On Polygonal Paths with Bounded Discrete-Curvature: The Inflection-Free Case

Sylvester Eriksson-Bique[1], David Kirkpatrick[2](✉), and Valentin Polishchuk[3]

[1] Courant Institute, New York, USA
seb522@nyu.edu
[2] University of British Columbia, Vancouver, Canada
kirk@cs.ubc.ca
[3] University of Helsinki, Helsinki, Finland
polishch@cs.helsinki.fi

Abstract. A shortest path joining two specified endpoint configurations that is constrained to have mean curvature at most ς on every non-zero length sub-path is called a ς-*geodesic.* A seminal result in non-holonomic motion planning is that (in the absence of obstacles) a 1-geodesic consists of either (i) a (unit-radius) circular arc followed by a straight segment followed by another circular arc, or (ii) a sequence of three circular arcs the second of which has length at least π [Dubins, 1957]. Dubins' original proof uses advanced calculus; Dubins' result was subsequently rederived using control theory techniques [Sussmann and Tang, 1991], [Boissonnat, Cérézo, and Leblond, 1994], and generalized to include reversals [Reeds and Shepp, 1990].

We introduce and study a discrete analogue of curvature-constrained motion. Our overall goal is to show that shortest polygonal paths of bounded "discrete-curvature" have the same structure as ς-geodesics, and to show that properties of ς-geodesics follow from their discrete analogues as a limiting case, thereby providing a new, and arguably simpler, "discrete" proof of the Dubins characterization. Our focus, in this paper, is on paths that have non-negative mean curvature everywhere; in other words, paths that are free of inflections, points where the curvature changes sign. Such paths are interesting in their own right (for example, they include an additional form, not part of Dubins' characterization), but they also provide a slightly simpler context to introduce all of the tools that will be needed to address the general case in which inflections are permitted.

1 Introduction

Curvature-constrained paths are a fundamental tool in planning motion with bounded turning radius. Paths that are *smooth* (continuously differentiable) have the advantage that they may look more appealing and realistic than *polygonal* (piecewise-linear) paths. Nevertheless, polygonal paths are a much more common model in geometry, exactly because of their discrete nature, and for this same reason they have the potential of providing simpler and more intuitive proofs

J. Akiyama et al. (Eds.): JCDCGG 2013, LNCS 8845, pp. 44–64, 2014.
DOI: 10.1007/978-3-319-13287-7_5

of properties of their smooth counterparts. Furthermore, from an applications perspective, polygonal paths are more natural to plan, describe and follow. For instance, in air traffic management—one of our motivating applications—an aircraft flight plan (a list of "waypoints") is represented on the strategic level by a polygonal path whose vertices are the GPS waypoints. The actual smoothly turning trajectory at a waypoint is decided by the pilot on the tactical level when executing the turn (see [25] for more on curvature-constrained route planning in air transport). We are thus motivated to formulate a *discretized* model of curvature-constrained motion.

Smooth paths of bounded curvature. In studying smooth paths of bounded curvature, L. E. Dubins [15] observed that if one restricts attention to paths whose curvature is defined at every point then there are situations in which shortest paths do not exist. On the other hand, if one only requires that paths are everywhere differentiable (that is their slope is well-defined at every point) then their *mean* curvature is well-defined on every non-zero length sub-path. Thus, Dubins chose to define a path to have bounded curvature if its mean curvature is bounded everywhere. Specifically, let γ be a smooth path, parameterized by its arclength. For any t in the domain of γ, let $\gamma'(t)$ denote the derivative of γ at t – the unit vector tangent to γ at t. The path γ is said to have mean curvature at most one if $\angle st \leq s - t$, for any $t < s < t + \pi$, where $\angle st$ denotes the angle between the directions of $\gamma'(s)$ and $\gamma'(t)$. (In other words, γ', viewed as mapping from the domain of γ to the unit circle, is 1-Lipschitz.) Furthermore, there is no loss of generality in restricting attention to the case where the mean curvature bound is one, since the general case can be reduced to the unit case by suitable scaling. Accordingly, we hereafter use the term "curvature-constrained" as a shorthand for "has mean curvature bounded by one", and we refer to paths with this property as *cc-paths*.

Dubins' characterization, a seminal result in curvature-constrained motion planning, states that, in the absence of obstacles, shortest curvature-constrained paths in the plane, are one of two types: (i) a circular arc followed by a line segment followed by another arc, or (ii) a sequence of three circular arcs, the second of which has length at least π.

Discrete circular arcs. While several possibilities suggest themselves as ways to formulate a discrete analogue of unit-bounded curvature[1], it seems that all such formulations are based on a natural notion of discrete circular arcs. Let $0 < \theta \leq \pi/2$ be a given angle. We say that a polygonal chain forms a *θ-discrete circular arc* (or simply a discrete circular arc if θ is understood) if (i) its vertices belong, in sequence, to a common circle of radius one, and (ii) successive edges have length at most $d_\theta = 2\sin\frac{\theta}{2}$ (that is they subtend a circular arc of length at most θ). Any portion of regular polygon with $k \geq 3$ sides, inscribed in a unit circle, provides a prototypical $(2\pi/k)$-discrete circular arc.

[1] In an earlier draft [18], the authors proposed an alternative definition which had some deficiencies that are resolved by the definition used in this paper.

Polygonal paths of bounded discrete-curvature. Discrete circular arcs not only satisfy our intuitive notion of polygonal path with bounded discrete-curvature but, like their smooth counterparts, they seem to capture the extreme case. It is interesting to ask what properties a path P with bounded discrete-curvature should have in general. To ensure that such a path does not turn "too sharply" it seems natural to require that, like discrete circular arcs, P should turn by at most θ at each of its interior vertices. However, such a restriction alone does not guarantee that P will serve as a bona fife discrete analogue of a bounded-curvature path: many short successive edges of P, each turning only slightly, can simulate a sharp turn. Taking further inspiration from discrete circular arcs, that permit short edges with correspondingly gentle turns, we define:

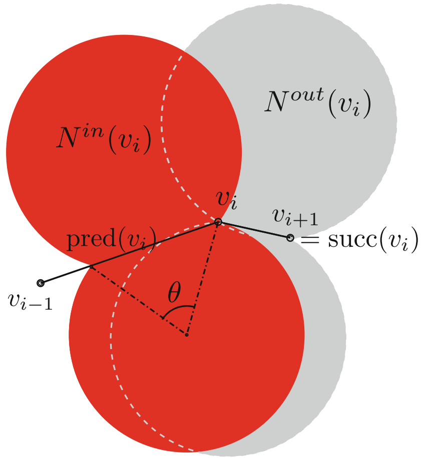

Fig. 1. Local conditions for θ-discrete curvature-constrained paths.

Definition 1. *A polygonal path $\langle v_1, v_2, \ldots, v_n\rangle$ has θ-discrete-curvature (or just discrete-curvature, if θ is understood) at most one if, for $1 < i < n$, the turn at vi, the difference between the angles of the rays $v_{i-1}v_i$ and v_iv_{i+1}, is no more than* $\sin^{-1}(\frac{\min\{d_\theta, |v_{i-1}v_i|\}}{2}) + \sin^{-1}(\frac{\min\{d_\theta, |v_iv_{i+1}|\}}{2})$.

Remark 1. As with its smooth counterpart, we hereafter use the term "θ-discrete-curvature-constrained" (frequently abbreviated as "θ-dcc") to mean "has θ-discrete-curvature bounded by one". It is easy to confirm that (i) any discrete circular arc is a dcc-path, and (ii) for every vertex v_i on a dcc-path, the point $\mathrm{succ}(v_i)$ at distance $\min\{d_\theta, |v_iv_{i+1}|\}$ from v_i along the edge v_iv_{i+1}, does not lie in the interior of the region $N^{in}(v_i)$ formed by the union of the unit circles passing through v_i and $\mathrm{pred}(v_i)$, the point at distance $\min\{d_\theta, |v_{i-1}v_i|\}$ from v_i along the edge $v_{i-1}v_i$. We will refer to $N^{in}(v_i)$ (respectively, $N^{out}(v_i)$, the region formed by union of the unit circles passing through v_i and $\mathrm{succ}(v_i)$ as the *in-neighbourhood* (respectively, *out-neighbourhood*) of v_i (see Fig. 1). Note that, if $|v_iv_{i+1}| \le d_\theta$ then $N^{out}(v_i) = N^{in}(v_{i+1})$.

Relating smooth and discrete curvature-constrained paths. Let γ be a smooth path, parameterized by its arclength. We say that a polygonal path $\hat{\gamma} = \langle v_1, v_2, \ldots, v_n\rangle$ is a *θ-discretization* of γ if (i) $v_i = \gamma(t_i)$, for $1 \le i \le n$, where (ii) $t_1 = 0$, $t_n = |\gamma|$, and $0 \le t_{i+1} - t_i \le \theta$, for $1 \le i < n$.

By definition, a θ-discrete circular arc is a θ-discretization of an arc of a (smooth) unit circle. In fact (cf. Theorem 1 below) *every* θ-discretization $\hat{\gamma} = \langle\gamma(t_1), \gamma(t_2), \ldots, \gamma(t_n)\rangle$ of *every* cc-path γ forms a θ-dcc-path.

Lemma 1. *For any r and s, $r < s < r + \pi$, in the domain of γ, $|\gamma(s) - \gamma(r)| \geq 2 \sin \frac{s-r}{2}$.*

Proof. Assume, without loss of generality, that $r = 0$ and that $\gamma(r)$ is at the origin O (i.e. $\gamma(0) = (0, 0)$); the lemma is then equivalent to $|\gamma(s)| \geq 2 \sin \frac{s}{2}$ (Fig. 2). We prove this by lower-bounding the derivative of $|\gamma(s)|^2$:

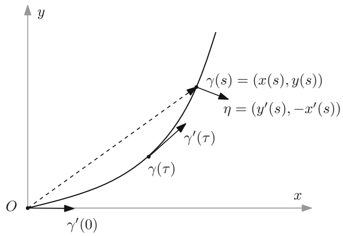

Fig. 2. $|O\gamma(s)| \geq 2 \sin \frac{s}{2}$, the slope of $O\gamma(s)$ is at most $\tan \frac{s}{2}$.

$$(|\gamma(s)|^2)' = 2\gamma(s) \cdot \gamma'(s) = 2 \int_0^s \gamma'(\tau) \cdot \gamma'(s)\, d\tau = 2 \int_0^s \cos \angle s\tau \, d\tau \geq 2 \int_0^s \cos(s - \tau)\, d\tau = 2 \sin s$$

Hence $|\gamma(s)|^2 \geq 2(1 - \cos s) = 4 \sin^2 \frac{s}{2}$. □

Corollary 1. *For all i, $1 \leq i < n$, $|\gamma(t_{i+1}) - \gamma(t_i)| \geq |t_{i+1} - t_i| \frac{\sin(\theta/2)}{\theta/2}$.*

Proof. It suffices to observe that, since $0 \leq t_{i+1} - t_i \leq \theta$, $\frac{\sin((t_{i+1}-t_i)/2)}{(t_{i+1}-t_i)/2} \geq \frac{\sin(\theta/2)}{\theta/2}$. □

Lemma 2. *For any r and s, $r < s < r + \pi$, in the domain of γ, the angle between $\gamma'(r)$ and the ray $\gamma(r)\gamma(s)$ is at most $\frac{s-r}{2}$.*

Proof. Assume again that $r = 0$ and that $\gamma(0) = O$; also assume w.l.o.g. that $\gamma'(0)$ is horizontal ($\gamma'(0) = (1, 0)$). Let $\gamma(s) = (x(s), y(s))$, and let $k(s) = y(s)/x(s)$ be the slope of the ray $O\gamma(s)$ (Fig. 2, left). Then the lemma is equivalent to $k(s) \leq \tan \frac{s}{2}$, which we will prove by showing that $k' \leq \frac{1-\cos s}{\sin^2 s} = \frac{1}{2\cos^2(s/2)} = (\tan \frac{s}{2})'$.

By definition, for any $\tau < s$ the angle between $\gamma'(\tau)$ and $\gamma'(s)$ is at most $s - \tau$; in particular $\gamma'(s) \leq s$. It follows that $x'(s) = (x'(s), y'(s)) \cdot (1, 0) = \gamma'(s) \cdot \gamma'(0) = \cos \angle s0 \geq \cos s$, and thus $x(s) \geq \sin s$. Next, consider the unit vector $\eta = (y'(s), -x'(s))$, orthogonal to $\gamma'(s)$ (Fig. 2). Since the angle between $\gamma'(\tau)$ and η is at least $\pi/2 - (s - \tau)$, it follows that $(x'(\tau), y'(\tau)) \cdot \eta \leq \cos(\pi/2 - (s - \tau))$, or $x'(\tau)y'(s) - y'(\tau)x'(s) \leq \sin(s - \tau)$. Integrating over τ from 0 to s, we get $x(s)y'(s) - y(s)x'(s) \leq 1 - \cos s$. Combining this with $x(s) \geq \sin s$, we obtain what we need: $k' = (y/x)' = \frac{y'x - x'y}{x^2} \leq \frac{1-\cos s}{\sin^2 s}$. □

Corollary 2. *The angle between the ray $\gamma(t_{i-1})\gamma(t_i)$ and $\gamma'(t_i)$ is at most* $\sin^{-1}$ $(\frac{\min\{d_\theta, |\gamma(t_{i-1})\gamma(t_i)|\}}{2})$

Proof. By the lemma, the angle between the ray $\gamma(t_{i-1})\gamma(t_i)$ and $\gamma'(t_i)$ is at most $\frac{t_i - t_{i-1}}{2}$, which is always at most $\theta/2 = \sin^{-1}(\frac{d_\theta}{2})$. So, it suffices to consider the case where $|\gamma(t_{i-1})\gamma(t_i)| < d_\theta$. But in this case, $\frac{t_i - t_{i-1}}{2} \leq \sin^{-1}(\frac{|\gamma(t_{i-1})\gamma(t_i)|}{2})$, by Lemma 1. □

In summary, we have shown the following:

Theorem 1. *If γ is any cc-path and $\hat{\gamma}$ any θ-discretization of γ, then (i) $\hat{\gamma}$ is a θ-bcc-path, and (ii) $|\gamma|\frac{\sin(\theta/2)}{\theta/2} \leq |\hat{\gamma}| \leq |\gamma|$.*

Proof. (i) That $\hat{\gamma}$ is a θ-bcc-path is an immediate consequence of Corollary 2, since the angle between the ray $\gamma(t_{i-1})\gamma(t_i)$ and the ray $\gamma(t_i)\gamma(t_{i+1})$ is just the sum of the angles formed by these rays with $\gamma'(t_i)$.
(ii) It is clear that the length of any θ-discretization of a smooth curve γ is no greater than the length of γ. On the other hand, it follows immediately from Corollary 1 that its length cannot be less than $|\gamma|\frac{\sin(\theta/2)}{\theta/2}$. □

The fact that the bounds on $|\hat{\gamma}|$ coincide in the limit as θ approaches zero, will be used to obtain properties of shortest smooth paths as a limit of the properties of their discrete counterparts.

Remark 2. It is worth noting at this point that our definition of θ-dcc-path, because of its "local" nature, rules out some paths that may be seen as having bounded curvature. For example, a "sawtooth" approximation of a straight line (see Fig. 3) does not qualify as a θ-dcc-path if the pitch of the teeth (turn angle), no matter how small, is too sharp relative to the size of the teeth (edge length).

Configurations. A *configuration* is a pair (p, ϕ), where p is a point and ϕ is a direction (unit vector). We say that a polygonal path $P = \langle v_1, v_2, \ldots, v_n\rangle$ satisfies endpoint configurations (v_1, ϕ_1) and (v_n, ϕ_n) if (i) the difference between the angle of the ray v_1v_2 and ϕ_1, is no more than $\sin^{-1}(\frac{\min\{d_\theta, |v_1v_2|\}}{2})$, and (ii) the difference between the angle of the ray $v_{n-1}v_n$ and ϕ_n, is no more than $\sin^{-1}(\frac{\min\{d_\theta, |v_{n-1}v_n|\}}{2})$.

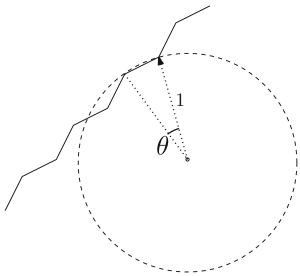

Fig. 3. A sawtooth path that does *not* qualify as a θ-dcc-path.

Remark 3. This is equivalent to asserting that the path $\langle v_0, v_1, v_2, \ldots, v_n, v_{n+1}\rangle$, formed from P by adding an arbitrarily short edge of direction ϕ_1 (respectively, ϕ_n) to the start (respectively, end) of P, has bounded discrete-curvature.

Remark 4. It is easy to confirm that, for *any* intermediate configuration (v_i, ϕ_i), the composition of any dcc-path $\langle v_1, v_2, \ldots, v_i\rangle$ that satisfies endpoint configurations (v_1, ϕ_1) and (v_i, ϕ_i) with any dcc-path $\langle v_i, v_{i+1}, \ldots, v_n\rangle$ that satisfies endpoint configurations (v_i, ϕ_i) and (v_n, ϕ_n) produces a dcc-path $\langle v_1, v_2, \ldots, v_n\rangle$ that satisfies endpoint configurations (v_1, ϕ_1) and (v_n, ϕ_n). Furthermore, if $P = \langle v_1, v_2, \ldots, v_n\rangle$ is any dcc-path that satisfies endpoint configurations (v_1, ϕ_1) and (v_n, ϕ_n), then, for all i, $1 < i < n$, there exists a direction ϕ_i such that (i) the sub-path $\langle v_1, v_2, \ldots, v_i\rangle$ is a dcc-path with endpoint configurations (v_1, ϕ_1) and (v_i, ϕ_i), and (ii) the sub-path $\langle v_i, v_{i+1}, \ldots, v_n\rangle$ is a dcc-path with endpoint configurations (v_i, ϕ_i) and (v_n, ϕ_n). On the other hand, breaking P at an arbitrary point in the interior of one of its edges may produce a path that no longer has bounded discrete-curvature (at the breakpoint).

We will be interested in characterizing shortest dcc-paths that satisfy specified endpoint configurations:

Definition 2. *A* discrete-geodesic *joining endpoint configurations* (v_1, ϕ_1) *and* (v_n, ϕ_n)*, is a dcc-path that (i) satisfies the endpoint configurations* (v_1, ϕ_1) *and* (v_n, ϕ_n)*, and (ii) has minimum total length among paths satisfying (i).*

Remark 5. It is by no means obvious that discrete-geodesics exist for all endpoint configurations. Dubins' proof [15] of the existence of smooth geodesics makes use of tools from functional analysis (in particular, Ascoli's theorem); in Sect. 4, we describe an alternative approach to establishing the existence of discrete-geodesics, having established a suitable characterization of the form the discrete-geodesics must take (if they exist).

Remark 6. It is straightforward to confirm that if a path contains a pair of successive edges $v_{i-1}v_i$ and v_iv_{i+1} whose *shortcut*, edge $v_{i-1}v_{i+1}$, has length at most d_θ, then the path can be made both shorter and smaller (fewer edges) by replacing the edges by their shortcut, without violating the curvature constraint. It follows that for the purposes of characterizing discrete-geodesics, we can restrict our attention to paths with finitely many edges formed by maximal[2] discrete circular arcs connected by (possibly degenerate) line segments. This assertion, which follows immediately from our definition, is the discrete analogue of a non-trivial property of smooth curvature-bounded geodesics, proved as Proposition 13 in [15].

1.1 Related Work

The books [26,27] provide general references that discuss curvature-constrained path planning in the broader context of non-holonomic motion planning. We note that study of curvature-constrained path planning has a rich history that long predates and goes well beyond robot motion planning, see for example the work of Markov [29] on the construction of railway segments.

The Dubins characterization of smooth geodesics has been rederived using techniques from optimal control theory in [7,40]. Variations and generalizations of the problem are studied in [6,8–10,12–14,19,28,30–34,37–39]. In addition, Dubins' characterization plays a fundamental role in establishing the existence as well as the optimality of curvature-constrained paths. Jacobs and Canny [23] showed that even in the presence of obstacles it suffices to restrict attention to paths of Dubins form between obstacle contacts and that if such a path exists then the shortest such path is well-defined. Fortune and Wilfong [20] give a super-exponential time algorithm for determining the existence of, but not actually constructing, such a path. Characterizing the intrinsic complexity of the existence problem for curvature-constrained paths is hampered by

[2] We ignore for the present the fact that successive maximal discrete circular arcs of opposite orientation could share an edge. In this case we are free to impose disjointness of arcs by assigning the shared edge to just one of the two arcs.

the fact that there are no known bounds on the minimum length or *intricacy* (number of elementary segments), expressed as a function of the description of the polygonal domain, of obstacle-avoiding paths in Dubins form. In a variety of restricted domains polynomial-time algorithms exist that construct shortest bounded-curvature paths [1,2,4].

A discretization of curvature-constrained motion was studied by Wilfong [42,43]. However, his setting is different from ours since he considered discretized *environment*, and not discrete paths. A practical way of producing paths with length and turn constraints is presented in [41]. For some other recent work on bounded-curvature paths see [3,5,11,16,17,21,22,24].

1.2 Our Approach

We study properties of discrete-geodesics that are free of inflections, as well as their smooth counterparts. In Sect. 2 we motivate the study of this restricted class of paths by proving that unrestricted discrete-geodesics never need more than two internal inflections; i.e. all discrete geodesics are formed by the concatenation of at most three inflection-free discrete-geodesics.

Section 3 establishes the central result of the paper: a precise characterization of the form of all discrete-geodesics (if they exist). In Sect. 4 we use this characterization to outline a proof of a characterization of smooth inflection-free geodesics (establishing, by a simple limiting argument, this interesting variant of the Dubins characterization). We also include a simple geometric proof of the existence of one important special case of discrete-geodesics that illustrates the strength of our characterization.

We note that similar methods, proving properties of smooth curves using discretization, were already used by Schur in his paper of 1921; interestingly, exactly these problems, considered by Schur [36] and Schmidt [35], led Dubins to his result.

2 Inflections in Discrete-Geodesics

We now start our investigation of the structure of discrete-geodesics. An edge e of a polygonal path is an *inflection* edge if the edges adjacent to e lie on the opposite sides of (the supporting line of) e. Such an edge is said to have *positive inflection* if the path makes a left turn into and a right turn out of e (and *negative inflection*, otherwise). Note that, in accordance with our interpretation of endpoint conditions as a zero-length edge of specified orientation, the first and last edges of a path are possible inflection edges. When we want to distinguish such edges, we refer to them as *endpoint inflection edges*; other inflection edges are referred to as *internal inflection edges*.

It is not hard to see that a dcc-path can have arbitrarily many inflection edges (of arbitrary lengths). However, minimum length such paths, can have no more than two internal inflection edges (of any length) *in total*.

2.1 More than Two Internal Inflection Edges is Impossible

Our first observation is that in any discrete-geodesic there can be at most one internal inflection edge of each turn type.

Lemma 3. *Any θ-dcc-path containing two or more internal inflection edges of the same type can be replaced by another θ-dcc-path, with fewer edges, whose total length is no longer than the original. In fact, if the inflection edges are non-parallel, the replacement path is strictly shorter.*

Proof. Let $P = \langle v_1, \ldots, a, b, c, d, \ldots, w, x, y, z, \ldots, v_n \rangle$ and suppose that both bc and xy are internal inflection edges with positive inflection (i.e. P turns left at both b and x and right at both c and y; see Fig. 4). Note that since edges bc and xy are internal inflection edges, the edges ab, cd, wx and yz all have strictly positive length.

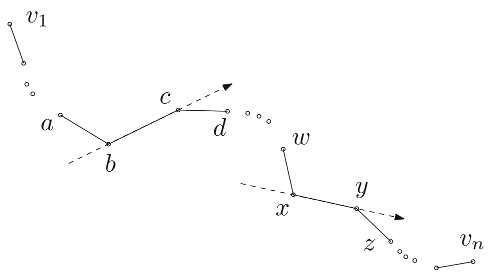

Fig. 4. A discrete curvature-constrained path with two positive inflection edges bc and xy.

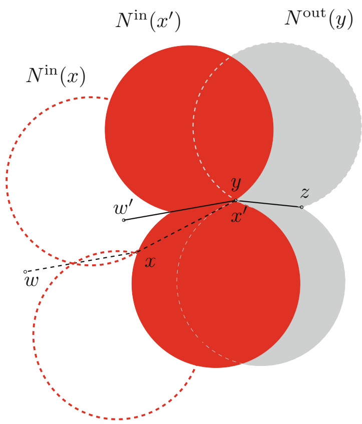

Fig. 5. Full translation

We assume, without loss of generality, that the ray from b through c and the ray from x through y are either parallel or diverge (as illustrated). Then, any transformation of P that results from a translation of the sub-path between c and x along the vector xy (taking $c, \ldots x$ to $c', \ldots x'$) reduces the length of P (except when the inflection edges are parallel, in which case the length of P is preserved) and maintains the dcc-path property at b and c (since the edge bc lengthens, while the turns at both b and c are not increased).

Consider the situation when x has been translated all the way to y (see Fig. 5). The dcc-path property holds for the resulting path as long as pred(x') lies in $N^{out}(y)$ (or, equivalently, succ(y) lies in $N^{in}(x')$), where x' denotes the translation of x, etc. In this case, there is nothing left to prove since the path has one fewer edge (namely xy) than P. Hence we can assume that, after this full translation, pred(x') lies in $N^{out}(y)$.

If pred(x') lies in the right component of $N^{out}(y)$ (see Fig. 6) then it must lie in the segment of this circle cut off by the line through x and y, which implies that $|\text{pred}(x')x'| < d_\theta$ and so pred$(x') = w'$. It follows that when the translation is taken just to the point where w' lies on the boundary of $N^{out}(y)$, at which point y must still lie outside $N^{\text{out}}(\text{pred}(x'))$ (see Fig. 7), we have $|w'y| < d_\theta$ and thus if we replace the edges $w'x'$ and $x'y$ at this point by the edge $w'y$, we must have a path that satisfies the dcc-path property.

Similarly, if pred(x') lies in the left half of $N^{out}(y)$ (see Fig. 6) then succ(y) must lie in the left half of $N^{in}(x')$ and in fact in the segment of this circle cut off by the line through x and y. As before, this implies that $|\text{succ}(y)y| < d_\theta$ and so succ$(y) = z$. It follows that when the translation is taken just to the point where z lies on the boundary of $N^{in}(x')$, at which point x' has not yet entered the interior of $N^{out}(y)$ (see Fig. 7), we have $|x'z| < d_\theta$ and thus if we replace the edges $x'y$ and yz at this point by the edge $x'z$, we must have a path that satisfies the dcc-path property.

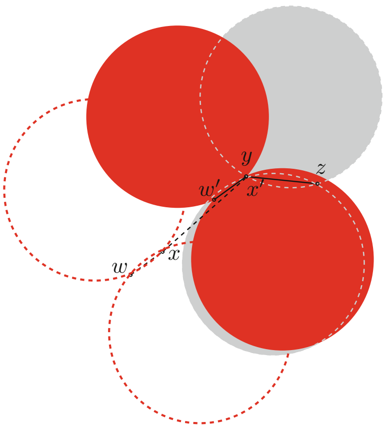

Fig. 6.

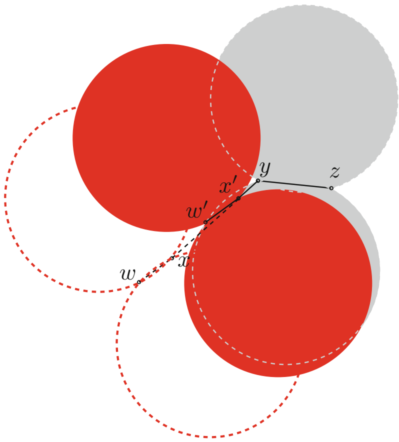

Fig. 7.

Since in both of these remaining cases the resulting path has one fewer edge than P, the result follows. Note that the only situation where the path has not had its length strictly reduced is where the inflection edges are parallel (and so the translation is length preserving). □

Remark 7. It is worth observing at this point that Lemma 3 applies as well to the case in which the inflection edge xy is an endpoint inflection. In this case, we are not able to conclude that the transformed path has one fewer edge (since edge yz has length zero), but it does follow from our argument that if P cannot be shortened then the endpoint inflection edge must be a chord of the circle that defines the endpoint configuration at y.

It follows from Lemma 3 that, ignoring the length zero edges at the path endpoints, any discrete-geodesic is the concatenation of at most three inflection-free

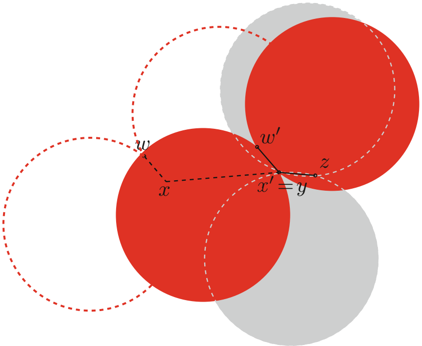

Fig. 8.

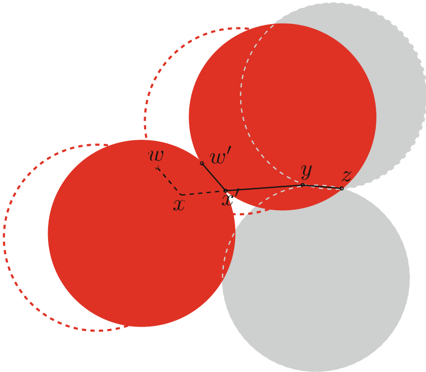

Fig. 9.

sub-paths. This motivates our next focus on the form of inflection-free discrete-geodesics.

3 A Normal Form for Inflection-Free Geodesics

We have already noted that any shortest dcc-path of minimum size consists of a finite number of maximal discrete circular arcs connected by (possibly degenerate) line segments that we refer to as *bridges*. Here we include the (possibly degenerate) circular arcs (e.g. vertex k in Fig. 10) supported by the circles (shown as dashed) that define the endpoint configurations of the path. A bridge vw is *degenerate* if $v = w$ (e.g. the first bridge, vertex d, in Fig. 10). Of course, if a given path has no inflection edges, all of the discrete circular arcs have the same orientation; without loss of generality we will assume that they are all clockwise oriented (see Fig. 10).

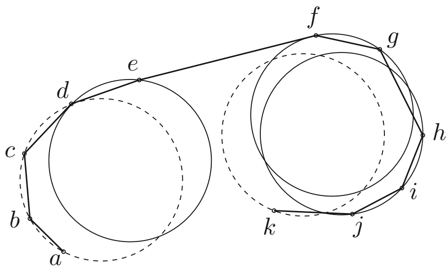

Fig. 10. A discrete-curvature-constrained path with five discrete arcs (including one degenerate arc) and no inflections.

The main result of this section is the following theorem. It amounts to a special case of a general characterization theorem for discrete-geodesics and is a fundamental building block for the proof of that theorem.

Theorem 2. *Any inflection-free discrete-geodesic, joining two specified endpoint configurations, is composed of a sequence of at most four discrete circular arcs, at most two of which are non-degenerate, joined by (possibly degenerate) bridges.*

The proof of Theorem 2 proceeds by induction on the number of discrete circular arcs in the path. Since the basis of the induction is obvious, it suffices to prove the following:

Lemma 4. *Any dcc-path, joining two specified endpoint configurations, that consists of three non-degenerate discrete circular arcs of the same orientation, joined by (possibly degenerate) bridges, can be replaced by a shorter path, joining the same endpoint configurations, that consists of at most two non-degenerate discrete circular arcs.*

Proof. Suppose we are given a sub-path consisting of three non-degenerate discrete circular arcs joined by two (possibly degenerate) line segments that we refer to as the first and second bridge. We will refer to the circles supporting the three arcs as the first, second and third circles (coloured blue, red and green, respectively, in all of our figures). We consider several cases depending on the nature of the two bridges (degenerate or not) and the total turn of the second arc (essentially whether or not it exceeds π). With only one explicitly noted exception, we use one of two continuous *shortening transformations* both of which involve moving all (or most) of the vertices on the second arc while keeping the other arcs fixed: (i) a *pivot* rotates all (except possibly the opposite endpoint) of the vertices on the second arc about one of the arc endpoints; and (ii) a *slide* translates all (except possibly the opposite endpoint) of the vertices on the second arc along one of the non-degenerate bridges. Since both transformations move the second arc in a rigid fashion, we need only consider vertices in the neighbourhood of the two bridges to confirm that the transformations preserve the dcc-path property.

To simplify the analysis that follows we will assume throughout that if a transformation leads to a *co-linearity event*: one of the bridges becomes co-linear with one of its adjacent edges (equivalently, the turn at some bridge endpoint becomes zero), then we will stop the transformation at this point and combine the two co-linear edges into a new bridge. Obviously, this results in a simpler path (with one fewer edge) and a possible degeneration of one of the discrete circular arcs. With this exception, all of our transformations terminate with either (i) the second arc becoming co-circular with the first or third (a *co-circularity event*), in which case a bridge has been eliminated, (ii) a formerly non-degenerate bridge becoming degenerate (a *bridge degeneration event*), or (iii) a bridge intersecting one of its associated circles in a chord of length d_θ (a *maximal chord event*). In the second event, the resulting path is simpler in the sense that a path with a narrow second arc (total turn less than π) gets measurably narrower, and one with a wide second arc (total turn greater than π) gets measurably wider. The third event is treated differently, depending on the intersection of the bridge with it second associated circle. In all cases, the transformations are easily seen to not only *shorten* the path, they also arguably leave it in a form that is *simpler* than it was to start, from which it immediately follows that the full reduction consists of only finitely many transformation steps.

Case I: both bridges are non-degenerate. We begin by considering the case where both bridges are non-degenerate. As we shall see, if one or more of the bridges is degenerate, a shortening transformation exists that will bring us back to this case.

There are two sub-cases to consider. In the first sub-case the turn from the first bridge edge to the second is less than or equal to π (see Fig. 11). First note that if we slide the middle discrete arc (vertices c through x) along the first bridge edge (taking c towards b) we maintain the discrete bounded curvature property at b and c as long as b (respectively, c) lies outside the second circle (respectively, first) circle (i.e. until the bridge bc becomes degenerate). Meanwhile, the discrete bounded curvature property is maintained at the endpoints of the second bridge edge (xy) as long as the predecessor (r) of the outer point (y) lies outside of the third circle (because of the direction of the translation the successor of the other bridge endpoint point (x) cannot enter the second circle). If this point r meets the third circle (a maximal chord event) while outside of the second circle, then point r can replace y as the outer point of the second bridge (leading to a shortening of the second discrete arc), and we can continue in Case I.

By symmetry the analogous properties hold if we slide the middle discrete arc along the second bridge edge (taking x to y). Since both of these translations serve to shorten the curve, we can assume that they have been done until either or both of the bridges have degenerated (taking us to Case II or Case III below) or we are left with unresolved maximal chord events on both bridges. In the latter case, the successor point of b (illustrated by p in Fig. 12) must lie on the first circle, the predecessor point of y (illustrated as r) must lie on the third circle, and both p and r must lie inside the second circle.

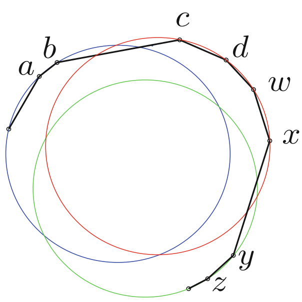

Fig. 11.

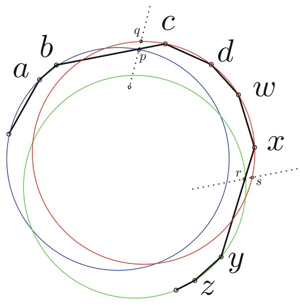

Fig. 12.

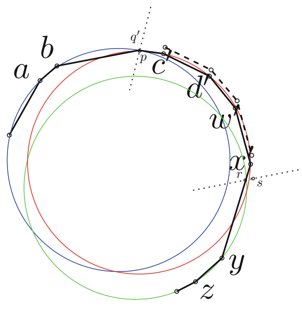

Fig. 13.

To deal with this last situation, we observe that either (i) the distance $|rx|$ must be at least the distance from p to the point q on the second circle intersected by the line through p with the slope of the second bridge edge, or (ii) the distance $|pc|$ must be at least the distance from r to the point s on the second circle intersected by the line through r with the slope of the first bridge edge. (It is easily confirmed that if neither of these hold, we get a contradiction of the fact that the slope of edge qc must exceed the slope of edge xs.) Without loss of generality, we assume that the first of these holds. Then, if we break edge bc at point p and translate the second arc along the second bridge edge, the second circle (specifically point q) must meet point p before x reaches r (note that, since $|ry| = d_\theta$, x reaches r before y enters the second circle). If we stop the translation at this point (see Fig. 13) we see that the first bridge has been replaced by two edges (bp and pc) which become part of the first and second discrete arc respectively; i.e. vertex p is a degenerate bridge, taking us to Case II. (It is worth noting here that as the translation takes q to p the discrete bounded curvature property is initially violated at c. It is restored just when q coincides with p. This is the reason why we need to ensure that the bridge remains feasible for vertices x and y until q reaches p. It is not at all clear that a transformation exists that is guaranteed to shorten the path in the situation under consideration, while preserving the discrete bounded curvature property throughout.)

The second sub-case, where the turn from the first bridge edge to the second is greater than π (see Fig. 14), is treated in a very similar fashion. As in the first sub-case, we note that if we slide the middle discrete arc along the first bridge edge (taking c towards b) we maintain the discrete bounded curvature property at b and c as long as b (respectively, c) lies outside the second (respectively, first) circle (i.e. until the first bridge becomes degenerate). Meanwhile the discrete bounded curvature property is maintained at the endpoints of the second bridge edge (xy) as long as the successor (r) of the inner point (x) lies outside of the second circle (because of the direction of the translation the predecessor of the other bridge endpoint (y) cannot enter the third circle). If this point r meets the second circle (a maximal chord event) while outside of the third circle, then point r can replace x as the inner point of the second bridge (leading to a lengthening of the second discrete arc), and we can continue in Case I.

By symmetry the analogous properties hold if we slide the middle discrete arc along the second bridge edge (taking x to y). Since both of these translations serve to shorten the curve, we can assume that they have been done until either or both of the bridges have degenerated (taking us to Case II or Case III below) or we are left with unresolved maximal chord events on both bridges. In the latter case, the predecessor point of c (illustrated by p in Fig. 15) must lie on the first circle, the successor point of x (illustrated as r) must lie on the third circle, and p (respectively, r) must lie inside the first (respectively, second) circle.

To deal with this last situation, we observe that either (i) the distance $|ry|$ must be at least the distance from p to the point q on the circle associated with the first arc intersected by the line through p with the slope of the second bridge edge, or (ii) the distance $|pb|$ must be at least the distance from r to the point s on the circle associated with the third arc intersected by the line through r with the slope of the first bridge edge. (As before, it is easily confirmed that if neither of these hold, we get a contradiction of the fact that the slope of edge bq must exceed the slope of edge ys.) Assume, without loss of generality that

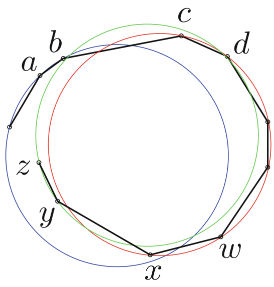

Fig. 14.

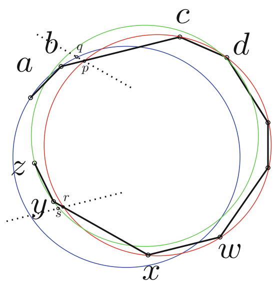

Fig. 15.

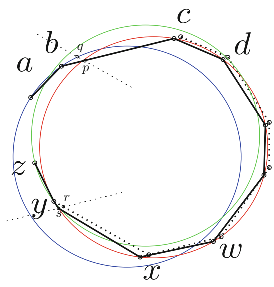

Fig. 16.

the second of these holds. Then, if we break edge xy at point r and translate the second arc, including the point r, along the first bridge edge, the point r must encounter the circle associated with the third arc before p reaches b. If we stop the translation at this point (see Fig. 16) we see that the second bridge has been replaced by two edges (xr and ry) which become part of the second and third discrete arc respectively; i.e. vertex r is a degenerate bridge, taking us to Case II. (As before, we note that the translation produces a path that violates the discrete bounded curvature property initially, but it is restored just when s coincides with r.)

Case II: one bridge is degenerate and the other is not. We assume, without loss of generality, that the first bridge is degenerate. There are two sub-cases again that depend on the span of the second arc. In the first sub-case (see Fig. 17) the total turn from the first edge (bc) after the degenerate bridge (b) to the second bridge edge (xy) is less than or equal to π. If we translate the middle discrete arc, excluding the first bridge point, (i.e. the vertices c through x) along the second bridge edge (taking x towards y) (see Fig. 18), we maintain the discrete bounded curvature property until the first of two events occurs: (i) x (respectively, y) joins the third (respectively, second) circle, or (ii) the successor (c) of the degenerate bridge (b) joins the first circle. The first event coincides with the degeneration of the second bridge, while the second event leaves us with a new degenerate first bridge (vertex c), and hence a new instance of Case II with a smaller second arc.

In the second sub-case (see Figs. 19 and 21), where the total turn from the first edge (bc) after the degenerate bridge (b) to the second bridge edge (xy) is greater than π, we again slide the middle discrete arc, this time including the first bridge point, (i.e. the vertices b through x) along the second bridge edge (taking x towards y). The discrete bounded curvature property is maintained at x and y unless x (respectively, y) joins the third (respectively, second) circle (i.e. the second bridge becomes degenerate). Meanwhile, if b moves outside of the first circle (see Fig. 20), then the discrete bounded curvature property is maintained at a and b until a joins the second circle, at which point a replaces b as a degenerate bridge, so we can continue in Case II. Alternatively, if b moves inside the first circle, we maintain feasibility of the transformed path by (continuously)

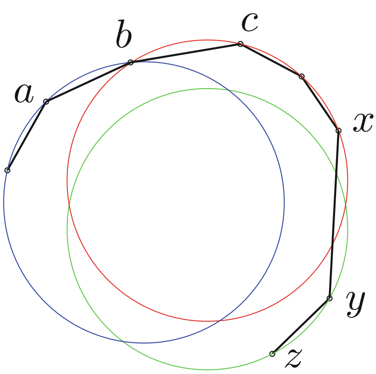

Fig. 17.

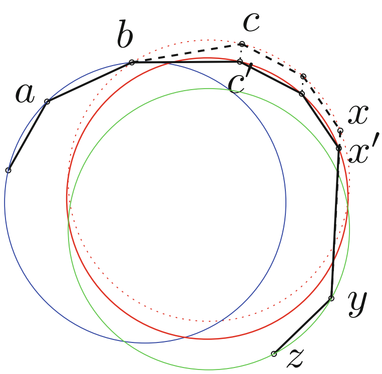

Fig. 18.

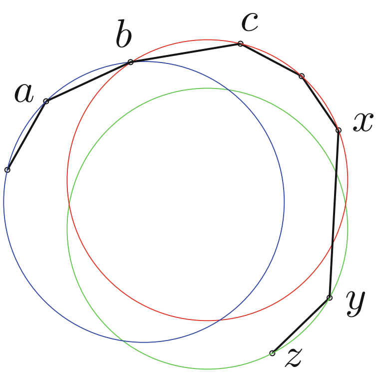

Fig. 19.

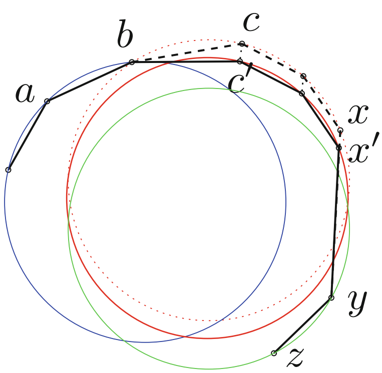

Fig. 20.

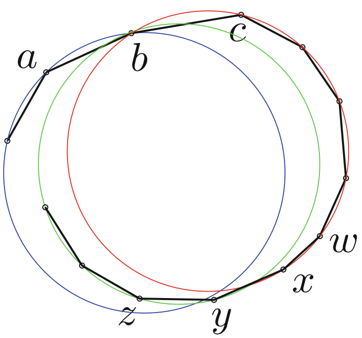

Fig. 21.

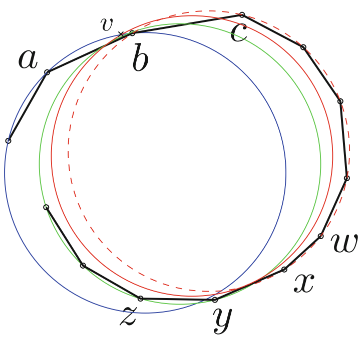

Fig. 22.

replacing the bridge point by the intersection point v of the first two circles (see Fig. 22). It is straightforward to confirm that both $|av| < |ab|$ and $|vc'| < |bc|$. Thus, the transformation can continue until v coincides with either a or c, at which point either a or c becomes a degenerate bridge, and we can continue in Case II, with one fewer edge.

Case III: both bridges are degenerate. As in both previous cases, there are two sub-cases depending on the span of the middle arc. If the middle arc spans less than a half circle (refer to Fig. 23) then we can transform the path by rotating the second arc, excluding the first bridge point, counterclockwise about the second bridge point (see Fig. 24). Of course, if this rotation continues long enough the second and third circle will coincide, at which point the second bridge disappears. Prior to this, the transformation preserves the length of all edges except for the first edge after the first bridge point (bc in Fig. 24) which shortens, since the distance from both endpoints of this edge to the second bridge point (y) is unchanged but the angle they form with the second bridge point decreases. The transformation continues until the first vertex on the second arc (c) meets the first circle, at which point it becomes a new degenerate first bridge. We then continue in Case III, with a smaller second arc.

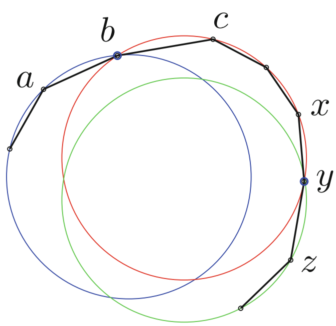

Fig. 23.

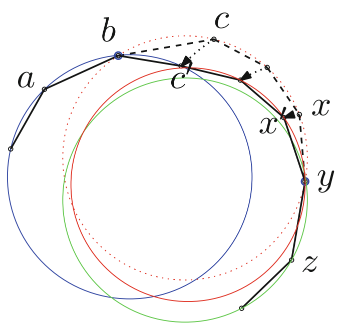

Fig. 24.

In the second sub-case, the middle arc spans at least a half circle (refer to Figs. 25 and 27). Then we can transform the path by rotating the second arc, this time including both bridge points, counterclockwise about the second bridge point. As in the previous sub-case, if this rotation continues long enough the second and third circle will coincide at which point the second bridge disappears. Prior to this, there are two cases to consider depending on the trajectory of the first bridge point (b).

If the first bridge point (b) moves outside the first circle (see Fig. 26), the discrete bounded curvature property is maintained at a and b until a joins the second circle, at which point either a or c replaces b as a degenerate bridge, so we can continue in Case III.

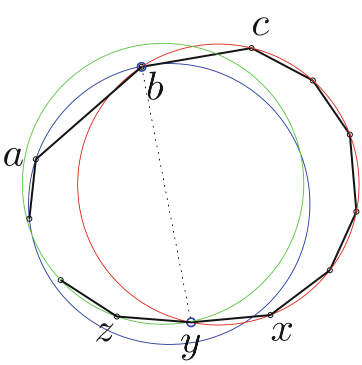

Fig. 25.

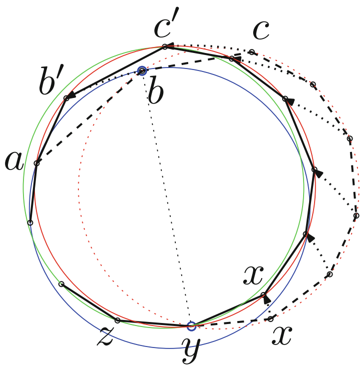

Fig. 26.

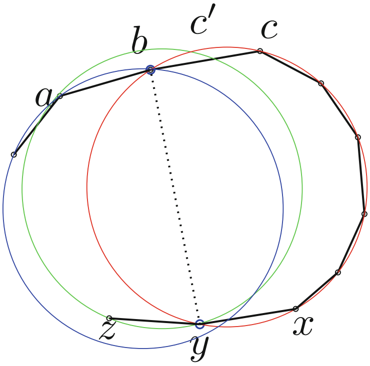

Fig. 27.

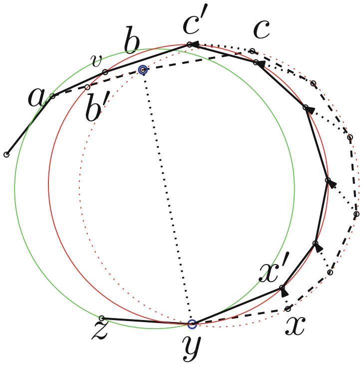

Fig. 28.

Alternatively, if the first bridge point (b) moves inside the first circle (see Fig. 28), we maintain feasibility of the transformed path by (continuously) replacing the bridge point by the intersection point v of the first two circles. It is straightforward to confirm that both $|av| < |ab|$ and $|vc'| < |bc|$. Thus, the rotation can continue until v coincides with either a or c, at which point either a or c becomes a degenerate bridge, and we can continue in Case III, with one fewer edge. □

4 Existence and Uniqueness of ς-geodesics

Careful inspection of the proof of Lemma 4 shows that it applies even when the first or third discrete circular arc is degenerate (i.e. arises from an endpoint constraint), provided the second discrete circular arc spans at most a half circle. Furthermore, if the second circular arc in some locally shortest path spans more than a half circle, then if the first (or third) arc is degenerate, it must be the case that the path starts (or ends) with an edge that (i) is an extension of

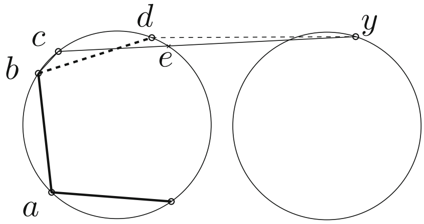

Fig. 29.

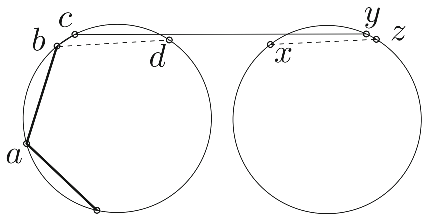

Fig. 30.

the corresponding endpoint configuration and (ii) cuts the middle circle with a maximal chord.

Taking this into consideration, Theorem 2 can be strengthened to provide a very tight characterization of the *form* of inflection-free θ-discrete-geodesics, if they exist: they are formed by two (or fewer) θ-discrete circular arcs of the same orientation, joined by a (possibly degenerate) bridge, and preceded and followed by (possibly degenerate) edges that are extensions of the endpoint configurations. Furthermore, when the extension of one endpoint configuration is non-degenerate, either (i) there is only a single non-degenerate θ-discrete circular arc, or (ii) the extension of the other endpoint configuration must be degenerate, and the adjacent θ-discrete circular arc must span more than a half circle.

We note that, as θ goes to zero, this refined characterization of inflection-free θ-discrete-geodesics describes a family of smooth geodesics, including the sole inflection-free geodesic specified by Dubins' general characterization. In this way, we can derive an analogue of Dubins' result for inflection-free geodesics. Uniqueness, in the smooth case, is a direct consequence of the uniqueness of their discrete counterparts, together with our discretization theorem (Theorem 1).

Clearly paths of the form specified by Theorem 2, joining specified endpoint configurations, always exist. To argue the existence of discrete-geodesics, it remains to argue that the infimum of the lengths of paths of this form is always realized by a path of this form. It would suffice to use a compactness argument (of the style used by Dubins), but it turns out to be both simpler and more revealing to argue this geometrically. We will do so for general (not necessarily inflection-free) discrete-geodesics in a companion paper. To give some sense of the kind of arguments involved, we consider just one special case here: an inflection-free dcc-path is *endpoint-anchored* if it is formed by two θ-discrete circular arcs respecting the two endpoint constraints, joined by a bridge of length at least $2d_\theta$. (Note that such paths correspond to the unique inflection-free paths in Dubin's characterization, in the limit as θ goes to zero.)

To this end, we say that a θ-discrete arc consisting of a sequence of edges of length exactly d_θ is *perfect*, and a θ-discrete arc consisting of a sequence of edges all but one of which have length exactly d_θ is *near-perfect*. With this, we can assert that:

Claim 3. *Endpoint-anchored geodesics exist and are composed of two perfect discrete circular arcs joined by a non-degenerate bridge.*

Proof. The key to establishing the existence of discrete-geodesics is the observation that any minimum length discrete arc spanning an angle Ψ is made up of $\lfloor \Psi/\theta \rfloor$ segments of length d_θ, and one additional edge spanning an arc of length $\Psi \bmod \theta$, i.e. it is near-perfect.

Suppose we have a dcc-path that consists of a near-perfect, but non-perfect, discrete arc followed by a non-degenerate bridge, followed by another near-perfect discrete circular arc. It remains to argue that such a path is not a discrete-geodesic, i.e. it can be shortened. There are two cases to consider. The first, shown in Fig. 29, the chord bd of length d_θ from b, the last vertex on the initial perfect arc, crosses the bridge cy. In this case, it is straightforward to show (since $|bd| + |de| \leq |bc| + |ce|$ and $|dy| < |de| + |ey|$) that replacing sub-path bcy by bdy must produce a shorter dcc-path with the same endpoints.

Alternatively, we can assume (see Fig. 30) that neither bd nor the corresponding chord xz on the second circle cross the bridge cy. In this case, if we pivot the bridge cy about its endpoint y then the dcc-path property is preserved until the bridge hits the first of b, d or x. But since this transformation leads to a shorter path in all three situations. □

5 Conclusion

We introduced a discrete model of curvature-constrained motion and studied some of its properties, in particular the structure of geodesics in this model. Our focus here has been primarily on inflection-free paths, which we have demonstrated constitute an essential component of unrestricted geodesics. We have also illustrated the utility of our characterization in relating properties of smooth geodesics as the limiting case of our discrete geodesics.

In a subsequent paper we will extend our characterization of inflection-free discrete-geodesics to the general case, including a re-derivation of the full Dubins characterization of smooth geodesics, using similar limiting arguments. We believe that discrete versions of curvature-constrained motions that include reversals (cf. [31]) can be formulated in the same way.

Acknowledgements. We thank Sergey Bereg, Stefan Foldes, Irina Kostitsyna and Joe Mitchell for discussions.

References

1. Agarwal, P.K., Biedl, T., Lazard, S., Robbins, S., Suri, S., Whitesides, S.: Curvature-constrained shortest paths in a convex polygon. In: SoCG (1998)
2. Agarwal, P.K., Raghavan, P., Tamaki, H.: Motion planning for a steering-constrained robot through moderate obstacles. In: SToC, pp. 343–352 (1995)
3. Ahn, H.-K., Cheong, O., Matousek, J., Vigneron, A.: Reachability by paths of bounded curvature in a convex polygon. Comput. Geom. **45**(1–2), 21–32 (2012)
4. Bereg, S., Kirkpatrick, D.: Curvature-bounded traversals of narrow corridors.In: SoCG, pp. 278–287 (2005)

5. Bitner, S., Cheung, Y.K., Cook IV, A.F., Daescu, O., Kurdia, A., Wenk, C.: Visiting a sequence of points with a bevel-tip needle. In: López-Ortiz, A. (ed.) LATIN 2010. LNCS, vol. 6034, pp. 492–502. Springer, Heidelberg (2010)
6. Boissonnat, J.-D., Bui, X.-N.: Accessibility region for a car that only moves forwards along optimal paths. Research report 2181, INRIA Sophia-Antipolis (1994)
7. Boissonnat, J.-D., Cérézo, A., Leblond, J.: Shortest paths of bounded curvature in the plane. Int. J. Intell. Syst. **10**, 1–16 (1994)
8. Bui, X.-N., Souères, P., Boissonnat, J.-D., Laumond, J.-P.: Shortest path synthesis for Dubins nonholonomic robot. In: IEEE International Conference Robotics Automation, pp. 2–7 (1994)
9. Chang, A., Brazil, M., Rubinstein, J., Thomas, D.: Curvature-constrained directional-cost paths in the plane. J. Global Optim. **53**(4), 663–681 (2011)
10. Chitsaz, H., LaValle, S.: Time-optimal paths for a Dubins airplane. In: 46th IEEE Conference on Decision and Control (2007)
11. Chitsaz, H., Lavalle, S.M., Balkcom, D.J., Mason, M.T.: Minimum wheel-rotation paths for differential-drive mobile robots. Int. J. Rob. Res. **28**, 66–80 (2009)
12. Chitsaz, H.R.: Geodesic problems for mobile robots. Ph.D. thesis, University of Illinois at Urbana-Champaign, Champaign, IL, USA, AAI3314745 (2008)
13. Djath, K., Siadet, A., Dufaut, M., Wolf, D.: Navigation of a mobile robot by locally optimal trajectories. Robotica **17**, 553–562 (1999)
14. Dolinskaya, I.: Optimal path finding in direction, location and time dependent environments. Ph.D. thesis, The University of Michigan (2009)
15. Dubins, L.E.: On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents. Am. J. Math. **79**, 497–516 (1957)
16. Duindam, V., Jijie, X., Alterovitz, R., Sastry, S., Goldberg, K.: Three-dimensional motion planning algorithms for steerable needles using inverse kinematics. Int. J. Rob. Res. **29**, 789–800 (2010)
17. Edison, E., Shima, T.: Integrated task assignment and path optimization for cooperating uninhabited aerial vehicles using genetic algorithms. Comput. Oper. Res. **38**, 340–356 (2011)
18. Eriksson-Bique, S.D., Kirkpatrick, D.G., Polishchuk, V.: Discrete dubins paths. CoRR, abs/1211.2365 (2012)
19. Foldes, S.: Decomposition of planar motions into reflections and rotations with distance constraints. In: CCCG'04, pp. 33–35 (2004)
20. Fortune, S., Wilfong, G.: Planning constrained motion. Ann. Math. AI **3**, 21–82 (1991)
21. Furtuna, A.A., Balkcom, D.J.: Generalizing Dubins curves: minimum-time sequences of body-fixed rotations and translations in the plane. Int. J. Rob. Res. **29**, 703–726 (2010)
22. Giordano, P.R., Vendittelli, M.: Shortest paths to obstacles for a polygonal dubins car. IEEE Trans. Rob. **25**(5), 1184–1191 (2009)
23. Jacobs, P., Canny, J.: Planning smooth paths for mobile robots. In: Li, Z., Canny, J.F. (eds.) Nonholonomic Motion Planning, pp. 271–342. Kluwer Academic Pubishers, Norwell (1992)
24. Kim, H.-S., Cheong, O.: The cost of bounded curvature. CoRR, abs/1106.6214 (2011)
25. Krozel, J., Lee, C., Mitchell, J.S.: Turn-constrained route planning for avoiding hazardous weather. Air Traffic Control Q. **14**, 159–182 (2006)
26. Latombe, J.-C.: Robot Motion Planning. Kluwer Academic Publishers, Boston (1991)

27. Li, Z., Canny, J.F. (eds.): Nonholonomic Motion Planning. Kluwer Academic Pubishers, Norwell (1992)
28. Ma, X., Castan, D.A.: Receding horizon planning for Dubins traveling salesman problems. In: Proceedings of the 45th IEEE Conference on Decision and Control, San Diego, CA, USA (2006)
29. Markov, A.A.: Some examples of the solution of a special kind of problem on greatest and least quantities. Soobshch. Kharkovsk. Mat. Obshch. **1**, 250–276 (1887). In Russian
30. Morbidi, F., Bullo, F., Prattichizzo, D.: On visibility maintenance via controlled invariance for leader-follower dubins-like vehicles. In: IEEE Conference on Decision and Control, (CDC), pp. 1821–1826 (2008)
31. Reeds, J.A., Shepp, L.A.: Optimal paths for a car that goes both forwards and backwards. Pac. J. Math. **145**(2), 367–393 (1990)
32. Reif, J., Wang, H.: Non-uniform discretization for kinodynamic motion planning and its applications. In: Laumond, J.-P., Overmars, M. (eds.) Algorithms for Robotic Motion and Manipulation, pp. 97–112. A.K. Peters, Wellesley, MA (1997); Proceedings of 1996 Workshop on the Algorithmic Foundations of Robotics, Toulouse, France (1996)
33. Robuffo Giordano, P., Vendittelli, M.: The minimum-time crashing problem for the Dubins car. In: International IFAC Symposium on Robot Control SYROCO (2006)
34. Savla, K., Frazzoli, E., Bullo, F.: On the dubins traveling salesperson problems: novel approximation algorithms. In: Sukhatme, G.S., Schaal, S., Burgard, W., Fox, D. (eds.) Robotics: Science and Systems II. MIT Press, Cambridge (2006)
35. Schmidt, E.: Über das extremum der bogenlänge einer raumkurve bei vorgeschreibenen einschränkungen ihrer krümmung. Sitzber. Preuss. Akad. Berlin, pp. 485–490 (1925)
36. Schur, A.: Über die schwarzsche extremaleigenschaft des kreises unter den kurven konstanter krümmung. Math. Ann. **83**, 143–148 (1921)
37. Shkel, A.M., Lumelsky, V.J.: Classification of the dubins set. Robot. Auton. Syst. **34**(4), 179–202 (2001)
38. Sigalotti, M., Chitour, Y.: Dubins' problem on surfaces ii: nonpositive curvature. SIAM J. Control Optim. **45**(2), 457–482 (2006)
39. Sussman, H.J.: Shortest 3-dimensional paths with a prescribed curvature bound. In: Proceedings of 34th IEEE Conference Decision Control, pp. 3306–3311 (1995)
40. Sussmann, H.J., Tang, G.: Shortest paths for the Reeds-Shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control. Research report SYCON-91-10, Rutgers University, New Brunswick, NJ (1991)
41. Szczerba, R.J., Galkowski, P., Glickstein, I.S., Ternullo, N.: Robust algorithm for real-time route planning. IEEE Trans. Aerosp. Electron. Syst. **36**, 869–878 (2000)
42. Wilfong, G.: Motion planning for an autonomous vehicle. In: Proceedings of IEEE International Conference Robotics Automation, pp. 529–533 (1988)
43. Wilfong, G.: Shortest paths for autonomous vehicles. In: Proceedings of 6th IEEE International Conference Robotics Automation, pp. 15–20 (1989)