corresponding to* (θ_i, θ_{i+1}) refers, for simplicity, to the (or any) shortest *CSC*-path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) . 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. Instead of trying to prove the reverse inclusions directly, we argue in the third step, that the union of the arcs forms a simple closed curve 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.

Step 3a. Suppose for a contradiction that there is a point (θ_i, θ_{i+1}) on the boundary of $L_i^{i+1}(\alpha)$ such that a shortest *CSC* path from (p_i, θ_i) to (p_{i+1}, θ_{i+1}) has both its circular arcs strictly shorter than $\alpha \leq \pi$. In a neighborhood of (θ_i, θ_{i+1}) , this *CSC*-path changes continuously while both its circular arcs remains strictly shorter than $\alpha \leq \pi$. By Lemma 3.1, these paths are shortest *CSC*-paths and thus, a neighborhood of (θ_i, θ_{i+1}) is included in $L_i^{i+1}(\alpha)$, contradicting the assumption.

Step 3b. We show that if the path corresponding to $(\theta_i, \theta_{i+1}) \in (0, 2\pi)^2$ has type $L_\alpha SL_\beta$ or $L_\alpha SR_\beta$, with $\beta \leq \alpha$, then (θ_i, θ_{i+1}) belongs to the curves of equations

$$\begin{aligned} \mathcal{C}_{L_\alpha SL_{\leq \alpha}} : \quad & \theta_{i+1} = \theta_i + \alpha + \arccos(\cos \alpha - d \sin(\theta_i + \alpha)) \\ & \text{for } \pi - \alpha \leq \theta_i \leq \pi - \alpha + \arcsin\left(\frac{1 - \cos \alpha}{d}\right), \\ \mathcal{C}_{L_\alpha SR_{\leq \alpha}} : \quad & \theta_{i+1} = \theta_i + \alpha - \arccos(2 - \cos \alpha + d \sin(\theta_i + \alpha)) \\ & \text{for } \pi - \alpha + \arcsin\left(\frac{1 - \cos \alpha}{d}\right) \leq \theta_i \leq \pi - \alpha + \arcsin\left(\frac{2 - 2 \cos \alpha}{d}\right). \end{aligned}$$

The key idea is to show that for paths of type $L_\alpha SL_\beta$, $\theta_{i+1} = \theta_i + \alpha + \beta$ and $\cos \beta = \cos \alpha - d \sin(\theta_i + \alpha)$; plugging the latter equation into the former, while carefully monitoring modulo effects, yields the equation of $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ (see [15, Section B2] for details). Paths of type $L_\alpha SR_{\leq \alpha}$ are handled similarly and lead to the equation of $\mathcal{C}_{L_\alpha SR_{\leq \alpha}}$.

The curves $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leq \alpha}}$ meet in one of their endpoints, at $\theta_i = \pi - \alpha + \arcsin\left(\frac{1 - \cos \alpha}{d}\right)$ and $\theta_{i+1} = \pi + \arcsin\left(\frac{1 - \cos \alpha}{d}\right)$. Their union τ is thus a connected curve, and that curve is simple since $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leq \alpha}}$ are graphs of functions whose domains share only one point. The endpoints of τ are $(\pi - \alpha, \pi + \alpha)$ and (t, t) with $t = \pi - \alpha + \arcsin\left(\frac{2 - 2 \cos \alpha}{d}\right)$, which belong to, respectively, the lines $\theta_{i+1} = 2\pi - \theta_i$ and $\theta_{i+1} = \theta_i$. The slope of τ is everywhere less than -1 so τ lies in the quadrant Q ; a straightforward calculation⁵ actually shows that this slope is everywhere less than -7 , a fact used in Step 4. Altogether, the

⁵The derivative of the function of $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ minus -7 is $8 + \frac{d \cos(\theta_i + \alpha)}{\sqrt{1 - (\cos(\alpha) - d \sin(\theta_i + \alpha))^2}}$. Noting that the term in the square root is non-negative and that $\cos(\theta_i + \alpha) \leq 0$ (since $\pi \leq \theta_i + \alpha \leq \frac{3\pi}{2}$), the expression is negative if and only if $64 \left(1 - (\cos(\alpha) + d \sin(\theta_i + \alpha))^2\right) - d^2 \cos^2(\theta_i + \alpha) < 0$. This is a degree-two polynomial in $\sin(\theta_i + \alpha)$ whose leading coefficient is $-63d^2 < 0$ and discriminant is $4d^2(-63d^2 + 4032 + 64 \cos^2 \alpha)$, which is negative for any $d \geq 8.6 > \sqrt{\frac{4032 + 64}{63}} \geq \sqrt{\frac{4032 + 64 \cos^2 \alpha}{63}}$. Hence, for any θ_i and $d \geq 8.6$, the slope of $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ is less than that of $\theta_{i+1} = 2\pi - \theta_i$. The calculation is similar for $\mathcal{C}_{L_\alpha SR_{\leq \alpha}}$.

union of $\mathcal{C}_{L_\alpha SL_{\leq \alpha}}$ and $\mathcal{C}_{L_\alpha SR_{\leq \alpha}}$ is thus a simple curve τ contained in the quadrant Q and with endpoints on the two diagonals of the square $[0, 2\pi]^2$.

Step 3c. Any boundary point of $L_i^{i+1}(\alpha)$ with corresponding path of type $L_\alpha SL_{\leq \alpha}$ or $L_\alpha SR_{\leq \alpha}$ must belong to τ . Step 2 then implies that any boundary point of $L_i^{i+1}(\alpha)$ must belong to τ or one of its image under the central symmetry with respect to (π, π) , the reflection with respect to the line $\theta_{i+1} = \theta_i$, and their composition. Let σ denote the union of τ and its images under these transformations. Since τ is a simple arc of curve that lies in the quadrant Q with its endpoints on the lines $\theta_{i+1} = \theta_i$ and $\theta_{i+1} = 2\pi - \theta_i$, σ is a simple closed curve in $[0, 2\pi]^2$.

Let $\hat{\sigma}$ denote the projection of σ on $\mathbb{S}^1 \times \mathbb{S}^1$. If $\alpha < \pi$, σ lies in the open square $(0, 2\pi)^2$ and $\hat{\sigma}$ is thus a simple closed curve. If $\alpha = \pi$, the two points $(0, 2\pi)$ and $(2\pi, 0)$ of σ are mapped to the same point of $\hat{\sigma}$, and $\hat{\sigma}$ then consists of two simple loops meeting in exactly the point $(0, 0)$. When $\alpha \leq \pi$, the boundary of $L_i^{i+1}(\alpha)$ cannot be a proper subset of a simple closed curve since both that set and its complement have interior points. It follows that when $\alpha < \pi$, the boundary of $L_i^{i+1}(\alpha)$ is the whole curve $\hat{\sigma}$; the same holds when $\alpha = \pi$ as, in that case, both loops forming $\hat{\sigma}$ contain some point of the boundary of $L_i^{i+1}(\pi)$.

Step 4. The segment that bounds D_i^{i+1} in the quadrant Q lies on the line $\theta_{i+1} = 2\pi - 7\theta_i$ with θ_i ranging from 0 to $\frac{\pi}{4}$. We have shown that the boundary of $L_i^{i+1}(\alpha)$ in Q is $\mathcal{C}_{L_\pi SL_{\leq \pi}} \cup \mathcal{C}_{L_\pi SR_{\leq \pi}}$ which slope is less than -7 everywhere. Since the leftmost point of $\mathcal{C}_{L_\pi SL_{\leq \pi}} \cup \mathcal{C}_{L_\pi SR_{\leq \pi}}$ is $(0, 2\pi)$ which lies on the line $\theta_{i+1} = 2\pi - 7\theta_i$, the boundary of $L_i^{i+1}(\pi)$ is strictly below that of D_i^{i+1} , in Q , except for their endpoint $(0, 2\pi)$. On the other hand, a simple calculation also shows that the rightmost point of $\mathcal{C}_{L_{\frac{3\pi}{4}} SL_{\leq \frac{3\pi}{4}}} \cup \mathcal{C}_{L_{\frac{3\pi}{4}} SR_{\leq \frac{3\pi}{4}}}$ is strictly above the line $\theta_{i+1} = 2\pi - 7\theta_i$, which implies that the boundary of $L_i^{i+1}(\frac{3\pi}{4})$ is strictly above that of D_i^{i+1} , in Q , and concludes the proof. \square

n-dimensional diamond \mathcal{D} . We extend the construction of the two-dimensional diamond D_i^{i+1} to an arbitrary number of points by defining the *diamond of p_1, \dots, p_n* as

$$\mathcal{D} = \bigcap_{i=1}^{n-1} (\mathbb{S}^1)^{i-1} \times D_i^{i+1} \times (\mathbb{S}^1)^{n-i-1}, \quad (5.1)$$

with the convention that $(\mathbb{S}^1)^0 \times A = A \times (\mathbb{S}^1)^0 = A$. The similarity with Equation (3.4),

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

and the inclusions $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$ from Lemma 5.1 yield:

COROLLARY 5.2. *If p_1, \dots, p_n satisfy the $(D_{8,6})$ condition then $\mathcal{L}(\frac{3\pi}{4}) \subset \mathcal{D} \subset \mathcal{L}(\pi)$.*

Lifting \mathcal{D} to convex polyhedra. Recall that ν_{i+1}^i denotes the polar angles of $\overline{p_{i+1}p_i}$. For $i = 1, \dots, n-1$, 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

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

We lift \mathcal{D} to \mathbb{R}^n through the lifting $(\mathbb{S}^1)^n \rightarrow \Lambda$. See Figure 5.2 for an illustration and note that, through this lifting, the image of one diamond D_i^{i+1} (or, more formally, of $(\mathbb{S}^1)^{i-1} \times D_i^{i+1} \times (\mathbb{S}^1)^{n-i-1}$) is in general not connected. For $1 < i < n$, Λ_i contains one point from $\nu_{i+1}^{i-1} + 2\pi\mathbb{Z}$ 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 $\Lambda_n^+ = \Lambda_n^- = \Lambda_n$. Finally, consider the family of hyperplanes in $(\mathbb{S}^1)^n$

$$\mathcal{H} = \{\theta_i = \nu_{i+1}^i \mid i = 1, \dots, n-1\} \cup \{\theta_i = \nu_i^{i-1} \mid i = 2, \dots, n\}.$$

LEMMA 5.3. *\mathcal{D} does not intersect the hyperplanes of \mathcal{H} , and the lifting $(\mathbb{S}^1)^n \rightarrow \Lambda$ 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 $4n$ linear inequalities in \mathbb{R}^n , each involving 2 variables.*




Fig. 5.2: Illustration of the images of D_{i-1}^i and D_{i+1}^{i+1} through the lifting $(\mathbb{S}^1)^n \rightarrow \Lambda$. The hyperplanes of \mathcal{H} that are not on the boundary of Λ (through the lifting) are shown in dashed.

Proof. Since D_{i+1}^{i+1} intersects none of the lines $\theta_i = \nu_{i+1}^i$ and $\theta_{i+1} = \nu_{i+1}^i$, in $(\mathbb{S}^1)^2$, the diamond \mathcal{D} intersects none of the hyperplanes of \mathcal{H} .

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 $\Xi = \prod_{1 \leq i \leq n} \Lambda_i^{\varepsilon_i}$ where $\varepsilon_i \in \{-, +\}$. By definition of D_{i+1}^{i+1} , its pre-image under the quotient map $\mathbb{R}^2 \rightarrow (\mathbb{S}^1)^2$ is a family of disjoint convex quadrilaterals, one in each open square defined by the lines $\theta_i = \nu_{i+1}^i + 2\pi\mathbb{Z}$ and $\theta_{i+1} = \nu_{i+1}^i + 2\pi\mathbb{Z}$, in \mathbb{R}^2 . The interior of the rectangle $\Lambda_i^{\varepsilon_i} \times \Lambda_{i+1}^{\varepsilon_{i+1}}$ does not intersect those lines and therefore meets at most one of the convex quadrilaterals of the pre-image. It follows that the image of D_{i+1}^{i+1} through the partial lift $(\mathbb{S}^1)^2 \rightarrow \Lambda_i^{\varepsilon_i} \times \Lambda_{i+1}^{\varepsilon_{i+1}}$ is a convex polygon defined by at most 8 inequalities: the 4 defining the quadrilateral and the 4 defining the boundary of the rectangle.

The image of \mathcal{D} through the (partial) lift $(\mathbb{S}^1)^n \rightarrow \Xi$ is thus a convex polyhedron defined by at most $8(n-1)$ inequalities, in one or two variables each. However, $4(n-1)$ of these inequalities are sufficient: since the closure of \mathcal{D} intersects no hyperplane of \mathcal{H} in a face of dimension more than $n-2$, we can drop, for each D_{i+1}^{i+1} , the 4 inequalities defining the boundary of $\Lambda_i^{\varepsilon_i} \times \Lambda_{i+1}^{\varepsilon_{i+1}}$. \square

Convex optimization sub-problems. We can now prove that computing a global shortest path through a sequence of points reduces to solving (possibly exponentially many) convex optimizations sub-problems.

PROPOSITION 5.4. *Let p_1, \dots, p_n be a sequence of points in the plane such that any two consecutive points are at least distance 8.6 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^{n-2} disjoint convex polyhedra in \mathbb{R}^n , each defined by at most $4n$ linear inequalities, in two variables each. Moreover, through this lifting, F is strictly convex over each of these polyhedra.*

Proof. The global minimum of F is realized over \mathcal{D} as it is realized over $\mathcal{L}(\frac{3\pi}{4})$ (by Proposition 4.2) and $\mathcal{L}(\frac{3\pi}{4}) \subset \mathcal{D}$ (by Corollary 5.2). For each cell c in $(\mathbb{S}^1)^n \setminus \mathcal{H}$, the lifting $(\mathbb{S}^1)^n \rightarrow \Lambda$ maps $\mathcal{D} \cap c$ to a convex polyhedron in \mathbb{R}^n defined by $4n$ linear inequalities (by Lemma 5.3). Moreover, the length function F is convex on each such polyhedron since it is convex on $\mathcal{L}(\pi)$ (by Proposition 3.3) and $\mathcal{D} \subset \mathcal{L}(\pi)$ (by Corollary 5.2). Since $(\mathbb{S}^1)^n \setminus \mathcal{H}$ has at most 2^{n-2} cells, the statement follows. \square

Note that the length function F may actually have up to 2^{n-2} local minima as the example of Figure 5.3 illustrates. We refer to the technical report [15, Corollary 10] for a proof. Since distinct local minima necessarily belong to distinct connected components of \mathcal{D} , the constant 2^{n-2} in Proposition 5.4 is best possible. The next two sections improve Proposition 5.4 by relaxing the distance condition from 8.6 to 4




Fig. 5.3: A sequence of n points $\{(2i, (-1)^i d) \mid i = 1, \dots, n\}$ such that, for d large enough, the sequences of polar angles $(\theta_2, \dots, \theta_{n-1})$ in $\{0, \pi\}^{n-2}$ define 2^{n-2} paths that are arbitrarily close to distinct local optima (two of which are shown).




Fig. 6.1: (Left and center) The two elementary transformations used to prove Proposition 4.2 under the (D_4) condition. (Right) A combination of the two transformations that shortens the path through p_i .


(§ 6) and by narrowing down the set of sub-problems sufficient to consider (§ 7).

6. Relaxing the distance conditions. We now review and sharpen our chain of arguments so that our reduction to convex optimization holds under the (D_4) condition.

All the arguments regarding the local convexity of the length function over $\mathcal{L}(\pi)$ holds under the (D_4) condition (in fact, the (D_4) condition is only used there to evacuate the need to consider *CCC*-paths). There is thus nothing to add to § 3.

In § 4, the fact that any global minimum of F lies in $\mathcal{L}(\frac{3\pi}{4})$ (Proposition 4.2) follows from the property that, if the length of a circular arc preceding or following p_i is more than $\frac{3\pi}{4}$, then the path can be globally shortened between p_{i-1} and p_{i+1} . The $(D_{2+2\sqrt{2}})$ condition allowed for a simple proof of this property because the segments preceding and following p_i have length at least 2. Under the (D_4) condition, this property can still be proven by applying adequate local transformations (illustrated in Figure 6.1). A more intricate case analysis is needed and we refer the interested reader to the technical report [15, Lemma 11] for more details.

In § 5, the construction of the “nice” domain \mathcal{D} , the $(D_{8.6})$ condition is used in the proof of the inclusions $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$ (Lemma 5.1). It turns out that these inclusions may fail under a weaker distance condition. We address this issue by making the diamond D_i^{i+1} sensitive




to the distance between the points. Specifically, with $d_i = |p_i p_{i+1}|$ and $\xi_i = 2\pi/(d_i - 1/d_i)$, 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 $(\nu_{i+1}^i, \nu_{i+1}^i)$. The inclusions $L_i^{i+1}(\frac{3\pi}{4}) \subset D_i^{i+1} \subset L_i^{i+1}(\pi)$ under the (D_4) condition, for the new D_i^{i+1} , are then proved in a similar way as in Lemma 5.1: we compare the analytical expressions of the boundaries of $L_i^{i+1}(\alpha)$ and D_i^{i+1} (in a single quadrant of the domain since the symmetries of $L_i^{i+1}(\alpha)$ and D_i^{i+1} remain unchanged). The arguments are standard though long and non-trivial, as we compute the sign of functions at their local extrema by bounding their values using partial sums of power series and Descartes' rule of signs. We refer the interested reader to the technical report [15, Lemma 15] for details. The n -dimensional diamond is then defined as before (Equation (5.1)) and the extension of the inclusions from the 2 to the d dimensional setting remain unchanged.

Altogether, this improves Proposition 5.4 as follows.

THEOREM 6.1. *Let p_1, \dots, p_n be a sequence of points in the plane 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^{n-2} disjoint convex polyhedra in \mathbb{R}^n , each defined by at most $4n$ linear inequalities, in two variables each. Moreover, through this lifting, F is strictly convex over each of these polyhedra.*

7. Reducing the number of sub-problems. Proposition 5.4 reduces the computation of a global minimum of F to the resolution of up to 2^{n-2} convex optimization sub-problems. To choose a sub-problem to consider means to fix one of the cells $\prod_i \Lambda_i^{\varepsilon_i}$ as a candidate cell for containing the global minimum of F . This choice has a simple geometric interpretation which can be used to considerably reduce the number of sub-problems to consider (Theorem 7.2).



“Class” of a shortest path. Let γ stand for the curvature-constrained shortest path that visits the configurations $(p_1, \theta_1) \dots (p_n, \theta_n)$ in order. Let $\varepsilon_i(\gamma) = -$ if the tangent to γ at p_i lies in the positive cone of the vectors $\overrightarrow{p_i p_{i-1}}$ and $\overrightarrow{p_{i+1} p_i}$ and $+$ otherwise. It is clear that $\varepsilon_i(\gamma) = +$ if and only if $\theta_i \in \Lambda_i^+$.⁶ Thus, if we can compute the “class” $(\varepsilon_1(\cdot), \dots, \varepsilon_n(\cdot))$ of the globally shortest path visiting p_1, \dots, p_n , then computing that path reduces to a *single* convex optimization problem. Figure 5.3 shows two paths in distinct classes.

Sharp turns. We now argue that unless the polygonal path $p_{i-1} p_i p_{i+1}$ makes a “sharp turn” at p_i , the tangent to the globally shortest path visiting p_1, \dots, p_n in order must belong to the positive cone of $\overrightarrow{p_i p_{i-1}}$ and $\overrightarrow{p_{i+1} p_i}$; in other words the “class” of the globally shortest path in p_i is trivial. Formally, 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 the triangle $p_{i-1} p_i p_{i+1}$ is acute at p_i and if the distance from p_{i-1} to the segment $p_i p_{i+1}$, or the distance from p_{i+1} to the segment $p_i p_{i-1}$, is at most 4.

LEMMA 7.1. *Let γ be a globally shortest path visiting p_1, \dots, p_n in order and assume that the (D_4) condition holds. If p_i is not a sharp turn then $\varepsilon_i(\gamma) = +$.*

Proof. Let γ be a globally shortest path. 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 which allows a global shortening of γ , contradicting the assumption that this path is globally shortest. To keep the notations simple we consider without loss of generality that $i = 2$.

Existence of a self-intersection. Refer to Figure 7.1(a). 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 left⁷ of ℓ ; since $\theta_2 \in \Lambda_2^-$, p_3 is to the left of ℓ and to the right of $(p_1 p_2)$. Moreover, if p_3 is on $(p_1 p_2)$ then p_2 is a sharp turn, so p_3 is *strictly* to the right of $(p_1 p_2)$ and, similarly, p_1 is strictly to the left of $(p_3 p_2)$.

⁶If $\overrightarrow{p_i p_{i-1}}$ and $\overrightarrow{p_{i+1} p_i}$ are opposite then $\varepsilon_i(\gamma)$ is defined so as to be consistent with the definition of Λ_i^\pm .

⁷Unless specified otherwise, the constraint to be to the left (or to the right) of an oriented line is considered non-strict.




Fig. 7.1: (a) If $\theta_2 \in \Lambda_2^-$ and p_2 is not a sharp turn, the path self-intersects. (For clarity the vectors are not drawn to scale.) (b) Shortening a path that self-intersects.

For $j = 1, 3$ let γ_j denote the portion of γ between p_j and p_2 . By Proposition 4.2 (or, rather, its generalization under the (D_4) condition mentioned in § 6), 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_2p_3) , 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} (shown in gray in Figure 7.1(a)). 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} .

Consider finally the intersection between γ_3 and ρ . Let γ'_3 be γ_3 minus its endpoint p_2 . γ'_3 intersects the line (p_1p_2) in exactly one point, m_3 , at which it traverses (p_1p_2) . 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 . 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 .

The self-intersection is between the two line segments. Let C_i^- and C_i^+ denote the circular arcs of γ that precede and follow p_i , respectively, and S^- and S^+ the line segments that precede and follow p_2 . The self-intersection identified above is an intersection between two elements in $\{C_1^+, S^-, C_2^-, C_2^+, S^+, C_3^-\}$. We discuss the various situations in turn. We assume here that the distance $|p_i p_{i+1}| > 4$; by continuity, the result will still hold under the non-strict (D_4) condition. Note also that $|p_1 p_3| > 4$ since p_2 is not a sharp turn.

The self-intersection cannot be between two circular arcs. Indeed, if $i \neq j$, C_i^\pm and C_j^\pm cannot intersect since $|p_i p_j| > 4$. Furthermore, since γ is locally shortest, Proposition 4.2 (and § 6) ensures that C_2^- and C_2^+ have length at most $\frac{3\pi}{4}$ and they cannot intersect (other than at p_2).

The self-intersection cannot be between a circular arc and a segment either. Any point in S^+ is within distance at most 2 from the segment $p_2 p_3$. Since p_2 is not a sharp turn, the distance from p_1 to the segment $p_2 p_3$ is more than 4 so C_1^+ cannot intersect S^+ . Similarly, C_3^- and S^- do not intersect. Since both S^+ and S^- are tangent to the circle supporting C_2^- and C_2^+ and to, respectively, C_3^- and C_1^+ , no other intersection between a circular arc and a segment is possible.

We therefore know that S^- and S^+ intersect in some point I .

Global shortening of the path. Let L^- and L^+ denote the two lines supporting S^- and S^+ . We consider the four unit circles C_1, \dots, C_4 tangent to both lines, and label the line/circle contact points as shown in Figure 7.1(b). 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 $\frac{3\pi}{4}$.

Let Q denote the endpoint of S^+ other than T_1^+ . If $T_3^+ \notin S^+$ then Q lies between T_3^+ and I , and 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_1 and p_3 . If that ray meets the disk centered in p_1 first, then p_1 is distance at most 4 from the segment $p_2 p_3$, and a similar argument holds in the symmetric case. Hence, p_2 is a sharp turn, a contradiction. It must therefore be that T_3^+ lies on S^+ . A symmetric argument shows that T_2^- lies on S^- . We can thus shorten γ using arcs of C_2 and C_3 (see Figure 7.1(b)), which concludes the proof. \square

We can now improve Theorem 6.1 as follows.

THEOREM 7.2. *Let p_1, \dots, 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 $4n$ linear inequalities, in two variables each. Moreover, through this lifting, F is strictly convex over each of these polyhedra.*

Proof. Let γ be a globally shortest path visiting p_1, \dots, p_n . If p_i is not a sharp turn then, the polar angle θ_i of γ at p_i belongs to Λ_i^+ by Lemma 7.1. Thus all global minima of F belong to the 2^k cells of $(\mathbb{S}^1)^n \setminus \mathcal{H}$ defined as $\prod_i \Lambda_i^{\varepsilon_i}$ where $\varepsilon_i = +$ if p_i is not a sharp turn, and $\varepsilon_i = \pm$ otherwise. The result then follows as in Proposition 5.4 and Theorem 6.1. \square

8. Extensions and perspectives. We conclude this paper by an overview of various refinements and extensions of our results.

Complexity analysis. 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 sum of minima of expressions using inverse trigonometric functions, for which computing *self-concordant barriers* currently seems out of reach. For the ellipsoid method, the number of steps needed to achieve an additive error of at most ε on the solution depends on the *fatness* of the domain [4, Theorem 5.2.1], which we only managed to analyze for the sub-problem $\mathcal{D} \cap \prod_i \Lambda_i^+$ under the $(D_{4.3})$ condition. As a consequence, under that distance condition and in the absence of sharp turn, we can compute a path (of curvature at most 1 that visits the p_i in order) whose length exceeds that of an optimal path by at most an *additive* constant ε in time⁸ $O(n^4 \log \frac{n}{\varepsilon})$, which is the usual complexity for the ellipsoid method in dimension n . We refer to the technical report [15, Section 8] for the details of the analysis.


⁸The complexity is considered in 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.

Fixed initial and final directions. If the initial and/or the final directions of the path are fixed, a globally shortest path may have a circular arc of length more than π incident to p_1 or p_n . This invalidates our sufficient condition for a *CSC*-path to be the shortest Dubins path between two configurations (Lemma 3.1) which was crucial for establishing the local convexity of F . A way around this difficulty is to consider four *fixed-type length functions* where the orientations (L or R) of the circular arcs following p_1 and preceding p_n are imposed. Our previous proof structure yields, mutatis mutandis, that each fixed-type length function is locally strictly convex at any point $(\theta_2, \dots, \theta_{n-1})$ such that every circular arc preceding or following a waypoint has length less than π , except the initial and/or final ones. From there, the rest of the proof can be easily adapted to show that the fixed-type length functions realize their global minima over some number of polyhedra over which they are locally strictly convex. The minimum of all fixed-type functions over all these polyhedra yields the minimum of the original length function F . See the technical report [15, Section 9.1] for a more detailed discussion.

Obstacles. Our technique 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 extremal configurations or contact points on the boundary of the obstacles [14, 18]. It can be shown 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; in other words, comparatively, “connecting the dots” is now relatively easy, and the difficult task appears to be the discrete sub-problem of computing the contact points and the anchored circular arcs. 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. We, again, refer the interested reader to the technical report [15, Section 9.2] for details.

Distance condition. The requirement that the points p_1, \dots, 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 $\mathcal{L}(\frac{3\pi}{4})$ (Proposition 4.2). Finally, it is instrumental for proving that the minima of F can be searched for in convex components over which F is convex (Lemma 5.1). 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, a task we did not undertake.

An alternative approach?. Lemma 4.1 and Proposition 4.2 imply that, under the (D_4) condition, if there is a globally shortest path starting from p_1 with a given direction \vec{u} and visiting p_1, \dots, p_n in order, that path is unique and can easily be computed. Indeed, there is at most one



circular arc of length less than $\frac{3\pi}{4}$ that is tangent to the ray $p_1 + \mathbb{R}\vec{u}$ and ends in p_2 ; the circular arc leaving p_2 must have the same length and orientation (the case where both arcs sum to 2π in Lemma 4.1 being ruled out by Proposition 4.2) so the ray supporting the segment following p_2 is fixed. An immediate induction determines the path upto the ray following p_{n-1} and this ray must contain p_n for the path to be globally shortest. Of course, if at some point, the ray following p_i is too far away from p_{i+1} then the path cannot be continued and \vec{u} is not the initial direction of a globally shortest path. The problem then reduces to finding a zero of the signed distance from the final ray to p_n , seen as a function of the polar angle θ_1

of \vec{u} . It seems plausible that this function is well-behaved (e.g. piecewise monotonic in θ_1) and that simple methods such as binary search could perhaps be used. However, this function can have at least 2^{n-2} zeros, as shown by the lower-bound example depicted in Figure 5.3, so this approach is likely to face the same combinatorial explosion as our reduction to convex optimization.

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, Otfried Cheong for helpful discussions and the two anonymous referees for their many valuable suggestions and in particular the second referee for pointing out the alternative approach sketched in the last section.

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 (2002), pp. 1814–1851.
- [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, New York, NY, USA, 1995, ACM Press, pp. 343–352.
- [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, Dec 2008, pp. 628–643.
- [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, New York, NY, USA, 2005, ACM Press, pp. 278–287.
- [6] J.-D. BOISSONNAT, A. CÉRÉZO, AND J. LEBLOND, *Shortest paths of bounded curvature in the plane*, Journal of Intelligent and Robotic Systems, 11 (1994), pp. 5–20.
- [7] J.-D. BOISSONNAT AND S. LAZARD, *Convex hulls of bounded curvature*, in Proc. 8th Canad. Conf. Comput. Geom., 1996, pp. 14–19.
- [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 (2003), pp. 189–229.
- [9] X.-N. BUI, *Planification de trajectoire pour un robot polygonal non-holonyme 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, San Diego, CA, 1994, pp. 2–7.
- [11] G. DESAULNIERS, *On shortest paths for a car-like robot maneuvering around obstacles*, Les Cahiers du GERAD G-94-35, École 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 (1957), pp. 497–516.
- [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 (1991), pp. 21–82.
- [15] X. GOAOC, H.-S. KIM, AND S. LAZARD, *Bounded-curvature shortest paths through a sequence of points*, Research report RR-7465, Inria, 2010.
- [16] D. HALPERIN, L. E. KAVRAKI, AND J.-C. LATOMBE, *Robotics*, in CRC Handbook of Discrete and Computational Geometry, CRC Press, 2004, ch. 41, pp. 1065–1093.
- [17] 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.
- [18] P. JACOBS AND J. CANNY, *Planning smooth paths for mobile robots*, in Nonholonomic Motion Planning, Z. Li and J.F. Canny, eds., Kluwer Academic, 1992, pp. 271–342.
- [19] J.-C. LATOMBE, *Robot Motion Planning*, Kluwer Academic Publishers, Boston, 1991.
- [20] S. M. LAVALLE, *Planning Algorithms*, Cambridge University Press, 2006.
- [21] 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, 2000, pp. 314–325.
- [22] X. MA AND D. A. CASTAÑÓN, *Receding horizon planning for Dubins traveling salesman problems*, in 45th IEEE Conference on Decision and Control, Dec 2006.
- [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] ———, *The curvature-constrained traveling salesman problem for high point densities*, in Proceedings of the 46th IEEE Conference on Decision and Control, 2007, pp. 5985–5990.
- [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 (2007).
- [26] J. A. REEDS AND L. A. SHEPP, *Optimal paths for a car that goes both forwards and backwards*, Pacific J. Math, 145 (1990), pp. 367–393.
- [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, Natick, MA, USA, 1998, A. K. Peters, Ltd., pp. 49–57.
- [28] K. SAVLA, E. FRAZZOLI, AND F. BULLO, *Traveling salesperson problems for the Dubins vehicle*, IEEE Transactions on Automatic Control, 53 (2008).
- [29] J. T. SCHWARTZ AND M. SHARIR, *Algorithmic motion planning in robotics*, in Algorithms and Complexity, J. van Leeuwen, ed., vol. A of Handbook of Theoretical Computer Science, Elsevier, Amsterdam, 1990, pp. 391–430.
- [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, Houston, Texas, USA, 1998, pp. 59–67.
- [31] H. J. SUSSMANN, *Shortest 3-dimensional paths with a prescribed curvature bound*, in Proceedings of the 34th IEEE Conference on Decision and Control, 1995, pp. 3306–3311.
- [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 (2001), pp. 1739–1772.